Current Biology *16*, 814–820, April 18, 2006 ©2006 Elsevier Ltd All rights reserved DOI 10.1016/j.cub.2006.03.027

Report

Discrimination Strategies of Humans and Rhesus Monkeys for Complex Visual Displays

Kristina J. Nielsen,[1] Nikos K. Logothetis,[1] and Gregor Rainer[1,*]
[1] Max-Planck Institute for Biological Cybernetics
D-72076 Tübingen
Germany

Summary

By learning to discriminate among visual stimuli, human observers can become experts at specific visual tasks. The same is true for Rhesus monkeys, the major animal model of human visual perception. Here, we systematically compare how humans and monkeys solve a simple visual task. We trained humans and monkeys to discriminate between the members of small natural-image sets. We employed the "Bubbles" procedure [1] to determine the stimulus features used by the observers. On average, monkeys used image features drawn from a diagnostic region covering about 7% ± 2% of the images. Humans were able to use image features drawn from a much larger diagnostic region covering on average 51% ± 4% of the images. Similarly for the two species, however, about 2% of the image needed to be visible within the diagnostic region on any individual trial for correct performance. We characterize the low-level image properties of the diagnostic regions and discuss individual differences among the monkeys. Our results reveal that monkeys base their behavior on confined image patches and essentially ignore a large fraction of the visual input, whereas humans are able to gather visual information with greater flexibility from large image regions.

Results and Discussion

We investigated the performance of monkey observers trained to discriminate among natural images. Natural images contain structure at many spatial scales distributed nonhomogenously across the image and are thus good examples of complex, redundant visual forms. The monkeys were trained to discriminate between three natural images by performing a saccade task with these stimuli (see Figure 1). Every stimulus presentation was followed by the presentation of three response targets, each of which was associated with one of the stimuli. Upon presentation of a particular stimulus, a saccade to the associated target was rewarded with a drop of juice.

After the monkeys reached a performance level of at least 80% correct for a particular stimulus set, "Bubbles" was used to identify the diagnostic regions for each stimulus in the set. In Bubbles, stimuli are sampled from a parametric search space. Here, we search the image space by presenting the stimuli behind occluders, which consist of a mid-gray mask punctured by a number of randomly located windows ("bubbles") through which the occluded image was visible (see Experimental Procedures for details). Unique occluders were generated on every trial by randomly placing the bubbles. The monkeys continued to perform their discrimination task on the partially visible images. Whether they could identify the partially visible stimuli depended on whether the occluder uncovered image parts critical for task performance. For quantitative analysis, we grouped the occluders from trials in which a stimulus was correctly identified. We similarly grouped occluders from incorrect trials, and we determined diagnostic regions by comparing these two groups. At each pixel, the distributions of occluder values for correct and incorrect trials were compared with the Kolmogorov-Smirnov test. Image pixels at which occlusion systematically influenced performance should show a different distribution of occluder values in correct and incorrect trials, whereas similar distributions should arise for pixels with no influence of occlusion. The p values of the Kolmogorov-Smirnov test were Bonferoni corrected for the number of image pixels, and diagnostic regions were taken to be image regions where the corrected p values were below 0.01.

Because monkeys had not been tested with Bubbles before, we first established that the technique is suitable for the study of visual recognition in monkeys. For this purpose, we used a custom-designed set of geometrical shapes for which we a priori determined the diagnostic regions. The results of this experiment are reported in the Supplemental Data available online. We then proceeded to use Bubbles to study visual-information use in a task that required the discrimination among the members of natural-image sets. The diagnostic regions for two image sets and both monkeys are shown on the left side of Figure 2. Diagnostic regions covered on average 7% ± 2% of each image and were similar in size for the two monkeys [G00: 4.5%; B98: 9.3%; paired t test: $t(5) = -1.34$, $p = 0.24$]. Note that no diagnostic region could be determined for one of the images in monkey B98, suggesting that the monkey used no region consistently to identify the image.

Whereas the size of the diagnostic regions was similar for the two monkeys, their diagnostic regions tended to contain different amounts of spatial structure and cover different image locations. For monkey B98, diagnostic regions were located closer to the image border and contained image regions that largely lacked spatial structure. In contrast, G00 used image regions with more spatial structure, located near the center of the images. Accordingly, diagnostic regions for the two monkeys differed in terms of their distance from the image center [average distance for B98: 2.93° of visual angle; G00: 1.99°; paired t test: $t(4) = -5.67$, $p < 0.01$]. Image structure was characterized by luminance and amount of edges. Both parameters were computed at four decreasing levels of spatial scale (see Experimental Procedures). The mean

*Correspondence: gregor.rainer@tuebingen.mpg.de

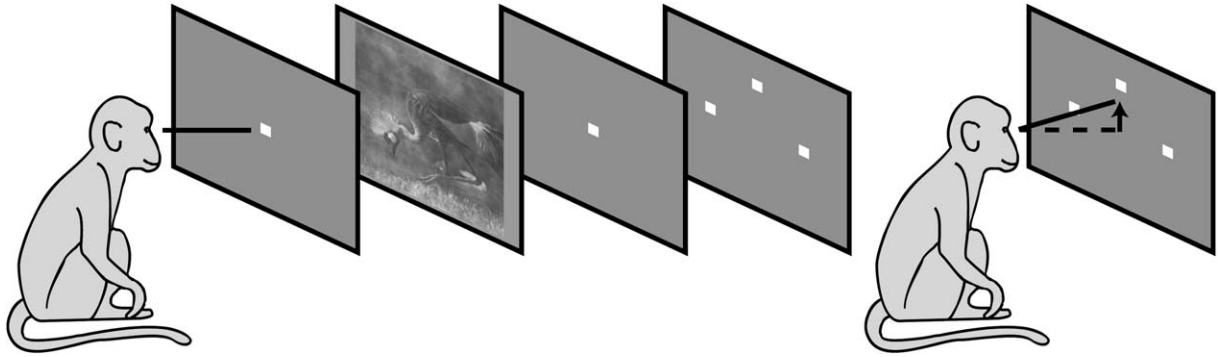

Figure 1. Task Design for the Monkey Observers

Each task began with the presentation of a central fixation spot, which the monkeys had to fixate. While the monkey continued to fixate, the fixation spot was replaced by the stimulus for 300 ms, after which time the fixation spot reappeared. Finally, three targets appeared in the periphery, each of which was associated with one of the stimuli. The monkey had to make a saccade to the correct target to receive a reward.

luminance of the diagnostic regions was similar for both monkeys, independent of spatial scale [paired t test: $t(4) \leq 0.75$, $p \geq 0.5$ for the four scales]. However, the diagnostic regions of B98 contained significantly fewer edges at the finest resolution [paired t test: scale 1: $t(4) = 3.33$, $p = 0.03$; scale 2: $t(4) = 2.63$, $p = 0.06$; scale 3: $t(4) = 2.05$, $p = 0.11$; scale 4: $t(4) = 2.63$, $p = 0.06$], confirming that the diagnostic regions of monkey B98 contained less spatial structure. Thus, diagnostic-region size, but not its location or spatial structure, was consistent across both monkeys. During task performance, we introduced catch trials on which the unoccluded images were shown to ensure that monkeys were maintaining high performance discriminating the unoccluded images. Both monkeys performed equally well on these catch trials [G00: 95% correct, B98: 98% correct, paired t test: $t(5) = -1.24$, $p = 0.27$].

Do the diagnostic regions identified by Bubbles bear any relation to performing the task outside the Bubbles paradigm? To address this question, we investigated whether monkeys could correctly identify images when presented with their diagnostic regions alone. For this purpose we constructed "diagnostic" stimuli that consisted of image regions with high diagnosticity by revealing the 10%, 30%, and 50% most diagnostic pixels (see Figure 3). Similarly, we constructed "nondiagnostic" stimuli consisting of the 10%, 30%, and 50% least diagnostic pixels. A unique stimulus set was generated for each monkey on the basis of that monkey's Bubbles results. All six stimuli thus constructed were matched to the original image in terms of mean luminance and contrast. Monkeys performed the discrimination task with the modified stimuli with no additional behavioral training.

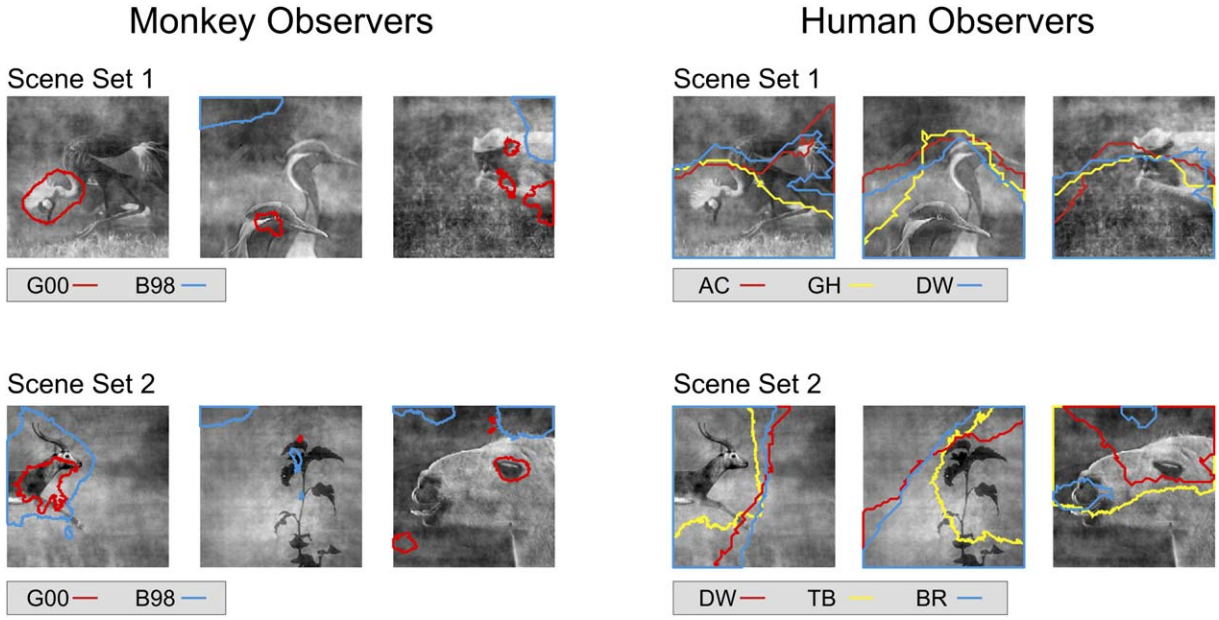

Figure 2. Diagnostic Regions in Natural Scenes for Monkeys and Humans

The left side shows the results for the two monkey observers, the right side for the human observers. Lines encircle the diagnostic regions, with each color corresponding to one observer (observer identity is given in the legend below each plot).

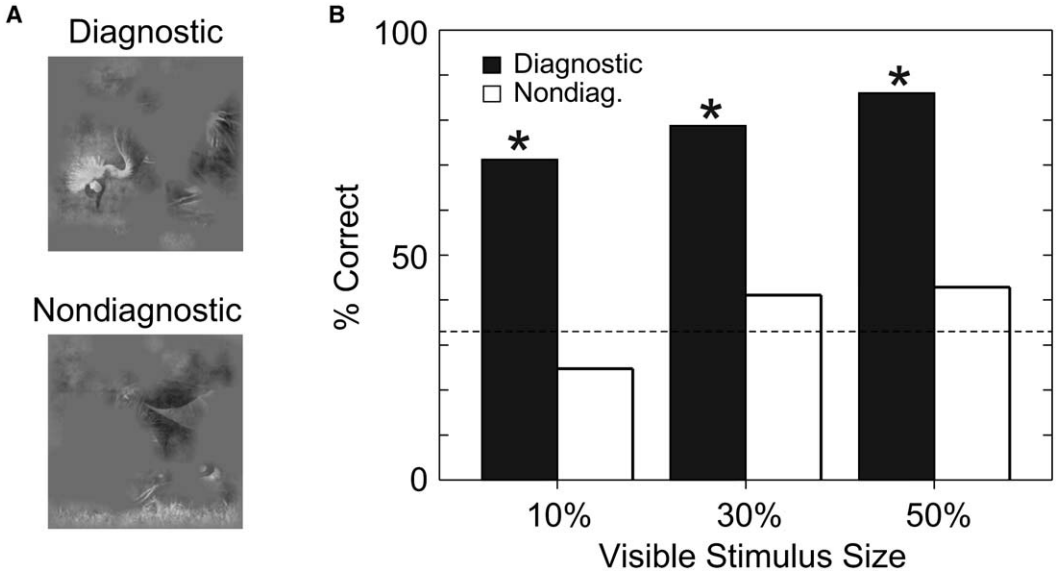

Figure 3. Verification of the "Bubbles" Results

(A) Exemplar stimuli showing a "diagnostic" and a "nondiagnostic" stimulus. These stimuli were generated based on the Bubbles results for the first scene in the first image set for monkey G00. In both stimuli, 50% of the original stimulus are exposed.

(B) Performance of the monkeys with the control stimuli. Data are averaged across both monkeys and stimulus sets. Black bars correspond to the performance for diagnostic stimuli, open bars to the performance for nondiagnostic stimuli. The dashed line shows the chance level (33.3%). Stars indicate deviations in the performance from chance level significant at $p < 0.05$, as assessed by a χ^2 test.

Performance levels were averaged across all six images in the two sets and both monkeys (see Figure 3). All performance levels were compared against the chance level of 33% correct responses with a χ^2 test. Monkeys performed significantly better than chance for all diagnostic stimuli (χ^2 tests: $\chi^2 \geq 37.6$, $p < 10^{-8}$ for the three stimulus sizes). However, monkeys performed at chance level for the three nondiagnostic stimuli ($\chi^2 \leq 2.48$, $p > 0.16$ for the three tests). This indicates that when monkeys were confronted with image regions of high diagnosticity, they treated these as the unoccluded images and were able to perform the task. In the absence of high-diagnosticity regions, monkeys were unable to perform above chance.

To compare the visual information use of monkeys with that of humans, we tested human observers with Bubbles on the identical image sets. The diagnostic regions for human observers are shown on the right side of Figure 2. With an average size of 51% ± 4% of the full image, diagnostic regions for human observers were an order of magnitude larger than the diagnostic regions determined for the monkeys. A t test showed this difference to be significant [$t(28) = -9.44$, $p < 0.001$]. These results are summarized in Figure 4A, which contrasts the diagnostic-region size for the two species.

On most of the trials, the diagnostic regions were not completely exposed, but only a small portion of them was visible. To analyze how much of the diagnostic region was visible on an average correct trial, we focused on the last trials of each testing session. Performance levels were similar for monkeys and humans for this data set [t test on the performance levels, $t(25) = 0.05$, $p = 0.96$]. For each trial, we then computed what fraction of the diagnostic region was visible through the occluder. These data, averaged across correct and incorrect trials separately, are plotted in Figure 4B, showing that monkeys needed to see more of their diagnostic regions to identify an image than human observers. On average, 42.9% of a diagnostic region needed to be visible for the monkeys to correctly identify a scene, whereas on incorrect trials only 32.9% of the diagnostic region was visible. This difference was statistically significant [$t(8) = 2.67$, $p = 0.03$]. For human observers, only an average of 4.2% of the diagnostic region was visible on correct trials, compared to 2.1% on incorrect trials. This difference was also statistically significant [$t(17) = 8.60$, $p < 0.001$], as were the differences between monkeys and human observers [correct trials: $t(25) = 12.65$, $p < 0.001$; incorrect trials: $t(25) = 11.98$, $p < 0.001$].

The results imply that monkeys needed to see more of the diagnostic image regions for a correct identification. However, the diagnostic regions of monkeys cover a smaller extent of the full image. Considering of each occluder only the bubbles that fall into the diagnostic regions, we found that for the monkeys, on average 2.0% of an image was visible on correct trials. This value was similar to the result for humans, for which 2.2% of an image was visible on an average correct trial [$t(25) = -0.41$, $p = 0.69$]. On the incorrect trials performed by monkeys, only 1.4% of the images was visible, whereas 1.1% of the images was visible on incorrect trials for human observers. Again, the two values were not significantly different [$t(25) = 0.91$, $p = 0.37$]. In conclusion, when only the diagnostic image regions are considered, monkeys and humans required the same amount of the full stimulus to be exposed for a similar performance.

Finally, we examined individual differences among human and monkey observers. We estimated the degree to which the diagnostic regions of different observers overlapped. Averaged across the two image sets, the diagnostic regions of the two monkeys overlapped in 1.2% of the full image, or 17.3% of an average

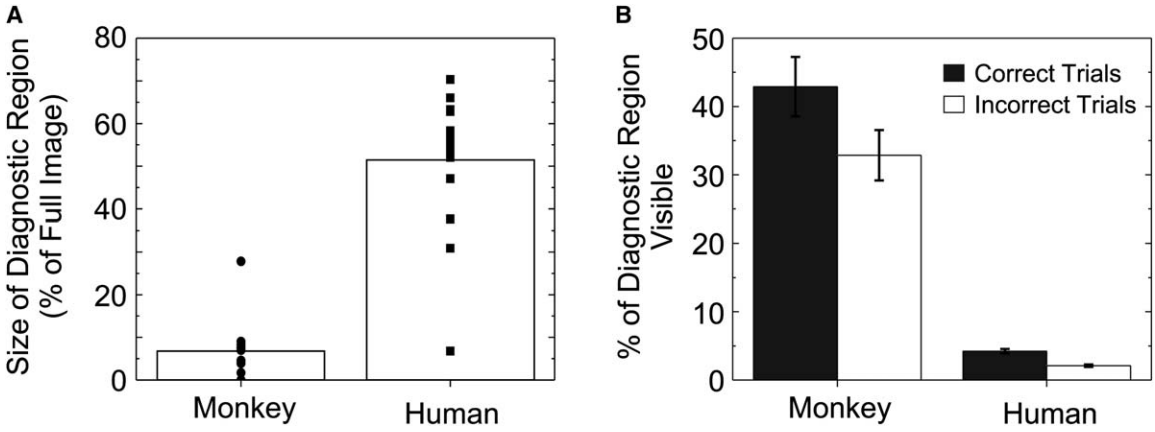

Figure 4. Comparison of the "Bubbles" Results between Monkeys and Humans

(A) Size of the diagnostic regions in percentage of original image size. The bars indicate the mean across all observers and images. Symbols indicate the values for individual diagnostic regions.

(B) Percentage of the diagnostic region visible on an average correct (black bar) or incorrect (open bar) trial. Bars show the average across all observers and images. The error bars correspond to the standard error of the mean.

diagnostic region. In contrast, the diagnostic regions of human observers overlapped on average in 35.6% of the full image, i.e., in 69.2% of the diagnostic regions. Thus, diagnostic regions of individual human observers tended to be more similar to each other than the diagnostic regions of the two monkeys. This raised the question of to what extent the behavior of human observers can be used to predict the monkey observers' behavior. We computed the overlap between the human observers' consensus diagnostic region and a monkey observer's diagnostic region. Across the two scene sets, the common diagnostic region for the human observers overlapped with 77.2% of a diagnostic region of monkey G00. For monkey B98, the data from human observers could be used to predict about 19.3% of the monkey's diagnostic region. This indicates that diagnostic regions estimated in human observers are not in general a good predictor for diagnostic regions in monkeys.

We have trained human and monkey observers to discriminate between natural images. Such images contain a wealth of features that observers might use to discriminate among them. Are observers using all available features in the images equally, or are they preferentially relying on certain features to guide their behavior? To answer this question, we used the Bubbles technique to determine diagnostic regions for each image. These diagnostic regions delineate the spatial location of the features that significantly contributed to observer performance in the discrimination task. Generally, diagnostic regions covered only a fraction of the entire visual stimulus, suggesting that observers were not drawing information equally from the entire image, but sampling preferentially from restricted image portions. Unavailability of diagnostic regions due to occlusion was associated with observers' inability to perform the discrimination task.

We observed robust differences in diagnostic-region size between monkey and human observers. Whereas monkey diagnostic regions covered only a small fraction of the images, they were approximately an order of magnitude larger for human observers, where they covered around half of the images. Diagnostic regions represent the image parts from which observers draw information, but how much visual information do observers need to see on an individual trial to enable them to perform correctly? Intriguingly, taking into account the diagnostic-region size, we found that around 2% of the entire image was visible on the average correct trial in both monkeys and humans. Thus, although the actual amount of visual information required for correct performance was similar for monkeys and humans, humans were able to gather this information from a much larger region. This suggests that human observers could extract task-relevant information from the visual environment with greater flexibility. A recent study applied the Bubbles technique to a face-classification task in human observers and pigeons and found general agreement between diagnostic regions in terms of size and location for those two species [2]. To what degree this is a result of the different visual stimuli used in that and our study remains to be determined.

Although diagnostic-region size was similar for both monkeys, they showed significant individual differences in location and image statistics of the diagnostic regions. Diagnostic regions in one monkey were located close to the image center and contained lots of spatial structure, whereas in the other monkey they were located close to the image border and contained little spatial structure. These differences in the diagnostic regions cannot be explained by different training histories, because the two monkeys received the same training. Our results therefore imply that the monkeys' individual biases led them to choose different strategies. Note that both monkeys performed the discrimination task with unoccluded images at similarly high levels. As our findings suggest, they achieved this performance by using very different strategies and focusing on different image regions. There is no way to infer this rather striking difference in visual information use on the basis of performance data on the discrimination task, and it can be detected only with a method, such as Bubbles, that directly visualizes information use.

We demonstrate that trained observers use particular spatial regions in complex scenes to perform

a discrimination task. Other forms of perceptual learning are based on enhancing sensitivity for orientation, spatial frequency, or other stimulus dimensions [3]. The spatial version of Bubbles employed here does not tell us whether observers are relying specifically on certain spatial frequency or orientation channels; however, in principle Bubbles can readily be adapted for the study of such effects. Indeed, it has already been used to identify how the performance of a task depends on different spatial-frequency channels [1], as well as on their phase [4]. The version of Bubbles we used relies on occlusion to study the contribution of image features to behavioral performance, raising the question how relevant our results are to real-world vision. Occlusion is common in everyday life, and we are generally able to recognize objects despite the fact that they are partially occluded. Thus, Bubbles can be thought of as a parametric unbiased method for simulating the occlusion that occurs in many real-world situations. In addition, several behavioral studies have provided evidence for the idea that chimpanzees [5], as well as macaque monkeys [6, 7], are able to recognize familiar stimuli despite partial occlusion.

Our findings reveal which features of a set of learned visual stimuli observers actually use during the performance of a task. In brain regions such as area TE of the inferior temporal cortex, learning has been associated with long-lasting modifications in neural activity to represent task-relevant attributes of visual stimuli [8]. After training, TE neurons become tuned to features diagnostic for a categorization or discrimination task [9, 10], or to the trained views of three dimensional objects [11]. Most previous studies on the effects of training on the perception of complex stimuli have used stimuli with predefined dimensions (for example, [12–14]). In these studies, stimuli are assigned to different classes according to experimenter-defined parameter ranges. For example, observers learned to sort Greeble stimuli into different classes based on the shapes of their components [15], or to sort face and fish stimuli into categories based on dimensions such as nose height or fin size [16]. Investigators have generally inferred that after learning, observer performance must be based on acquired expertise about which aspects of the stimuli are diagnostic. However, most complex visual stimuli are not parametrically defined according to simple generative models and contain many elements that are at different spatial scales and could be employed by observers. Direct methods such as Bubbles may thus prove particularly useful for understanding how such complex stimuli are encoded in the brain.

Experimental Procedures

Subjects

Two adult male Rhesus monkeys (*Macaca mulatta*) weighing 10 and 13 kg participated in the experiments. Before the experiments, a metal head post and a scleral search coil [17] were implanted under aseptic conditions [18]. Monkeys received their daily amount of liquid during the experimental sessions and were provided with dry food ad libitum. The monkeys were tested daily and performed between 500 and 1000 trials per day. About 20 sessions were collected per monkey for each stimulus set. All studies involving the monkeys were approved by the local authorities (Regierungspräsidium Tübingen) and were in full compliance with the guidelines of the European Community (EUVD, European Union directive 86/609/EEC) and the National Institutes of Health for the care and use of laboratory animals.

A total of eight human observers (3 males, 5 females) were tested. All subjects were naive as to the purpose of the experiments. Informed consent was obtained from all subjects. Subjects had normal or corrected-to-normal vision. Testing sessions usually lasted between 1 and 3 hr, with subjects completing between 1000 and 2000 trials in this time. Subjects returned to the lab for additional sessions, until a total of 3000 to 6000 trials had been collected.

Task and Stimuli

Two stimulus sets of three natural scenes each were used. All stimuli had a size of 256 × 256 pixels, corresponding to 6° × 6° of visual angle. The natural scenes were taken from Corel PhotoCDs and normalized to have equal Fourier amplitude spectra [19]. All stimuli were presented centrally. Both monkey and human observers worked with one stimulus set at a time. During each trial, one of the stimuli was randomly chosen and presented to the observer. Observers had to indicate which of the three stimuli they had just seen.

For monkeys, each trial began with the presentation of a yellow fixation spot in the center of the screen, combined with the sounding of a tone. After 100 ms fixation time, the spot was turned off and the stimulus was presented for 300 ms. During stimulus presentation, the monkeys had to maintain fixation at the center of the screen in a window with a radius of 3°. After another 100 ms of central fixation, three small white squares (the targets) were presented at 6° eccentricity. Each of the three members of a stimulus set was associated with one of the targets. A saccade to the correct target was rewarded by a drop of juice.

For human observers, trials began with the presentation of a yellow fixation spot for 500 ms, followed by one of the stimuli for 500 ms. Observers responded after the presentation of the stimulus by pressing designated keys on the numerical keypad of a standard computer keyboard. Each of the images in a stimulus set was associated with a specific response key. No constraints were imposed on reaction time. For the data described above, no fixation constraints were imposed because of the brief presentation time, and observers were not given feedback about the correctness of their answer to prevent learning and behavioral nonstationarity during "Bubbles." To ensure that these factors did not significantly contribute to the differences between humans and monkeys, we performed control experiments during which three human observers were required to maintain fixation within 3° of the center of the screen. The control experiments were performed with the second natural-scene set. The tested subjects did not participate in the previous experiments. During these control experiments, stimuli were presented for 300 ms, and we provided human observers with performance feedback such that after each trial a "+" or "−" sign on the screen indicated correctness of the response. These conditions thus recreated the exact parameters we used for the monkeys. We found that the characteristics of the diagnostic regions were not changed by these additional controls. In particular, the diagnostic-region size was 33.5% ± 5% and thus statistically indistinguishable from that obtained in the original experiments [46.1% ± 7%, t test, $t(16) = -1.54$, $p = 0.14$]. There was a 73% overlap between diagnostic regions determined in these control experiments and the regions obtained in the original experiment.

All observers were initially trained to associate each of the unoccluded stimuli with its assigned saccade target or button press. For monkeys, this was done by introducing a brightness cue in the saccade targets, with the correct target being brighter than the incorrect targets. This cue was gradually removed as the monkeys' performance improved. Monkeys were always trained with the entire stimulus set. Monkeys initially learned to associate visual stimuli with saccade directions prior to the stimulus sets reported here. This initial training lasted for a period of several months. Once the monkeys had learned the rules of the task, they quickly acquired new stimulus sets. Humans were provided with a printout that showed both the stimuli and their associated buttons in order to inform them about the mapping between stimuli and response buttons. They were then given a training period of 20 trials, in which they could use the printout to guide their responses. After these training trials, the printout was removed. All subjects performed

better than 90% correct on the original stimuli after these training trials.

After observers had acquired the task with unoccluded stimuli, we additionally introduced stimuli with occluders. The presentation of unoccluded images was maintained (10% of trials for human observers, 40% for monkeys) as a baseline control of the performance. The occluders were constructed as described in [1]. In brief, each occluded image appeared to be shown behind a surface punctured by round windows ("bubbles"), through which parts of the image were visible. Bubbles had the profile of a 2D Gaussian, so that they smoothly merged into the nontransparent background. Bubbles were randomly positioned, with the restriction that the center of each bubble fell onto an image pixel, and two bubbles could not have identical center coordinates. All bubbles had the same size, which was determined by setting the standard deviation of the 2D Gaussian profile to 14 pixels. For the human subjects, bubbles numbers were adapted to each subject's performance by a staircase protocol. Staircases were run independently for each image in a stimulus set and converged to a performance of 75% correct. After every fourth trial of an image, the bubbles number was updated. The number was decreased by three if the image had been identified correctly in the last four trials, and it was increased by two if fewer than three trials had been correct. For the monkeys, we employed the same staircase procedure in most sessions. As an additional control, we used a modified staircase procedure for one dataset (second set of natural images for monkey B98). During these sessions, the bubbles numbers were identical for all images, rather than adapted to each stimulus independently as in the original staircase procedure. The modified procedure thus showed all stimuli through the same bubbles number, preventing the number of bubbles itself from serving as a potential cue to the stimulus. They were initialized to a value at which the monkeys could perform the task at ceiling performance. After 15 trials, the bubbles numbers were successively decreased by a fixed amount until the monkey's performance dropped below 70% correct. At this point, the numbers of bubbles were reset to the initial value and the cycle was restarted.

Setup

Monkeys performed experiments in acoustically shielded chambers. Eye movements were monitored with the scleral-search-coil technique [20] and digitized at 200 Hz. Stimuli were presented on a 21" monitor (Intergraph 21sd115, Intergraph Systems, Huntsville) with a resolution of 1024 by 768 pixels and a refresh rate of 75 Hz. Background luminance of the monitor was set to 41 cd/m^2, and the monitor was γ corrected. The monitor was placed at a distance of 95 cm from the monkey. Stimuli were generated in an OpenGL-based stimulation program under Windows NT. Similar equipment was used for human observers, who were seated 85 cm from the monitor (background luminance of 27 cd/m^2). When eye movements of human observers were measured, the head position of the observers was restrained by using a chinrest. Eye movements were measured with iView 1.1 (SensoMotoric Instruments GmbH, Teltow, Germany).

Data Analysis

Analyses were carried out in Matlab (The Mathworks, Natick). To determine how much of the diagnostic regions was visible through the occluder on any trial, we analyzed the last 40 trials of each staircase session for each stimulus. Bubbles numbers were stable throughout these trials. Monkey B98 was tested on the second set of natural scenes with a method of constant stimuli; therefore, this data set was excluded from the analysis. Because only four to six sessions were run for human subjects, we used only the last four testing sessions for the monkeys. An image pixel was considered to be visible when the occluder value for this pixel was equal to or larger than 0.5.

Physical properties of an image were characterized as the distribution of luminance, as well as edges across the image [21]. Both parameters were computed at four spatial resolutions, which were generated through progressively low-pass filtering and subsampling the image. The four resolutions corresponded to horizontal and vertical image-reduction factors of 1, 0.5, 0.25, and 0.125. Luminance information was computed at each resolution by convolution of the image with a 2D Gaussian with a kernel size of 20 by 20 pixels and a standard deviation of 4 pixels. Edges of four different orientations (0º, 45º, 90º, 135º) were detected at each resolution. They were extracted by applying quadrature filter pairs to the images, i.e., pairs of similarly oriented sine and cosine Gabor filters. The standard deviation of the filters was set to 4 pixels, and the frequency to 1/10 pixels. Artifacts at the image borders were avoided by appending copies of an image to its borders. These copies were only present while convolutions were computed. All computed luminance and edge maps were rescaled to half the size of the original image, i.e., to 128 by 128 pixels.

Supplemental Data

Supplemental Data include Supplemental Results, one figure, and one table and are available with this article online at: http://www.current-biology.com/cgi/content/full/16/8/814/DC1/.

Acknowledgments

This work was supported by the Max Planck Society. G.R. is a Deutsche Forschungsgemeinschaft Heisenberg investigator (RA 1025/1-1). We thank C. Kayser for help with the analysis of image structure and C. Wehrhahn for comments on the manuscript.

Received: October 16, 2005
Revised: March 2, 2006
Accepted: March 2, 2006
Published: April 17, 2006

References

1. Gosselin, F., and Schyns, P.G. (2001). Bubbles: A technique to reveal the use of information in recognition tasks. Vision Res. *41*, 2261–2271.
2. Gibson, B.M., Wasserman, E.A., Gosselin, F., and Schyns, P.G. (2005). Applying bubbles to localize features that control pigeons' visual discrimination behavior. J. Exp. Psychol. Anim. B. *31*, 376–382.
3. Fahle, M. (2005). Perceptual learning: specificity versus generalization. Curr. Opin. Neurobiol. *15*, 154–160.
4. McCotter, M., Gosselin, F., Sowden, P., and Schyns, P.G. (2005). The use of visual information in natural scenes. Vis. Cogn. *12*, 938–953.
5. Fujita, K. (2001). What you see is different from what I see: Species differences in visual perception. In Primate Origins of Human Cognition and Behavior, T. Matsuzawa, ed. (Tokyo: Springer), pp. 29–54.
6. Kovács, G., Vogels, R., and Orban, G.A. (1995). Selectivity of macaque inferior temporal neurons for partially occluded shapes. J. Neurosci. *15*, 1984–1997.
7. Sugita, Y. (1999). Grouping of image fragments in primary visual cortex. Nature *401*, 269–272.
8. Logothetis, N.K., and Sheinberg, D.L. (1996). Visual object recognition. Annu. Rev. Neurosci. *19*, 577–621.
9. Sigala, N., and Logothetis, N.K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. Nature *415*, 318–320.
10. Baker, C.I., Behrmann, M., and Olson, C.R. (2002). Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. Nat. Neurosci. *5*, 1210–1216.
11. Logothetis, N.K., Pauls, J., and Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. Curr. Biol. *5*, 552–563.
12. Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. J. Exp. Psychol. Gen. *115*, 39–57.
13. Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. Psychol. Rev. *94*, 115–147.
14. Freedman, D.J., Riesenhuber, M., Poggio, T., and Miller, E.K. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. Science *291*, 312–316.
15. Gauthier, I., Williams, P., Tarr, M.J., and Tanaka, J. (1998). Training 'Greeble' experts: A framework for studying expert object recognition processes. Vision Res. *38*, 2401–2428.

16. Sigala, N., Gabbiani, F., and Logothetis, N.K. (2002). Visual categorization and object representation in monkeys and humans. J. Cogn. Neurosci. *14*, 187–198.
17. Judge, S.J., Richmond, B.J., and Chu, F.C. (1980). Implantation of magnetic search coils for measurement of eye position: An improved method. Vision Res. *20*, 535–538.
18. Lee, H., Simpson, G.V., Logothetis, N.K., and Rainer, G. (2005). Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. Neuron *45*, 147–156.
19. Rainer, G., Augath, M., Trinath, T., and Logothetis, N.K. (2001). Nonmonotonic noise tuning of BOLD fMRI signal to natural images in the visual cortex of the anesthetized monkey. Curr. Biol. *11*, 846–854.
20. Robinson, D.A. (1963). A method of measuring eye movement using a scleral search coil in a magnetic field. IEEE Trans. Biomed. Eng. *BME-10*, 137–145.
21. Itti, L., Koch, C., and Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. IEEE Trans. Pattern Anal. Mach. Intell. *20*, 1254–1259.

Behavioral/Systems/Cognitive

Dissociation Between Local Field Potentials and Spiking Activity in Macaque Inferior Temporal Cortex Reveals Diagnosticity-Based Encoding of Complex Objects

Kristina J. Nielsen, Nikos K. Logothetis, and Gregor Rainer
Max Planck Institute for Biological Cybernetics, D-72076 Tübingen, Germany

Neurons in the inferior temporal (IT) cortex respond selectively to complex objects, and maintain their selectivity despite partial occlusion. However, relatively little is known about how the occlusion of different shape parts influences responses in the IT cortex. Here, we determine experimentally which parts of complex objects monkeys are relying on in a discrimination task. We then study the effect of occlusion of parts with different behavioral relevance on neural responses in the IT cortex at the level of spiking activity and local field potentials (LFPs). For both spiking activity and LFPs, we found that the diagnostic object parts, which were important for behavioral judgments, were preferentially represented in the IT cortex. Our data show that the effects of diagnosticity grew systematically stronger along a posterior–anterior axis for LFPs, but were evenly distributed for single units, suggesting that diagnosticity is first encoded in the posterior IT cortex. Our findings highlight the power of combined analysis of field potentials and spiking activity for mapping structure to computational function in the brain.

Key words: **monkey; visual cognition; object recognition; electrophysiology; psychophysics; behavior**

Introduction

Because we live in a three-dimensional world, distant objects are often only partially visible, and in part covered by closer objects. Under most circumstances, partially occluded objects are recognizable despite the lack of information about the occluded shape regions. However, it has been demonstrated previously that occlusion of specific, behaviorally relevant shape regions renders both humans and monkeys unable to perform tasks on partially occluded shapes (Biederman, 1987; Gosselin and Schyns, 2001; Nielsen et al., 2006). Occlusion of other shape regions leads to no behavioral impairments. In this study, our goal is to systematically examine how occlusion of visual shape regions of differing behavioral relevance impacts the neural representation of these shapes in the inferior temporal (IT) cortex of the macaque monkey.

The IT cortex is thought to play a major role in object recognition processes and contains many neurons that respond to ethologically relevant objects such as faces (Perrett et al., 1982; Desimone et al., 1984), but also to arbitrary shapes after the monkey has learned to identify them (Logothetis and Sheinberg, 1996; Tanaka, 1996). IT neurons have been shown to retain their shape selectivity despite occlusion of randomly selected shape portions (Kovács et al., 1995). However, it has not been tested whether occlusion effects depend on which parts of a shape are occluded, taking the behavioral relevance of the occluded shape parts into account. Yet, several studies have provided evidence that parts of objects can be sufficient to evoke responses from IT neurons (Tanaka et al., 1991; Tsunoda et al., 2001; Baker et al., 2002). Furthermore, it has been shown that learning modifies neural responses in the IT cortex. Learning of associations between different shapes (Sakai and Miyashita, 1991; Messinger et al., 2001) and learning of task-relevance of shape features (Baker et al., 2002; Sigala and Logothetis, 2002; Sigala, 2004) are both reflected in IT cortical neural activity. As different shape regions acquire behavioral relevance because of training on a task, and are thus the outcome of a learning process, it is likely that the effects of occlusion will depend on the behavioral relevance of the occluded shape parts.

Materials and Methods

Behavioral and electrophysiological methods. Two adult male monkeys (*Macaca mulatta*) participated in the experiments. All studies were approved by the local authorities (Regierungspräsidium, Tübingen, Germany) and were in full compliance with the guidelines of the European Community (European Union directive 86/609/EEC) for the care and use of laboratory animals. Stimuli were presented on a γ-corrected 21 inch monitor, placed at a distance of 97 cm from the monkeys. Each image subtended 6 by 6° of visual angle. Stimuli were generated as described in a previous study (Nielsen et al., 2006). Of the six natural scenes used in the previous study, we chose four scenes for each monkey. The average gray-scale value of each stimulus was set to the same value to control overall luminance. Furthermore, all modifications of an image had the same overall contrast as the original image (measured as the SD of the gray-scale values). For occluded images, only unoccluded image

Received May 29, 2006; revised July 21, 2006; accepted Aug. 10, 2006.

This work was supported by the Max Planck Society. G.R. is a Deutsche Forschungsgemeinschaft Heisenberg investigator (RA 1025/1-1). We thank N. Sigala and A. Tolias for comments on this manuscript.

Correspondence should be addressed to Dr. Gregor Rainer, Max Planck Institute for Biological Cybernetics, Spemannstrasse 38, D-72076 Tübingen, Germany. E-mail: gregor.rainer@tuebingen.mpg.de.

DOI:10.1523/JNEUROSCI.2273-06.2006

 0270-6474/06/269639-07$15.00/0

parts were considered when computing the mean and SD of the gray-scale values.

During the recording sessions, the monkeys performed a fixation task. Each trial began when the monkeys acquired fixation on a central fixation point. After a variable baseline duration of at least 100 ms, a stimulus was presented for 500 ms. The monkeys were required to maintain fixation within 1° of the center of the screen for the whole trial. Fixation was monitored with a scleral search coil and sampled at 200 Hz (CNC Engineering, Enfield, CT). Successful fixation was rewarded with a drop of juice delivered 1 s after stimulus offset. The monkeys completed at least 10 repetitions for each condition during a recording session.

Single-cell activity and the local field potential (LFP) were recorded from a recording chamber consisting of a ball-and-socket joint with an 18-gauge stainless-steel tube passing through its center (Schiller and Koerner, 1971). Horsley–Clark coordinates for the chambers were anteroposterior (AP), 18.1, mediolateral (ML), 17.7 for monkey 1, and AP, 15.4, ML, 16.8 for monkey 2. Neural signals were recorded using a five-channel electrode drive (Thomas Recording, Giessen, Germany), and platinum/tungsten electrodes coated with quartz glass with an impedance between 1 and 2 MΩ (ESI2ec; Thomas Recording). The recorded signal was divided into multiunit activity (band-passed signal between 500 Hz and 10 kHz) and LFPs (band-passed signal between 1 Hz and 100 Hz). From the multiunit activity, the activity of single neurons was extracted using standard spike-sorting techniques (Offline Sorter; Plexon, Dallas, TX). To ensure an unbiased estimate of neural activity, we made no attempt to select neurons based on task selectivity. Instead, we advanced each electrode until the activity of one or more neurons was well isolated and then began collecting data. The position of each electrode in terms of AP and ML coordinates and distance from the superior temporal sulcus was noted. We sampled different AP positions in a systematic manner in both monkeys. In monkey 1, initial recording positions were anterior; over the course of the experiments, the recording positions were moved more and more posterior to minimize structural brain damage caused by guide tube movement. In monkey 2, we proceeded in the opposite way, and recording locations were moved from posterior to anterior locations.

Data analysis. Single-unit activity was analyzed in a 300 ms time window beginning 100 ms after stimulus onset to account for neural latency. Baseline activity was determined in a 100 ms time window preceding stimulus onset. Spike density functions were computed by convolution of the spike trains with a Gaussian kernel ($\sigma = 10$ ms), using a resolution of 1 ms. Spike density functions of different modifications of the same scene were normalized by the maximal value observed across all modifications. The diagnostic variance was computed as $V_{\text{Group}}/V_{\text{total}} \times 100\%$, where V_{total} is the total firing rate variance,

$$V_{\text{group}} = \tfrac{1}{2} \cdot [(\bar{f}_{\text{diag}} - \bar{f})^2 + (\bar{f}_{\text{ndiag}} - \bar{f})^2],$$

where $\bar{f}_{\text{diag}}$, $\bar{f}_{\text{ndiag}}$, and $\bar{f}$ represent the diagnostic, nondiagnostic, and overall firing rate mean, respectively (Bortz, 1993).

The LFP, which was originally sampled at a rate of 4.46 kHz, was first downsampled to 1 kHz. A bandpass filter (first order Butterworth filter, bandpass between 5 and 80 Hz) was applied to remove slow drifts. Finally, each LFP channel was z-transformed using the mean LFP amplitude and SD of the channel in the 100 ms baseline period preceding stimulus onset. Visual evoked potentials (VEPs) were computed by stimulus-locked averaging of the LFP data. Individual sites were identified as responsive to a particular stimulus if the absolute VEP amplitude was larger than 1.5 SD at three consecutive time bins during the stimulus presentation. Computation of the variance explained by diagnosticity was based on the mean LFP amplitude in an interval of 20 ms duration centered on the maximum of a positive VEP peak at ~140 ms (P140). The same formula was used as for the single units, but replacing mean firing rate with mean LFP amplitude. The P140 latency depended on the visible stimulus size; it also differed between monkeys. We therefore used a different interval for each condition and monkey. Because the visible stimulus size seemed to be the major determinant of the peak latency, the same interval was used for diagnostic and nondiagnostic conditions of the same visible stimulus size. Intervals were always 20 ms long. Their placement was determined by computing the grand average VEP over all responsive LFP cases from one monkey for one particular stimulus size (either full, 10, 30, or 50%). The latency of the peak of the P140 was determined, and used as the center for the 20 ms interval.

Results

In a previous study, two Rhesus monkeys (*Macaca mulatta*) were trained to discriminate between the members of small sets of natural scenes. We used natural scenes because they are good examples of complex visual stimuli, and contain information at many different spatial scales. After training, we systematically determined the parts of each scene that the monkeys relied on to perform the discrimination task (Nielsen et al., 2006). To investigate how occlusion of different scene parts influenced neural responses in the IT cortex, we used these results to split each scene into parts with and without behavioral relevance. By constructing appropriate masks, we generated three occluded versions of each scene which revealed only the behaviorally relevant parts of the scene (diagnostic conditions). We similarly constructed three occluded versions in which only behaviorally irrelevant scene parts were visible (nondiagnostic conditions). Across the three diagnostic versions of each scene, and similarly across the three nondiagnostic versions, we varied how much of the original scene remained visible (visible stimulus size: 10, 30, or 50%). Exemplar stimuli are shown in Figure 1. To avoid low-level differences between conditions, all stimuli were adjusted to have the same mean luminance, as well as the same overall contrast. Because each of the monkeys relied on different image regions to perform the discrimination task, each monkey had its own stimulus set. We verified that the monkeys could correctly identify a scene when presented with any of the diagnostic, but not when presented with any of the nondiagnostic conditions (Nielsen et al., 2006).

Using these behaviorally defined stimuli, we recorded the activity of well isolated single neurons in area TE in the two monkeys. During the recording sessions, the monkeys viewed a set of 28 stimuli, consisting of four scenes and the corresponding diagnostic and nondiagnostic conditions. Activity of 423 neurons was recorded from both monkeys. Neural responses to the different scenes were treated independently. For each neuron, the responses to all versions of a scene were included in the analysis if at least one version (either the original scene or one of the modifications) evoked significant excitatory responses from the neuron (t test vs baseline activity, $p < 0.05$ corrected for the 28 comparisons). Thus, each neuron could contribute between one and four "cases" (the responses to all versions of one scene) to the group analysis. By these criteria, 220 cases generated by 135 neurons were selected for additional analysis.

The activity of an exemplar neuron from this group is shown in Figure 1. Presentation of the full natural scene elicited a visual response from the neuron (t test vs baseline activity, $p < 10^{-6}$), as did all diagnostic conditions ($p = 0.0004$, $p < 10^{-4}$, and $p < 10^{-4}$ for the three conditions). In contrast, only the largest nondiagnostic condition triggered a significant response from the neuron ($p = 0.54$, $p = 1.0$, and $p = 0.003$ for these conditions). The visible stimulus size also influenced the neural firing rates, with larger responses to stimuli that revealed more of the original natural scene. However, responses to the diagnostic conditions were always larger than to the nondiagnostic conditions (ANOVA with factors diagnosticity and size: main effect size, $p < 0.001$; main effect diagnosticity, $p < 0.001$; interaction, $p = 0.29$).

Similar effects were seen across the whole population of neurons. The population spike density function (Fig. 2*a*) showed a response to the full stimulus that began with the typical latency of

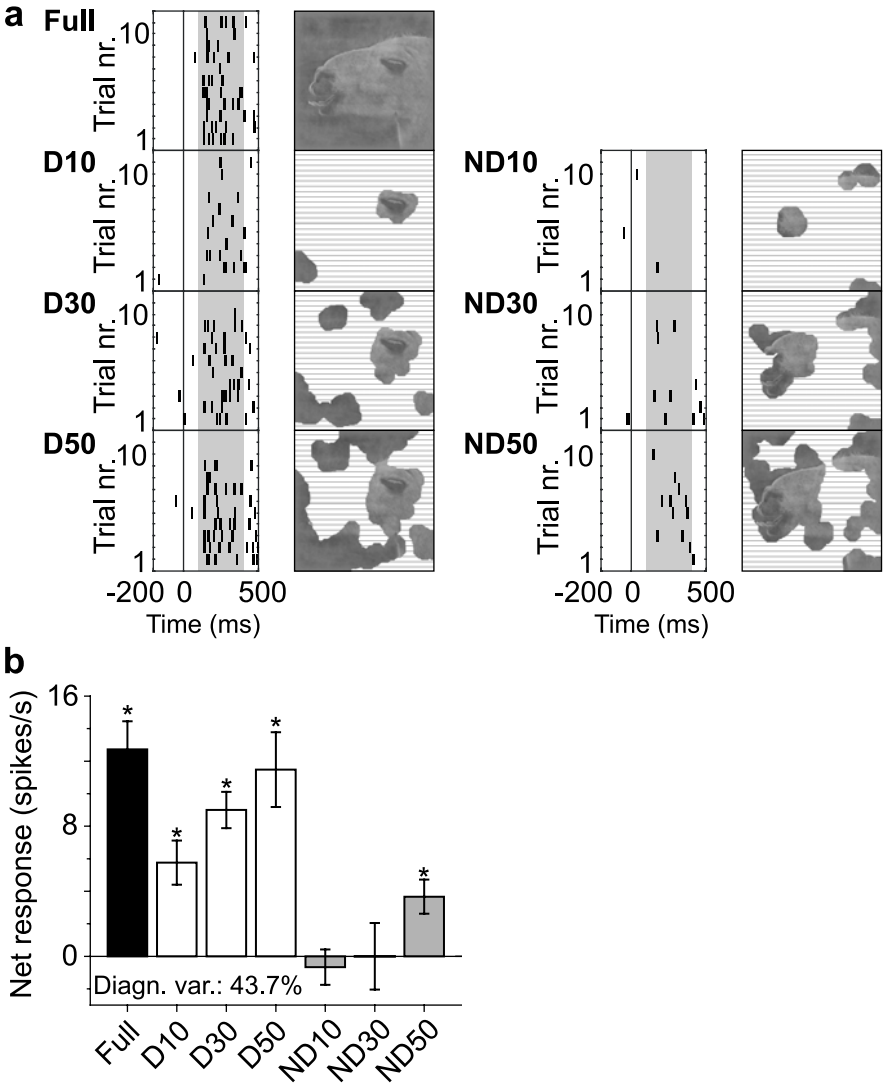

Figure 1. Responses of an exemplar single unit. ***a***, Raster plots. In these plots, each line denotes the occurrence of an action potential generated by the selected neuron (stimulus onset at 0 ms). Each plot summarizes the responses in one of the seven conditions, using the stimuli shown next to each raster plot (occluded image parts are shown as hatched regions; they were gray in the actual stimuli). The labels next to each raster plot indicate the type of condition (D, diagnostic; ND, nondiagnostic; numbers correspond to the visible stimulus size). The gray region in each raster plot corresponds to the time window used for computing the stimulus evoked firing rate. ***b***, Average net firing rate for the selected neuron. Error bars denote SEM, asterisks indicate conditions for which a *t* test between stimulus and baseline firing rate yielded $p < 0.05$. As a reference, the diagnostic variance for this case is given in the plot.

TE neurons of ~100 ms (Baylis et al., 1987; Tamura and Tanaka, 2001) and lasted throughout the stimulus presentation period. As shown in Figure 2*b*, the corresponding diagnostic and nondiagnostic conditions evoked significantly less activity than the full stimulus (paired *t* tests between the full and other conditions, $p < 0.001$ in all cases). Thus, occlusion of parts of the original scenes in general reduced the response rate of TE neurons. Similar effects of occlusion on responses of TE neurons have been reported previously (Kovács et al., 1995). In addition, we found that responses to diagnostic stimulus parts were greater than responses to nondiagnostic parts. Diagnostic conditions resulted in larger mean firing rates than nondiagnostic conditions at all visible stimulus sizes (paired *t* tests between conditions of the same visible stimulus size, $p \leq 0.02$ in all three cases). Interestingly, responses in the diagnostic conditions were on average independent of the visible stimulus size (one-way repeated measures ANOVA, $p = 0.9$). Revealing 10% of the image with high behavioral relevance triggered the same responses as revealing 50% with high behavioral relevance. The differences between diagnostic and nondiagnostic conditions were seen in many individual neurons. For 90 cases, firing rates in at least one diagnostic condition were significantly different from the matching nondiagnostic condition (*t* test, $p < 0.05$ adjusted for the three comparisons). Furthermore, we plotted the net firing rate of each case to a diagnostic stimulus condition against the net firing rate for the matching nondiagnostic condition (for the visible stimulus size of 10%, see Fig. 2*c*). At all visible stimulus sizes, more cases had higher firing rates in the diagnostic than in the nondiagnostic condition (χ^2 test, visible stimulus size of 10%, 147 vs 70 cases, $p < 0.001$; 30%, 129 vs 87, $p = 0.004$; 50%, 127 vs 90, $p = 0.01$). Diagnostic regions were determined in experiments in which the monkeys performed a discrimination task. In contrast, we used a passive fixation paradigm for the neurophysiological recordings. We verified in a separate control experiment that our findings were not influenced by the different tasks (supplemental experiment 1, Fig. 1, available at www.jneurosci.org as supplemental material). In conclusion, our results indicate that the behavioral relevance of a scene part is a major determinant for the influences of occlusion in area TE. Note that many neurons also responded to nondiagnostic parts, suggesting that learning the visual discrimination task led to a relative reweighting of neural representations of parts according to their diagnosticity, but did not completely abolish responses to nondiagnostic regions.

Given the robust influences of diagnosticity on occlusion effects on the population level, we investigated how different subregions of TE were influenced by diagnosticity. A subdivision of TE into smaller subregions has been suggested based on anatomical data (Seltzer and Pandya, 1978; Iwai and Yukie, 1987; Yukie et al., 1990), but also because a functional specialization of neurons has been observed in different parts of TE (Hasselmo et al., 1989; Janssen et al., 2000; Perrett et al., 1991, 1992; Tamura and Tanaka, 2001). To map the influences of diagnosticity across TE, we quantified the effect of stimulus diagnosticity on each case by computing how much of the total trial-by-trial variance in firing rate could be explained by diagnosticity (the "diagnostic variance"). If the firing rate of a neuron for the occluded conditions was solely determined by the diagnosticity of the visible parts, then the diagnostic variance equals 100%, whereas a diagnostic variance of 0% would indicate equal responses to diagnostic and nondiagnostic conditions.

Cases with high diagnostic variance values responded preferentially to diagnostic scene parts. To show this, we selected the

cases for which the diagnostic variance value was above the 90th percentile of all diagnostic variance values. The average spike density function for these 22 "high-diagnosticity cases" is plotted in Figure 3*a*, showing that the neurons with high diagnostic variance values indeed responded more strongly to diagnostic than to nondiagnostic conditions (repeated measures ANOVA on the mean firing rates with factors diagnosticity and size; significant main effect of diagnosticity, $p = 0.001$, with no interaction with size, $p = 0.92$). Similarly, across all cases, higher diagnostic variance values were associated with increasingly larger responses to diagnostic than to nondiagnostic conditions (supplemental Fig. 2, available at www.jneurosci.org as supplemental material).

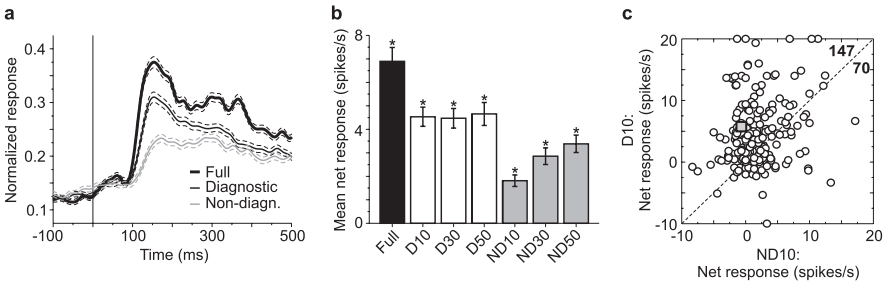

Figure 2. Population response. ***a***, Average normalized spike density function for the 220 visually responsive cases. Spike density functions are averaged across the three diagnostic and nondiagnostic conditions. Dashed lines correspond to the SEM. The stimulus onset occurs at time 0 ms. ***b***, Mean net firing rate for the complete population. Errors bars show the SEM and asterisks indicate conditions with a mean significantly different from zero (*t* test, $p < 0.05$). Labels for conditions are as in Figure 1. ***c***, Net firing rate in a diagnostic condition (visible stimulus size 10%) versus the net firing rate in the matching nondiagnostic condition. Each point represents one case. A minority of cases had firing rates higher than 20 spikes/s or lower than −10 spikes/s. These cases are plotted overlying the corresponding axis. The square represents the example neuron depicted in Figure 1. The dashed line indicates equal responses in the diagnostic and nondiagnostic condition. The numbers list the cases above and below this line.

To study the influences of diagnosticity across TE, we plotted the recording locations of the high diagnosticity cases. Figure 4*a* shows their distribution along the AP dimension of the recording region for one of the monkeys. The results for the other monkey were similar (supplemental Fig. 3, available at www.jneurosci.org as supplemental material). For each monkey, we divided the recording region into two halves with an equal extent along the AP axis. We found that high-diagnosticity cases were evenly distributed across the posterior and anterior half of the recording region (monkey 1, χ^2 test, $p = 1$; monkey 2, χ^2 test, $p = 1$). As a second step, we plotted the diagnostic variance of each case as a function of its AP recording location (Fig. 4*c*). There was no influence of the recording location on the diagnostic variance. This was the case for each monkey individually, as well as for the data of both monkeys combined (Pearson correlation coefficients not significantly different from zero; monkey 1, $p = 0.5$; monkey 2, $p = 0.9$; combined data, $p = 0.6$). We similarly tested for differences along the ML axis (data not shown). Again, no consistent influence of recording location on the diagnostic variance could be observed (Pearson correlation coefficients not significantly different from zero; monkey 1, $p = 0.1$; monkey 2, $p = 0.8$; combined data, $p = 0.8$). Furthermore, we investigated whether cases in the lower bank of the superior temporal sulcus (STS) and ventral TE were differently influenced by stimulus diagnosticity. These TE regions have been shown previously to be differently involved in the encoding of three-dimensional objects; they also differ in their connection pattern with other brain regions (Janssen et al., 2000). Here, no differences were found between cases located in the lower bank of the STS or ventral TE. Critically, the dependency of the diagnostic variance on AP position remained the same in both regions. The correlation coefficients computed between diagnostic variance and AP position were not significantly different between the lower bank of the STS and ventral TE, both for the combined data as well as both monkeys individually (monkey 1, $p = 0.7$; monkey 2, $p = 0.6$; combined data, $p = 0.3$). Thus, across the tested region, TE neurons were homogeneously influenced by stimulus diagnosticity.

Spike counts capture local processing as well as long range outputs of neurons in a brain region. However, the LFP is a mass signal that is influenced by currents originating from axons, somata, and dendrites around the electrode (Mitzdorf, 1987; Logothetis, 2002; Logothetis and Wandell, 2004), and thus reflects local neural processes as well as the inputs from other brain regions to the region under study. It has been shown previously that

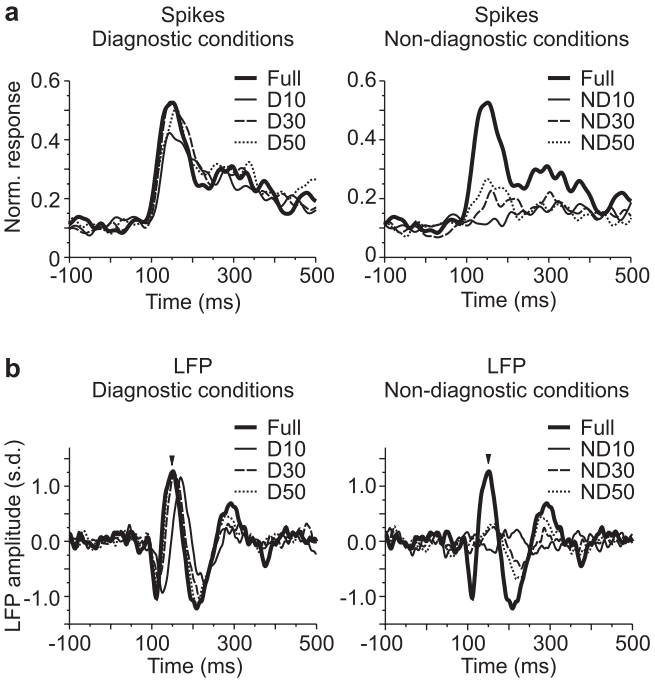

Figure 3. Population average for the high diagnosticity cases. ***a***, Average normalized spike density function, computed across the 22 single-unit cases for which the diagnostic variance exceeded the 90th percentile of all cases. ***b***, Average VEP for the 23 LFP cases for which the diagnostic variance exceeds the 95th percentile. The VEP component labeled with an arrow is the P140. For spiking activity and LFP, the diagnostic and nondiagnostic conditions are plotted separately; the response to the full condition is repeated in the two plots. Labels indicate the condition (D, diagnostic; ND, nondiagnostic; numbers correspond to the visible stimulus size).

the LFP recorded from individual sites in the IT cortex carries object-selective information (Kreiman et al., 2006). Here, we study the influence of diagnosticity on the LFP. If task-related neural signals can be observed at the level of spiking activity but not LFP, it suggests that these signals are locally computed rather than relayed to the region under study from other brain areas. The relation between spiking activity and LFP can thus provide useful information about the localization of particular computations. We subjected LFP signals recorded concurrently with the spiking activity discussed above to an analysis that was similar to the previous analysis, but took into account the continuous and time-varying nature of the LFP. We first selected LFP sites exhib-

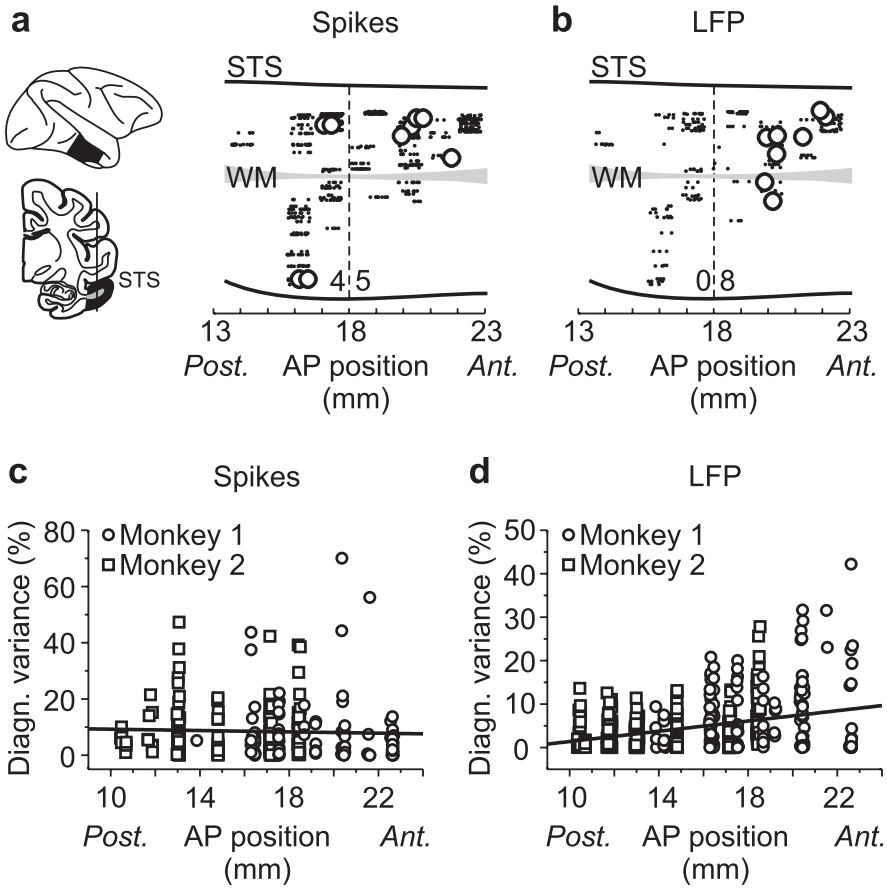

Figure 4. Influence of recording position on the properties of single units and the LFP. ***a–b***, Location of high diagnosticity cases in monkey 1, shown on a sagittal view of parts of the temporal lobe. In ***a***, the two small brain pictures on the left indicate the location of the selected brain region. This region is indicated in black in the upper image; it is generated by slicing along the line depicted in the lower image. The right side in ***a*** shows the distribution of single-unit high-diagnosticity cases. ***b***, Distribution of the high-diagnosticity LFP cases. In these plots, each dot corresponds to one case recorded from this monkey. Large dots show the location of the diagnosticity cases; small dots show the locations of the rest of the cases. To allow a better separation of different cases, the AP position of each case was randomly jittered by a small amount for display purposes only. The dashed line divides the recording region into a posterior and anterior half with equal extent along the AP axis; numbers list the diagnosticity cases in each half. Thick black lines indicate the location of the STS and the ventral end of the brain. The white matter (WM), which separates the lower bank of the STS from ventral TE in the selected slice, is shown by the gray region. The position of these landmarks is plotted as estimated during recordings. ***c–d***, Diagnostic variance as a function of recording location. In these plots, the diagnostic variance of each case is plotted as a function of its AP position. Symbols indicate the monkey in which a case was recorded; the thick line plots the regression computed between diagnostic variance and AP position. ***c***, Single unit data. ***d***, LFP data. In all plots, Post and Ant label the posterior and anterior end of the recording region, respectively.

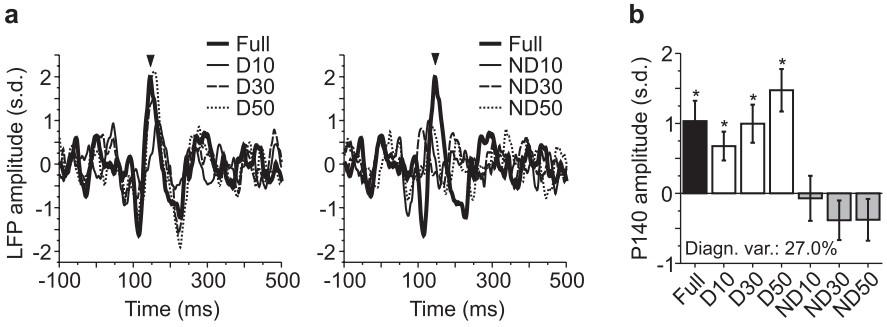

Figure 5. Responses of an exemplar LFP site. ***a***, Visual evoked potentials. The P140 is labeled with an arrow. Diagnostic and nondiagnostic conditions are plotted separately; the response to the full condition is repeated in both plots. Conditions are labeled as in Figure 3. ***b***, Amplitude of the P140, averaged across trials. Error bars denote the SEM; asterisks indicate the conditions for which a *t* test of the peak amplitude against 0 yielded $p < 0.05$. As a reference, the diagnostic variance for this site is also given.

iting a visual response to at least one stimulus, as indicated by the VEPs (Materials and Methods). Responses to different scenes were again treated as separate cases, and a total of 458 cases from 214 LFP sites were analyzed further.

The grand average VEP of these cases showed three prominent peaks in the time interval from 100 to 200 ms after stimulus onset (data not shown). A negative deflection ~100 ms after stimulus onset (N100) was followed by a positive peak at ~140 ms (P140), and finally, a second negative deflection at 200 ms (N200). To characterize activity at each LFP site, we analyzed the LFP amplitude in a 20 ms bin around the peak time of the most prominent component, which was the P140 (Materials and Methods). An analysis of the N200 amplitudes yielded similar results, whereas the N100 amplitude exhibited no systematic effects (supplemental Fig. 4, available at www.jneurosci.org as supplemental material). To illustrate the behavior of individual LFP sites, the VEP and the P140 amplitude of an exemplar site are plotted in Figure 5.

The P140 amplitude of this LFP site clearly distinguished between diagnostic and nondiagnostic conditions. All diagnostic conditions generated a P140 with an amplitude significantly larger than zero (*t* test, $p = 0.006$, $p = 0.003$, and $p < 0.001$ for the three conditions), as did the full condition ($p = 0.003$). However, none of the nondiagnostic conditions evoked a P140 with an amplitude larger than the baseline level (*t* test against 0, $p = 0.83$, $p = 0.20$, and $p = 0.23$ for the three conditions). The visible stimulus size had no influence on the P140 amplitudes for this site (ANOVA with factors diagnosticity and size: main effect size, $p = 0.61$; main effect diagnosticity, $p < 0.001$; interaction, $p = 0.15$).

Based on the P140 amplitudes, we computed the percentage of trial-by-trial variance in the LFP, which could be explained by influences of diagnosticity. As was the case for spiking activity, LFP cases with high diagnostic variance values responded preferentially to diagnostic conditions. Figure 3*b* plots the average VEP for the 23 LFP cases with high diagnostic variance values (above the 95th percentile). For these cases, the amplitude of the P140 was significantly greater in diagnostic than in nondiagnostic conditions (repeated measures ANOVA with factors diagnosticity and size; significant main effect of diagnosticity, $p < 0.001$, with no interaction with size, $p = 0.39$). In addition to the influences of diagnosticity, the visible stim-

ulus size also had an effect on the P140, as its latency depended on the visible stimulus size. In all diagnostic conditions, the latency of the P140 (computed as the latency of the positive peak between 100 and 200 ms after stimulus onset) was significantly longer than in the full condition (paired t tests, $p < 0.001$, $p = 0.001$, and $p = 0.005$ for the three conditions). We did not determine the latencies of the P140 in the nondiagnostic conditions, as the peak amplitudes in these conditions were too small to allow a reliable measurement of latencies. A high diagnostic variance was not only linked to larger P140 amplitudes in the diagnostic conditions for the high diagnosticity cases; a positive correlation between the diagnostic variance and stronger responses to diagnostic conditions was also obtained at the population level, suggesting that in general higher explained variance values were linked to stronger responses in diagnostic conditions (supplemental Fig. 1, available at www.jneurosci.org as supplemental material).

Using the diagnostic variance values, we then mapped across TE how occlusion effects observed in the LFP depended on stimulus diagnosticity. Figure 4*b* shows the locations of the LFP cases strongly influenced by diagnosticity in one monkey. It can be seen that high-diagnosticity LFP cases clearly clustered in the anterior half of the recording location (χ^2 test, $p = 0.008$). The same was the case for the other monkey (χ^2 test, $p = 0.005$). Figure 4*d* plots the diagnostic variance as a function of AP position for the entire population of LFP cases. Across all LFP cases, there was a significant correlation between the AP location of a LFP case and its diagnostic variance. This was the case for each monkey individually, as well as for the combined data from both monkeys (Pearson correlation coefficient, monkey1, $r = 0.16$, $p = 0.04$; monkey 2, $r = 0.31$, $p < 0.001$; combined data, $r = 0.33$, $p < 0.001$). This indicates that the influence of diagnosticity on LFP responses grew systematically stronger the more anterior in TE the LFP responses were recorded.

Learning effects can modify the responses of TE neurons during individual recording sessions (Messinger et al., 2001). The observed gradient in the LFP responses could thus have been generated because of a systematic sampling of the recording locations. Recording from posterior locations in the initial sessions of the experiment, and from anterior locations in the final sessions could have made a learning effect appear as a spatial gradient. Whereas the recording location was slowly moved from posterior to anterior locations across the different sessions for one monkey, we used the opposite direction for the other monkey. Because we find the same gradients for both monkeys, learning-dependent changes occurring during the recording sessions cannot account for the spatial gradient in LFP responses.

As for the single units, we tested whether this relationship was similarly present in the lower bank of the STS and ventral TE, and computed correlation coefficients separately for cases located in these two regions. The correlation coefficients were not significantly different between the two regions (monkey 1, $p = 0.7$; monkey 2, $p = 0.4$; combined data, $p = 1$). We also tested whether the position of an LFP case along the ML axis had an influence on the diagnostic variance of the case. Across the whole population of LFP cases, we observed a significant correlation between ML location and diagnostic variance ($r = 0.26$; $p < 0.001$). However, this effect was caused by a strong correlation observed for monkey 2 ($r = 0.27$; $p < 0.001$). In monkey 1, there was no significant correlation between ML position and diagnostic variance ($r = -0.04$; $p = 0.6$).

Discussion

We have experimentally determined which parts of natural scenes monkeys are using in a visual task. Occlusion of these diagnostic parts had a larger influence on neural responses in area TE of macaque cortex than occlusion of nondiagnostic parts, at the level of individual single neurons as well as at the level of local populations as measured by the LFP. This suggests that not all aspects of learned stimuli are encoded equally, but instead that those parts are preferentially represented which are diagnostic for the behaviors associated with these stimuli. Thus, we find signatures of a scene encoding in IT which is based on the diagnosticity of scene parts. In our case, monkeys learned to perform a particular saccadic eye movement associated with each member of a small set of natural images. In learning this task, each of the monkeys came to rely on particular features in each image and these behaviorally relevant or diagnostic features were preferentially encoded in the IT cortex. This lends support to the notion that the neural representation of objects in IT may be not be fixed but instead strongly influenced by the visual experience and viewing history of each observer.

Extensive studies of how single-cell responses in area TE to whole objects can be understood in terms of the responses to object parts have been performed in anesthetized monkeys (Tanaka et al., 1991; Tanaka, 2003). In these studies, experimenters used a reductive determination procedure to identify optimal features for each neuron under study. Our approach is different in that we rely on the monkeys' performance to systematically determine for each stimulus the parts that allow correct recognition behavior. Our results show that this behavioral relevance or diagnosticity is a major determining factor of how learned stimuli are encoded in memory. There has also been published work with behaving monkeys that more indirectly speaks to the issue of parts based representation. For example, TE neurons tended to show systematic tuning for dimensions in parametrically defined line-drawing stimuli that were important for the performance in a categorization task, but not for unimportant features (Sigala and Logothetis, 2002; Sigala et al., 2002). Similarly, TE neurons selectively represented feature conjunctions in visual stimuli composed of two parts, when these were relevant for correct task performance (Baker et al., 2002). In these studies, behaviorally relevant parts exerted a greater influence on neural responses consistent with our findings, but responses to whole stimuli were never directly compared with the responses to the parts alone.

We have observed effects of diagnostic parts-based encoding not only at the level of spiking activity, but also at the level of local populations as measured by the LFP. The LFP is a mass signal that originates from current flow in dendrites and somata in neural populations near the tip of the electrode. It is estimated that between 60 and 70% of excitatory connections of a given pyramidal cell remain local and only between 30 and 40% project to other cortical area (Braitenberg and Schüz, 1998; Binzegger et al., 2004). The LFP thus provides a combined measure of local processing, as well as the inputs from other brain regions. Spiking activity, however, can be considered to provide a combined measure of local processing and outputs to connected target areas. Our results reveal that diagnosticity-related spiking activity was found evenly along the posterior to anterior progression of area TE. Critically, diagnostic LFP activity was only observed in the anterior part of the recording area. To our knowledge this is the first such dissociation between spiking and LFP activity as a function of anatomical recording area. This finding has two major consequences. In posterior TE, diagnosticity was represented

only at single cell but not LFP level. Thus, absence of task-relevant signals in the LFP of a given brain region does not necessarily imply that no single neuron in that region shows such task-relevant signals. Posterior TE regions project strongly to the anterior TE, where we did find evidence of diagnostic parts-based encoding in both LFP and single cell activity. Our findings are consistent with the idea that diagnosticity is first represented by select populations of neurons in the posterior TE, and then transmitted to the anterior TE. We suggest that observation of task-relevant effects at the LFP level does not necessarily imply that the associated functions are performed in the region under study, but rather that they may be computed in brain areas that project to the region under study. Because the LFP is closely related to EEG signals recorded in human subjects, this has profound effects on the interpretation of related findings in humans. For example, EEG studies in humans have been used to link activity in the human lateral occipital complex to the perception of coherent objects from their isolated parts (Doniger et al., 2000; Murray et al., 2004). Our work suggests that this may underestimate the size of the computational network underlying this function. Thus, in general, brain areas where correlates of cognitive functions are actually computed should show effects at the single unit but not LFP level, whereas regions where this information is dynamically routed should show effects at both the single cell and LFP levels. Combined single-cell and LFP recordings thus provide more information than each kind of signal alone, and analysis of the relationship between these signals can provide a novel and powerful method for mapping structure to function in the brain.

References

Baker CI, Behrmann M, Olson CR (2002) Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. Nat Neurosci 5:1210–1216.

Baylis GC, Rolls ET, Leonard CM (1987) Functional subdivisions of the temporal lobe neocortex. J Neurosci 7:330–342.

Biederman I (1987) Recognition-by-components: a theory of human image understanding. Psychol Rev 94:115–147.

Binzegger T, Douglas RJ, Martin KA (2004) A quantitative map of the circuit of cat primary visual cortex. J Neurosci 24:8441–8453.

Bortz J (1993) Statistik für sozialwissenschaftler. Berlin: Springer.

Braitenberg V, Schüz A (1998) Cortex: statistics and geometry of neuronal connectivity. Heidelberg: Springer.

Desimone R, Albright TD, Gross CG, Bruce C (1984) Stimulus-selective properties of inferior temporal neurons in the macaque. J Neurosci 4:2051–2062.

Doniger GM, Foxe JJ, Murray MM, Higgins BA, Snodgrass JG, Schroeder CE, Javitt DC (2000) Activation timecourse of ventral visual stream object-recognition areas: high density electrical mapping of perceptual closure processes. J Cog Neurosci 12:615–621.

Gosselin F, Schyns PG (2001) Bubbles: a technique to reveal the use of information in recognition tasks. Vision Res 41:2261–2271.

Hasselmo ME, Rolls ET, Baylis GC (1989) The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. Behav Brain Res 32:203–218.

Iwai E, Yukie M (1987) Amygdalofugal and amygdalopetal connections with modality-specific visual cortical areas in macaques (*Macaca fuscata*, *M. mulatta*, and *M. fascicularis*). J Comp Neurol 261:362–387.

Janssen P, Vogels R, Orban GA (2000) Selectivity for 3D shape that reveals distinct areas within macaque inferior temporal cortex. Science 288:2054–2056.

Kovács G, Vogels R, Orban GA (1995) Selectivity of macaque inferior temporal neurons for partially occluded shapes. J Neurosci 15:1984–1997.

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.

Logothetis NK (2002) The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. Philos Trans R Soc Lond B Biol Sci 357:1003–1037.

Logothetis NK, Sheinberg DL (1996) Visual object recognition. Annu Rev Neurosci 19:577–621.

Logothetis NK, Wandell BA (2004) Interpreting the BOLD signal. Annual Rev Physiol 66:735–769.

Messinger A, Squire LR, Zola SM, Albright TD (2001) Neuronal representations of stimulus associations develop in the temporal lobe during learning. Proc Natl Acad Sci USA 98:12239–12244.

Mitzdorf U (1987) Properties of the evoked potential generators: Current source-density analysis of visually evoked potentials in the cat cortex. Int J Neurosci 33:33–59.

Murray MM, Foxe DM, Javitt DC, Foxe JJ (2004) Setting boundaries: brain dynamics of modal and amodal illusory shape completion in humans. J Neurosci 24:6898–6903.

Nielsen KJ, Logothetis NK, Rainer G (2006) Discrimination strategies of humans and rhesus monkeys for complex visual displays. Curr Biol 16:814–820.

Perrett DI, Rolls ET, Caan W (1982) Visual neurones responsive to faces in the monkey temporal cortex. Exp Brain Res 47:329–342.

Perrett DI, Oram MW, Harries MH, Bevan R, Hietanen JK, Benson PJ, Thomas S (1991) Viewer-centered and object-centered coding of heads in the macaque temporal cortex. Exp Brain Res 86:159–173.

Perrett DI, Hietanen JK, Oram MW, Benson PJ (1992) Organization and functions of cells responsive to faces in the temporal cortex. Philos Trans R Soc Lond B Biol Sci 335:23–30.

Sakai K, Miyashita Y (1991) Neural organization for the long-term memory of paired associates. Nature 354:152–155.

Schiller PH, Koerner F (1971) Discharge characteristics of single units in superior colliculus of the alert rhesus monkey. J Neurophys 34:920–936.

Seltzer B, Pandya DN (1978) Afferent cortical connections and architectonics of the superior temporal sulcus and surrounding cortex in the rhesus monkey. Brain Res 149:1–24.

Sigala N (2004) Visual categorization and the inferior temporal cortex. Behav Brain Res 149:1–7.

Sigala N, Logothetis NK (2002) Visual categorization shapes feature selectivity in the primate temporal cortex. Nature 415:318–320.

Sigala N, Gabbiani F, Logothetis NK (2002) Visual categorization and object representation in monkeys and humans. J Cog Neurosci 14:187–198.

Tamura H, Tanaka K (2001) Visual response properties of cells in the ventral and dorsal parts of the macaque inferotemporal cortex. Cereb Cortex 11:384–399.

Tanaka K (1996) Inferotemporal cortex and object vision. Annu Rev Neurosci 19:109–139.

Tanaka K (2003) Columns for complex visual object features in the inferotemporal cortex: clustering of cells with similar but slightly different stimulus selectivities. Cereb Cortex 13:90–99.

Tanaka K, Saito HA, Fukada Y, Moriya M (1991) Coding visual images of objects in the inferotemporal cortex of the macaque monkey. J Neurophys 66:170–189.

Tsunoda K, Yamane Y, Nishizaki M, Tanifuji M (2001) Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. Nat Neurosci 4:832–838.

Yukie M, Takeuchi H, Hasegawa Y, Iwai E (1990) Differential connectivity of inferotemporal area TE with the amygdala and the hippocampus in the monkey. In: Vision, memory, and the temporal lobe (Iwai E, Mishkin M, eds), pp 129–135. New York: Elsevier.

Using spikes and local field potentials to reveal computational networks in monkey cortex

Kristina J. Nielsen[1] & Gregor Rainer[2]

[1] Salk Institute for biological studies, La Jolla,USA

[2] Max Planck Institute for biological cybernetics, Tübingen, Germany

Introduction

Traditionally, neurophysiological investigations in awake non-human primates have largely focused on the study of single unit activity (SUA), recorded extracellularly in behaving animals using microelectrodes. The general aim of these studies has been to uncover the neural basis of cognition and action by elucidating the relation between brain activity and behavior. This is true for studies in sensory systems such as the visual system, where investigators are interested in how SUA covaries with aspects of visually presented stimuli, as well as for studies in the motor system where SUA covariation with movement targets and dynamics are investigated. In addition to these SUA studies, there has been increasing interest in the local field potential (LFP), a signal that reflects aggregate activity across populations of neurons near the tip of the microelectrode. In this chapter, we will describe recent progress in our understanding of brain function in awake behaving monkeys using LFP recordings. We will show that the combination of recording the activity of single neurons and local populations simultaneously offers a particularly promising way to gain insight into cortical brain mechanisms underlying cognition and memory.

Measures of neural activity at the level of neurons and networks

The activity of single neurons (SUA) is estimated by amplifying and collecting the comprehensive broadband electrical signal, which can be detected in the brain by using microelectrodes. This signal is digitized at rates of 20kHz or higher, and high-pass filtering to remove its low frequency components at a typical cutoff frequency of 300Hz. Clustering methods are then used to extract the times of action potentials generated by one or more neurons near the electrode tip, by identifying and gathering the occurrence of waveforms with a particular, predefined shape corresponding to the signature of action

potentials emitted by the neuron being studied. The part of the comprehensive signal that has often been filtered away to identify and sort action potentials has received comparatively little attention, but recent accumulating evidence suggests that it carries potentially very useful information. When the unfiltered broadband signal picket up at the microelectrodes is considered, it becomes clear that it actually consists of several components (for a review, see (Logothetis 2002; Logothetis and Wandell 2004)). To extract these components, the broadband signal is usually split into two different frequency bands by high- and lowpass filtering, respectively. Highpass filtering is used to extract the multiunit activity (MUA), while lowpass filtering isolates the local field potential (LFP). The MUA, obtained by bandpass filtering the comprehensive signal in a frequency range of 400 Hz to about 3 kHz, represents the weighted average of the spiking activity within a sphere of about 200 – 300 μm around the electrode tip. This MUA bears close resemblance to SUA, but it differs in that it represents average action potential activity generated by neurons close to the electrode tip rather than those of an individual single neuron. Another difference between the signals is that there tends to be an overrepresentation of large excitatory pyramidal cells in SUA estimates (Henze et al. 2000) due to sampling biases related to the recording technique, whereas MUA is less susceptible to this limitation. According to current estimates, over 70% of excitatory synapses in the cortex remain local, and only about 30% target distant brain regions (Binzegger et al. 2004; Braitenberg and Schüz 1998). Since both MUA and SUA capture spiking activity, they thus represent local processing within a cortical column as well as the long-range output that targets distant brain regions. The LFP, on the other hand, is extracted from the broadband signal picked up at an electrode by lowpass filtering below about 300 Hz. It measures extracellular fields generated by membrane currents originating from axons, somata and dendrites surrounding the electrode tip(Logothetis 2002; Logothetis and Wandell 2004; Mitzdorf 1985). Synchronized dendritic activity is thought to have the largest contribution to the LFP (Mitzdorf 1985), making the LFP a measure of the local processing in a brain region, as well as of the inputs that the brain region receives. The LFP is a mass signal, representing the weighted average of the synaptic signals of a neuronal population within 0.5 – 3 mm of the electrode tip (Logothetis 2002). Accordingly, LFPs depend on temporal synchronization between

dendritic events in the sampled population, as well as on the spatial alignment of the constituting neurons. Furthermore, the LFP may not only reflect the activity of the large pyramidal neurons, but also the interneuronal activity within a cortical volume. In summary, MUA and SUA capture different aspects of neural processing than the LFP. On the one hand, MUA and SUA contain signals related to the spiking output of a brain region, along with signatures of local processing carried out in that region. On the other hand, the LFP reflects dendritic input to the region near the electrode tip, as well as local processing in that region. The spiking output of a single neuron is related to the synaptic input by a nonlinear transformation, and the same is true at the level of neuronal populations. Thus systematically comparing the input to a brain region to its output represents valuable information. Using appropriate experimental designs, one can directly compare neural activity at the level of LFP to MUA/SUA, and the differences between these measures can be thought of as representing the local operations carried out by a particular brain structure.

When comparing LFP and SUA, the different nature of the two signals has to be kept in mind. SUA consists of the occurrence of action potentials from a particular neuron, and is thus a discrete signal. In most cases, SUA is summarized by computing the mean firing rate of a neuron in a selected time interval. In contrast, the LFP is a continuous signal. The LFP is commonly analyzed either in the spectral or the temporal domain. Spectral analysis consists of computing the power spectrum of the LFP, possibly in a time-resolved manner by computation of the spectrogram. The power in selected frequency bands, or the overall shape of the spectrum is then analyzed further. Analysis in the temporal domain is usually performed by computing the averages of the LFP across multiple trials of the same condition. Before averaging, trials are aligned to a selected time point, such as the onset of a stimulus or a movement. The resulting evoked potentials usually show a series of positive and negative peaks, whose amplitudes can then be analyzed further. MUA is a continuous signal like the LFP, which is often analyzed by computing its average amplitude in a selected time interval.

Previous work on LFP and SUA

Throughout the last years, a number of studies have performed detailed comparisons between the LFP and SUA or MUA in the macaque monkey. Comparisons were carried out across a number of brain regions, as well as a number of tasks. Many of these studies have largely focused on describing similarities between these signals during various cognitive tasks. Such similarities emerged for example when studying the LFP responses in area V1, the initial cortical stage of visual processing. In agreement with the behavior of single neurons, LFPs recorded in area V1 display sensitivity to the orientation of a grating pattern, as well as the grating's contrast (Frien et al. 2000; Henrie and Shapley 2005). In both cases, it seems that LFP components with frequencies in the gamma band (i.e. above about 30 Hz) are most sensitive to stimulus parameters. Similar results have been obtained for area MT in the visual cortex, which is strongly implicated in motion perception based on the response properties of single MT neurons (for a review, see (Born and Bradley 2005)). The LFP recorded in MT, specifically its frequency components above 40 Hz, also carries information about the direction and speed of a moving stimulus (Liu and Newsome 2006). Agreement between LFP and MUA could be observed at individual recording sites, as the preferences for particular stimulus speeds and motion directions were highly correlated between LFP and MUA signals recorded from the same electrode. Furthermore, variability in both LFP and MUA correlated well with the monkey's trial-to-trial performance in a speed discrimination task. LFPs were also recorded in the inferotemporal (IT) cortex, which represents the final stage of the ventral visual processing stream. Neurons in this area show a strong preference for complex objects including faces (for reviews, see (Logothetis and Sheinberg 1996; Tanaka 1996)). Again, the LFP recorded in area IT is in agreement with the behavior established for individual neurons, and LFPs recorded from sites in IT show selectivity for complex objects (Kreiman et al. 2006). In further agreement with the behavior of single neurons, the LFP recorded in IT also shows tolerance to changes in an object's position in space, as well as the object's size. Object selectivity can be observed by analyzing the LFP either in the temporal or in the spectral domain. In the temporal domain, the range of the LFP signal (the difference between the maximum and minimum LFP amplitude) differs between objects, which correlates with a modulation of LFP

broadband power in the spectral domain. Finally, a number of studies have assessed the LFP signals related to the planning and execution of hand and eye movements. The LFP recorded in the arm regions of the primary motor area M1, as well as of the supplementary motor area (SMA) has been found to convey information about arm movements. The evoked potentials in these areas contain components that are influenced by the direction of an arm movement. Evoked potentials are furthermore sensitive to which hand is used for a task, and whether the monkey moves only one or both hands (Donchin et al. 2001; Mehring et al. 2003). The same response properties are displayed by single neurons recorded in M1 and SMA (Donchin et al. 1998). Further similarities between LFP and SUA have been documented for the parietal reach region (PRR) and the lateral intraparietal area (LIP) in the posterior parietal cortex. Analyses of single neuron responses have established that these areas contain maps for the direction of either arm or eye movements that the monkey is intending to perform (for a review, see (Andersen and Buneo 2002)). The LFP in areas PRR and LIP also encodes the direction of planned arm and eye movements (Pesaran et al. 2002; Scherberger et al. 2005). In area LIP, tuning widths for movement directions are similar for the LFP and SUA. Consistent with the results obtained for the visual cortex, planning either an eye or an arm movement is most strongly reflected in LFP components with frequencies in the gamma range.

In conclusion, many of the studies performed so far have revealed that the LFP in general shows response properties similar to that of the neurons recorded in the same brain region. However, experiments have also documented interesting differences between LFP and SUA or MUA behavior. A number of these differences could be due to the fact that the LFP pools signals over a larger neuronal population than the other two signals. This means that the neurons contributing to the LFP signals have more diverse response properties than the ones contributing to either SUA or MUA. For example, it has been demonstrated that the LFP recorded on an electrode is a poor predictor of the behavior of single neurons recorded from the same electrode. Instead, the LFP correlates better with the average signal of the neuronal population within a 3 mm radius around the electrode tip (Kreiman et al. 2006). An example of such a discrepancy between LFP and SUA is shown in Figure 1, which illustrates the lack of agreement between LFP and SUA recorded from the same electrode using data collected during an experiment performed

by our group. LFP and SUA were recorded in IT while monkeys were presented with four different natural scenes, which were matched in overall contrast and luminance (see (Nielsen et al. 2006b) for details on the stimuli). LFP responses were characterized by analyzing the amplitude of a positive peak occurring at about 140 ms after stimulus onset in the visual evoked potentials (VEP). For individual neurons, the mean firing rate in a 300 ms interval during stimulus presentation was used to compute responses to the natural scenes. To test for similarities between the LFP and SUA recorded at the same electrode, we determined the agreement in the stimulus preferences of the two signals. For this purpose, the natural scene evoking the largest response from a single neuron was determined. The scenes were also rank ordered according to their peak amplitudes determined from the VEP. The rank of the best single neuron stimulus was then determined for the LFP signal recorded at the same electrode. Figure 1 plots the distribution of ranks across the population of 54 IT neurons, which responded excitatory to at least one of the full scenes, and their corresponding LFP sites. LFP-SUA pairs were only considered if the LFP was also responsive to at least one stimulus. If stimulus preferences were similar for SUA and LFP, then the best SUA stimulus should rank first for the LFP signal. However, we observed an almost equal distribution of all ranks. There was no indication of a similar scene preference for SUA and LFP. Most importantly, the number of rank 1 cases (i.e. the number of LFP-SUA pairs with the same scene preference) was not significantly lager than chance (χ^2 test, p=.9). The stimulus preference of the LFP can thus in general not be inferred from the stimulus preference of a locally recorded single neuron. It is possible that this is a consequence of the fact that the LFP pools across substantial brain regions, and thus individual stimulus preferences of local neurons are lost in this averaging process.

Pooling signal from neurons with various orientation preferences may similarly explain why the LFP recorded in V1 is less orientation selective than the MUA (Frien et al. 2000). It may also be the reason for the fact that less LFP than MUA sites are tuned for stimulus speed (17.3% vs. 2.2%) and motion direction (22.1% vs. 3.4%) in MT (Liu and Newsome 2006), and that in IT less LFP sites than MUA sites are object selective (44% vs 72%) (Kreiman et al. 2006). It becomes more difficult to explain how averaging over a diverse population of neurons can be linked to the finding that about 20% of the

LFP sites in MT are not visually responsive, as almost all MUA sites respond visually (Liu and Newsome 2006). It is likely that the local three dimensional structure of cortex is responsible for this, since dendritic dipoles can cancel each other if oriented opposite one another. Yet, there are a number of additional results that cannot be easily explained by assuming that the LFP is simply an average over a larger neuronal population than MUA or SUA. These findings hint at the different sources generating LFP and SUA or MUA, and highlight the importance of performing combined analyses of these signals. First, both SUA and MUA recorded in area V1 show strong adaptation effects, with responses to a continuously presented stimulus ceasing after about 3 s. The LFP response, on the other hand, remains elevated above baseline level throughout the presentation duration (Logothetis et al. 2001). Another discrepancy between the SUA and LFP recorded in V1 pertains to their dependency on stimulus contrast. V1 neurons initially increase their responses to a grating whose contrast is increasing. After a threshold contrast has been reached, the responses of most V1 cells saturate. The LFP shows a similar dependency on stimulus contrast, with increases in responses with increasing contrast. However, the LFP keeps increasing at contrast levels at which the single cell responses have already saturated (Henrie and Shapley 2005). Furthermore, the comparison of LFP and SUA in M1 and SMA shows that correlations between LFP and SUA recorded from the same electrode may be absent in one brain area, but present in another (Donchin et al. 2001).

Differences between LFP and SUA and MUA have also turned up when these signals were used to predict monkey behavior. In the motor cortex, the LFP can be used to successfully decode a movement direction about 50 ms after this is possible based on SUA and MUA. Furthermore, combining LFP and SUA or LFP and MUA results in higher decoding accuracy than possible based on any signal alone, suggesting independent information in these signals (Mehring et al. 2003). In LIP, both SUA and LFP components with frequencies above 30 Hz can be used to predict the direction of an eye movement. However, only the LFP can be used to decode the transition from planning an eye movement to executing it. Interestingly, a different frequency band in the LFP carries this information, as decoding performance for the behavioral state is best when based on LFP components with frequencies below 20 Hz (Pesaran et al. 2002).

Differences between LFP and SUA were also observed in the parietal reach region (PRR). Here, SUA could predict better than the LFP which direction a planned eye or arm movement was going to have. In contrast, the LFP was better at distinguishing between eye and arm movements (Scherberger et al. 2005).

Most recently, a study has described interesting differences between LFP and MUA in V1 (Wilke et al. 2006), confirming earlier findings (Gail et al. 2004). Monkeys were trained to indicate the presence of a large circle by pulling a lever. By adding small dots in the periphery around the large circle, the circle could be made to disappear perceptually, despite its continuous physical presence on the screen. At the level of the MUA, effects of the perceptual stimulus disappearance were first observed in V4. MUA in V1 and V2 did not reflect the perceptual disappearance. A largely similar result was obtained when the gamma frequency components of the LFP were analyzed. However, the LFP power in the alpha-band range (9 – 14 Hz) was modulated by perception in V1, V2, and V4. Interestingly, the modulation in LFP power was observed later than the first influence of perceptual disappearance on the MUA in V4.

All these findings highlight the fact that LFP and SUA or MUA indeed reflect different brain processes. As mentioned at the beginning of the chapter, the LFP reflects input and local processing to a brain region, whereas SUA and MUA represent the output of that region. The similarities between LFP and SUA or MUA therefore suggest a high degree of similarity between the inputs to a brain region and its outputs. The listed discrepancies between LFP and SUA or MUA in contrast are instances where input and output are not this closely related. These are cases that may allow us to identify the unique contributions of individual brain regions. The data from Scherberger et al. for example indicate that while information about the movement type is already present in the input to PRR, the direction of the movement becomes more precisely defined in this area. On the other hand, the results reported by Wilke et al. represent a case in which information is initially absent in the output from a brain region, but is later relayed to the brain region from a different source – likely activated after additional processing. In conclusion, a combined analysis of LFP and SUA or MUA has the power to reveal how different brain regions interact with each other to process information.

Combining LFP to SUA to reveal computational networks across the brain

Recent results from our group provide a compelling example of how a comparison between LFP and SUA can be used to identify how different brain regions participate in the extraction of information from a visual scene (Nielsen et al. 2006b). In this study, two monkeys were trained to discriminate between natural scenes. We then determined the regions of each natural scene on which the monkeys relied to perform the discrimination task. These diagnostic regions were determined by using a behavioral paradigm called 'Bubbles' (Gosselin and Schyns 2001). During each trial of the Bubbles paradigm, the scenes appeared behind randomly constructed occluders, which consisted of a non-transparent surface punctured by randomly placed, round windows. The monkeys continued to perform the discrimination task on the occluded scenes. Depending on which scene regions were occluded on a trial, the monkeys could or could not identify a scene correctly. Occlusion of diagnostic scene regions rendered the monkeys unable to perform the task correctly, while occlusion of nondiagnostic scene regions did not. By collecting performance data for a large number of different occluders, and determining the scene regions that were systematically occluded during incorrect responses, we could determine how relevant each scene region was for the monkeys' behavior (Nielsen et al. 2006a).

Based on the behavioral results, we then constructed unique stimulus sets for each monkey. Stimulus sets consisted of four natural scenes, presented either as their original version, or as one of six modifications. Three modifications showed diagnostic scene regions (diagnostic conditions), and three modifications showed nondiagnostic scene regions (nondiagnostic conditions). All scene regions except the selected ones were covered by an occluder. The three diagnostic conditions varied in how much of the original scene remained visible (10, 30 or 50%); the same was the case for the non-diagnostic conditions. All stimuli were matched in luminance and overall contrast. We used this stimulus set to probe the influence of diagnosticity on the responses of single neurons and the LFP in the IT cortex. Monkeys passively viewed the stimuli during the recording sessions. Figure 2 shows the responses of a selected neuron, as well as a selected LFP site to the different versions of one natural scene. In this figure, the LFP responses are shown in form of the trial averaged visual evoked potential (VEP). As can

be seen, the VEP contains a prominent positive peak at about 140 ms after stimulus onset (the P140), whose amplitude differs between diagnostic and nondiagnostic conditions. We quantified single neuron responses to different stimuli by computing the mean firing rate in a 300 ms interval beginning 100 ms after stimulus onset. For the LFP, stimulus responses were determined as the mean LFP amplitude in a 20 ms bin centered on the P140 maximum.

We could establish that stimulus diagnosticity is represented at the level of SUA in IT as the average response to the diagnostic conditions was significantly larger than the average response to the nondiagnostic condition in the tested single neuron population (paired t tests between diagnostic and nondiagnostic conditions showing the same amount of the original scene, $p \leq .02$ in all three cases). We then mapped the influences of diagnosticity across IT. For each single neuron, we quantified the influence of diagnosticity on its responses by computing the amount of variance in the firing rate that could be attributed to differences in responses to diagnostic and nondiagnostic conditions (diagnostic variance). We performed the same analysis on the LFP responses. Using this measure, we could locate the single neurons that were most strongly influenced by diagnosticity. Figure 3A plots their location within the recording region for one of the monkeys. We similarly located the LFP sites most strongly influenced by diagnosticity. The location of these sites for the same monkey is also plotted in Figure 3A. As can be seen, diagnosticity strongly influences single neurons throughout the whole recording region. In contrast, the influence of diagnosticity on the LFP increases from posterior to anterior recording locations, as all LFP sites with strong diagnosticity influences cluster in the anterior half of the recording region. We confirmed these conclusions by plotting the diagnostic variance of all single neurons as a function of their anterior-posterior location (see Figure 3B). The diagnostic variance for the LFP sites was plotted in the same way (Figure 3C). While the diagnostic variance did not depend on the recording location for the single neurons, it increased from posterior to anterior locations for the LFP. Both findings were obtained by analyzing the data from each monkey individually, as well as combining their results (SUA: Pearson correlation coefficients r not significantly different from 0 for each monkey individually and the combined data, $p \geq .1$ for the three tests; LFP: $r \geq .16$, $p \leq .04$ for the three tests).

In summary, these data suggest that diagnosticity is encoded in the output of single neurons throughout IT. However, the encoding of diagnosticity at the input level – as indicated by the LFP – increases from posterior to anterior IT. Our findings highlight a novel way to combine LFP and SUA recordings to reveal computational networks underlying a particular cognitive function, in the present case the neural encoding of diagnostic elements in complex visual displays. Diagnosticity is first encoded by neurons in posterior IT, and then transmitted to more anterior regions. The extraction of diagnosticity is accomplished by neural networks in posterior IT, since its signatures can be seen in the spiking output (SUA) from that region, but not in their input (LFP). Thus even though anterior IT contains signals of diagnosticity at both LFP and SUA levels, extraction of diagnosticity is largely not accomplished in that region itself but is already present in its input signals. The assignment of computational functions to connected brain networks is possible only by having LFP and SUA signals available during a suitable experimental paradigm, which is designed in such a way that critical parametric task condition variations are detectable at the level of LFPs. The insights given by these joint LFP/SUA analyses go a step beyond traditional brain-behavior correlations based on each of the signals considered alone. Our approach can generalize to other behavioral tasks, and promises to allow a delineation of functional networks with far greater accuracy than has previously been possible.

Reference List

Andersen RA, Buneo CA (2002) Intentional maps in posterior parietal cortex. Annu Rev Neurosci 25: 189-220

Binzegger T, Douglas RJ, Martin KAC (2004) A quantitative map of the circuit of cat primary visual cortex. J Neurosci 24(39): 8441-8453

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

Braitenberg V, Schüz A (1998) Cortex: Statistics and geometry of neuronal connectivity.

Donchin O, Gribova A, Steinberg O, Bergman H, Cardoso de Oliveira S, Vaadia E (2001) Local field potentials related to bimanual movements in the primary and supplementary motor cortices. Exp Brain Res 140: 46-55

Donchin O, Gribova A, Steinberg O, Bergman H, Vaadia E (1998) Primary motor cortex is involved in bimanual coordination. Nature 395: 274-278

Frien A, Eckhorn R, Bauer R, Woelbern T, Gabriel A (2000) Fast oscillations display sharper orientation tuning than slower components of the same recordings in striate cortex of the awake monkey. European Journal of Neuroscience 12: 1453-1465

Gail A, Brinksmeyer HJ, Eckhorn R (2004) Perception-related modulations of local field potential power and coherence in primary visual cortex of awake monkey during binocular rivalry. Cereb Cortex 14: 300-313

Gosselin F, Schyns PG (2001) Bubbles: A technique to reveal the use of information in recognition tasks. Vision Res 41: 2261-2271

Henrie JA, Shapley R (2005) LFP power spectra in V1 cortex: the graded effect of stimulus contrast. J Neurophys 94: 479-490

Henze DA, Borhegyi Z, Csicsvari J, Mamiya A, Harris KD, Buzsaki G (2000) Intracellular features predicted by extracellular recordings in the hippocampus in vivo. J Neurophys 84: 390-400

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

Liu J, Newsome WT (2006) Local field potential in cortical area MT: Stimulus tuning and behavioral correlations. J Neurosci 26(30): 7779-7790

Logothetis NK (2002) The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. Philos Trans R Soc Lond B Biol Sci 357: 1003-1037

Logothetis NK, Sheinberg DL (1996) Visual object recognition. Annu Rev Neurosci 19: 577-621

Logothetis NK, Wandell BA (2004) Interpreting the BOLD signal. Annual Review of Physiology 66: 735-769

Logothetis NK, Pauls J, Augath M, Trinath T, Oelterman A (2001) Neurophysiological investigation of the basis of the fMRI signal. Nature 412: 150-157

Mehring C, Rickert J, Vaadia E, Cardoso de Oliveira S, Aertsen AMHJ, Rotter S (2003) Inference of hand movements from local field potentials in monkey motor cortex. Nat Neurosci 6(12): 1253-1254

Mitzdorf U (1985) Current source-density method and application in cat cerebral cortex: Investigation of evoked potentials and EEG phenomena. Physiological Reviews 65(1): 37-100

Nielsen KJ, Logothetis NK, Rainer G (2006a) Discrimination strategies of humans and rhesus monkeys for complex visual displays. Curr Biol 16: 814-820

Nielsen KJ, Logothetis NK, Rainer G (2006b) Dissociation between local field potentials and spiking activity in macaque inferior temporal cortex reveals diagnosticity-based encoding of complex objects. J Neurosci 26(38): 9639-9645

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(8): 805-811

Scherberger H, Jarvis MR, Andersen RA (2005) Cortical local field potential encodes movement intentions in the posterior parietal cortex. Neuron 46: 347-354

Tanaka K (1996) Inferotemporal cortex and object vision. Annu Rev Neurosci 19: 109-139

Wilke M, Logothetis NK, Leopold DA (2006) Local field potential reflects perceptual suppression in monkey visual cortex. Proc Natl Acad Sci USA 103(46): 17507-17512

Figure Legends

Figure 1: Agreement in stimulus selectivity between single IT neurons and the LFP signals recorded at the same electrode. Only visually responsive neurons and LFP sites were considered for this analysis. We first identified the best stimulus (out of a set of four natural scenes) for each single neuron. It was then determined how strong the LFP responses to this stimulus were in comparison to the other stimuli. For this purpose, the natural scenes were rank ordered according to the LFP responses that they evoked, and the rank of the best single neuron stimulus was determined for the LFP. Rank 1 cases are cases in which the best single neuron stimulus evokes the largest responses from the LFP. Rank 4 cases imply that the best single neuron stimulus was the worst stimulus for the LFP. The histogram shows the percentage of LFP sites for each rank, the dashed line indicates the chance level of 25%.

Figure 2: Responses of an exemplar single neuron and an exemplar LFP site. **A**, Exemplar stimuli. The three images show one of the natural scenes in its original version, the diagnostic version of this image as constructed for one of the monkeys, and the matching nondiagnostic scene version. In the latter two images, occluded image parts are indicated by hatched regions. The diagnostic and nondiagnostic stimulus show 30% of the original scene. **B,** Responses of an exemplar neuron. The two plots summarize the responses to the seven versions of one natural scene by plotting spike density functions (spike trains were smoothed with a Gaussian kernel with a standard deviation of 30 ms). The stimulus appeared at time 0 ms and stayed on the screen for 500 ms. Spike density functions were normalized so that the maximum across all conditions equaled 1. The plot on the left shows the responses to the original (full) scene and the three diagnostic versions. The right side plots the responses to the nondiagnostic stimulus versions, the response to the full stimulus is repeated as a reference. In the legend, 'D' indicates diagnostic, 'ND' nondiagnostic conditions; the numbers correspond to the amount of the original image visible in a condition. **C,** Responses of an exemplar LFP site. These plots show the VEP for each condition. VEPs are plotted in units of standard deviation. These units are computed by subtracting the mean LFP amplitude in a 100 ms window preceding stimulus onset from the LFP of each trial, and dividing this signal by the

standard deviation calculated from the same baseline period. The arrow points at the P140. The layout of these two plots as is in **B**.

Figure 3: Influence of recording position on the properties of single neurons and the LFP. **A**, Location of high diagnosticity cases in one monkey, shown on a sagittal view of parts of the temporal lobe. The two small brain pictures on the left indicate the location of the selected brain region. This region is indicated in black in the upper image; it is generated by slicing along the line depicted in the lower image. The right side shows the location of the single neurons (circles) and LFP sites (triangles) strongly influenced by diagnosticity. To allow a better separation of different cases, the AP position of each case was randomly jittered by a small amount for display purposes only. Thick black lines indicate the location of the superior temporal sulcus (STS) and the ventral end of the brain. The position of these landmarks is plotted as estimated during recordings. **B-C**, Diagnostic variance as a function of recording location. In these plots, the diagnostic variance of each case is plotted as a function of its AP position. Symbols indicate the monkey in which a case was recorded; the dashed line plots the regression computed between diagnostic variance and AP position. **B**, Single unit data. **C**, LFP data. In all three plots, "Post" and "Ant" label the posterior and anterior end of the recording region, respectively.

Figure 1:

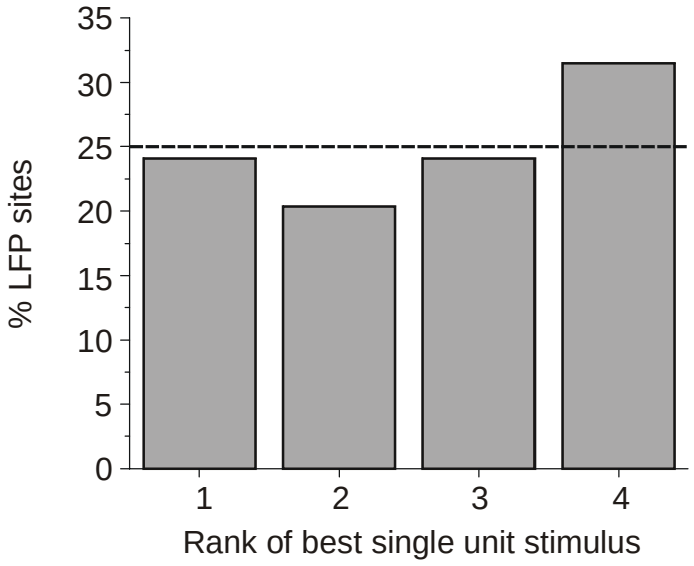
35
30
25
20
15
10
5
0
% LFP sites
1
2
3
4
Rank of best single unit stimulus

Figure 2:

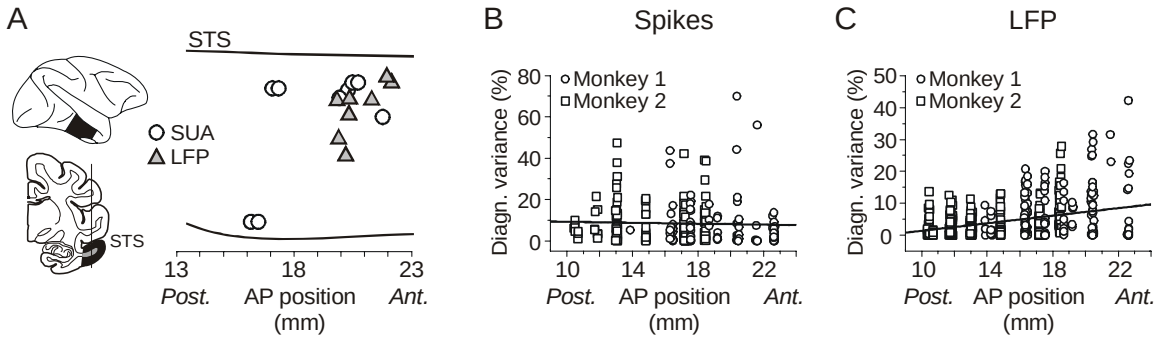
A
STS
SUA
LFP
STS
13
18
23
Post.
AP position
(mm)
Ant.
B
Spikes
Monkey 1
Monkey 2
Diagn. variance (%)
80
60
40
20
0
10
14
18
22
Post.
AP position
(mm)
Ant.
C
LFP
Monkey 1
Monkey 2
Diagn. variance (%)
50
40
30
20
10
0
10
14
18
22
Post.
AP position
(mm)
Ant.

Figure 3:

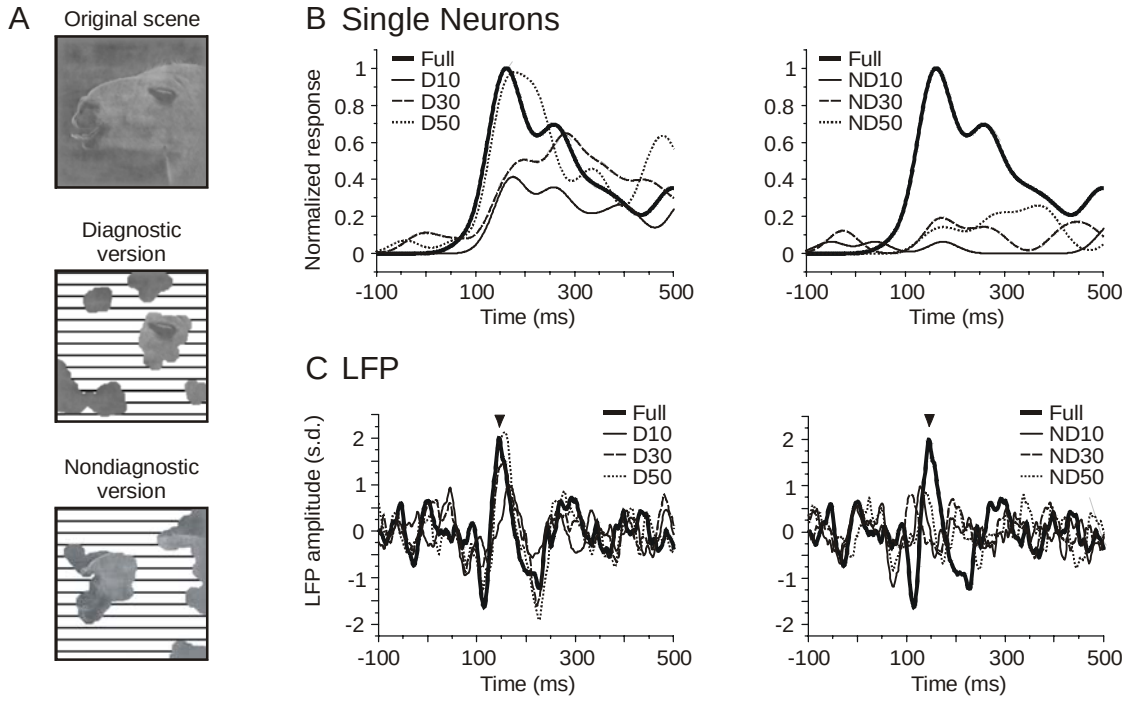
A
Original scene
Diagnostic version
Nondiagnostic version
B Single Neurons
Normalized response
Full
D10
D30
D50
ND10
ND30
ND50
Time (ms)
C LFP
LFP amplitude (s.d.)

Localizing Cortical Computations during Visual Selection

Gregor Rainer[1,*]
[1]Max Planck Institute for Biological Cybernetics, D-72076 Tübingen, Germany
*Correspondence: gregor.rainer@tuebingen.mpg.de
DOI 10.1016/j.neuron.2008.02.002

Local field potentials (LFPs) and spikes are two signals that can be recorded from the brain using extracellular microelectrodes. A study by Monosov et al. in this issue of *Neuron* using timing relations between these signals suggests that selection of a target from an array of distractors is a computational operation performed specifically and locally in the frontal eye field (FEF).

Much of what we know about the functional organization of the brain comes from single-unit activity recording studies, which have examined how spiking activity in given brain regions correlates with sensory, motor, or cognitive aspects of behavioral tasks. Spiking activity recorded using extracellular microelectrodes measures the output of neurons near the tip of the electrode. It is estimated that for cortical pyramidal cells well over 60% of connections remain local and generate synaptic activity in the same cortical region (Braitenberg and Schüz, 1998). The remaining connections are with remote cortical areas, to which the results of the cortical computations are transmitted (see Figure 1). The demonstration that a given task attribute is represented in spiking activity in a given brain region, however, does not mean that this task attribute is actually computed there. It could have been present already in the synaptic inputs, and therefore the computational work might have been performed elsewhere, and is merely being transmitted to further stages of processing. In principle, one might rule this out by simultaneously recording not only from the brain region of interest, but also from the major regions which provide its synaptic input. This is feasible for early sensory cortices, as demonstrated for example by joint recordings from the LGN and V1. For higher cortical areas, it cannot be done, mainly due to the large number of connections among cortical areas (Felleman and Van Essen, 1991). It is simply not feasible to record simultaneously from the required number of brain regions. It turns out, however, that considering the local field potential (LFP) together with spiking activity is an approach that can be used to resolve this question. Because LFPs reflect synaptic processing in a volume around the electrode tip, they are influenced by synaptic inputs arriving from other brain regions as well as by local processing (see Figure 1). In this scheme, spikes are related to local processing and output, whereas LFPs are related to local processing and input. Simultaneous measurement of spikes and LFPs in appropriate behavioral tasks could thus be used to compare these two measurements of neural activity, and thus provide estimates of which signals are already present in the inputs to the brain region under study and which signals are computed there de novo.

In this issue of *Neuron*, Monosov et al. (2008) report on joint LFP and spike recordings during spatial selection in the frontal eye fields (FEFs). During their spatial selection task, monkeys had to report the orientation or location of a target stimulus embedded in a visual array of distractors. They compared timing of LFP and spike responses in relation to two important events during each trial: array onset and selection time. In relation to the onset of the array, LFP latencies were shorter than spike latencies, as one expects since the visual inputs arriving from the sensory periphery first cause synaptic activity in the FEF before this synaptic activity is converted to spike output. In relation to selection time, i.e., the time the neural response first distinguishes the target from the distractors, they observed the opposite pattern: selection times occurred earlier in spikes than in LFPs. This suggests that information related to the target is not received from distant brain regions, but rather computed locally in the FEF. The results of the local computations are first visible in FEF spiking activity, and subsequently amplified in recurrent circuits such that they are later measurable also at the level of the LFP in the same brain area. This is a particularly exciting finding, because it allows us to consider a brain area as a computational unit that transforms incoming input signals to outputs, and to actually estimate both inputs and outputs directly from extracellular recordings. The implication is that timing differences between LFP and spike-related task-relevant neural signals can be used to distinguish whether the brain region under study is performing computational operations on incoming neural signals or simply acting to relay this activity to other brain regions.

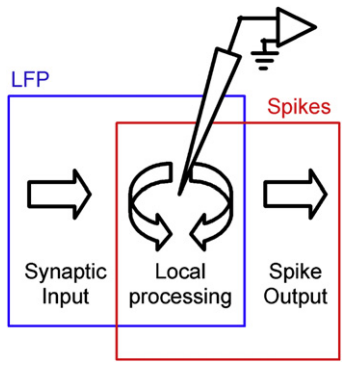

Figure 1. Simple Scheme Describing Relation of LFP and Spike Signals to Information Flow
LFPs are related primarily to local processing in synaptic inputs from other brain areas, whereas spikes are related to local processing and spike outputs.

The LFP is a mass signal that is related to synaptic activity in large populations of many thousands of neurons near the electrode tip. By contrast, spiking activity represents the output of a single neuron. This raises methodological concerns; for example, LFPs might be poorly selective to target location, and therefore selection

times estimated from LFP signals might be delayed compared with single-unit values. This was not the case in the study from Monosov and colleagues: although LFP tuning width for target location tended to be somewhat larger than those estimated from spiking activity, there was highly robust directional tuning present in both signals. In addition, directional preference was highly correlated across the two signal types, suggesting that the LFP measurements were not spatially broadly distributed and unselective, but rather precisely related to the properties of single neurons encountered at the same site. Related findings have been reported in area MT for speed and directional tuning (Liu and Newsome, 2006), and in the posterior parietal cortex for movement direction (Scherberger et al., 2005). In these cases authors examined oscillatory LFP activity and not evoked responses as in the present study, but found robust tuning for task parameters in the LFPs as well as correlations to unit activity. How does the FEF take nonselective inputs and convert them to directional signals for eye movement control? If the interpretation of the authors is correct, the neural signals entering the FEF upon array onset already contain information about the distractors and the target in their respective locations, as well as the current behavioral goals of the animal. Presumably, the job of the FEF in this context is to extract target information and convert it to an explicit representation suitable for control of action. How this conversion is achieved by FEF neural networks still needs to be worked out, but the general approach now provides a method that can be used to address this question.

In the context of understanding cortical information processing, it is of great interest whether particular influences on a given brain region have a bottom-up or top-down origin, because this would allow a dissociation between effects of sensory origin and those derived from internal representations of task demands. A recent study has presented evidence that top-down and bottom-up communication in the brain might be supported by different frequency bands of the LFP (Buschman and Miller, 2007). During an easy visual search task, the target was reflected first in the lateral interparietal (LIP) cortex and subsequently in a frontal cortical region that included the FEF, whereas the opposite was true during difficult visual search. At the same time, the authors found task-dependent differences in LIP-frontal LFP coherence, such that a 22–34 Hz LFP band showed greater activity during the difficult search task, whereas a 35–55 Hz band showed less activity. The authors suggest that LIP identifies the target first in the easy task and communicates this information to the frontal cortex in a bottom-up fashion, whereas during the hard task, the target is first reflected in frontal activity and then sent to LIP in a top-down manner. Joint LFP-spiking analyses of the kind employed by Monosov and colleagues could be used to directly test this idea; during the easy task, frontal target-related signals should appear first in the LFP, and later, in spiking activity. More generally, it is known that top-down and bottom-up projections tend to have different projection patterns, forming synapses preferentially in apical and proximal parts of dendrites, respectively. This makes it particularly appealing to extend electrical recordings of neural activity with imaging methods such as Ca^{+2} imaging (Stosiek et al., 2003), particularly if these can be further developed to examine layer-specific synaptic activity in populations of neurons. Such an approach would yet further refine our picture of information processing by including the measurement of top-down and bottom-up inputs into the brain region under study, in addition to electrical neuronal activity measurements.

The approach used by Monosov and colleagues is related to one used in a recent study of inferior temporal (IT) cortex (Nielsen et al., 2006). That study described a dissociation between spiking and LFP activity as a function of recording location in IT cortex. For posterior sites, learning-dependent object selectivity was seen in spiking activity, but not LFP activity, whereas at anterior sites this selectivity was present in both LFP and spiking activity. The interpretation of this result was that the learning-dependent signals were first generated in posterior IT and thus present in the output of that brain area and not in the input. Anterior IT sites already showed these signals in their inputs, consistent with receiving signals from posterior IT regions. This represents in a sense an orthogonal approach to one used in the study by Monosov and colleagues. The two studies have applied a similar logic to describe differences between LFP and spiking activity as a function of brain topography and temporal response dynamics. Both of these approaches can in principle also be applied together, and this combination promises substantial further advances in our understanding of computational and informational flow in cortex.

REFERENCES

Braitenberg, V., and Schüz, A. (1998). Cortex, Statistic and Geometry of Neuronal Connectivity (Berlin: Springer).

Buschman, T.J., and Miller, E.K. (2007). Science *315*, 1860–1862.

Felleman, D.J., and Van Essen, D.C. (1991). Cereb. Cortex *1*, 1–47.

Liu, J., and Newsome, W.T. (2006). J. Neurosci. *26*, 7779–7790.

Monosov, I.E., Trageser, J.C., and Thompson, K.G. (2008). Neuron *57*, this issue, 614–625.

Nielsen, K.J., Logothetis, N.K., and Rainer, G. (2006). J. Neurosci. *26*, 9639–9645.

Scherberger, H., Jarvis, M.R., and Andersen, R.A. (2005). Neuron *46*, 347–354.

Stosiek, C., Garaschuk, O., Holthoff, K., and Konnerth, A. (2003). Proc. Natl. Acad. Sci. USA *100*, 7319–7324.

Neuron
Article

Measurements of Simultaneously Recorded Spiking Activity and Local Field Potentials Suggest that Spatial Selection Emerges in the Frontal Eye Field

Ilya E. Monosov,[1,2] Jason C. Trageser,[1] and Kirk G. Thompson[1,*]
[1]Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Building 49, Room 2A50, Bethesda, MD 20892, USA
[2]Brown-NIH Graduate Partnership Program, Brown Department of Neuroscience, Providence, RI 02906, USA
*Correspondence: kgt@lsr.nei.nih.gov
DOI 10.1016/j.neuron.2007.12.030

SUMMARY

The frontal eye field (FEF) participates in selecting the location of behaviorally relevant stimuli for guiding attention and eye movements. We simultaneously recorded local field potentials (LFPs) and spiking activity in the FEF of monkeys performing memory-guided saccade and covert visual search tasks. We compared visual latencies and the time course of spatially selective responses in LFPs and spiking activity. Consistent with the view that LFPs represent synaptic input, visual responses appeared first in the LFPs followed by visual responses in the spiking activity. However, spatially selective activity identifying the location of the target in the visual search array appeared in the spikes about 30 ms before it appeared in the LFPs. Because LFPs reflect dendritic input and spikes measure neuronal output in a local brain region, this temporal relationship suggests that spatial selection necessary for attention and eye movements is computed locally in FEF from spatially nonselective inputs.

INTRODUCTION

Visual spatial selection describes the process that guides visual attention (Serences and Yantis, 2006) and selectively couples perception to action (Allport, 1987). Understanding the time course of this process is key to understanding the neural computations that underlie it. Typically, this question has been addressed by analyzing event-related brain potentials (ERPs) recorded from scalp electrodes in humans (Hillyard and Anllo-Vento, 1998; Luck et al., 2000) and neuronal spiking activity in behaving primates (Schall and Thompson, 1999). In visual search studies, in which subjects are required to discriminate a target among distractors, human ERPs (Luck and Hillyard, 1994) and single units recorded in primate frontal eye field (FEF) (Sato et al., 2001; Thompson et al., 1996), lateral intraparietal area (LIP) (Ipata et al., 2006; Thomas and Pare, 2007), and superior colliculus (McPeek and Keller, 2002) exhibit an initial period of nonselective activation followed by a discrimination process that identifies the location of the target in the search array.

Local field potentials (LFPs) are electrical potentials recorded with an electrode positioned in the brain. The LFP signal represents the summed synaptic activity occurring near the tip of the electrode. It is a combined measure of local processing and synaptic inputs from other brain regions regardless of whether or not spikes are generated (Chen et al., 2007; Cruikshank et al., 2002; Juergens et al., 1999; Kaur et al., 2004; Kreiman et al., 2006; Logothetis and Wandell, 2004; Mitzdorf, 1985, 1987; Nielsen et al., 2006). In contrast, spiking activity represents the results of local neural processing and is the output signal from the neurons near the tip of the electrode. Although both LFPs and spiking activity have been used to measure the time course of spatial attention processes, the relationship between these neurophysiological signals is still unclear. Analysis of concurrently recorded LFP and spiking activity can shed light on how sensory representations in dendritic input are transformed into cognitive signals (Kreiman et al., 2006; Nielsen et al., 2006).

The FEF is a brain area in monkeys and humans that participates in the visual spatial selection process (Awh et al., 2006; Pessoa et al., 2003; Schall and Thompson, 1999; Serences and Yantis, 2007). The spatial selection process localizes behaviorally important objects in a complex visual scene and is necessary for guiding visual attention and goal-directed behaviors. In a previous report we showed that spiking activity in monkey FEF reflects the locus of spatial attention during covert visual search tasks in the absence of eye movements (Thompson et al., 2005b). During the collection of these neuronal spiking data, LFPs were also recorded simultaneously from the same electrodes. The goals of this study were to determine whether LFP responses were spatially selective, and if so, to compare the time course and spatial tuning of the spatially selective signals in neuronal spiking activity with LFP responses.

We found that in the covert visual search task, both the LFPs and the spiking activity exhibited initial nonselective visual responses that evolved into significant spatial tuning in the time period before the monkeys' behavioral report. The directional tuning of the spatially selective responses in the visual search task matched the directional tuning of the visually evoked responses to a single visual stimulus in the memory-guided saccade task. Although the initial visual responses appeared first in the LFP signals in both tasks, the spatially selective responses in

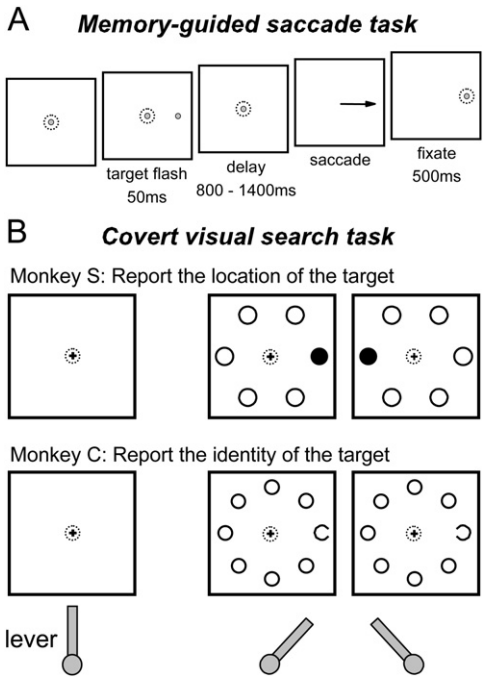

Figure 1. The Tasks

(A) The memory-guided saccade task. After the monkey fixated on a central spot, a peripheral stimulus identical to the fixation spot was flashed for 50 ms randomly at one of the six or eight locations matching the stimulus locations in the covert visual search task. After a delay, the fixation spot disappeared, and the monkey was rewarded for making a saccade to the remembered target location.

(B) The covert visual search tasks. After the monkey grasped the lever in the vertical position, a small fixation cross appeared. After fixating the central cross, a search array appeared in which one of the stimuli was different. Monkey S was rewarded for turning the lever in the same direction as a different-colored stimulus in relation to the fixation cross. Monkey C was rewarded for turning the lever in the same direction as the gap in the C target stimulus regardless of its location in the search array. The depiction of the lever at the bottom shows the correct behavioral responses for the example trials shown in the search displays.

the visual search task appeared first in the spiking activity. These results suggest that during visual search, spatial selectivity is generated in FEF from spatially nonselective inputs.

RESULTS

Spiking activity and LFP responses were recorded concurrently on single electrodes inserted into the FEF of two monkeys in 43 separate recording sessions. The monkeys performed a memory-guided saccade task (Figure 1A) and one of two covert visual search tasks (Figure 1B). In the covert visual search tasks, the monkeys made a manual lever turn as the behavioral report. Monkey S was required to report the location of the singleton target in the search array (20 recording sites) and monkey C was required to report the orientation of the C among Os in the search array (23 recording sites). Single-neuron activity recorded with this task was described previously (Thompson et al., 2005b). For this study we combined the activity from simultaneously recorded single neurons into a single representation of spiking activity at each recording site. The primary aim of this study was to compare the times that a spatially selective response first appeared in the LFPs with spikes in the covert visual search task. We refer to this time as the *selection time*. For the data collected at a recording site to be included in the study, there must have been measurable visual response onset latencies in both the LFPs and spikes, and a measurable selection time in the visual search task for either the LFP response or the spiking activity. In addition, the visual response latencies and selection times must have occurred before the average reaction time of the session. Over all sessions, lever turn reaction times averaged at 284 ms for monkey S and 297 ms for monkey C.

There were strong correlations between the directional tuning of the spatially selective responses in the LFPs and spikes within and across the visual search task and the memory-guided saccade task, which is consistent with a functional relationship between the LFPs and spikes (see Figure 8 below). But first we describe the results of the time course analysis, which is blind to the preferred target directions of the two signals.

Visual Response Latencies and Spatial Selection Times of LFPs and Spikes

The spiking activity and LFP signals recorded simultaneously at each recording site were analyzed using the same methods to obtain the visual response onset latencies and the time of spatial selection measured from the time of search array presentation. Figure 2, Figure 3, and Figure 4 illustrate the analysis applied to the data collected from a single recording site in monkey S (see Experimental Procedures for details). Briefly, selection time was defined as the first time following visual stimulus presentation that the response differed significantly across target locations based on an analysis of variance (ANOVA) at each millisecond (Figure 2 and Figure 3). In the memory-guided saccade task, spikes and LFPs exhibited initial responses that differed across target location. Therefore, for the memory-guided saccade task, selection time measures the initial visual response latency to a single stimulus. In the visual search task, however, a visual stimulus appears at all locations on every trial and the initial responses of spikes and LFPs did not vary with target position. Selection time in the visual search task, therefore, measures the first time that the response differentiates the target stimulus from the distractors. To get a measure of the visual response latency in the visual search task, we defined the visual latency as the first time following the visual search array presentation that the combined response across all trials differed from baseline (Figure 4).

Even though visual response latencies were measured using different visual stimuli and measurement methods in the memory-guided saccade and visual search tasks, the temporal relationship between initial visual response latencies measured in LFPs and spikes was the same across the two tasks. The initial visual response occurred earlier in the LFPs than in the spikes. For the memory-guided saccade task, the average ± standard error (SE) selection time was 63.4 ± 3.2 ms for LFPs, and 72.8 ± 4.3 ms for spikes (paired t test: $p < 0.001$). For the visual search task, the average ± SE onset latency was 56.5 ± 2.4 ms for LFPs, and 71.8 ± 4.0 ms for spikes ($p < 0.001$). There were also strong correlations between the selection times obtained from the

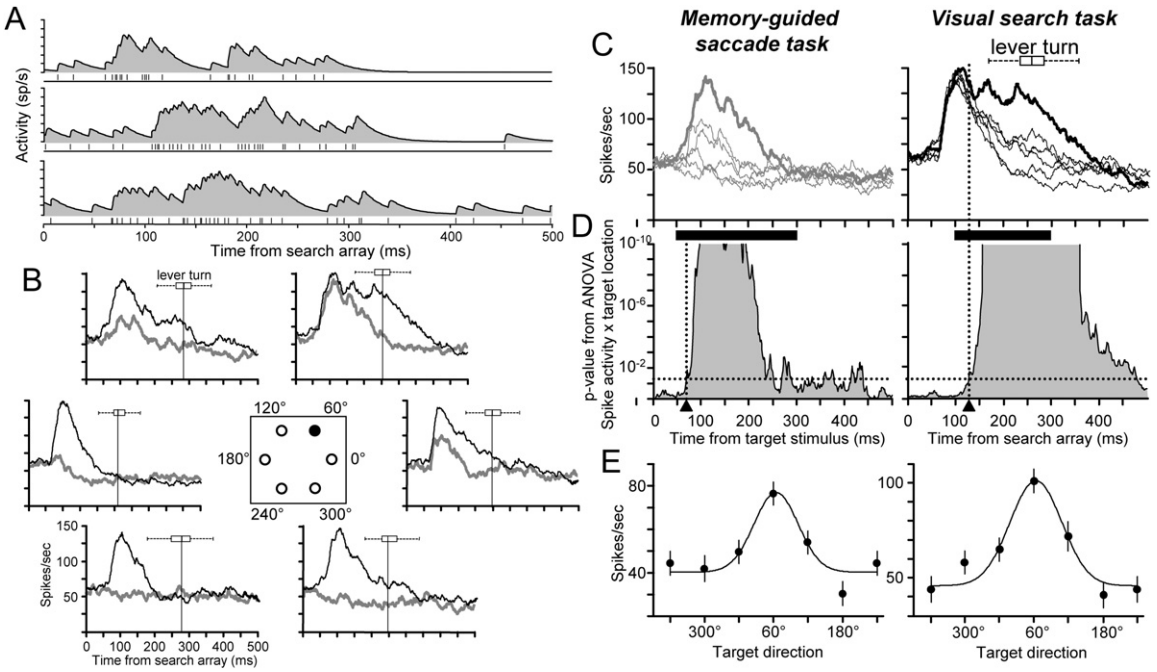

Figure 2. Spatial Tuning Analysis of Spiking Activity Recorded at the Same Site as the LFP Shown in Figure 3

(A) Spike density functions, derived from a filter resembling an EPSP, are plotted above tick marks representing times of action potentials for three representative trials.

(B) The average target-aligned spiking activity at each target location from the memory-guided saccade task (gray) and the covert visual search task (black). The box-whisker plot in each panel indicates the median, quartiles, and range of reaction times in the covert visual search task. The neuron's preferred target direction (60°) corresponds to the filled circle in the search array at the center.

(C) The superimposed average activity for each target position from the memory-guided saccade task (left) and the visual search task (right). The thick line represents the average activity on trials when the target was at the preferred spatial location.

(D) The p value (ANOVA) at each millisecond in the memory-guided saccade task (left) and in the visual search task (right) that estimates the probability that the spiking activity did not vary across target locations. The black triangle at the bottom of the plot marks the selection time (vertical dotted line: memory-guided = 70 ms, visual search = 128 ms), which was defined as the first millisecond that the p value crossed $p = 0.05$ (horizontal dotted line), but only if it continued past $p = 0.001$ and $p < 0.05$ for more than 20 of the next 25 ms.

(E) The spatially selective response measured from 50–300 ms following the target flash in the memory-guided saccade task (left), and from 100–300 ms following the time of search array presentation (right) as a function of target direction. The time ranges for measuring spatial tuning are indicated by black bars in (C). The points plot the average response within the time interval at each target location, and the error bars are the 95% confidence intervals. The parameters of the best-fit Gaussian curve from the memory-guided saccade task (left) are B = 40.26 spikes/second (sp/s), R = 36.54 sp/s, $\Phi = 64.42°$, and $T_{\varphi} = 38.25°$; and from the covert visual search task (right), they are B = 45.75 sp/s, R = 55.36 sp/s, $\Phi = 63.22°$, and $T_{\varphi} = 45.02°$.

memory-guided saccade task and the visual response onset latencies obtained from the visual search task at each recording site (LFPs: $r = 0.48$, $p = 0.001$; spikes: $r = 0.78$, $p < 0.001$). Because we were interested in comparing visual onset times to spatial selection times in visual search, in this study we will focus mostly on results obtained in the visual search tasks.

Cumulative distributions of onset latencies and selection times measured in the visual search task are shown separately for the two monkeys in Figures 5A and 5B. Visual response latencies were obtained for the spiking activity and the LFP response from all 43 recording sites. For spiking activity, the average ± SE onset latency was 68.4 ± 3.5 ms for monkey S, and 74.7 ± 3.3 ms for monkey C. For the LFP response, the average ± SE onset latency was 53.6 ± 1.0 ms for monkey S, and 59.0 ± 0.9 ms for monkey C. An ANOVA that factored the monkey and response measure revealed a significant difference in response latencies between the two monkeys ($p = 0.02$), and between spiking activity and LFP response ($p < 0.001$) with no interaction between monkey and activity measure ($p = 0.86$).

Selection times in the visual search task were obtained for spiking activity from 38 (88.4%) recording sites and for the LFP response from all 43 recording sites. For spiking activity, the average selection time was 124.6 ± 5.1 ms for monkey S, and 113.0 ± 6.2 ms for monkey C. For the LFP response, the average selection time was 155.2 ± 6.3 ms for monkey S, and 133.3 ± 7.1 ms for monkey C. An ANOVA revealed a significant difference in the selection times in the visual search task between the two monkeys ($p = 0.01$), and between spiking activity and LFP response ($p < 0.001$) with no interaction between monkey and the activity measure ($p = 0.43$).

The differences in visual response latencies and selection times between the two monkeys may be due to individual differences or to the different visual stimuli used in two different visual search tasks in the two monkeys. It has previously been shown that a search for a gap in a C among Os is very easy (Treisman and Gormican, 1988), and the visual system may be able to resolve a single gap in a circle faster than it can resolve a color difference in a search array. Nevertheless, the important result is the absence of significant interaction between monkeys

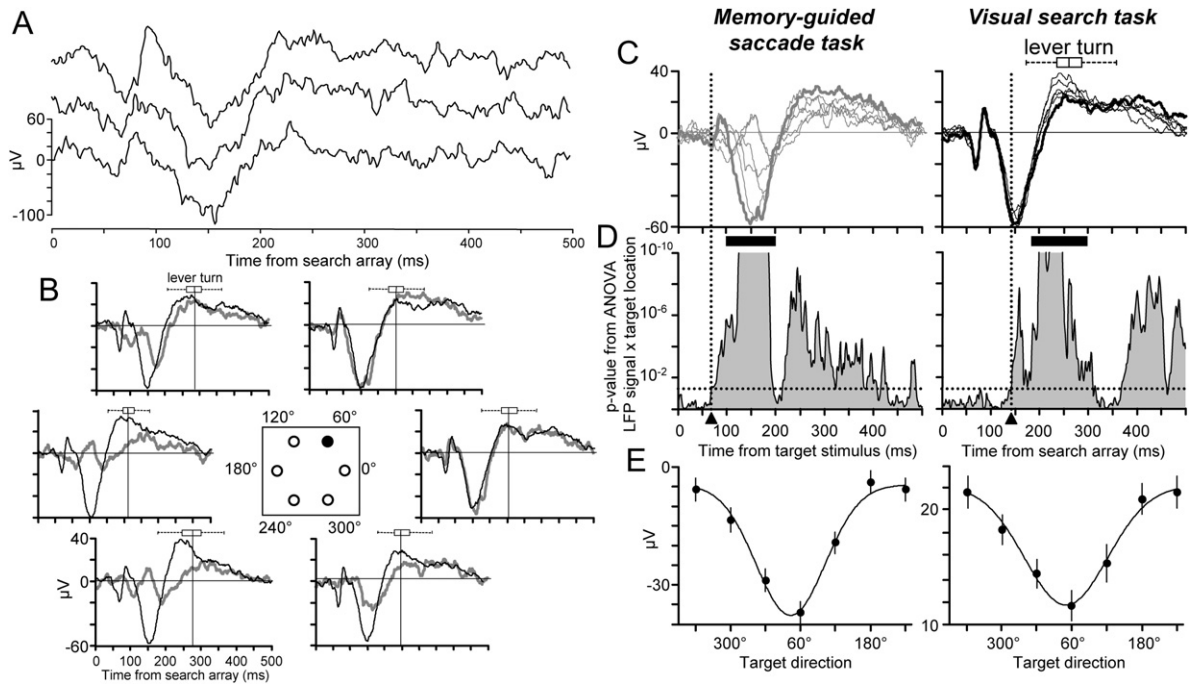

Figure 3. Spatial Tuning Analysis of the LFP Response Recorded at the Same Site as the Spiking Activity Shown in Figure 2

Conventions are the same as in Figure 2.

(A) The LFP responses on three representative visual search trials.

(B) The average target-aligned LFP response in the memory-guided saccade task (gray) and in the covert visual search task (black), sorted by target location.

(C) The superimposed average LFP response for each target position from the memory-guided saccade task (left) and the covert visual search task (right).

(D) The p value (ANOVA) at each millisecond that estimates the probability that the LFP response did not vary across target locations. The selection time of the LFP response at this recording site is 69 ms for the memory-guided saccade task and 142 ms for the visual search task.

(E) The spatially selective response measured from 100–200 ms following the target flash in the memory-guided saccade task (left), and from 180–300 ms following the time of search array presentation (right) as a function of target direction. The time interval used for determining the spatial tuning of the LFP response was the interval that exhibited the most variability in the ANOVA analysis shown in (D) (see Supplemental Data and Figure S1). The points plot the average response within the time interval at each target location, and the error bars are the 95% confidence intervals. The parameters of the best-fit Gaussian curve from the memory-guided saccade task (left) are $B = -4.59$, $R = -33.14$, $\Phi = 43.05°$, and $T_{\varphi} = 59.24°$; and from the covert visual search task (right), they are $B = 19.59$, $R = -16.33$, $\Phi = 64.05°$, and $T_{\varphi} = 62.91°$.

performing different visual search tasks and the measured timing differences between LFPs and spikes. This means that in spite of the individual differences, the temporal relationships between LFPs and spikes were the same in the two monkeys.

To summarize the results across the population, we plotted the percentage of recording sites showing significant modulation at each millisecond following the presentation of the search array, and did so separately for monkey S (Figure 5C) and monkey C (Figure 5D). These continuous measures of significant modulation across the population are another way to visualize the timing differences across LFPs and spikes, and they validate the results obtained from the calculations of initial visual response onset latencies and spatial selection times. For both monkeys, significant visual responses are evident in the LFPs before the spikes and significant spatially selective responses are evident in the spikes before the LFPs. Because the relationships between spiking activity and LFP responses were the same for both monkeys, the data from the two monkeys are combined in the following analyses.

We compared the response latencies and selection times measured from the spiking activity with LFP responses recorded simultaneously at individual recording sites during the visual search task (Figures 6A and 6B). Significant positive correlations between spiking activity and LFP responses for onset latencies ($r = 0.46$, $p = 0.002$) and for selection times ($r = 0.51$, $p = 0.001$) support the claim that spiking activity and LFP responses are related. Spiking activity and LFP response onset latencies for each recording site are plotted in Figure 6A, and selection times are plotted in Figure 6B. In both plots, the times from each site are sorted according to the time measured in the spiking activity, and a histogram shows the distribution of differences between the times obtained from the LFPs and spikes. For nearly all (41/43 = 95%) of the recording sites, the measured response onset latency was earlier in the LFP response than in the spiking activity. On average, the LFP visual response began 15.3 ± 2.2 ms earlier than the spike visual response. The visual latencies of the LFP responses varied less than the spike responses. As a consequence, the difference between visual onset latency measured in the spikes and in the LFP increased with increasing spike response latency. Nevertheless, even the recording sites with the earliest spike responses had LFP response latencies that were significantly earlier. For the quartile of recording sites with

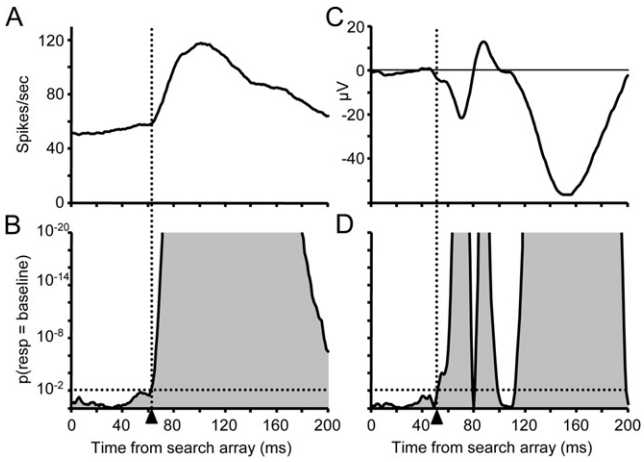

Figure 4. Visual Response Latency Analysis of the Spike (Left) and LFP (Right) Responses Recorded during the Covert Visual Search Task

The activity is from the same recording session as shown in Figure 2 and Figure 3.

(A) The average spike density function constructed by convolving each spike with a kernel that resembles an EPSP and averaging across all trials.

(B) The p value (paired t test) at each millisecond that estimates the probability that the spiking activity is equal to the baseline activity (measured from −50 to 0 ms). The visual response latency was defined as the first time that the p value crossed p = 0.01 (horizontal dotted lines), but only if it continued past p = 0.001 and p < 0.01 for more than 20 of the next 25 ms. The spiking visual response latency in the covert visual search task at this recording site is 63 ms (black triangles and vertical dotted lines).

(C) The average LFP signal across all trials.

(D) The p value (paired t test) at each millisecond that estimates the probability that the LFP signal is equal to the baseline signal (measured from −50 to 0 ms). The LFP visual response latency is 53 ms.

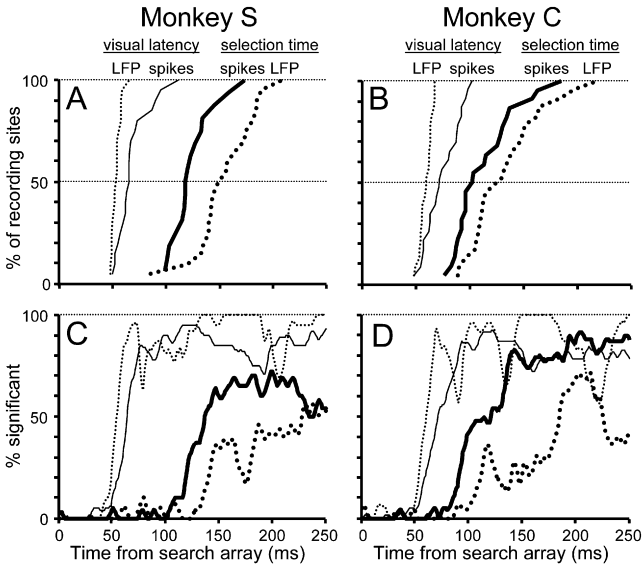

Figure 5. Population Results from the Covert Visual Search Tasks Shown Separately for the Two Monkeys

(A) Cumulative distributions of visual response latencies and spatial selection times for all recording sites in monkey S performing the "location" version of the covert visual search task. The average ± SE times, from left to right, were 53.6 ± 1.0 ms for LFP visual latencies (thin dotted line; median = 53.5), 68.4 ± 3.5 ms for spike visual latencies (thin solid line; median = 65), 124.6 ± 5.1 ms for spike selection times (thick solid line; median = 119), and 155.2 ± 6.3 ms for LFP selection times (thick dotted line; median = 152.5).

(B) The same as (A) but for monkey C performing the "identity" version of the covert visual search task. The average ± SE times, from left to right, were 59.0 ± 0.9 ms for LFP visual latencies (median = 59 ms), 74.7 ± 3.3 ms for spike visual latencies (median = 72 ms), 113.0 ± 6.2 ms for spike selection times (median = 102.5 ms), and 133.3 ± 7.1 ms for LFP selection times (median = 129 ms).

(C and D) The percentage of recording sites showing significant modulation at each millisecond following the presentation of the search array in monkey S (C) and monkey C (D). The plots were smoothed using a running window of 5 ms for easier viewing. The line types correspond to those in (A) and (B).

the earliest spike visual response (range: 48–60 ms), the LFP visual response began on average 2.3 ± 0.7 ms earlier than the spike visual response (paired t test, p = 0.01). The earlier initial visual onsets in the LFP signal are consistent with the expected result that feedforward visual inputs in postsynaptic potentials precede the visually evoked spiking activity (Schroeder et al., 1998).

Selection times in the visual search tasks were obtained for the LFP response from all 43 sites and for spiking activity from 38 sites. For the 38 recording sites with selection times from both measures, selection times occurred later in the LFP response than in the spiking activity for 84% (32/38) of the recording sites, and differed, on average, by 24.7 ± 5.0 ms (Figure 6B). However, at ten recording sites, the spatial tuning of the LFP response and spiking activity differed by more than 40° of visual angle (see Figure 8B); these are indicated in Figure 6B by the filled circles in the scatter plot and shaded bars in the histogram. It is possible that at these recording sites the LFP response and spiking activity were less related to each other than at the sites in which the spatial tuning of the two signals corresponds. When these ten sessions were removed from the analysis, the selection times occurred later in the LFP response than in the spiking activity at 93% (26/28) of the recording sites, and differed, on average, by 31.5 ± 5.1 ms.

We also compared the selection times for LFPs and spiking activity in the memory-guided saccade task (Figure 6C). Selection time in the memory-guided saccade task measures visual response latency because it identifies the first time that the responses differed across target locations for a single visual stimulus presented alone. It corresponds to the visual response latency measured in the covert visual search task, and across the recording sites the two measures were strongly correlated for both spikes (Pearson's r = 0.78, p < 0.001) and LFPs (r = 0.60, p < 0 001). Just like the visual response latencies measured in the visual search task (Figure 6A), the selection times measured in the memory-guided saccade task were earlier (9.9 ± 2.5 ms) for LFPs than for spikes (Figure 6C). The similarity in the results across the tasks and analysis methods adds to our confidence in the accuracy of our timing measurements (also see Supplemental Data).

Relationship of LFP Visual Response Latency to Selection Times

Studies have shown that the earliest visual response latencies of LFPs recorded in dorsal stream areas of visual cortex are in

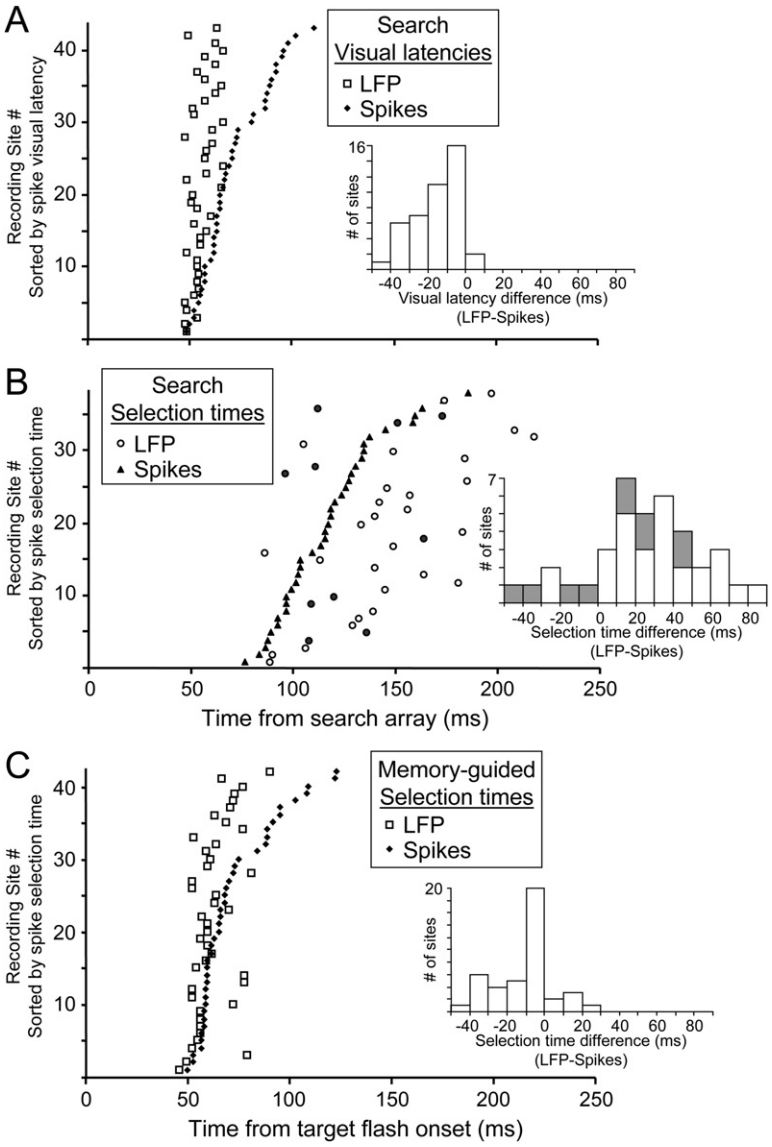

Figure 6. Population Results from the Covert Visual Search Task at Each Recording Site Combined across the Two Monkeys

(A) Visual response latencies of the LFP responses in the covert visual search task (open squares) and spikes (filled diamonds) at each recording site, sorted by the visual response latency of the spikes. LFP and spike visual response latencies were obtained from all 43 recording sites. The histogram shows the distribution of LFP visual response latency relative to spike visual response latency obtained across all recording sites (LFP – spikes; mean = −15 ± 2.2 ms). Similar results were obtained from the selection times measured in the memory-guided saccade task (see Figure 6C).

(B) Selection times in the covert visual search task of the LFP responses (open and filled circles) and spikes (filled triangles) at each recording site, sorted by the selection time of the spikes. LFP and spike selection times were obtained from 38 recording sites. The histogram shows the distribution of LFP selection time relative to spike selection time obtained across all recording sites (LFP – spikes; mean = 24.7 ± 5.0 ms). The filled circles in the scatter plot and filled bars in the histogram represent the ten recording sites in which the spatial tuning of the LFP and spikes differed by more than 40° of visual angle (see Figure 8B).

(C) Selection times in the memory-guided saccade task, measured from the LFP responses (open squares) and spikes (filled diamonds) at each recording site and sorted by the selection time of the spikes (n = 42). The histogram shows the distribution of LFP selection time relative to spike selection time across all recording sites (LFP – spikes; mean = −9.9 ± 2.5 ms). Compare with results in Figure 6A.

cortical layer 4, which corresponds to the feedforward projection of visual inputs (Chen et al., 2007; Schroeder et al., 1998). We hypothesized that if the inputs to FEF from visual cortex were spatially selective, they would be evident first at the recording sites with the earliest LFP visual response latencies. Therefore, we examined whether LFP visual response latencies were related to times of spatial selection (Figure 7). It should be noted that this analysis does not establish the cortical layer of the recording sites, but it is motivated by the assumption that recording sites in FEF with earlier visually evoked LFP activity are functionally closer to the feedforward visual input from visual cortex. To visualize the data we plotted how the spike visual response latencies and selection times, and LFP selection times, changed with increasing LFP visual response latency (Figure 7). For statistical analysis, the recording sites were divided into two groups based on LFP visual response latency measured in the visual search task. The sites with LFP visual response latencies between 48 and 55 ms were assigned to the "early" group (n = 22), and sites with latencies between 56 and 67 ms to the "late" group (n = 21). The large symbols in Figure 7 indicate the average ± SE of each group. The spike visual response latencies differed significantly across the early groups (65.0 ± 3.3 ms) and late groups (79.0 ± 2.9 ms) (t test, p = 0.003). This is consistent with the result that LFP and spike visual latencies were positively correlated. LFP response selection times did not differ significantly between the early (145.4 ± 6.9 ms) and late (141.5 ± 7.6 ms) groups (p = 0.7). For the spiking activity, the selection times of the early (109.5 ± 4.7 ms) and late (126.3 ± 6.6 ms) groups were marginally different (p = 0.04). The surprising result was that the recording sites with the earliest LFP visual response latencies, and therefore those functionally closest to the feedforward visual input to FEF, exhibited the earliest spike selection times and the latest LFP selection times. The difference between spike and LFP selection times in the early group was highly significant (paired t test, $p < 10^{-5}$). For the late group, the difference between the LFP and spike selection times did not reach statistical significance (p = 0.07). We also divided the recording sessions into early and late groups based on the selection times measured in the data collected from the memory-guided saccade task, which were recorded in a separate block of trials in each session. Note that selection time for the memory-guided saccade data is determined using the exact same analysis method as for visual search data, but actually measures the

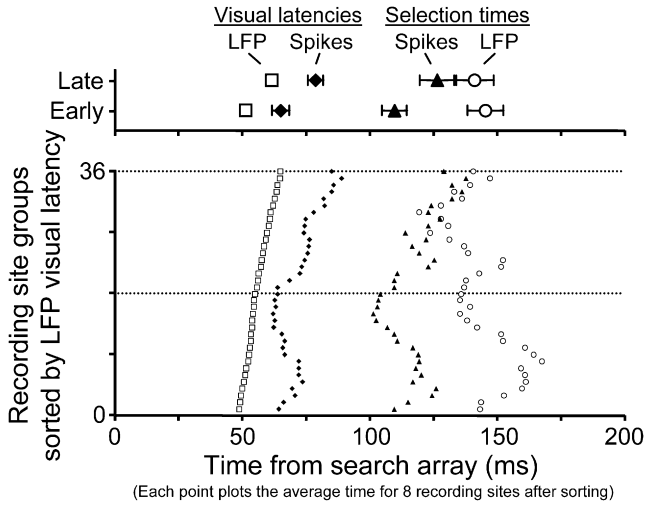

Figure 7. The Relationship of Selection Time in LFPs and Spikes to LFP Visual Response Latency

The symbols representing the different times are the same as in Figure 6. The visual response latencies and selection times across all recording sites are sorted by increasing LFP visual response latency. Each of the data points plots the average for a group of eight sorted recording sites. Consecutive data points represent the average of eight recording sites after shifting the averaging window by one. The statistical comparisons are shown at the top (large symbols). The averages ± SE of the response latencies and selection times are plotted after dividing the recording sites into two groups based on LFP visual response latency. The recording sites with LFP visual response latencies between 48–55 ms were assigned to the "early" group (n = 22; LFP visual response latencies = 51.7 ± 0.5 ms; spike visual response latencies = 65.0 ± 3.3 ms; spike selection times = 109.5 ± 4.7 ms; LFP selection times = 145.4 ± 6.9 ms), and recording sites with visual response latencies between 56–67 ms to the "late" group (n = 21; LFP visual response latencies = 61.5 ± 0.8 ms; spike visual response latencies = 79.0 ± 2.9 ms; spike selection times = 126.3 ± 6.6 ms; LFP selection times = 141.5 ± 7.6 ms).

visual onset latency to a single visual stimulus. The results were statistically identical to those shown in Figure 7.

Comparison of Directional Tuning

The variation of spatially selective LFP and spiking responses with target direction in the memory-guided saccade and visual search tasks was characterized with Gaussian functions (Figure 2E and Figure 3E). The spatial parameters of the best-fit Gaussian curves provide estimates of the preferred direction and spatial extent of the LFP and spiking response fields. Details of the spatial tuning analysis are provided in the Experimental Procedures. There were no differences in the directional tuning measures between the two monkeys. The preferred direction was provided by the optimum direction (Φ) parameter. The preferred tuning directions of the spiking activity and LFP responses in the memory-guided saccade and visual search tasks were compared by taking the angle difference between the two measures. Angle differences can range from −180° to +180°. Figure 8 shows the distributions of angle differences between the preferred target directions of the LFPs and spikes for the memory-guided saccade and visual search tasks (Figures 8A and 8B), and between the preferred target directions obtained from the

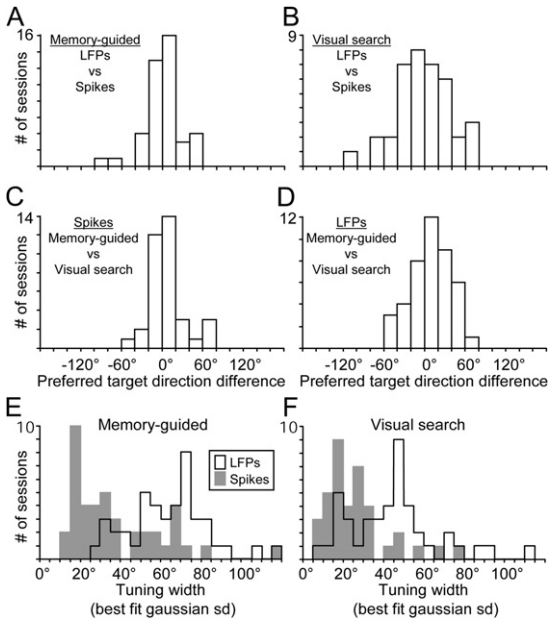

Figure 8. Comparisons of Spatial Tuning in Spiking Activity and LFP Responses Recorded in the Memory-Guided Saccade and Covert Visual Search Tasks

(A–D) The distributions of the differences in the preferred target directions measured from the spiking activity and LFP responses for the recording sites that exhibited significant spatial tuning in the spiking activity. LFP responses exhibited significant spatial tuning at all 43 recording sites in both the memory-guided saccade and the covert visual search tasks. Angle differences can range from −180° to +180°. All of the distributions are peaked near 0° (Rayleigh test, $p < 0.001$). A circular correlation analysis (Mardia and Jupp, 2000) showed that the preferred target directions are significantly correlated between (A) LFPs and spikes recorded in the memory-guided saccade task ($n = 42$; $p < 10^{-9}$); (B) LFPs and spikes recorded in the covert visual search tasks ($n = 38$; $p = 0.001$); (C) spikes recorded in the memory-guided saccade task and spikes recorded in the covert visual search tasks ($n = 37$; $p < 10^{-8}$); and (D) LFPs recorded in the memory-guided saccade task and LFPs recorded in the covert visual search tasks ($n = 43$; $p = 0.001$).

(E) The distribution of tuning widths of the LFP responses (open bars) and spike responses (filled bars) in the memory-guided saccade task. Tuning width was defined as the standard deviation (T_φ) parameter of the best-fit Gaussian curves. The average response field width is 38.7° ± 3.7° for spiking activity, and 64.0° ± 3.0° for LFP responses; and the two distributions differ significantly (paired t test, $p < 10^{-7}$).

(F) The distribution of tuning widths of the LFP and spike responses in the covert visual search tasks. The average response field width is 26.5° ± 2.6° for spiking activity, and 44.7° ± 3.3° for LFP responses (paired t test, $p < 10^{-4}$).

memory-guide saccade and visual search tasks for spikes and LFPs (Figures 8C and 8D). All the distributions are peaked near 0° (Rayleigh test, $p < 0.001$). An analysis that measures the correlation between two circular variables (Mardia and Jupp, 2000) showed that there were strong correlations between the preferred directions obtained from LFPs and spikes in the memory-guided saccade task (Figure 8A; $p < 10^{-9}$) and in the visual search task (Figure 8B; $p = 0.001$). There were also strong correlations between the preferred directions obtained across the two tasks for both spikes (Figure 8C; $p < 10^{-8}$) and LFPs (Figure 8D, $p = 0.001$). In summary, there were overall strong correlations between the directional tuning of the LFP and spike response

fields across the memory-guided saccade task, in which a visual stimulus is presented alone, and the covert visual search tasks, in which the target must be identified among distractors.

The tuning widths of the LFP and spike response fields were estimated by the standard deviation (T_φ) parameter of the best-fit Gaussian curves. The distributions of tuning widths (in polar angle coordinates) for the single visual stimulus in the memory-guided saccade task, and for the target among distractors in the search tasks, are shown in Figures 8E and 8F, respectively. For the memory-guided saccade task, the average response field width is 38.7° ± 3.7° for spiking activity, and 64.0° ± 3.0° for LFP responses. For the visual search tasks, the average response field width is 26.5° ± 2.6° for spiking activity, and 44.7° ± 3.3° for LFP responses. The results of an ANOVA showed that LFP tuning widths were significantly larger than spike tuning widths ($p < 0.001$), and the tuning widths of responses in the memory-guided saccade task were significantly larger than those in the visual search task ($p < 0.001$). When converted to visual field angles according to the law of cosines, the width of receptive fields for a single visual stimulus alone averaged 6.5° ± 0.4° for spikes and 10.4° ± 0.4° for LFPs, and receptive fields in the visual search task averaged 4.6° ± 0.4° for spikes and 7.4° ± 0.5° for LFPs. The sizes of receptive fields of the spiking activity and the narrower spatial tuning in the visual search task (as compared with that of single targets) are similar to previous reports that used comparable methods (Schall et al., 1995a; Schall et al., 2004).

DISCUSSION

We show that LFPs in FEF exhibit visually evoked responses that are spatially selective; they identify the location of a target presented alone in a memory-guided saccade task, and identify the location of a behaviorally important stimulus during covert visual search in the absence of eye movements. We compared the LFP responses to the single-unit activity recorded concurrently on the same electrodes (Thompson et al., 2005b). In the covert visual search task, both the LFPs and spikes exhibited a short-latency, spatially nonselective visual response followed by a selective response that identified the location of the behaviorally relevant stimulus that instructed the monkey to manually turn a lever to the left or right. The spatial selectivity for the behaviorally relevant target in the visual search task appeared in the spiking activity before the LFP response. This result is especially intriguing because it suggests that a cognitive representation identifying the location of behaviorally important visual stimuli is computed in the FEF from spatially nonselective inputs (Thompson and Bichot, 2005; Thompson et al., 2005a).

The spatial tuning for target location was consistent across tasks and across LFPs and spikes at each recording site, but was generally broader in the LFP signal than in the spikes. Previous spike versus LFP comparisons either used full-field visual stimulation (e.g., Chen et al., 2007; Logothetis et al., 2001) or placed visual stimuli based on the spatial extent of the spike receptive fields (Fries et al., 2001; Liu and Newsome, 2006; Pesaran et al., 2002). We are not aware of any study that compared the spatial extent of visual responses of LFPs with that of spikes recorded on the same electrode. But the broader spatial tuning in LFPs as compared with that of spikes is consistent with the view that LFPs reflect synaptic activity over a larger area of cortex than is reflected in the spiking output of a few localized neurons (Kreiman et al., 2006; Liu and Newsome, 2006; Logothetis et al., 2007; Logothetis and Wandell, 2004; Mitzdorf, 1985, 1987). Nevertheless, the overall strong correlations of spatial tuning between the LFP responses and spiking activity when a target was presented alone and when presented among distractors indicate that the LFP and spike signals originate from the same region of FEF.

In a recent study, Buschman and Miller (2007) compared the time course of spatially selective spiking activity recorded simultaneously in FEF and the LIP, an area that is interconnected with FEF, in monkeys performing visual search tasks. Their results suggest that spatial attention signals appear first in the FEF during top-down attention and first in LIP during bottom-up attention. The implication is that visually driven attention signals flow from LIP to FEF and cognitively driven attention signals flow from FEF to LIP. Although simultaneous spike recordings can be used to compare signals in interconnected areas, this experimental method does not address whether or how different brain areas influence each other or how synaptic inputs are transformed into spiking outputs in a given area. In addition, the results of Buschman and Miller (2007) have been called into question mainly due to the difficulty in knowing whether the neurons recorded in LIP and FEF in that study were those that received input from or influenced activity in the other brain area (Schall et al., 2007). The combined LFP-spike analysis described in this study may be able to address some of these unresolved issues.

Combined analysis of LFP and spiking activity can provide information about computations that cannot be obtained when these signals are considered separately (Kreiman et al., 2006; Nielsen et al., 2006). In the cerebral cortex, there is strong evidence that the LFP is a mass signal that is primarily influenced by the excitatory postsynaptic potentials (EPSPs) of dendrites (Chen et al., 2007; Cruikshank et al., 2002; Juergens et al., 1999; Kaur et al., 2004; Kreiman et al., 2006; Logothetis and Wandell, 2004; Mitzdorf, 1985, 1987; Nielsen et al., 2006), and thus reflects inputs from other brain regions as well as local neural processes mediated by interneurons. Spiking activity reflects local processing and the long range outputs of neurons to other brain regions. Simultaneous LFP and spike recordings provide a way to compare the dendritic input with the spiking output, which is required to understand the transformation of neural signals from one processing stage to the next. In general, brain areas where cognitive functions are computed should show response modulations in the spiking activity of single units before they appear in the LFP—whereas the brain areas that receive this information from other areas should show response modulations first in the LFP, or simultaneously in the LFP and spiking activity (Nielsen et al., 2006). In this study, we specifically examined the transformation of a nonselective visual representation of items in a search array into a cognitive signal that identifies the location of the behaviorally relevant target stimulus.

The FEF is an important site of convergence in the visual system (Jouve et al., 1998; Schall, 1997; Schall et al., 1995b; Vezoli et al., 2004). The FEF receives retinotopically organized input

from dorsal stream visual areas MT, MST, and LIP; ventral stream visual areas V4, TEO, and TE; and from the supplementary eye field and prefrontal areas 46 and 12. The dorsal stream innervation is most likely responsible for the fast nonselective initial visual responses we measured in the LFP and spikes (Bisley et al., 2004; Chen et al., 2007; Pouget et al., 2005; Schmolesky et al., 1998). The latencies of the initial visually evoked LFP and spike responses were correlated, appearing in the LFP signals about 15 ms before the spikes in the visual search tasks, and about 10 ms before the spikes in the memory-guided saccade task. At the recording sites with the earliest spike latencies, the LFP latency was about 2 ms earlier. The earlier visually evoked modulation in the LFP is consistent with studies in visual cortex (Logothetis et al., 2001; Schroeder et al., 1998), and with the hypothesis that the LFP signal reflects synaptic input and indicates that the initial visual response was relayed to the FEF from other brain areas.

The reverse temporal relationship was found in the visual search data when we compared the time course of spatial selectivity in the LFP response with spiking activity. Following the initial nonselective visual response, a spatially selective signal identifying the location of the search array target emerged first in the spiking activity, and then in the LFP signal about 30 ms later. The earlier spatially selective signal in the spiking activity suggests that the representation of the location of the behaviorally relevant target stimulus is computed within the FEF rather than relayed from other brain areas.

The alternative interpretation is that some modulations in synaptic activity cannot be detected in event-related LFPs using the methods we employed in this study. It is possible that FEF generates the strong spatially selective spiking signals by amplifying weak differences in the synaptic inputs. Although the exact nature of the input signals to FEF is currently unknown, they must contain information about the visual stimuli, and differences between them. Our results suggest that computations in FEF convert these differences into a strong categorical representation identifying the target location, regardless of the visual feature that differentiates the target from distractors. Consistent with this view, in our study we used two different classes of visual features, color and shape, and we obtained the same results.

The recording sites with the earliest LFP visual response latencies tended to have the earliest spatial selection times in the spiking activity. In dorsal stream visual areas of monkey cortex, LFPs recorded in lamina 4 have the shortest visual response latencies due to feedforward input from lower areas (Chen et al., 2007; Schroeder et al., 1998). We therefore made the reasonable assumption that the FEF recording sites with the earliest LFP visual response latencies were functionally closer to the feedforward inputs. Although we cannot identify the cortical layers we were recording from, the results depicted in Figure 7 suggest that spatial selectivity in FEF originates first in neurons near the feedforward input and then is distributed to the functionally more distant regions in FEF via local connections or feedback from other areas. Consistent with this view, at the recording sites with the latest LFP visual response latencies, the selection times measured in the LFP and spikes did not differ significantly. The results reported here provide evidence for such a functional architecture, though further studies are needed to test this hypothesis in greater detail.

Our results suggest that spatial selectivity during a pop-out covert visual task is generated in FEF from spatially nonselective inputs. A few studies have examined the relationships between LFP and spiking responses in other areas. In area MT, for example, Liu and Newsome (2006) found that tuning for motion direction and speed in LFP responses is highly correlated with that of spike activity. In inferotemporal (IT) cortex, Kreiman et al. (2006) showed a simultaneous time course of object selectivity in LFP responses and spiking activity. A study by Nielsen et al. (2006) showed that spikes and LFPs in IT exhibited learned object selectivity, and that the modulation of LFP responses, but not spiking activity, grew stronger from posterior to anterior IT. Because LFP modulation reflects the synaptic input, they concluded that learned object selectivity was encoded first in posterior IT and then transmitted to anterior IT. Only one study, conducted in area V4, has compared the spatial selection process measured in LFPs and spikes during visual search (Bichot et al., 2005). In that study, spatially selective responses appeared in the LFP and spikes at the same time. Although it was not specifically addressed in that study, the simultaneous modulation in LFP and spikes suggests that the spatial selectivity was present in the inputs.

The combined analysis of LFPs and spikes promises to provide useful information for understanding computations in the brain. Also, LFPs recorded in monkeys can be an important link between monkey single-unit data and human EEG and imaging data (Logothetis and Wandell, 2004; Woodman et al., 2007). For example, the spatially selective LFP response we report could be related to the attention-related modulations observed in human EEG recordings during visual search (Luck and Hillyard, 1994). Single units, LFPs, and EEG recordings provide high temporal resolution. It is more difficult, however, to localize the source of the computations reflected in EEG recordings than in the other two signals. EEGs recorded from scalp electrodes reflect the postsynaptic potentials summed over a large region of the brain that could include many areas that are related to spatial vision. The FEF is just one of the potential sources of the spatially selective signals necessary for spatial attention (Pessoa et al., 2003; Serences and Yantis, 2006). Further work is needed to determine the relationships between LFPs and spikes within and between the many regions of the brain involved in spatial attention.

EXPERIMENTAL PROCEDURES

Data Collection

The data were collected from two male monkeys (*Macaca mulatta*) weighing 8 kg (monkey S) and 6.5 kg (monkey C). All surgical and experimental protocols were approved by the National Eye Institute Animal Care and Use Committee and complied with the National Institutes of Health *Guide for the Care and Use of Laboratory Animals*.

The surgical procedures, behavioral control, and visual stimulation techniques have been described previously (Thompson et al., 2005b). The single-unit spiking activity analyzed in this study is the same as in the previous study (Thompson et al., 2005b). Often two or three units were recorded simultaneously on one electrode and sorted offline. For this study, all the single units recorded at each site were combined to represent the overall spiking activity at each recording site.

The LFPs were recorded simultaneously on the same glass-insulated tungsten electrodes as the spikes using a Plexon data acquisition system (Plexon

Inc.). The impedance of the head-stage was 40 MΩ at 1 kHz. Electrode impedance ranged from 0.5 to 1.5 MΩ. A stainless steel guide tube resting on the surface of the dura served as the reference. The signals were amplified and filtered between 154 Hz and 8.8 kHz to obtain spike data. LFP signals were digitized and sampled at 1 kHz after filtering the electrode signal between 3 Hz and 88 Hz. Analog eye position and lever position signals were digitized and sampled at 1 kHz. A test of the effects of the LFP signal filtering is provided in the Supplemental Data (see Figure S2 available online). This test showed that signal distortions from the data acquisition system did not affect the results.

Behavioral Tasks

At each recording site monkeys performed a memory-guided saccade task (Figure 1A) and one of two visual search tasks (Figure 1B) in separate blocks of trials as described in a previous report (Thompson et al., 2005b). In the memory-guided saccade task, after the monkey fixated on a 0.3° diameter gray spot for 400–800 ms, an identical gray spot was flashed for 50 ms at one of six or eight isoeccentric peripheral target locations spaced equally around the central fixation spot. The eccentricity was adjusted so that at least one of the stimulus locations was inside the receptive field of the neuron being recorded. The eccentricities of the stimuli ranged between 8° and 12° across recording sessions, depending on receptive field location. Monkeys were required to maintain fixation on the central fixation spot for a random period ranging from 800 to 1400 ms. After the fixation spot disappeared, the monkeys were rewarded for making a saccade to the remembered target location.

In the covert visual search tasks, monkeys initiated a trial by grasping a lever and holding it in a vertical position. Once the lever was within 10° of vertical, a small central yellow fixation cross (0.3°) appeared. After fixating the cross for 400 to 800 ms, a search array appeared that was made up of a target randomly placed at one of the locations used in the memory-guided saccade task and distractors at the remaining locations. Each of the search array stimuli subtended 1.5° of visual angle. The monkeys were rewarded for maintaining fixation on the central cross and making the correct lever turn (>15° from vertical) within 2 s after search array presentation; in practice, the monkeys nearly always turned the lever to the physical limit of 35° from vertical. If the monkey broke fixation on the central cross, released the lever, or made an incorrect lever turn the trial was aborted immediately. The reward was given after a correct lever turn; however, the search array remained on for an additional 250–500 ms to probe for latent saccade plans. The monkeys did not tend to make saccades to the target of the search array after obtaining the reward (Thompson et al., 2005b).

Monkey S was trained to report the location of the color singleton target of the search array (Figure 1B, upper). The stimuli were isoluminant green and red discs. The target could be either green or red, but within a block of trials the color of the target and distractors did not change. The singleton target appeared randomly at one of six stimulus locations, three to the left and three to the right of the fixation cross. A correct response was a lever turn to the left or right corresponding to the location of the target stimulus relative to the fixation cross.

Monkey C was trained to report the orientation of a C among O distractors (Figure 1B, lower). The stimuli were gray rings with one of them having a 0.5° gap randomly on the left or right. The C target appeared randomly at one of eight locations positioned around the fixation cross. A correct response was a lever turn to the left or right corresponding to the location of the gap in the C target regardless of its location in the search array.

Data Analysis

The LFP signal is a continuous measure of brain activity. A comparable measure of spiking activity was obtained by convolving each spike with a function that resembles an EPSP (Thompson et al., 1996). With this method, each spike exerts influence only forward in time and represents the postsynaptic consequences of spiking activity. The resulting spike density function reflects the onset of spiking activity at a 1 ms time resolution and is comparable to the onset of activity measured in the LFP signal. Examples of the EPSP spike density functions are shown in Figure 2A. Below we describe the analytical methods used to determine the time course of visual activation and spatial selection, and characterize the spatial tuning of spiking activity and LFP responses recorded during the memory-guided saccade and covert visual search tasks.

Selection Time

The time course of spatial selectivity in the LFP and spiking activity was determined with an ANOVA at each millisecond following the target flash in the memory-guided saccade task and the presentation of the search array in the visual search tasks (Figure 2 and Figure 3). The running ANOVA estimated the probability at each millisecond that the response did not vary across target locations. Figure 2 and Figure 3 illustrate the time course analysis for the spiking activity (Figure 2) and the LFP response (Figure 3) recorded concurrently at a single site. The selection times of the spiking activity and the LFP response were determined separately and were defined as the first millisecond that the p value dropped below the 0.05 level before continuing past the 0.001 level and remaining below the 0.05 level for more than 20 of the next 25 ms. To obtain the earliest possible selection times, a threshold of p = 0.05 was used. However, a threshold of p = 0.01 did not alter the temporal relationship between the selection times of the LFP and spiking activity. Again, the important point is that the same statistical analysis and threshold was used to determine selection times in the LFP and spiking activity in the memory-guided saccade task and in the visual search task. In Figure 2D and Figure 3D, p values obtained from the running ANOVA are plotted as a function of time on a log axis from 1 to 10^{-10} for spikes (Figure 2D) and LFPs (Figure 3D) recorded concurrently at a single site during the memory-guided saccade task (left) and the visual search task (right).

It is important to note that selection time measured in the memory-guided saccade task is qualitatively different from that measured in the visual search task. In the memory-guided saccade task, a single target stimulus is presented alone and evokes a different initial response across target locations. Therefore, selection time in the memory-guided saccade task corresponds to the initial visual response latency to a single visual stimulus. In the visual search task, however, selection time measures the first time that the responses to the target of the search array are different from the responses to the distractors. As previously shown for spiking activity (Thompson et al., 1996), and as we now demonstrate for LFPs, the initial visually evoked responses in FEF during visual search do not distinguish the target from the distractors. Therefore we used a different method to determine visual response latency in the visual search task.

Visual Response Latency during Visual Search

Figure 4 illustrates how we measured the initial visual response latencies of the spiking activity and the LFP response recorded simultaneously during the visual search task. A paired t test was performed across all correct trials comparing the average activity during the 50 ms preceding the appearance of the search array on each trial to the activity at each millisecond following the appearance of the search array. Reliable results were obtained when the visual response latency was defined as the first time that the p value dropped below the 0.01 level, but only if it continued past the 0.001 level and remained below the 0.01 level for more than 20 of the next 25 ms. When the p value threshold was 0.05, the results were about the same, except that the results from a few of the recording sites were obviously false. Therefore, a more strict threshold of p = 0.01 was used to determine visual response latency, as opposed to the less strict value used to determine selection time (above). The important point is that the same threshold was used for determining visual response latencies in the LFP and spiking activity recording during the visual search task.

Spatial Tuning

To describe the variation in the spiking and LFP responses with the location of the singleton target, the response averaged over a time interval was fit with a Gaussian function of the form

$$A(\varphi) = B + R \cdot \exp\left(-1/2[(\varphi - \Phi)/T_{\varphi}]^2\right),$$

where activation (A) as a function of meridional direction (φ) depends on the baseline response (B), peak response (R), optimum direction (Φ), and tuning width (T_{φ}). Previous reports have shown that this function effectively characterizes the spatial pattern of FEF spiking activity (Bruce and Goldberg, 1985; Schall et al., 1995a, 2004).

The best-fit Gaussian curve was obtained for the average activity measured over a time range following visual stimulus presentation. For spiking activity,

the time range was from 50 ms to 300 ms for the memory-guided saccade task, and from 100 ms to 300 ms for the visual search task. These time intervals were used because they encompassed the period of spatial selectivity observed across the data (Thompson et al., 2005b). For the LFP response in the memory-guided saccade task, the time range was from 100 ms to 200 ms because this interval encompassed a strong, spatially selective, negative-going deflection observed across all the LFP recordings (see Figure 3C, left panel). For the LFP response in the visual search task, it was necessary to determine the appropriate time interval individually for the different recording sites. This is because a spatially selective response could emerge in a positive or a negative difference in the LFP signal. Therefore, to determine the spatial tuning of the LFP signal, we made the reasonable assumption that the preferred direction was in the visual hemifield contralateral to the brain hemisphere in which the LFP signals were recorded. In some of the LFP recordings, spatial tuning was evident in positive tuning during one time interval and in negative tuning during another time interval that was separated by a nonselective period during which time the polarity of the spatial tuning switched (see Figure S1). The time interval we used for determining the directional tuning was the interval that exhibited the strongest spatial selectivity in the running ANOVA analysis described above because it was most reliable. In the Supplemental Data we show that the spatial tuning during the two time intervals was essentially the same across the population (see Figure S1). For the recording site shown in Figure 3, the strongest spatial selectivity was in the interval between 180 and 300 ms (Figures 3C and 3D), and during this interval the preferred direction corresponded to the most negative LFP signal (Figure 3E).

SUPPLEMENTAL DATA

The Supplemental Data for this article can be found online at http://www.neuron.org/cgi/content/full/57/4/614/DC1/.

ACKNOWLEDGMENTS

This work was supported by the Intramural Research Program of the NIH, National Eye Institute. We thank P. Pouget, J.D. Schall, D.L. Sheinberg, E. Bromberg-Martin, R.E. Mruczek, and B. Anderson for helpful discussions and valuable comments.

Received: June 26, 2007
Revised: October 2, 2007
Accepted: December 28, 2007
Published: February 27, 2008

REFERENCES

Allport, A. (1987). Selection for action: Some behavioral and neurophysiological considerations of attention and action. In Perspectives on Perception and Action, H. Heuer and A.F. Sanders, eds. (Hillsdale, NJ: Lawrence Erlbaum Associates), pp. 395–419.

Awh, E., Armstrong, K.M., and Moore, T. (2006). Visual and oculomotor selection: links, causes and implications for spatial attention. Trends Cogn. Sci. *10*, 124–130.

Bichot, N.P., Rossi, A.F., and Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. Science *308*, 529–534.

Bisley, J.W., Krishna, B.S., and Goldberg, M.E. (2004). A rapid and precise on-response in posterior parietal cortex. J. Neurosci. *24*, 1833–1838.

Bruce, C.J., and Goldberg, M.E. (1985). Primate frontal eye fields: I. Single neurons discharging before saccades. J. Neurophysiol. *53*, 603–635.

Buschman, T.J., and Miller, E.K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science *315*, 1860–1862.

Chen, C.M., Lakatos, P., Shah, A.S., Mehta, A.D., Givre, S.J., Javitt, D.C., and Schroeder, C.E. (2007). Functional anatomy and interaction of fast and slow visual pathways in macaque monkeys. Cereb. Cortex *17*, 1561–1569.

Cruikshank, S.J., Rose, H.J., and Metherate, R. (2002). Auditory thalamocortical synaptic transmission in vitro. J. Neurophysiol. *87*, 361–384.

Fries, P., Reynolds, J.H., Rorie, A.E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. Science *291*, 1560–1563.

Hillyard, S.A., and Anllo-Vento, L. (1998). Event-related brain potentials in the study of visual selective attention. Proc. Natl. Acad. Sci. USA *95*, 781–787.

Ipata, A.E., Gee, A.L., Goldberg, M.E., and Bisley, J.W. (2006). Activity in the lateral intraparietal area predicts the goal and latency of saccades in a free-viewing visual search task. J. Neurosci. *26*, 3656–3661.

Jouve, B., Rosenstiehl, P., and Imbert, M. (1998). A mathematical approach to the connectivity between the cortical visual areas of the macaque monkey. Cereb. Cortex *8*, 28–39.

Juergens, E., Guettler, A., and Eckhorn, R. (1999). Visual stimulation elicits locked and induced gamma oscillations in monkey intracortical- and EEG-potentials, but not in human EEG. Exp. Brain Res. *129*, 247–259.

Kaur, S., Lazar, R., and Metherate, R. (2004). Intracortical pathways determine breadth of subthreshold frequency receptive fields in primary auditory cortex. J. Neurophysiol. *91*, 2551–2567.

Kreiman, G., Hung, C.P., Kraskov, A., Quiroga, R.Q., Poggio, T., and DiCarlo, J.J. (2006). Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. Neuron *49*, 433–445.

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

Logothetis, N.K., and Wandell, B.A. (2004). Interpreting the BOLD signal. Annu. Rev. Physiol. *66*, 735–769.

Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. Nature *412*, 150–157.

Logothetis, N.K., Kayser, C., and Oeltermann, A. (2007). In vivo measurement of cortical impedance spectrum in monkeys: implications for signal propagation. Neuron *55*, 809–823.

Luck, S.J., and Hillyard, S.A. (1994). Electrophysiological correlates of feature analysis during visual search. Psychophysiology *31*, 291–308.

Luck, S.J., Woodman, G.F., and Vogel, E.K. (2000). Event-related potential studies of attention. Trends Cogn. Sci. *4*, 432–440.

Mardia, K.V., and Jupp, P.E. (2000). Directional Statistics (Chichester: John Wiley & Sons Ltd).

McPeek, R.M., and Keller, E.L. (2002). Saccade target selection in the superior colliculus during a visual search task. J. Neurophysiol. *88*, 2019–2034.

Mitzdorf, U. (1985). Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. Physiol. Rev. *65*, 37–100.

Mitzdorf, U. (1987). Properties of the evoked potential generators: current source-density analysis of visually evoked potentials in the cat cortex. Int. J. Neurosci. *33*, 33–59.

Nielsen, K.J., Logothetis, N.K., and Rainer, G. (2006). Dissociation between local field potentials and spiking activity in macaque inferior temporal cortex reveals diagnosticity-based encoding of complex objects. J. Neurosci. *26*, 9639–9645.

Pesaran, B., Pezaris, J.S., Sahani, M., Mitra, P.P., and Andersen, R.A. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. Nat. Neurosci. *5*, 805–811.

Pessoa, L., Kastner, S., and Ungerleider, L.G. (2003). Neuroimaging studies of attention: From modulation of sensory processing to top-down control. J. Neurosci. *23*, 3990–3998.

Pouget, P., Emeric, E.E., Stuphorn, V., Reis, K., and Schall, J.D. (2005). Chronometry of visual responses in frontal eye field, supplementary eye field, and anterior cingulate cortex. J. Neurophysiol. *94*, 2086–2092.

Sato, T., Murthy, A., Thompson, K.G., and Schall, J.D. (2001). Search efficiency but not response interference affects visual selection in frontal eye field. Neuron *30*, 583–591.

Schall, J.D. (1997). Visuomotor areas of the frontal lobe. In Cerebral Cortex, K. Rockland, J.H. Kaas, and A. Peters, eds. (New York: Plenum Press), pp. 527–638.

Schall, J.D., and Thompson, K.G. (1999). Neural selection and control of visually guided eye movements. Annu. Rev. Neurosci. *22*, 241–259.

Schall, J.D., Hanes, D.P., Thompson, K.G., and King, D.J. (1995a). Saccade target selection in frontal eye field of macaque. I. Visual and premovement activation. J. Neurosci. *15*, 6905–6918.

Schall, J.D., Morel, A., King, D.J., and Bullier, J. (1995b). Topography of visual cortex connections with frontal eye field in macaque: Convergence and segregation of processing streams. J. Neurosci. *15*, 4464–4487.

Schall, J.D., Sato, T.R., Thompson, K.G., Vaughn, A.A., and Juan, C.H. (2004). Effects of search efficiency on surround suppression during visual selection in frontal eye field. J. Neurophysiol. *91*, 2765–2769.

Schall, J.D., Pare, M., and Woodman, G.F. (2007). Comment on "Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices". Science *318*, 44.

Schmolesky, M.T., Wang, Y., Hanes, D.P., Thompson, K.G., Leutgeb, S., Schall, J.D., and Leventhal, A.G. (1998). Signal timing across the macaque visual system. J. Neurophysiol. *79*, 3272–3278.

Schroeder, C.E., Mehta, A.D., and Givre, S.J. (1998). A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque. Cereb. Cortex *8*, 575–592.

Serences, J.T., and Yantis, S. (2006). Selective visual attention and perceptual coherence. Trends Cogn. Sci. *10*, 38–45.

Serences, J.T., and Yantis, S. (2007). Spatially selective representations of voluntary and stimulus-driven attentional priority in human occipital, parietal, and frontal cortex. Cereb. Cortex *17*, 284–293.

Thomas, N.W., and Pare, M. (2007). Temporal processing of saccade targets in parietal cortex area LIP during visual search. J. Neurophysiol. *97*, 942–947.

Thompson, K.G., and Bichot, N.P. (2005). A visual salience map in the primate frontal eye field. Prog. Brain Res. *147*, 251–262.

Thompson, K.G., Hanes, D.P., Bichot, N.P., and Schall, J.D. (1996). Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. J. Neurophysiol. *76*, 4040–4055.

Thompson, K.G., Bichot, N.P., and Sato, T.R. (2005a). Frontal eye field activity before visual search errors reveals the integration of bottom-up and top-down salience. J. Neurophysiol. *93*, 337–351.

Thompson, K.G., Biscoe, K.L., and Sato, T.R. (2005b). Neuronal basis of covert spatial attention in the frontal eye field. J. Neurosci. *25*, 9479–9487.

Treisman, A., and Gormican, S. (1988). Feature analysis in early vision: evidence from search asymmetries. Psychol. Rev. *95*, 15–48.

Vezoli, J., Falchier, A., Jouve, B., Knoblauch, K., Young, M., and Kennedy, H. (2004). Quantitative analysis of connectivity in the visual cortex: extracting function from structure. Neuroscientist *10*, 476–482.

Woodman, G.F., Kang, M.S., Rossi, A.F., and Schall, J.D. (2007). Nonhuman primate event-related potentials indexing covert shifts of attention. Proc. Natl. Acad. Sci. USA *104*, 15111–15116.

The Journal of Neuroscience, May 1, 2001, 21(9):3175–3183

Gating of Human Theta Oscillations by a Working Memory Task

Sridhar Raghavachari,[1] Michael J. Kahana,[1,2] Daniel S. Rizzuto,[1] Jeremy B. Caplan,[1] Matthew P. Kirschen,[1] Blaise Bourgeois,[2] Joseph R. Madsen,[1,2] and John E. Lisman[1]

[1]Volen Center for Complex Systems, Brandeis University, Waltham, Massachusetts 02454, and [2]Department of Surgery, Harvard Medical School and Children's Hospital, Boston, Massachusetts 02115

Electrode grids on the cortical surface of epileptic patients provide a unique opportunity to observe brain activity with high temporal–spatial resolution and high signal-to-noise ratio during a cognitive task. Previous work showed that large-amplitude theta frequency oscillations occurred intermittently during a maze navigation task, but it was unclear whether theta related to the spatial or working memory components of the task. To determine whether theta occurs during a nonspatial task, we made recordings while subjects performed the Sternberg working memory task. Our results show event-related theta and reveal a new phenomenon, the cognitive "gating" of a brain oscillation: at many cortical sites, the amplitude of theta oscillations increased dramatically at the start of the trial, continued through all phases of the trial, including the delay period, and decreased sharply at the end. Gating could be seen in individual trials and varying the duration of the trial systematically varied the period of gating. These results suggest that theta oscillations could have an important role in organizing multi-item working memory.

Key words: theta oscillations; working memory; Sternberg; intracranial EEG; brain waves; human

Oscillations in the theta frequency band (4–9 Hz) have been extensively studied in rats (Vanderwolf, 1969; Bland, 1986; O'Keefe and Recce, 1993; Skaggs et al., 1996), where they are especially prominent during spatial exploration. These oscillations can be seen in the field potential and in the potentials recorded from individual pyramidal cells (Leung and Yim, 1986; Fox, 1989; Ylinen et al., 1995; Kamondi et al., 1998). An important observation that sheds light on the function of theta is that hippocampal place cells systematically change their phase of firing relative to theta as the rat moves through a place field (O'Keefe and Recce, 1993; Skaggs et al., 1996; Jensen and Lisman, 2000). This suggests that one function of theta is to provide a reference frame for a neural code in which different spatial information is represented at different phases of the theta cycle. It remains controversial whether theta oscillations in the rat are specialized for the organization of spatial information in the hippocampus or are more generally involved in other functions (O'Keefe and Burgess, 1999).

Given the importance of theta oscillations in the rat, it has been of interest to determine whether similar oscillations occur in humans. Theta band energy can be detected in humans by both MEG and EEG methods and is evident during working memory tasks (Gevins et al., 1997; Sarnthein et al., 1998; Klimesch, 1999; Tesche and Karhu, 2000). It has recently become possible to observe large-amplitude (>100 μV) theta oscillations in humans by intracranial EEG (iEEG), a method that uses electrode arrays to record the EEG directly from the cortical surface (Kahana et al, 1999a, b; Caplan et al., 2000). These electrodes are implanted in epileptic patients to determine the location of seizure foci. The high signal-to-noise ratio of these recordings makes it possible to detect large-amplitude oscillations with a clear spectral peak in the theta frequency range and to study the dynamics of these oscillations during individual trials. This is not generally possible with the smaller MEG or EEG (1–10 μV) signals recorded from the scalp. The iEEG study of Kahana et al. (1999b) showed that theta oscillations occurred in intermittent bouts during a maze navigation task and that the probability of their occurrence was related to task difficulty. However, it remains unclear whether theta was related to the memory or spatial components of the task.

To determine whether large-amplitude theta can occur in a task that lacks a spatial component, we have recorded from intracranial electrode arrays while subjects performed the Sternberg task, a classic test of nonspatial, multi-item, verbal working memory (Sternberg, 1966). We found that theta oscillations occur during this task and have investigated its properties. The Sternberg task is particularly well suited for examining the temporal properties of theta because each trial has a well defined period over which working memory must be maintained. Thus, it was possible to investigate the timing of changes in theta with respect to the period of working memory.

MATERIALS AND METHODS

Subjects

Our four subjects had normal range of personality and intelligence and were all able to perform the task within normal limits. Subject 1 (male, age 23), subject 2 (male, age 18), and subject 3 (female, age 22) had implanted electrode arrays, whereas subject 4 (male, age 19) had bilateral depth electrodes in the temporal lobe. The research protocol was approved by the institutional review board at Children's Hospital (Boston, MA), and informed consent was obtained from the subjects.

Received Nov. 13, 2000; revised Jan. 10, 2001; accepted Jan. 26, 2001.

This work was supported by National Science Foundation Grant IBN-9723466, National Institutes of Health Grant MH-55687, and the Alfred P. Sloan Foundation.

We thank Larry Abbott, Xiao-Jing Wang, Marc Howard, and Adam Kepecs for helpful comments on a previous version of this manuscript. We acknowledge the enthusiastic cooperation of colleagues in the Children's Hospital Epilepsy Program, including Dr. Peter M. Black and Lewis Kull. Finally, we are most grateful to the patients and their families for their participation and support.

Correspondence should be addressed to John E. Lisman, Volen Center for Complex Systems, Brandeis University, 415 South Street, Waltham, MA 02454-9110. E-mail: lisman@brandeis.edu.

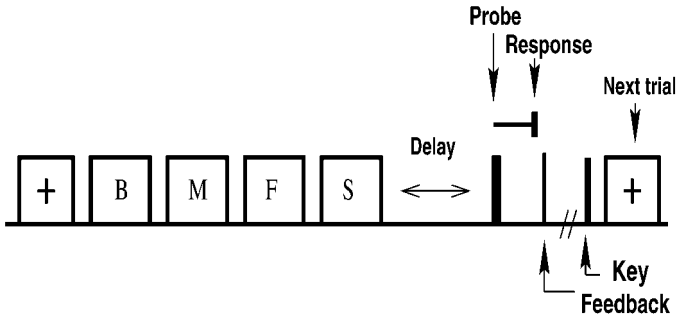

Figure 1. Schematic of the Sternberg task illustrating a four-item list using the "yes/no" procedure. A series of letters was presented after an orienting cue (+). After a delay period, a probe item was shown. Subjects indicated whether the probe was on the list, and RT was measured. After the response, the probe was turned off, and subjects received feedback on their performance and initiated the next trial by a key press.

Intracranial EEG recording

iEEG signal was recorded from arrays (grids or strips) containing multiple platinum electrodes (3 mm diameter) with an interelectrode spacing of 1 cm. Grids varied in size but covered several square centimeters of the cortical surface. The location of the electrodes was determined using coregistered postoperative computed tomograms and preoperative MRIs by an indirect stereotactic technique (Talairach and Tournoux, 1988). The iEEG signal was amplified, sampled at 200 Hz (Telefactor Corporation apparatus; band-pass filter: 0.5–100 Hz) for subjects 1 and 2, and at 256 Hz (Biologic Corp. apparatus; bandpass filter, 0.3–70 Hz) for subjects 3 and 4. Because of clock time discrepancies between the recording and experimental computers, our clock calibration was accurate to only ±200 msec.

Sternberg protocol

Lists of 1–4 consonants were presented sequentially on a computer screen. Although items were presented visually, this form of the Sternberg task is nevertheless considered a verbal working memory task because the stimuli are meaningful linguistic units (Baddeley, 1986). To start each trial, a visual orienting cue was presented 1 sec before the first list item (Fig. 1). Items were presented for 1.2 sec each with a 200 msec interval between items. The termination of the last item in the list was followed by a delay period of either 0.9 sec (subjects 1 and 2) or 2 sec (subjects 3 and 4), after which the probe was presented. The probe consisted of two letters for subjects 1 and 2 (forced choice variant), with one letter drawn from the presented list. The subject responded by pressing the left Control key if the first probe item was on the list and the right Control key if the second probe item was on the list. Subjects 3 and 4 were tested using the standard "yes"/"no" version of the Sternberg task, with a single probe item (Fig. 1*a*). The subjects responded by pressing the left Control key if the probe item was on the list and the right Control key otherwise. After each response, subjects received accuracy feedback (correct, incorrect) and latency feedback (very fast, fast, good response time, slow) via a screen message and then initiated the next trial by pressing a key. The subsequent trial began 1.6 sec after this key press. The mean interval between the response for one trial and the start of the next trial was ~2.5 sec. During each session, trials of each list length were randomly interleaved. We obtained 50, 96, 140, and 140 trials at each list length for subjects 1–4, respectively. Only correct trials with RTs <2.5 sec were used for analysis. Because there was no significant difference in our results for correct "yes" and "no" trials, data were pooled across these trial types.

Exclusion criteria

Subjects were excluded from analysis if their behavioral performance was poor (mean response times >2 sec or had high error rates). Approximately half the subjects (four of a total of nine subjects) that were tested were able to perform the task satisfactorily. Sites that were located over known lesions (determined from clinical records) or were involved in seizure onsets (identified by examining the seizure records) were excluded as were sites that showed epileptiform spiking (interictal spikes or spike-and-waves) activity. A total of 73 such sites (of 320) were rejected.

Data analysis

Power spectra. Because the oscillatory nature of the iEEG data was of interest, data analysis was done in the frequency domain. The power spectrum is the Fourier transform of the autocorrelation function. A simple estimate of the power spectrum, the magnitude-squared Fourier transform of the data has poor "bias" (the power at nearby frequencies contribute to the power at any given frequency, distorting the estimate) and variance (the estimate of the spectrum does not converge to the true value even if the data length increases) properties (Thomson, 1982). Multitaper techniques (Thomson, 1982; Mitra and Pesaran, 1999) provide a formal method to obtain estimates of the spectrum with optimal bias and variance properties. Briefly, the data set is windowed (tapered) using a set of special windows (Slepian windows), which are maximally concentrated in a time duration, T, and a bandwidth in frequency, W (Thomson, 1982). The time and frequency resolution of the windows thus fixes the number of windows, $K = 2TW - 1$, that can be used. The windowed data is then transformed to the frequency domain by calculating the Discrete Fourier transform, resulting in K estimates of the spectrum, $S_k(f)$. Averaging these estimates reduces the variance of the spectrum by $\sqrt{K}$. Our typical choices for T and W were 1 sec and 2 Hz, respectively. The averaged power spectra were obtained by averaging the single trial estimates.

Spectrograms. The spectral properties of stationary data sets do not change over time, i.e., the power spectrum of any stretch of data is statistically similar to any other stretch. If however, the spectrum varies over time, the data set is nonstationary. One method to quantify nonstationarity is to compute a time-varying spectrum, or spectrogram. Spectrograms were computed using the squared modulus of the complex demodulates [projection of the iEEG data onto different frequency bands using filters (1 sec duration, 4 Hz bandwidth) constructed from the Slepian windows (Mitra and Pesaran, 1999)]. Estimates from different Slepian windows were averaged together to obtain the spectrogram for each trial. The spectrograms for each trial were aligned with the onset of the first list item and averaged together. Only oscillatory activity with high signal-to-noise ratio will be apparent in averaged spectrograms (Tallon-Baudry et al., 1996).

Test for gating. Gating of theta was tested by comparing the energy in the average spectrogram during the trial to the energy in the 1 sec before the orienting stimulus. Because the distribution of energies in the spectrogram is non-Gaussian, a nonparametric method (Mann–Whitney U test; $p < 0.05$) was used to compare the average energy in each 250 msec epoch during the trial with the intertrial energy. Because the analysis windows were 1-sec-long, adjacent 250 msec bins are not independent. Multiple comparisons (for the number of electrodes, frequencies and bins) were corrected for by a Bonferroni correction.

Nonstationarity test. A second method to quantify nonstationarity is to expand the spectrogram, $S(f, t)$, along an orthogonal set of basis functions $A_l(t)$ such that:

$$S(f, t) = \sum_{l=0}^{L-1} a_l(f) A_l(t),$$

results in coefficients, $a_l(f)$ that are functions of the frequency alone, with L denoting the number of terms retained in the expansion. Quadratic-inverse theory (Thomson, 1990, 2001) can be used to pick an appropriate basis set, $A_l(t)$, such that the number of terms in the expansion, L is fixed to be $4TW$, where T and W are defined above, fixing the time and frequency resolution. Coefficients of higher order are identically zero. The coefficients, $a_l(f)$, of for the quadratic-inverse basis then take on special meaning. The 0th order coefficient, $a_0(f)$ is approximately $S(f)$, or the time-averaged spectrum. The first order coefficient, $a_1(f)$, is the time-derivative of the spectrum and so on. Thus, features such as sharp or gradual changes in power, frequency drifts etc. can be readily identified in a noisy background.

For a constant amplitude signal of a single frequency, the expansion coefficients vanish for all orders ≥1. For a stationary process, the ratio:

$$\Xi(f) = \left(\sum_{l=0}^{L-1} \frac{a_l(f)}{S(f)} \right)^2,$$

where $S(f)$ is the mean power, is χ^2_{L-1}-distributed. If the signal is systematically nonstationary at a given frequency f across several trials, the ratio will be significantly different from the expected value $L - 1$ for

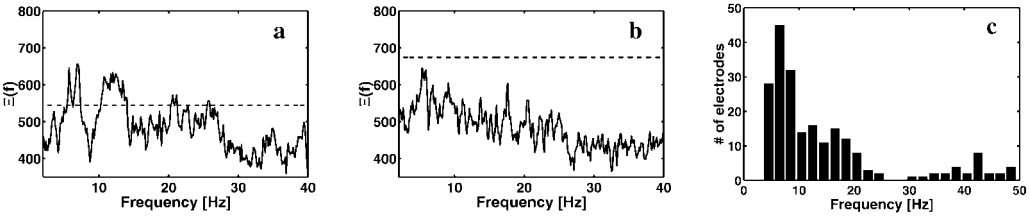

Figure 2. iEEG data show significant task-related nonstationarity, predominantly in the theta frequency band. To be considered significant, the nonstationarity index must be higher than the horizontal line, which denotes the 99.999% point of the χ^2 distribution with the appropriate degrees of freedom (see Materials and Methods). Nonstationarity index $\Xi(f)$ shown for two representative electrodes. *a*, Subject 1, Talairach coordinates (left–right, anterior–posterior, inferior–superior) are +44, +11, +38. This electrode exhibits significant nonstationarity in the theta and beta bands, as well as some peaks in the 20–30 Hz range. *b*, Subject 3, Talairach coordinates are −42, +8, +42. This shows an electrode that has no significant peaks. In both cases, the average power spectrum had peaks in the theta frequency range. For a majority of frequencies, the measure is distributed between the 5 and 95% confidence intervals with most values around the theoretical mean value, indicating that the method is appropriate. *c*, Summary plot of the number of electrodes which showed significant nonstationarity as a function of frequency. A given electrode could show nonstationarity at several different frequencies.

a χ^2_{L-1} process. This results in a single number representing the amount of nonstationarity at each frequency. Consider a stretch of iEEG data around a trial of the Sternberg task. If the spectral characteristics of the iEEG change because of the onset and offset of the task, or within the task itself, the nonstationary index $\Xi(f)$, will be significantly different from that expected by chance. The degrees of freedom, L for a single trial is equal to the highest order term retained in the expansion, L_{max}. For multiple trials, this becomes $L_{max} \times N_{trials}$. We considered the signal to be nonstationary at a given frequency f if the ratio exceeded a percentile threshold (typically 99.999% or $p < 0.00001$) of the χ^2_{L-1} distribution at that frequency. It is appropriate in our case to use a high value of significance given the large number of frequencies (256) and sites that were tested. This test allows a classification of iEEG data as stationary or non-stationary at a given frequency.

Test for continuity. To assess the continuity of theta within a trial, the baseline level of theta was first established by calculating the spectrum for a 1 sec interval after the response for each trial, and the individual estimates of the spectrum were log-transformed. The jack-knife variance (Mitra and Pesaran, 1999) was computed from individual estimates by leaving out one trial at a time and averaging over the remaining estimates. The resulting jack-knife statistic is *t*-distributed, and a threshold value was chosen as the 99.999% point of this distribution. The continuity of theta within a trial was now assessed as the fraction of trials (from all trials of all list lengths) at which the narrow-band power dipped below the threshold for any interval of >0.25 sec. The jack-knife statistic is used because it is a robust measure that does not make any assumptions about the underlying distribution of the data.

RESULTS

Figure 1 illustrates the structure of each trial of the Sternberg task. We visually presented lists of one to four consonants. After a delay period, the subjects' task was to indicate as rapidly as possible whether a probe item was on the list. We quantified the speed of the response by measuring the response time (RT). This task was administered to three subjects who had intracranial electrode arrays and one with depth electrodes. Each of these subjects performed the Sternberg task with very high accuracy; for subjects 1–4, accuracy was 86, 98, 97, and 96%, respectively. RT increased significantly with list length (LL) for all subjects ($p < 0.005$). This increase, approximated by the equation $RT = a \times LL + b$ msec, had coefficients $(a, b) = (89, 817)$, (95, 1008), (40, 463), and (37, 353) for subjects 1–4, respectively. The differences in RTs between the first two and last two subjects was most likely a consequence of differences in the design of the trial structure (see Materials and Methods).

To examine the oscillations occurred during the Sternberg task, it was desirable to have an unbiased algorithm to detect consistent task-related changes in several frequency ranges. Because the data set obtained was extensive (~200 trials/subject; total of 247 sites), examination of the entire data set by eye was impossible. We adapted a test developed by Thomson (2001) to detect task-related changes in different frequency bands (see Materials and Methods). The nonstationarity index, $\Xi(f)$, identifies electrodes in which the task produces transient or maintained changes in spectral power at any frequency f. No assumptions are made about the timing, duration, or sign of the changes. The only requirement for detection is that the changes be consistent across trials.

This nonstationarity test was applied to all sites that were not rejected for epileptic artifacts (see Materials and Methods). Figure 2, *a* and *b*, shows $\Xi(f)$ for two representative electrode locations. The broken line in the two panels represents the 99.999% confidence level for the statistic. Figure 2*a* shows an electrode for which $\Xi(f)$ exceeds this level in the theta, beta, and gamma frequency bands, indicating consistent task-related changes at these frequencies. Figure 2*b* shows an electrode in which $\Xi(f)$ did not exceed the required significance level at any frequency, suggesting little or no task-related activity at this site. We detected a total of 74 electrodes (of a total of 247 across all subjects) for which $\Xi(f)$ exceeded the 99.999% significance level at one or more frequencies. The electrode locations of these sites were widely dispersed over the cortex (24 in the temporal lobe, 18 in the occipital lobe, 18 in the parietal/motor/premotor areas, and 14 in the frontal lobe). Figure 2*c* shows a plot of the number of nonstationary electrodes at different frequencies. The majority of these (60) had significant nonstationarity in the theta frequency range (4–9 Hz), most prominently between 6 and 8 Hz. We therefore conclude that there are widespread task-related changes in theta during a memory task that lacks a spatial component. Task-related changes were also observed in the gamma frequency range, but these will be analyzed elsewhere.

To determine how theta changed during the task, we computed trial-averaged spectrograms for the nonstationary sites. An examination of these spectrograms revealed an interesting pattern of task-related activity at some sites: theta power increased at the beginning of the trial, was elevated through item presentation and the delay period, and decreased after the response. Figure 3*a* shows the averaged spectrograms from sites that display such a pattern in each of the four subjects. The average spectrograms for these sites show a clear peak in the 4–9 Hz range, the theta frequency band. Although peak frequencies and overall levels of activity varied across subjects, the general pattern at these gated sites is similar. Because the pseudocolor plots of the spectrograms emphasize certain transitions in power while making others less

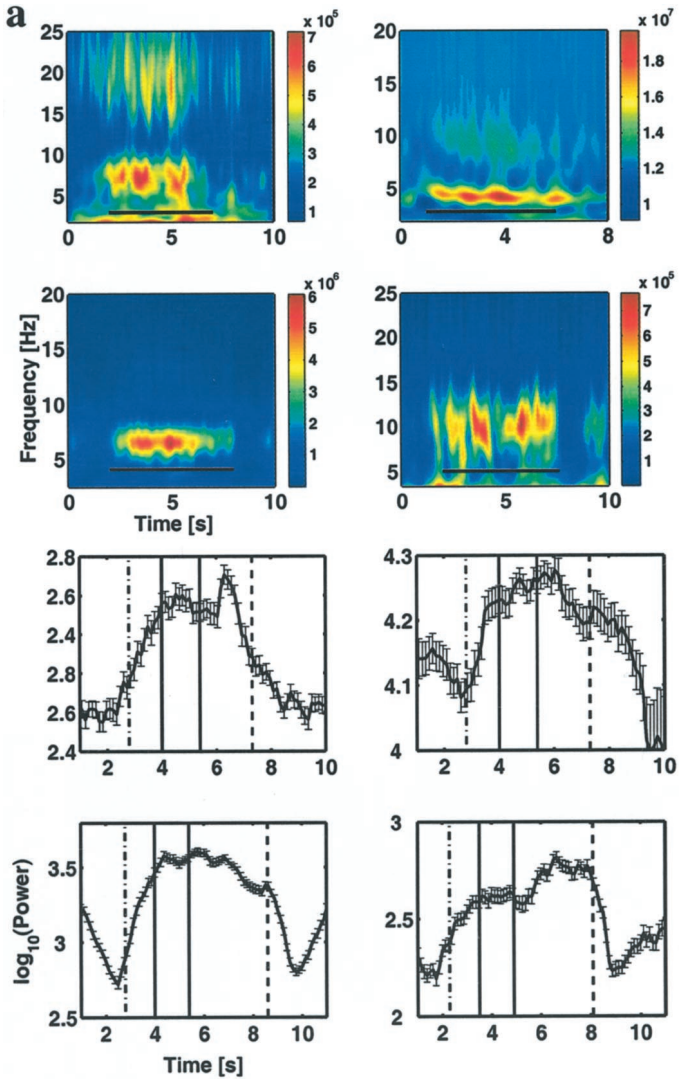

Figure 3. Theta is gated during the Sternberg task. *a*, Time-frequency energy averaged over all two-item lists shows sustained theta activity. An example is shown from each of the four subjects. The data illustrated were obtained from a right frontal site in subject 1 (Talairach coordinates are +44, +11, +38) (*top left*); a left temporal site in subject 2 (Talairach coordinates are −42, −9, −10) (*top right*); a right frontal site (Talairach coordinates are −52, −34, +38) in subject 3 (*bottom left*); and a depth electrode in left temporal lobe (Talairach coordinates are −25, −72, +6) in subject 4 (*bottom right*). *Black bars* in the spectrograms denote the trial duration from the orienting cue to the mean response time for two item lists. Because of limitations in our synchronization techniques and methods of time-frequency analysis, the determination of the onset and offset of theta has a precision of ±200 msec. The color scale represents power in square microvolts. In subject 1, we also observed similar gating centered around 18 Hz. However, because this finding was not duplicated across subjects, we did not analyze it further. *b*, Evolution of average theta power in time for the above four electrodes for a bandwidth of 4 Hz around the peak frequency. The *dot–dashed vertical line* marks the orienting cue, the *two solid lines* denote the list items, and the *dashed line* denotes the probe. Theta power is elevated throughout the trial, with fluctuations within the trial. The error bars denote the 95% confidence intervals.

visible, it is important to graphically plot changes in theta power as a function of time. Figure 3*b* shows the evolution of narrowband power at the same four sites averaged over all trials with two-item lists (4 Hz bandwidth around the peak frequency in the spectrogram). In all cases, theta power was elevated during the trial relative to the intertrial period. Note that the falling phase before trial onset in the bottom left panel occurs because the average intertrial interval for this subject was unusually brief (2 sec): because this interval was of the same order as the data window for the spectrogram (see Materials and Methods), the falling phase can be attributed to the offset of the previous trial. Note also that the shifts in theta power during the trial were relatively small. The most prominent feature at these sites is the gating "on" of theta at the onset of the trial and the gating "off" at the end of the trial. The average changes in power were high, increasing by a factor of 2 (top panels, subjects 1 and 2) or 8 (bottom panels, subjects 3 and 4). One question that remains unclear is whether theta is activated by the cue initiating the trial or the presentation of the first memory item. Technical limitations (see Materials and Methods) prevent us from determining the onset of gating with a precision better than ±200 msec. It is therefore unclear whether theta turns on with the orienting cue in anticipation of the need for engaging working memory or whether it turns on with the presentation of the first memory item. Experiments with longer delays between the orienting cue and the first list item would be useful in clarifying this issue.

It was desirable to develop a statistical test to determine whether this gating was statistically significant and whether it could be seen at a large number of sites. We therefore adopted a test for gating: that the average theta power (across trials) in every overlapping 250 msec epoch within the trials exceed the power during the intertrial period at the 95% confidence level. Thirty sites (of the 74 classified as nonstationary) met this criterion ($p < 0.01$, by a Bonferroni corrected, two-tailed, Mann–Whitney U test). One or more gated sites was detected in each of our subjects. It should be emphasized that sites that pass the gating test necessarily have an increase in theta power during the "pure" memory period, i.e., the interval after the offset of the last list-item and the onset of the probe (0.9 sec in the forced-choice variant and 2 sec otherwise) compared with the baseline power immediately after the response and before the onset of the subsequent trial (t test; $p < 0.01$). This observation indicates that theta is engaged during the pure working memory period without possible confounds of item presentation. The remainder of the electrodes that showed significant nonstationarity in the theta range typically had elevated theta power during a fraction of the trial duration, and thus were not classified as gated sites. We will not discuss these sites further.

To determine whether gating was dependent on the duration of the task, we examined responses to trials of different list lengths (and consequently trial duration). Figure 4 illustrates the change in gating with list length. Two examples are shown, one from a recording site on the surface of the left parietal lobe and one from a depth electrode in the left temporal lobe. In both cases, the duration of sustained theta closely followed the duration of the trial. It can also be seen that the maximum of the average theta power at these sites did not vary significantly with list length. Similarly, the frequency of theta did not change as additional items were presented (Fig. 3*a*). The pattern of gating at other sites was similar. We conclude that theta oscillations of relatively stereotyped frequency and power were gated by each trial of the task and that the period of gating coincided well with the duration of the trial.

Although Figures 3 and 4 indicate that the average theta power is continuous at gated sites, the possibility remains that theta is not continuous during individual trials. In fact, this seemed likely, because previous iEEG recordings (Kahana et al., 1999b) showed

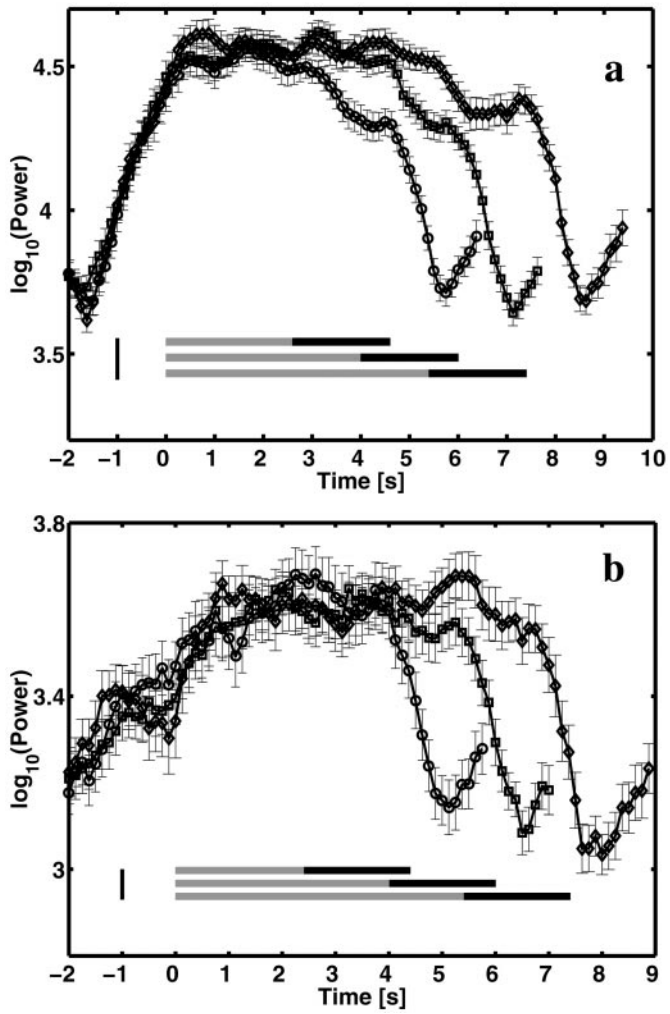

Figure 4. Gating varies systematically with list length. Averaged theta power (5–9 Hz) as a function of time shows that theta is elevated for the entire duration of the trial. The three different traces are averages over trials with two-, three- and four-item lists (*circles*, *squares*, and *diamonds*, respectively). *Gray bars* mark the presentation of the list items, and the *black bars* mark the delay interval until the presentation of the probe for the two-, three-, and four-item lists. The large tick at −1 sec marks the onset of the orienting cue. *a*, Recording from a subdural electrode in the parietal cortex (subject 3, Talairach coordinates are −52, −34, +38). *b*, Recording from a depth electrode in left temporal lobe (subject 4, Talairach coordinates are −25, −72, +6). The rise subsequent to the end of the trial is attributable to the onset of the next trial.

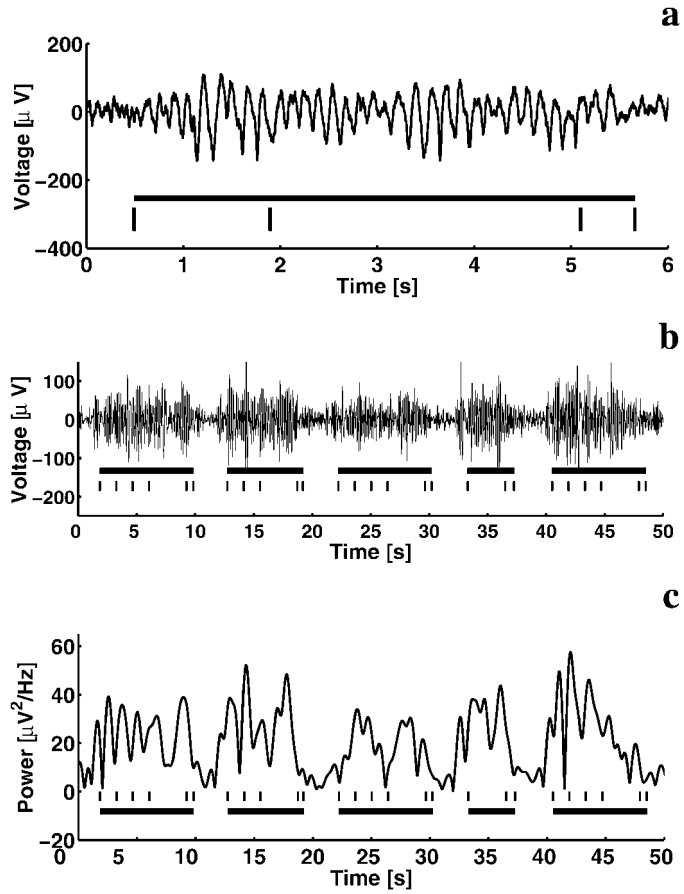

Figure 5. Gating of theta oscillations is evident in single trials of the raw iEEG signal. *a*, Sample raw iEEG trace recorded from an electrode in the parietal cortex (subject 3, Talairach coordinates are −52, −34, +38) during a two-item list. The *black bar* below the trace marks the task duration, whereas the *ticks* denote the presentation of the list items, probe, and response, respectively. *b*, A 50 sec iEEG trace with five consecutive trials from the same electrode shows clear enhancement in theta activity for the duration of each trial. *Bars* and *tick marks* are as above. *c*, Narrow-band power (7 ± 1 Hz) for the 50 sec trace above shows clear enhancement during trials relative to intertrial intervals.

that theta occurs intermittently during a spatial maze navigation task. As seen in Fig. 5*a*, which shows an unfiltered trace, theta appears to be continuously elevated during a trial of the Sternberg task. Indeed, theta oscillations were similarly gated during each of five consecutive trials (Fig. 5*b*). Also shown (Fig. 5*c*) is the time evolution of the narrow band power (2 Hz bandwidth) at the peak frequency (7 Hz) over the course of these successive trials. This plot shows that theta power during the task was greater than the level during the intertrial periods for a large fraction of each trial. In a more rigorous analysis of the ten gated sites with the largest amplitude theta (central region in subject 3; depth electrodes in subject 4), we calculated the fraction of trials for which there was a return to baseline theta power (see Materials and Methods) for any interval >0.25 sec during individual trials. The fraction of such trials was very low (ranging from 0.05 to 0.1 over all trials for all three list lengths). We conclude that there are many sites at which theta is continuous or nearly so during individual trials.

Several interleaved controls indicate that the signals at theta-gated sites were not directly related to sensory stimulation or to the execution of a motor response. Between successive trials, subjects were given visual feedback regarding their performance on the previous trial. This information was presented on the same monitor as the list items. However, as illustrated in Figure 6, this sensory stimulus did not evoke theta activity at gated sites. A second issue concerns the possibility that theta might occur in preparation for motor responses. However, Figure 6 shows that theta did not occur in anticipation of the motor response (key press) by which subjects initiated the next trial. More quantitatively, we compared the theta power in the 1 sec before the response at the end of the trial to the theta power 1 sec before the key press (75 trials of all list lengths in each subject with a 1 sec interval between the response and the key press). The theta power before the key press was significantly smaller ($p < 0.01$; t test) than before the response. We conclude that theta activation

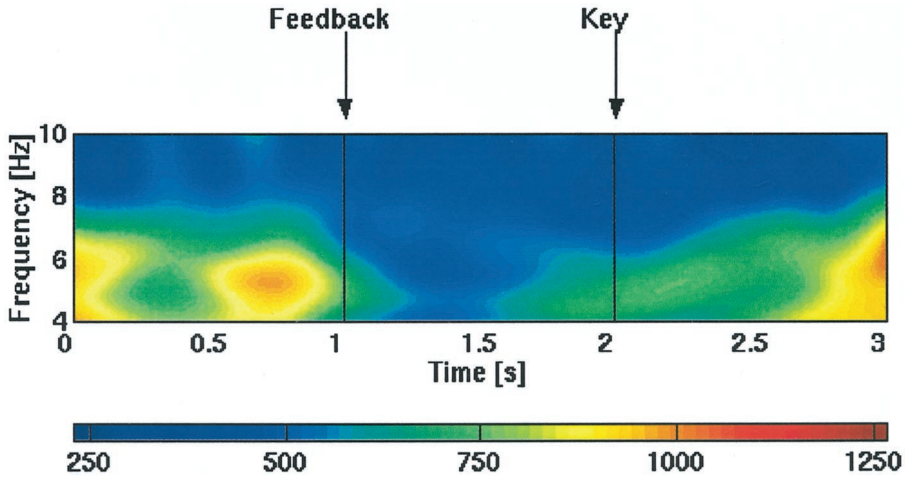

Figure 6. Theta activity not caused by sensory stimulus or motor responses. Averaged spectrogram for a site showing gated theta activity (subject 1, Talairach coordinates are +44, +11, +38). The trials were aligned to the response. The bars marked *Feedback* and *Key* denote the sensory stimulus (visual feedback after the response) and the mean time of the motor response to initiate the subsequent trial, respectively. One hundred trials of all list lengths with a mean delay of 1 sec between the response and the key press (to initiate the next trial) were used to compute the spectrogram. Theta activity (~6 Hz peak frequency) has a sharp offset after the response and stays off until the beginning of the next trial (1.6 sec after the key press). The small increase in theta activity after the key press is caused by averaging trials of different mean intervals between the feedback and key press. The color scale represents power in square microvolts per Hertz.

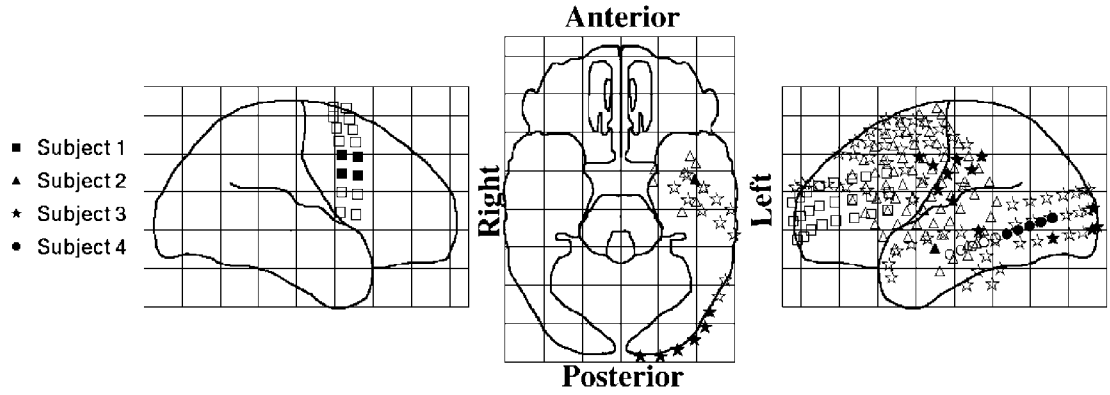

Figure 7. Views of standardized brain showing electrode locations of gated sites. *Filled symbols* indicate electrodes in which theta was gated "on" by the task. *Open symbols* indicate electrodes that did not meet our significance threshold. *Different symbols* indicate different subjects.

cannot be simply explained as a sensory or motor preparatory response.

Although we have found a large number of gated sites (30), it is difficult to make any strong conclusions about the distribution of these sites on the cortical surface because of the sparse sampling. It is important to understand that the electrode arrays were placed in candidate seizure loci. While providing details within the coverage area, the electrodes only covered a small fraction of any one lobe. Thus, our failure to detect activity in a given lobe of a subject does not imply that it was not present in that lobe. Despite these limitations, it is important to document the location of gated sites (Fig. 7). These appear to be distributed widely over the cortex in frontal, temporal, parietal, and occipital lobes. Furthermore, in two patients in which theta was detected with large arrays, we found that many of the gated sites were clustered near each other. However, closely spaced sites did not necessarily show similar activity. Figure 8 shows average spectrograms (*left panels*) and average power spectra (*right panels*) from each of three nearby electrodes (1–2 cm separation; subject 3). Sites with gated theta activity (*middle row*) sometimes occurred near other sites with no clear task-related theta activity (*top row*). Furthermore, in this subject (subject 3), there were sites (*bottom row* in Fig. 8) where theta activity was gated off by the task (i.e., the theta power was suppressed throughout a trial and rebounded after the response). Because such "off" gating was only detected in one subject, we describe it here only because it provides further evidence that large-amplitude theta can be very different at closely spaced sites.

DISCUSSION

Theta in rats has been most reliably elicited by movement, and it has therefore been suspected that theta may have a special role in spatial processing. It was thus of considerable interest that the first observation of large-amplitude theta in humans was during a spatial task (Kahana et al., 1999b). However, this task also had a memory component, leaving the possibility that theta might also occur in memory tasks that lack a spatial component. We therefore obtained iEEG data from subjects performing a verbal working memory task to test whether this nonspatial, working memory task also elicited large theta frequency oscillations. Using an objective test for nonstationarity, we showed that the Sternberg task evokes clear task-related changes in the iEEG in the theta frequency band (Fig. 2) at some cortical and subcortical sites. The power spectra showed a theta peak, the amplitude of which increased markedly during the task compared with baseline levels. Our finding that theta occurred during a task that lacked a spatial component strongly argues against the view that human theta is uniquely specialized for spatial computations. This con-

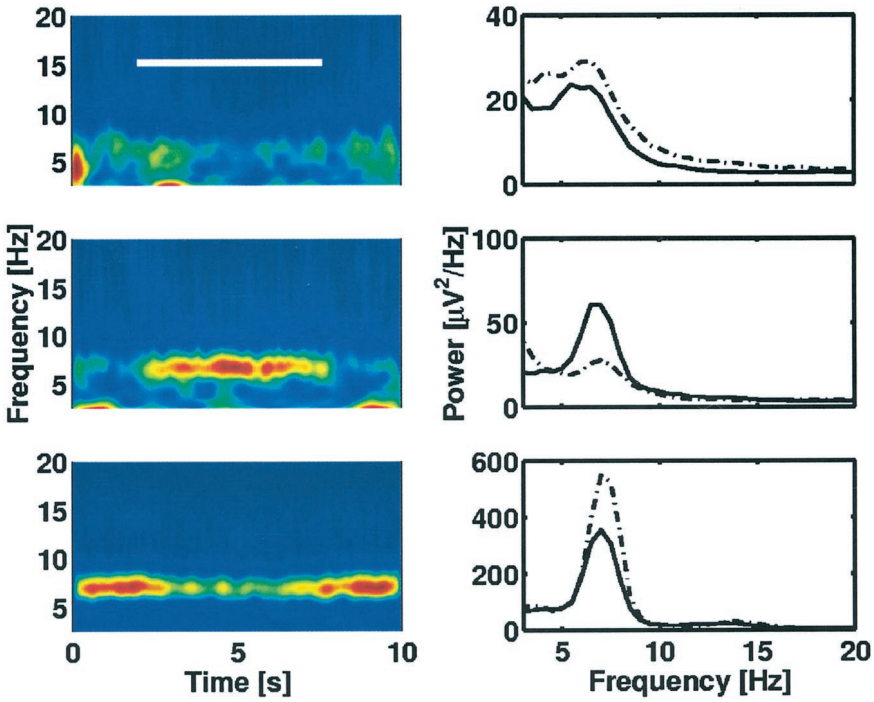

Figure 8. Nearby regions (spacing ~1–2 cm) can have dramatically different patterns of theta activity. *Left panels* show averaged spectrograms. *Right panels* show averaged power spectra for the corresponding electrodes (*solid lines* denote in-task power, *dot–dashed lines* denote out-of-task). The *white bar* in the *top left panel* marks the duration of the trial from the onset of the orienting cue until the presentation of the response. All three electrodes show theta activity evidenced by the peaks in the power spectra. The *top panel* shows an electrode with continuous, theta activity that is weakly modulated by the task. Significant task-gated theta increase in activity is evident for the electrode in the *middle panel.* The *bottom panels* show suppression of theta activity during the trial. (Talairach coordinates from *top* to *bottom* are −50, −14, +41; −54, −16, +29; −56, −24, +23.)

clusion is consistent with several observations in rat (Macrides et al., 1982; Givens and Olton 1990, 1995; Givens, 1996) and humans (Gevins et al., 1997; Sarnthein et al., 1998; Klimesch, 1999; Tesche and Karhu, 2000), indicating that theta can occur in nonspatial contexts.

A second major finding of this study is that we detected a large number of sites (30 in four subjects) in which the amplitude of theta oscillations increased at the beginning of the trial, stayed elevated through the entire trial, and decreased at the end (Fig. 3). We term this phenomenon "gating." When the duration of the trial was changed, theta gating changed accordingly (Fig. 4). Although previous EEG studies indicated that theta could occur during working memory tasks (Gevins et al., 1997; Klimesch, 1999), the timing of the involvement of theta was not investigated because the structure of the tasks was not suited for the study of timing issues. In contrast, the working memory component of the Sternberg task has a well defined onset and offset, which allowed us to detect a direct correlation of theta with task duration.

Although average theta power (across trials) was gated at these sites, we confirmed that individual trials also exhibited this gating. The high signal-to-noise ratio of the iEEG allowed us to determine that theta is continuous, relative to a baseline, at electrodes with high-amplitude, gated theta (Fig. 5). This is in contrast to the intermittent nature of theta observed in a spatial task (Kahana et al., 1999b). However, it is possible that if we knew when the memory demands occurred in the latter task, it would be continuous during those periods. The continuous nature of theta has important implications for models of working memory that are based on oscillatory activity (see below).

Several findings indicate that activity observed at theta gated sites cannot be a simple consequence of the sensory and motor components of the task. In interleaved controls (Fig. 6), we found that a sensory input or a motor response that was unrelated to working memory did not evoke theta at gated sites. These results, along with the tight temporal linkage of theta gating to the onset and offset of the period of working memory, suggest that theta oscillations may play an important role in human verbal working memory.

Spatial organization of theta

Our findings indicate that gated theta, although common, is not uniformly present (Fig. 7). The locations of gated sites were widely dispersed over the cortex. Although iEEG is uniquely well suited to give a fine grained, high temporal resolution view of theta, it is not well suited to establishing the regional localization of theta, because the grid placement is sparse and determined solely by clinical considerations. It is nevertheless tempting to try to relate the limited data available to brain regions implicated in working memory by fMRI methods (Ungerleider, 1995; Goldman-Rakic, 1995; Smith and Jonides, 1998). However, we caution against this for several reasons. First, recent work on theta in rats indicates that periods of high and low theta have nearly the same overall rates of firing (Csicsvari et al., 1999). Thus, changes in theta amplitude may not be detected by hemodynamic methods. A second point concerns the special methods that are used in fMRI studies to isolate the brain regions that are specifically involved in an aspect of brain function by subtracting the activation evoked by a simpler task that controls for sensory and other processes. In our case, all areas engaged by the task, including purely sensory areas, might be expected to show task-related oscillations. Indeed, the activation we observe in the occipital cortex might be related to sensory processes rather than memory processes. Our finding that widely distributed brain regions generate theta during a working memory task is consistent with EEG studies also show increased theta synchronization between posterior and frontal regions (Sarnthein et al., 1998) during a working memory task.

Oscillatory basis of working memory

Electrophysiological studies of working memory indicate that persistent firing of cells underlies working memory (Goldman-Rakic, 1995). Our results suggest that this firing may have an oscillatory character. Oscillatory single unit activity has not generally been reported in the delayed response tasks in monkeys (but see Nakamura et al., 1992), but it is not clear how to relate animal electrophysiological studies on single-item nonverbal

working memory (Goldman-Rakic, 1995) to the multi-item verbal working memory that we have studied in humans. It is possible that verbal working memory is more complex than the simpler forms used in animal studies (Baddeley, 1986) and that this may explain why oscillatory activity has not generally been seen in single units during simple working memory tasks in monkeys.

Relevance to models

Memory performance in the Sternberg task has been extensively studied, and the behavioral results strongly constrain possible models. Jensen and Lisman (1998) have proposed several variants of oscillatory models that account for the details of response time distributions in the Sternberg task. Their models were inspired by the observation (in rat) that different spatial information is encoded at different phases of the theta cycle (see introductory remarks). They propose that similar phase coding may be important in multi-item working memory (Lisman and Idiart, 1995; Jensen and Lisman, 1998) with different memory items active at different phases of the theta cycle. The continuous nature of theta during individual trials of the Sternberg task (Fig. 6) provides support for such models. In one of the variants of the Jensen–Lisman model, the frequency of theta oscillations was assumed to decrease as a function of the number of items being held in working memory. This model would seem to be ruled out by our finding that theta frequency does not vary significantly with memory load (Fig. 3*a*) at gated sites. A second model was based on the assumption that the phase of theta is reset by the arrival of the probe, an assumption that is supported by recent MEG results (Tesche and Karhu, 2000).

Although we have focused here on the possible role of theta in multi-item working memory, there are other possible roles, none of which are mutually exclusive. One possibility, for which studies of long-term potentiation provide some evidence (Pavlides et al., 1988; Huerta and Lisman, 1993), is that theta is used to rapidly encode information directly into long-term memory by synaptic modification. Another possible function of theta is to synchronize different regions of the cortex that participate in the task (Sarnthein et al., 1998). Analysis of the synchronization of theta at different sites during the Sternberg task is currently underway.

Concerns about the validity of data derived from epileptic patients

Intracranial recordings from epileptic patients are increasingly being used to study brain activity during cognitive tasks (Fried et al., 1997; Fernandez et al., 1999; Kahana et al., 1999b; Caplan et al., 2000; Kreiman et al., 2000). In such studies the possible contribution of epilepsy to the conclusions needs to be addressed. A number of observations suggest that the presence of theta activity during the Sternberg task is not a result of seizure activity. First, the precision of the gating (Fig. 3) and the high degree of spatial localization (Fig. 7) are exactly the opposite of what would be expected from an uncontrolled process like epilepsy. Second, because the location of the seizure origin is not known before the electrode implantation, iEEG from many regions is sampled to identify the focus. Thus, many of the sampled sites are far from the clinically determined epileptogenic foci. We observed task-related theta in each of our subjects at sites that were distant from the seizure foci, sometimes even in different hemispheres (in subjects 1 and 4). Third, direct examination of seizure activity in these subjects showed it to have a much higher amplitude (>1 mV peak-to-peak) with different spectral structure than even the largest amplitude theta signal during the task (200 μV peak-to-peak). Fourth, recent work using MEG has detected theta activity in normal subjects during the Sternberg task (Tesche and Karhu, 2000). Finally, the patients in this study had behavioral performance similar to normals on the Sternberg task. Thus, the task-related theta activity does not appear to be a consequence of the pathology of the subjects.

Concluding remarks

Elucidation of the properties of theta in rats progressed rapidly because spatial exploration is such a reliable task for eliciting rat theta. The present findings indicate that a working memory task is as good at eliciting theta in humans as spatial exploration is in rats. The highly reliable way in which human theta can be elicited by the Sternberg task and the ability to precisely control the cognitive demands of the task make this an ideal experimental system for the further study of the role of theta in memory and cognition.

REFERENCES

Baddeley AD (1986) Working memory. Oxford, UK: Clarendon.

Bland BH (1986) The physiology and pharmacology of hippocampal formation theta rhythms. Prog Neurobiol 26:1–54.

Caplan JB, Kahana MJ, Sekuler R, Kirschen M, Madsen JR (2000) Task dependence of human theta: the case for multiple cognitive functions, in press.

Csicsvari J, Hirase H, Czurko A, Mamiya A, Buzsáki G (1999) Fast network oscillations in the hippocampal CA1 region of the behaving rat. J Neurosci 19:RC20:1–4.

Fernandez G, Effern A, Grunwald T, Pezer N, Lehnertz K, Dumpelmann M, Van Roost D, Elger CE (1999) Real-time tracking of memory formation in the human rhinal cortex and hippocampus. Science 285:1582–1585.

Fox SE (1989) Membrane potential and impedance changes in hippocampal theta rhythm. Exp Brain Res 77:283–294.

Fried I, MacDonald KA, Wilson CL (1997) Single neuron activity in the human hippocampus and amygdala during recognition of faces and objects. Neuron 18:753–765.

Gevins A, Smith ME, McEvoy D, Yu L (1997) High-resolution EEG mapping of cortical activation related to working memory: effects of task difficulty, type of processing, and practice. Cereb Cortex 7:374–385.

Givens B (1996) Stimulus-evoked resetting of the dentate theta rhythm: relation to working memory. NeuroReport 8:159–163.

Givens BS, Olton DS (1990) Cholinergic and GABAergic modulation of medial septal area: effect on working memory. Behav Neurosci 104:849–855.

Givens B, Olton DS (1995) Bidirectional modulation of scopolamine-induced working memory impairments by muscarinic activation of the medial septal area. Neurobiol Learn Mem 63:269–276.

Goldman-Rakic P (1995) Cellular basis of working memory. Neuron 14:477–485.

Huerta PT, Lisman JE (1993) Heightened synaptic plasticity of hippocampal CA1 neurons during a cholinergically induced rhythmic state. Nature 364:723–725.

Jensen O, Lisman JE (1998) An oscillatory short-term memory buffer model can account for data on the Sternberg task. J Neurosci 18:10688–10699.

Jensen O, Lisman JE (2000) Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding. J Neurophysiol 83:2602–2609.

Kahana MJ, Caplan JB, Sekuler R, Madsen JR (1999a) Using intracranial recordings to study theta. Response to O'Keefe J and Burgess N (1999). Trends Cognit Sci 3:406–407.

Kahana MJ, Sekuler R, Caplan JB, Kirschen M, Madsen JR (1999b) Human theta oscillations exhibit task dependence during virtual maze navigation. Nature 399:781–784.

Kamondi A, Acsady L, Wang X-J, Buzsaki G (1998) Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: activity-dependent phase-precession of action potentials. Hippocampus 8:244–261.

Klimesch W (1999) EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. Brain Res Brain Res Rev 29:169–195.

Kreiman G, Koch C, Fried I (2000) Category-specific visual responses of single neurons in the human medial temporal lobe. Nat Neurosci 3:946–953.

Leung LS, Yim CY (1986) Intracellular records of theta rhythm in hippocampal CA1 cells of the rat. Brain Res 367:323–327.

Lisman JE, Idiart MA (1995) Storage of 7 ± 2 short-term memories in oscillatory subcycles. Science 267:1512–1515.
Macrides F, Eichenbaum HB, Forbes WB (1982) Temporal relationship between sniffing and the limbic θ rhythm during odor discrimination reversal learning. J Neurosci 2:1705–1717.
Mitra PP, Pesaran B (1999) Analysis of dynamic brain imaging data. Biophys J 76:691–708.
Nakamura K, Mikami A, Kubota K (1992) Oscillatory neuronal activity related to visual short-term memory in monkey temporal pole. Neuro-Report 3:117–120.
O'Keefe J, Burgess N (1999) Theta activity, virtual navigation and the human hippocampus. Trends Cognit Sci 3:403–406.
O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3:317–330.
Pavlides C, Greenstein YJ, Grudman M, Winson J (1988) Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of theta-rhythm. Brain Res 439:383–387.
Sarnthein J, Petsche H, Rappelsberger P, Shaw GL, von Stein A (1998) Synchronization between prefrontal and posterior association cortex during human working memory. Proc Natl Acad Sci USA 95:7092–7096.
Skaggs WE, McNaughton BL, Wilson MA, Barnes C (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus 6:149–172.
Smith E, Jonides J (1998) Neuroimaging analyses of human working memory. Proc Natl Acad Sci USA 95:12061–12068.
Sternberg S (1966) High-speed scanning in human memory. Science 153:652–654.
Talairach J, Tournoux P (1988) Co-planar stereotaxic atlas of the human brain. Stuttgart, Germany: Verlag.
Tallon-Baudry C, Bertrand O, Delpuech C, Pernier J (1996) Stimulus specificity of phase-locked and non-phase-locked 40 hz visual responses in human. J Neurosci 16:4240–4249.
Tesche C, Karhu J (2000) Theta oscillations index human hippocampal activation during a working memory task. Proc Natl Acad Sci USA 97:919–924.
Thomson DJ (1982) Spectrum estimation and harmonic analysis. Proc IEEE 70:1055–1096.
Thomson DJ (1990) Quadratic-inverse spectrum estimates: applications to plaeoclimatology. Philos Trans R Soc Lond [A] 332:539–597.
Thomson DJ (2001) Multitaper analysis of nonstationary and nonlinear time series data. In: Nonlinear and nonstationary signal processing (Fitzgerald WJ, Smith RL, Walden AT, and Young PC, eds) Cambridge, UK: Cambridge UP, in press.
Ungerleider L (1995) Functional brain imaging studies of cortical mechanisms for memory. Science 270:769–775.
Vanderwolf CH (1969) Hippocampal electrical activity and voluntary movement of the rat. Electroencephalogr Clin Neurophysiol 26:407–418.
Ylinen A, Soltesz I, Bragin A, Pentonnen M, Sik A, Buzsaki G (1995) Intracellular correlates of hippocampal theta rhythm in identified pyramidal cells, granule cells, and basket cells. Hippocampus 5:78–90.

Neuron, Vol. 45, 147–156, January 6, 2005, Copyright ©2005 by Elsevier Inc. DOI 10.1016/j.neuron.2004.12.025

Phase Locking of Single Neuron Activity to Theta Oscillations during Working Memory in Monkey Extrastriate Visual Cortex

Han Lee,[1] Gregory V. Simpson,[2] Nikos K. Logothetis,[1] and Gregor Rainer[1,*]
[1]Max Planck Institute for Biological Cybernetics
Spemannstrasse 38
D-72076 Tübingen
Germany
[2]Department of Radiology
University of California, San Francisco
185 Berry Street, Suite 350
San Francisco, California 94143

Summary

Working memory has been linked to elevated single neuron discharge in monkeys and to oscillatory changes in the human EEG, but the relation between these effects has remained largely unexplored. We addressed this question by measuring local field potentials and single unit activity simultaneously from multiple electrodes placed in extrastriate visual cortex while monkeys were performing a working memory task. We describe a significant enhancement in theta band energy during the delay period. Theta oscillations had a systematic effect on single neuron activity, with neurons emitting more action potentials near their preferred angle of each theta cycle. Sample-selective delay activity was enhanced if only action potentials emitted near the preferred theta angle were considered. Our results suggest that extrastriate visual cortex is involved in short-term maintenance of information and that theta oscillations provide a mechanism for structuring the recurrent interaction between neurons in different brain regions that underlie working memory.

Introduction

Working memory—the short-term maintenance of behaviorally relevant information—is an essential component for any higher cognitive function. It provides a mechanism that allows the brain to delay action and gather further relevant sensory evidence or query long-term memory stores to prepare an optimal behavioral response, rather than having to act immediately in reflex-like fashion to sensory stimuli. Working memory thus facilitates the coordination of multiple neural systems, the integration of task-relevant information from a variety of sources, and the selection of appropriate actions. It also provides a mechanism by which current plans or behavioral goals and contingencies can be maintained (Asaad et al., 2000; Rainer et al., 1998) and if necessary rapidly modified (Asaad et al., 1998). Since the description of neurons showing sustained elevated discharge of action potentials during delay periods (Fuster and Alexander, 1971; Fuster and Jervey, 1981; Kubota and Niki, 1971) several decades ago, this so-called "delay activity" has been considered to play a major role in the short-term maintenance of memories. Many studies since then have provided support for this view and greatly advanced our knowledge of the effects of stimulus type and modality on delay activity and its temporal dynamics (Funahashi et al., 1993; Fuster et al., 2000; Romo et al., 1999).

In humans, working memory has also been a subject of intense investigation using scalp and intracranial electroencephalography (EEG, iEEG) as well as magnetoencephalography (MEG), which provide estimates of local population activity. The published findings include reports of systematic changes in signal amplitude during working memory in the theta (Raghavachari et al., 2001; Tesche and Karhu, 2000), alpha (Gevins et al., 1997; Jensen et al., 2002), beta (Tallon-Baudry et al., 1999), and gamma (Howard et al., 2003; Lutzenberger et al., 2002; Tallon-Baudry et al., 1998) bands. There have also been reports of working memory-specific increases in synchronization between distinct brain regions, including elevation of synchronization in the beta band between signals recorded using iEEG at two occipital sites (Tallon-Baudry et al., 2001) and elevation of synchronization in the theta band between frontal and occipital EEG electrode signals (Sarnthein et al., 1998; Stam et al., 2002).

The relationship between these oscillatory phenomena in the EEG and the effects seen in the discharge of single neurons remains largely unexplored, yet it represents an important link between large-scale activation across neural populations and action potentials that are the fundamental computational elements that neurons use to communicate.

In this study, we investigate the relation between oscillatory phenomena and single neuron activity during working memory by simultaneously recording the local field potential (LFP)—a signal similar to the iEEG recorded in human patients—and single unit activity (SUA) from awake behaving monkeys performing a working memory task. The task required monkeys to retain information about a briefly presented visual stimulus over the course of a delay period. We presented visual stimuli at different contrast levels, chosen such that monkeys performed well at high contrast but were unable to exceed chance performance at low contrast. This allowed us to compare conditions in which monkeys were, or were not, holding information in short-term memory. We recorded neural activity simultaneously from multiple electrodes placed in the occipital visual cortex (extrastriate area V4) and focused our analyses largely on the theta frequency (4–8 Hz) band. We did this because occipital visual cortex is an area where the iEEG is characterized by high-amplitude oscillatory activity in the theta range during working memory (Raghavachari et al., 2001) and because synchronization between occipital and prefrontal EEG in the theta range has been implicated in working memory (Sarnthein et al., 1998). Below, we perform a frequency analysis of the occipital visual LFPs and describe their relationship to the activity of single neurons during the maintenance of information

*Correspondence: gregor.rainer@tuebingen.mpg.de

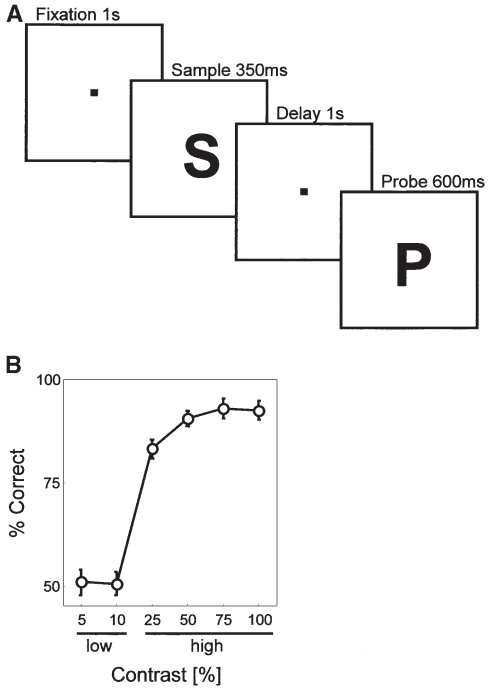

Figure 1. Behavioral Task and Performance

(A) The sequence of trial events and timing for the delayed-matching-to-sample (DMS) task. After a fixation period, a sample stimulus *S* (one of eight natural images at one of six contrast levels) was briefly presented. After a delay period, a full-contrast probe stimulus *P* was presented. In half of the trials, the probe matched the sample. Monkeys were rewarded for making a lever response on these trials and for withholding a response for nonmatching probe stimuli.

(B) Behavioral performance on the DMS task is shown as a function of sample stimulus contrast. Monkeys' performance was at chance (50% correct) on low (5% and 10%) contrast trials (t tests, $p > 0.1$) and significantly better than chance at contrasts of 25% and higher (t tests, $p < 0.001$). The error bars show the SEM.

in working memory, in order to better understand the functional significance and the mechanisms by which oscillations contribute to working memory.

Results

Behavioral Performance

Monkeys performed a delayed-matching-to-sample (DMS) task summarized in Figure 1A. They were required to view a sample stimulus and retain its identity over the course of a 1 s delay period. After the delay, a probe stimulus was presented, and monkeys had to release a lever if the probe matched the sample. Probe stimuli were always presented at full contrast, whereas sample stimuli were presented at one of six contrast levels. Behavioral performance is shown as a function of contrast in Figure 1B. For low contrasts (5% and 10%), performance was at the chance level of 50% correct (t tests, $p > 0.1$), whereas performance was significantly above chance (at 83% correct or better) for high contrasts (25%, 50%, 75%, and 100%; t tests, $p < 0.001$). This suggests that monkeys were unable to identify the sample stimulus and/or maintain it in working memory when it was presented at low (5% and 10%) contrast but had little or no difficulty for high (25% and above) contrast stimuli. Because sorting according to behavioral performance separated trials into the two categories of low and high contrast, we grouped trials into these two classes for some of our analyses of neural data.

Spectral Analyses

To examine whether there were any systematic variations in energy at different frequencies in the LFP signal, we computed trial-averaged spectrograms of the LFP. These are shown for an example recording site in Figure 2A for each contrast value separately. An increase in low-frequency power during the delay period following high-contrast sample stimuli is apparent in these spectrograms. To quantify this effect for this recording site, we computed the trial-averaged magnitude of the Fourier transform for the last 700 ms of the fixation (baseline) and delay periods. We used the fixation period as a baseline and compared it to corresponding values during the subsequent delay period as a function of contrast. The fixation period provides a good baseline condition, because the physical characteristics of the visual stimulation are identical to those during the delay period. While monkeys need to maintain a memory of the sample stimulus during the delay period, there is no working memory requirement during the fixation period. Thus, the main difference between these two periods is working memory demand. Comparing energy in the theta band between fixation and delay period, we observed an average 76% increase in theta energy at each individual contrast value of 25% or higher (t tests, $p < 0.001$), a reduction for 5% contrast (t test, $p < 0.001$), and no change at 10% contrast (t test, $p = 0.12$). Grouping the trials into low versus high contrast, we observed a significant increase at theta frequency for high contrast and a significant decrease at low contrast (t tests, $p < 0.01$). We conclude that a robust characteristic of the LFP at this recording site was an increase in signal amplitude in the theta range (4–8 Hz) during the delay period, when information about the sample stimulus needed to be maintained.

Working memory-specific theta frequency oscillations were observed at many of the recording sites. In Figure 2B, we show LFP energy as a function of frequency during the delay period relative to fixation averaged across all 37 recording sites. A peak in the theta (4–8 Hz) range is apparent in these group data. Across this population of recording sites, there was a significant 20% increase in theta frequency amplitude during the delay (t test, $p << 0.001$). This enhancement in theta amplitude was significant in each of the monkeys when analyzed separately (monkey A: t test, $p < 0.01$; monkey B: t test, $p < 0.05$). In addition, a significant majority of recording sites showed increases in theta amplitude as opposed to decreases (28/37 increases versus 9/37; χ^2 test, $p < 0.05$). In the alpha band (8–12 Hz), on the other hand, we observed—on average—decreases in amplitude for high-contrast sample stimuli at 6 out of 7 sites at which there was significant modulation (t tests, $p < 0.01$) and an overall significant 5% average reduction in amplitude (t test, $p < 0.05$). In the gamma band (30–75 Hz), we observed a significant average 11% increase in amplitude (t test, $p < 0.001$), and significant

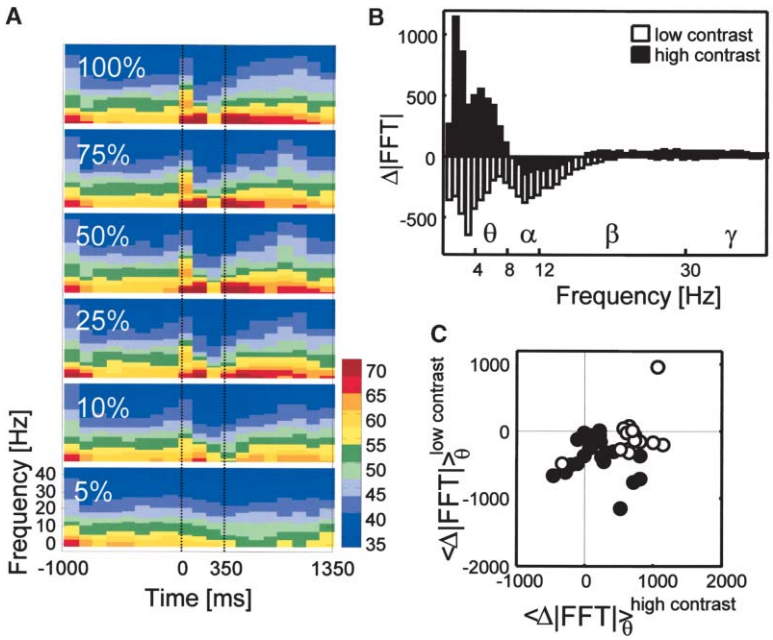

Figure 2. Frequency Analysis of Local Field Potential

(A) Trial-averaged spectrograms for frequencies from 0 to 45 Hz for an example recording site are shown for each contrast level separately averaged across all objects. The fixation period, the sample stimulus presentation from 0 to 350 ms, and the subsequent delay period are shown.

(B) The change in energy between fixation and delay periods for low- and high-contrast sample trials averaged across all recording sites (n = 37) is shown. This reveals energy increases in the delta (0–4 Hz), theta (4–8 Hz), and gamma (30–70 Hz) bands; decreases in the alpha (8–12 Hz) band; and mixed effects in the beta band (12–30 Hz) at high contrast. Low-contrast sample stimuli by contrast led to broad band decreases in energy.

(C) The theta frequency energy change from fixation to delay period for high contrast (x axis) is plotted against the corresponding low contrast value (y axis) for all recording sites. Empty symbols denote sites that (n = 14) individually showed significant theta frequency changes at high contrast.

modulations were apparent at 15 recording sites, with 12/15 sites showing increases (t tests, $p < 0.01$). The beta band appeared to be split into low beta showing working memory-dependent decreases and high beta showing increases, so this band is not further analyzed here. Elevations were also observed in these group data in the delta band (below 4 Hz), but these are also not further analyzed, because their cycle duration are of the same order of magnitude as our analysis window of 700 ms. In Figure 2C, we directly compare for each recording site how theta frequency LFP energy changed from fixation to the delay period for low- versus high-contrast sample stimuli. The empty symbols mark sites for which, individually, the increase in theta frequency energy was significant for high contrast (25% and greater) sample stimuli (t tests, $p < 0.01$). This occurred at 14 sites, and at all but one (13/14, or 93%) we observed more theta frequency energy during the delay compared to the fixation period. We found that a similar proportion (χ^2 test, $p > 0.1$) of sites showed theta enhancement in each of the monkeys (monkey A, 7 of 19 sites; monkey B, 6 out of 18 sites). The elevated theta band frequency energy represented primarily oscillatory activity that varied from trial to trial and was largely not tightly phase locked to the presentation time of the sample stimulus. Accordingly, the theta peak was attenuated in a frequency decomposition of the trial-averaged LFP. To demonstrate this, we performed a Fourier analysis of the trial-averaged LFP (i.e., FFT<LFP>) and computed how the magnitude of this quantity changed from fixation to delay period. We then subtracted this quantity from the averaged frequency energy across trials [i.e., <FFT(LFP)>] as analyzed above. The difference between these quantities [i.e., <FFT(LFP)> − FFT<LFP>] exhibited a significant increase in the theta range (t test, $p < 0.05$). This indicates that elevated frequency energy in the theta band was indeed largely non-phase locked. Thus, our findings are consistent with a working memory-dependent theta process initiated during the delay period, rather than simple upregulation of theta amplitude. Taken together, these findings suggest that holding information in working memory is accompanied by oscillatory changes in several frequency bands of the LFP in extrastriate area V4, with increases in amplitude observed mainly in the theta and gamma band and decreases in the alpha band relative to a baseline of identical sensory input but absent memory load.

Theta Oscillations and Behavior

We now focus on oscillatory activity in the theta band and proceed with a detailed analysis of the time course of energy in the theta band during the trial. A continuous estimate of theta frequency energy during the trial $E(LFP_\theta)$ as a function of contrast is shown in Figure 3A for a representative recording site. For sample stimuli at 5% contrast, theta energy was reduced throughout the delay relative to values during the fixation period. For samples at high contrasts (25% or greater), theta energy increased during the delay period and tended to peak 300–400 ms before the presentation of the probe stimulus. This pattern is similar to observations in humans using intracranial EEG recordings, where theta energy during a multiple-item memory task also decreased toward the end of the delay (Raghavachari et al., 2001). Whereas theta energy early in the delay varied systematically with sample contrast, this was not the case during the later part of the delay around the time of the peak. The trial-by-trial variability of theta frequency energy changes from fixation to delay period is shown in Figure 3B for 100% contrast sample stimuli, comparing the last 700 ms of the two periods. More data points fall above the diagonal, indicating greater theta energy on average during the delay (t test, $p < 0.01$) on the majority of trials. The opposite was true for 5% contrast sample stimuli shown in Figure 3C, where the majority of data points lie below the diagonal (t test, $p < <0.01$). On a trial-by-trial basis, the amount of theta energy was not correlated with behavioral performance on that trial,

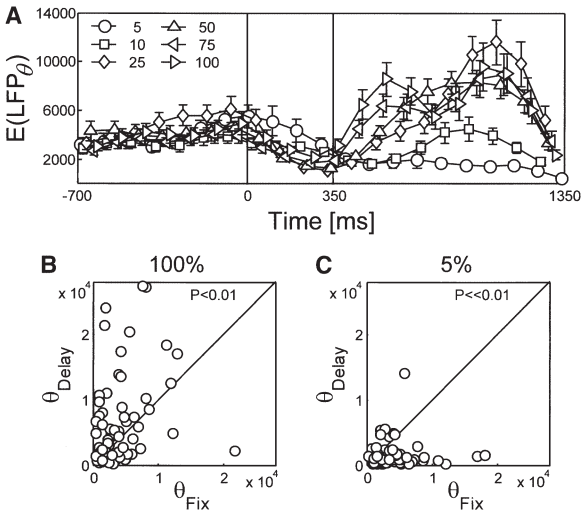

Figure 3. Time Course of Theta Energy during the Trial

(A) The time course of theta frequency energy is shown as a function of contrast at an example recording site. Relative to levels during the fixation period, this shows that theta frequency energy decreased during sample stimulus presentation (0–350 ms) and remained low throughout the delay (350–1350 ms) at 5% contrast and behaved similarly at 10% contrast. At higher contrasts, it began to rise during the delay period and peaked around 300 ms before the end of the delay. The trial-by-trial variability of (B) increased theta energy at 100% contrast and (C) decreased theta energy at 5% contrast during the delay is visualized by directly comparing the average levels of theta oscillations during the fixation versus the delay period. Each symbol represents one trial.

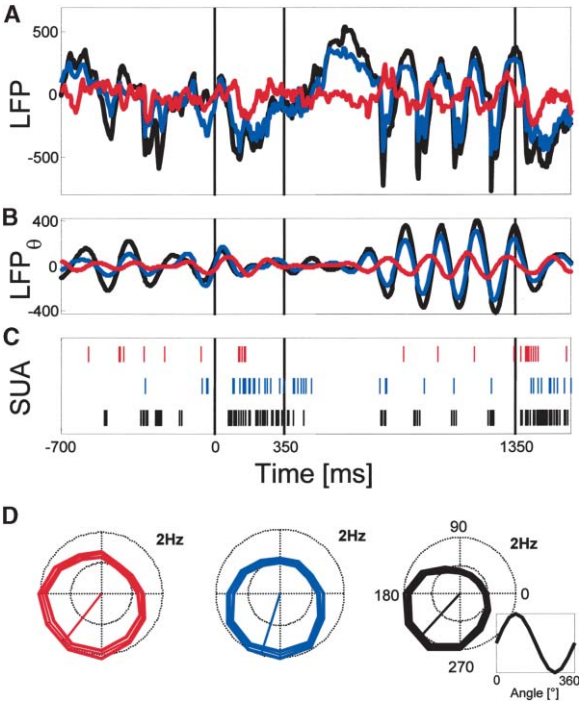

Figure 4. Single Trial Local Field Potential and Single Unit Activity at Three Sites

(A) The raw LFP signals recorded simultaneously at three sites in extrastriate area V4 are shown for fixation, sample, delay, and probe period, along with (B) the theta band-filtered versions of these signals (LFP$_\theta$) and (C) action potentials from single units (SUA) isolated on each of these recording channels. The relation between action potential timing for each of the neurons to the angle of their local LFP$_\theta$ signal during the delay period averaged across all high-contrast trials is shown in (D). Each of these neurons tended to emit action potentials preferentially in a particular range of theta angles. The polar plots represent how many action potentials were emitted at each phase angle of the theta cycle. The bold number on the upper right of each plot marks how many action potentials are represented by the outer circle gridline. The colored radial line represents the preferred theta angle for each neuron. The legend shows an example theta cycle for reference.

neither in terms of reaction time nor correctness. Accordingly, there was no significant difference between the distributions of theta energy for correct versus incorrect trials at any of the contrast levels (t tests, $p > 0.1$), and reaction times for correctly executed match trials were uncorrelated with the amount of theta energy preceding the probe presentation on that trial (t test, $p > 0.1$). Note that there were very few correctly executed match trials on low-contrast trials (three trials at low contrast in this session), because monkeys almost never responded to the probe stimulus after low-contrast samples and thus made mostly false-negative errors (i.e., omissions of required responses).

Theta Oscillations and Single Neuron Activity

Having established that theta frequency oscillations are common and robust characteristics of the LFP during working memory, we asked whether they had any systematic influence on the activity of single neurons recorded at those sites. Our results are based on a total of 72 recorded neurons (monkey A, 44; monkey B, 28). We examined the relationship between the LFP and action potentials emitted by single units isolated on the same channels. An example of three raw LFP signals acquired simultaneously for a high-contrast sample is shown in Figure 4A. Four cycles of oscillation in the theta range are apparent in each of these signals. The blue and black traces were recorded at a horizontal distance across the cortical surface of 1 mm and at similar depth, and the signals are accordingly highly correlated for this trial (Pearson $r = 0.94$ for the entire raw LFP signal) as well as across all trials in the recording session (t test, $p << 0.001$; $\langle r \rangle = 0.89$). Such high correlation of the raw LFP implies a similarly high correlation of theta phase for nearby sites. The red trace also exhibits several cycles of theta frequency oscillation, which are phase shifted relative to the black/blue channel by about 90°. This trace was recorded from an electrode at a distance of about 3 mm away, and average Pearson correlation between the signals was still significant but dropped sharply to $\langle r \rangle = 0.25$ and $\langle r \rangle = 0.35$ (for blue versus red and black versus red channels, respectively, t tests, $p << 0.001$). In Figure 4B, the theta band filtered signal (LFP$_\theta$) is shown for the three channels, demonstrating that the oscillatory activity during the delay is reliably extracted.

To study the relationship between theta oscillations and SUA shown in Figure 4C, we estimated the angle of the theta band oscillation LFP$_\theta$ during the delay period and generated histograms in polar format of the phase angles at which action potentials from isolated single

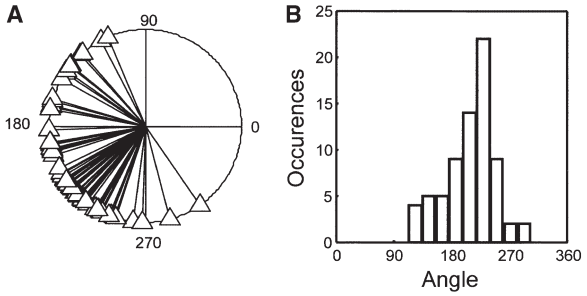

Figure 5. Distribution of Single Unit Activity Alignment to Theta Oscillations

The distribution of preferred theta angles for each of the 72 neurons in the population is shown in (A) polar and (B) histogram format, revealing a strong dependence of SUA on the phase of theta frequency oscillations.

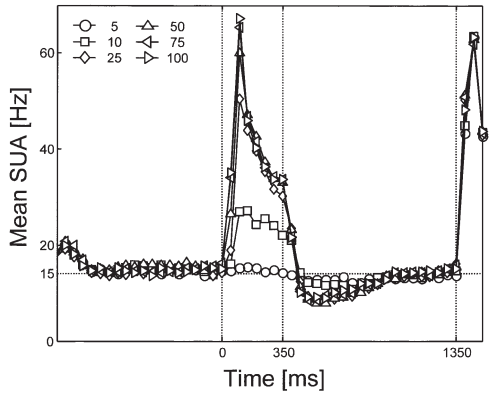

Figure 6. Mean Population Single Unit Activity during the Trial

The peristimulus time histogram averaged across all sample stimuli and all 72 neurons in the population as a function of contrast, revealing robust visual responses during the sample and probe periods at contrast levels higher than 10%. There was little difference in terms of mean firing rate comparing the fixation to the delay period. The average fixation period activity level is marked by a horizontal line at 15 Hz.

neurons occurred relative to the theta-filtered LFP at the same cortical site. Such polar histograms, based on all high-contrast trials, are shown for the three single neuron examples in Figure 4D. It is clear that action potentials were not evenly distributed as a function of theta angle, but tended to occur between angles of 90° to 300°. Thus, action potentials tended to occur at times when the LFP was decreasing from the peak of each cycle at 90° to the cycle trough at 270°. We confirmed that this phase-locking relationship between LFP_θ angle and action potential timing was a highly significant and robust effect. Each unit on average emitted about twice as many action potentials near its preferred theta angle compared to its antipreferred angle, and statistical tests confirmed that data were nonuniformly distributed (Rayleigh tests, $p < 0.01$). Indeed, action potential timing varied systematically with LFP_θ angle for a great majority of the recorded neurons (65/72, or 90%, Rayleigh test, $p < 0.01$). Considering each monkey separately, we confirmed that a significant majority of neurons (χ^2 test, $p << 0.01$) showed this systematic variation of action potential timing with respect to LFP_θ angle (monkey A, 43 of 44 neurons, or 98%; monkey B, 22 of 28 neurons, or 79%). Note that variation in LFP_θ phase among sites can lead to neurons emitting action potentials at different absolute times, even though they show significant alignment to their local LFP_θ. Accordingly, activity of the red neuron shows similar theta phase dependence as the blue or black neuron in Figure 4D but tends to emit action potentials that are unsynchronized with those of the blue or black neuron during the delay, as shown in Figure 4C. We estimated the preferred angle for each of the 72 neurons in the population with respect to the LFP_θ recorded on the same electrode and show the results in polar and histogram format in Figures 5A and 5B, revealing that, although there was some variability among neurons in terms of their preferred angle, with values ranging from around 100° to 290°, the distribution of preferred angles was unimodal exhibiting a peak near a theta angle of 220°. Note that the increased theta power in the LFP during the delay period and SUA phase locking appear to be closely related phenomena.

It is interesting to note that phase locking of action potentials to the LFP_θ signal occurred largely in the absence of modulations in mean firing rate. The mean peristimulus time histogram averaged across all recorded neurons and all sample stimuli as a function of sample contrast shown in Figure 6 demonstrates that there was little overall change in mean firing rate when comparing fixation and delay periods. In fact, for high-contrast trials, activity during the last 700 ms of the fixation period, the average activity was 15.0 Hz and did not significantly differ from the average activity of 14.5 Hz during the last 700 ms of the subsequent delay (paired t test, $p = 0.08$). Of the 25 neurons that did show systematic changes in mean firing rate (t tests, $p < 0.01$), about half showed increases from fixation to delay (15 neurons), with the rest showing decreases (10 neurons). So, unlike neurons in prefrontal or inferior temporal regions (Fuster and Alexander, 1971; Fuster and Jervey, 1981), V4 neurons do not exhibit robust elevated discharge of action potentials during working memory delays. Consistent with this observation, selectivity for the information that is being held in mind (i.e., the sample stimulus) was observed in relatively few V4 neurons. We asked for how many neurons activity during the last 700 ms of the delay depended upon which stimulus had been presented during the sample period. Considering high-contrast trials, we found selective delay period activity in 17/72 (24%) neurons (one-way ANOVA, $p_{SAMPLE} < 0.01$). This number is slightly larger than the figure of 11% previously reported (Chelazzi et al., 2001) at the $p < 0.05$ significance level. However, our figure is based on about 40 repetitions per stimulus condition. Basing the ANOVA on similar trial numbers by restricting the analysis to 100% contrast trials only (about ten repetitions), we in fact obtain a statistically identical estimate (significant delay activity in 5/72 neurons [6%]; χ^2 test, $p = 0.34$) to that reported by Chelazzi and colleagues (11/81 neurons, or 11%).

The fact that SUA tended to vary with theta angle raises the possibility that stimulus-selective signals during the delay period might also vary with respect to theta angle. For a given single neuron, selective activity might

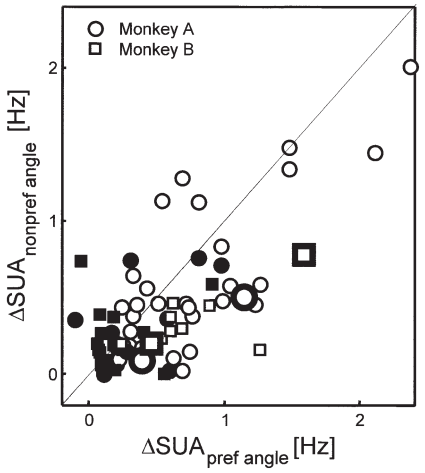

Figure 7. Comparison of Sample Selectivity of Single Unit Activity as a Function of Theta Angle

This scatter plot shows the trial-averaged difference in neural activity for the preferred versus the nonpreferred sample stimulus (ΔSUA) of each neuron, when considering only action potentials near the preferred (x axis) and nonpreferred (y axis) theta angle. Each symbol represents a single neuron; the empty symbols represent neurons that were significantly modulated by both sample stimulus identity and theta angle. The four large symbols represent the example neurons shown in Figure 8. Circle and square markers represent neurons recorded in monkey A or B, respectively.

preferentially occur near its preferred angle of the theta cycle. To evaluate if this occurs, we performed a two-way ANOVA on the activity of each neuron during the last 700 ms of the delay period. One factor was the identity of the sample stimulus (SAMPLE) being held in working memory—as in the one-way ANOVA above—and the other factor was theta angle (ANGLE). This analysis replicated the finding that the majority of neurons showed delay activity that varied systematically with theta angle, as demonstrated by the fact that 54/72, or 75%, of neurons showed a main effect of theta angle (two-way ANOVA, $p_{ANGLE} < 0.01$), with only 1/72 neurons showing a significant interaction ($p_{INTER} < 0.01$). Over two-thirds of the neurons (51/72, or 71%) showed a main effect of sample stimulus (two-way ANOVA, $p_{SAMPLE} < 0.01$), and the activity of the majority of these neurons (42/72, or 58%) varied significantly as a function of both sample stimulus and theta angle (two-way ANOVA, $p_{SAMPLE} < 0.01$ and $p_{ANGLE} < 0.01$). Thus, taking into account the angle of the LFP_θ significantly increases (χ^2 test, $p < 0.001$)—and in fact more than doubles—our estimate of how many V4 neurons contribute to working memory from 17/72 (24%) to 42/72 (58%).

The above statistical analysis revealed that over half of the recorded V4 neurons were significantly modulated by both theta angle and sample stimulus. We examined the relation between these two influences on neural activity by assessing the difference in activity for its preferred and nonpreferred sample stimulus in two angular windows of the LFP_θ signal. The first window was chosen to be $\pm 60^\circ$ around each neuron's preferred theta angle, and the second window shifted by 180° relative to the first window. The results are shown in Figure 7 for all 72 neurons, where each symbol represents a single neuron, and the difference in activity to its preferred compared to its nonpreferred sample stimulus during its preferred theta angle range (x axis) is plotted against the corresponding value during its nonpreferred theta angle range (y axis). The values are always positive, since they represent increases in activity for the preferred over the nonpreferred sample stimulus. The data points tend to lie below the diagonal (paired t test, $p < 0.001$), indicating that stimulus-selective signals—that is, activity increases for the preferred over the nonpreferred sample stimulus—were greater near the preferred theta angle than near the nonpreferred theta angle. We confirmed that this was also the case when the data of each monkey were analyzed alone (monkey A: t test, $p < 0.01$; monkey B: t test $p < 0.05$). In additional analyses, we confirmed that sample-selective signals in the range of $\pm 60^\circ$ near the preferred theta angle were also greater than sample-selective signals measured for the entire 700 ms delay without considering LFP_θ (paired t test, $p < 0.001$). The 42 neurons jointly modulated by theta angle and sample stimulus (two-way ANOVA, $p_{SAMPLE} < 0.01$ and $p_{ANGLE} < 0.01$) are shown as empty symbols in Figure 7. For these neurons, sample selectivity was also greater near the preferred theta angle (paired t test, $p < 0.001$), and sample selectivity was greater near the preferred theta angle than the nonpreferred angle for a greater fraction of these neurons (33/42, or 79%, compared to 9/42, or 21%; χ^2 test, $p < 0.01$). The four large symbols in Figure 7 represent four single neurons for which the dependence of neural activity on theta angle is shown in Figures 8A–8D. The left panels in each row show the theta angle dependence of the SUA averaged across all high-contrast sample stimuli, and the radial line represents the resultant vector or preferred theta angle for each neuron. All four of these neurons showed significant theta angle dependence (Rayleigh test, $p < 0.01$), with mean vectors between 180° and 270°. The right panels show the theta angle dependence of SUA for the preferred and nonpreferred sample stimulus and demonstrate that robust sample selectivity tended to occur near the preferred theta angle. Thus, more action potentials are emitted by V4 neurons during working memory delays around their preferred angle of the theta oscillation, with the consequence of an enhancement of object-selective neural signals near this preferred theta angle. Taking into account the angle of the theta oscillation leads to an enhancement in stimulus-selective activity during the delay period, and stimulus-selective signals could thus be extracted or decoded more accurately by a downstream brain region if it takes into account only action potentials emitted near the preferred theta angle.

Discussion

We have described oscillatory activity in the local field potential (LFP) during working memory in occipital visual cortex of the monkey. Holding a stimulus "in mind" over the course of a brief delay was associated with elevations at theta (4–8 Hz) and gamma (30–75 Hz) frequencies and reduced activity at alpha (8–12 Hz) frequencies. We observed a relationship between theta oscillations and the timing of action potential discharge of single

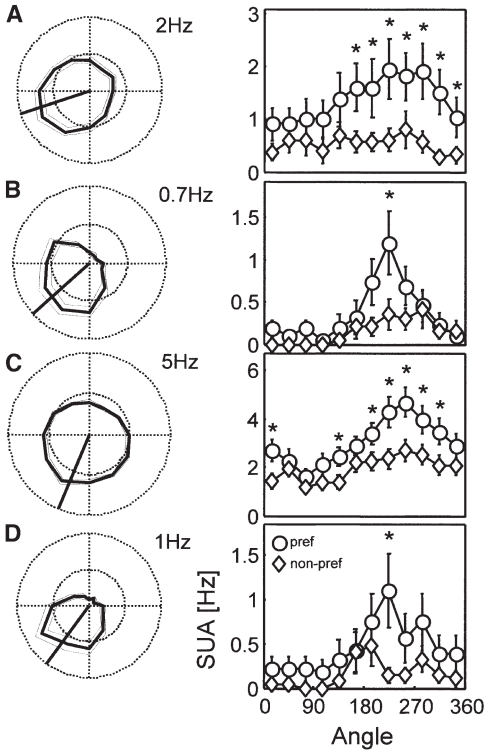

Figure 8. Single Neuron Examples Showing Sample Selectivity as a Function of Theta Angle

For four example neurons (A–D) corresponding to the large symbols in Figure 7, theta angle dependence of neural activity and sample stimulus selectivity are shown. The modulation of activity averaged across sample stimuli as a function of phase angle is shown on the left panel in each row. Dark lines represent means, thin lines represent standard errors for all high-contrast trials; radial thick lines represent the preferred theta angle of each neuron. The number on the top right of each panel represents the firing rate designated by the outer circle gridline. The right panel in each row shows the activity for the preferred and nonpreferred sample stimulus for each neuron, along with the SEM. Asterisks denote statistical significance (t tests, $p < 0.05$). The upper and lower two panels represent data collected from monkeys A or B, respectively.

neurons during the delay period, such that single unit activity varied systematically with the angle of the LFP theta oscillation. The observation that SUA was phase locked to the theta oscillation allowed the definition of a preferred theta angle for each neuron. SUA reflecting the identity of the remembered stimulus during the delay was greatest near this preferred theta angle, and for many neurons, stimulus-selective signals occurred only near the preferred theta angle. The alignment of SUA to theta oscillations and the associated effects on stimulus selectivity occurred largely in the absence of overall changes in mean firing rate of neurons during the delay period. These findings suggest that theta oscillations observed in the LFP, which reflects aggregate activity of the local population of neurons, are accompanied by working memory-related signals in the activity of single neurons in the local population. Thus, theta oscillations during working memory can be interpreted as reflections of stimulus-selective signals occurring at the level of single neurons.

Our findings of increases in theta and gamma frequency oscillations during working memory are largely in agreement with results obtained in humans. For example, a relationship between theta frequency power and working memory has been documented with high-resolution EEG (Gevins et al., 1997), and coherence analyses between scalp EEG signals have revealed increases in synchronization between prefrontal and occipital sites in the theta frequency range (Sarnthein et al., 1998; Stam et al., 2002) during working memory. This finding is of particular relevance to the results reported here because it provides evidence for cooperation between prefrontal and occipital sites (Klimesch, 1996) in the maintenance of information in short-term memory and implicates specifically theta frequency oscillations in occipital cortex in this maintenance. Recent work using intracranial EEG (Kahana et al., 2001) and independent component analysis applied to scalp EEG (Makeig, 2004) has provided further evidence implicating theta frequency oscillations in cognitive processing. In particular, one study reported theta frequency oscillations at occipital cortical sites while subjects were remembering lists of one to four items over the course of a delay period (Raghavachari et al., 2001)—findings which share many of the characteristics seen in our data. For example, we note that the time course of the theta elevation observed in that study was similar to our results in the monkey, peaking in the mid to late delay period and decreasing in anticipation of the probe stimulus toward the end of the delay. In our study, the delay was always followed by a high-contrast probe stimulus. Since we observed no elevations—in fact, reductions relative to the fixation period levels—of theta frequency energy for low-contrast sample stimuli, it follows that increased theta frequency oscillations are unlikely to be due to anticipation of the probe stimulus. We suggest, in agreement with Raghavachari and colleagues, that theta frequency oscillations reflect a form of reverberation of activity related to the memory or maintenance of the sample stimulus. These similarities suggest that theta frequency modulations recorded in monkeys and humans may be functionally equivalent.

Similar to the effects at theta, we observed increases in gamma frequency (30–75 Hz) energy during working memory, replicating findings in humans. Using a delayed-matching-to-sample paradigm, elevations at gamma frequencies have been observed at both frontal and occipital sites in the scalp EEG (Lutzenberger et al., 2002; Tallon-Baudry et al., 1998, 1999), MEG (Kaiser et al., 2003), and iEEG (Howard et al., 2003). Published findings in the alpha frequency (8–12 Hz) band appear to be somewhat more variable, with reports of both increases (Jensen et al., 2002) and decreases (Gevins et al., 1997). In our study, we observed decreases at alpha frequencies relative to baseline levels acquired while monkeys were waiting for the sample stimulus to be presented. Alpha frequency oscillations have been reported to be correlated to attention and mental effort (Gundel and Wilson, 1992; Worden et al., 2000), so it appears at least possible that effects observed in the alpha band reflect task contingencies other than the working memory demands during the delay period.

Studies of oscillatory phenomena in monkeys have mostly focused on periods of sensory stimulation in

occipital visual areas and described effects predominately in the gamma band, for example, as a function of attention (Fries et al., 2001), figure-ground segmentation (Gail et al., 2000), visual stimulus type (Rols et al., 2001), or feature integration (Singer and Gray, 1995). There has been far less work investigating oscillations during memory periods in monkeys. In one recent study, LFP recordings were made from monkey V4 using an implanted recording grid (Tallon-Baudry et al., 2004), documenting a correlation between behavioral performance and coherence in the beta frequency band. Another study has reported direction-selective increases in gamma frequency oscillations during a delayed saccade task in parietal cortex (Pesaran et al., 2002). Generally consistent with this, we observed elevations at gamma frequencies at some sites. No changes in the theta band were reported in that study. This might reflect differences between occipital sensory and parietal cortex, such as the low firing rates we observed during the delay period in V4, or might be related to motor planning of eye movements in the oculomotor task. Indeed, previous studies have documented LFP oscillatory activity in the 20–40 Hz range and SUA synchronization in motor cortical areas (Murthy and Fetz, 1996; Riehle et al., 1997).

At the level of single neurons, working memory has been associated with delay activity—the elevated discharge of action potentials that is commonly observed in the prefrontal cortex (Funahashi and Takeda, 2002; Goldman-Rakic, 1996; Miller and Cohen, 2001; Rainer and Ranganath, 2002). One interesting observation regarding delay activity is that single neurons rarely show stationary elevated discharge that is initiated by the sensory stimulus and lasts throughout the delay. On the contrary, more detailed analysis of the delay activity time course indicates that delay activity is subject to dynamics. One robust finding is that delay activity tends to increase toward the end of the delay, reflecting anticipation of the expected stimulus, action, or reward (Asaad et al., 1998; Rainer et al., 1999; Watanabe, 1996), whereas activity early during the delay tends to reflect retrospective reverberation related to the previously seen sample stimulus (Rainer et al., 1999). Together with the results of the present study, this suggests that information held in working memory may initially be stored in reverberating circuits that include prefrontal and occipital cortices and in which theta frequency oscillations play an important role. Computational modeling work suggests that such reverberating activity might indeed be important for maintenance of elevated neural activity during working memory (Seung et al., 2000; Wang, 2001). The information stored in reverberatory networks may then be converted into a prospective code in a task-dependent fashion as the delay ends, in concert with prospective coding in the prefrontal cortex and consistent with the observed decrease in theta frequency oscillations ahead of probe presentation observed in humans (Raghavachari et al., 2001) as well as in our study.

Our demonstration of theta-locked SUA in occipital visual cortex raises the question of whether a similar relationship might hold in the prefrontal cortex. A recent study addressing this issue by examining SUA in the prefrontal cortex found little evidence for oscillatory structure in the timing of action potentials (Compte et al., 2003), suggesting that oscillations may not play a prominent role in prefrontal delay activity. However, for a number of reasons, including the stochastic nature of SUA, low firing rates, and the variability in terms of individual cycle duration of oscillations in a given frequency band, oscillations are not easily detectable from SUA alone (Csicsvari et al., 1999). It thus remains a possibility that oscillations at theta and/or other frequencies might play a role during working memory also in the prefrontal cortex.

A systematic relation between theta oscillations and SUA is a well-known characteristic of neurons in the rat hippocampus (Buzsaki, 2002; O'Keefe and Recce, 1993), where pyramidal neurons that are sensitive to the animal's location in the environment also discharge action potentials systematically near a particular phase angle of the theta oscillation. Our study demonstrates a similar relationship in monkey occipital cortex. It has been suggested that multiple item memory might be supported by information packets aligned to different angles of the theta rhythm (Lisman and Idiart, 1995). We have demonstrated that working memory for a single stimulus is indeed associated with SUA aligned to the angle of the theta oscillation, a key prediction of that model. Whether multiple item memory is indeed stored at distinct theta angles, however, is an open but experimentally addressable question.

Although it remains to be shown that downstream brain regions take advantage of the variation in fidelity of neural signals as a function of the theta cycle, our findings demonstrate that behaviorally relevant signals contained in action potential discharges of single neurons are considerably enhanced when large-scale activity in the local population is taken into account. This provides evidence for the idea that the timing of action potentials—in our case relative to the theta cycle—may play an important role in the encoding of information (Abeles, 1982; Hopfield, 1995; Singer and Gray, 1995) and structuring information flow among different brain regions (Diesmann et al., 1999; Salinas and Sejnowski, 2001).

Experimental Procedures

Subjects

Two adult male rhesus monkeys (*Macaca mulatta*) weighing 12 and 13 kg were used in this study. Anatomical modified driven equilibrium Fourier transform (MDEFT) pulse sequence slices were obtained from each monkey (0.5 × 0.5 mm in-plane resolution, 0.5 mm slice thickness, 256 × 256 matrix with FOV 12.8 × 12.8 mm, 160 slices, TR = 20 ms, TE = 4 ms, FA = 20°) on a 4.7 T vertical MRI scanner with 40 cm bore and 50 mT/m gradient coil (BIOSPEC 47/40v, Bruker, Rheinstetten, D). Three-dimensional models of the skull and brain were extracted from these scans using ANALYZE (Mayo Foundation, Rochester, MN) software. These were used to locate the desired position of the recording chamber and headpost. Titanium implants were machined to fit the individual monkey heads using a CNC machine (Willemin-Macodel W428, Bassecourt, CH). The lower parts of the implants were then coated (Medipure 22-00 and 22-15, Medicoat, Mägenwil, CH) to ensure tissue compatibility. The headpost, recording chamber, and a scleral search coil (Robinson, 1963) were surgically implanted under aseptic conditions without the use of dental acrylic. After injections of Robinul (0.01 mg/kg) and Ketavet (15 mg/kg), we prepared monkeys for intubation by administering the analgesic Fentanyl (0.003 mg/kg), the barbiturate anesthetic Trapanal (5 mg/kg), and the paralytic Lysthenon (3 mg/kg) intravenously. For the duration of the surgery, monkeys were maintained in anesthesia using Isoflurane. Postoperatively, monkeys re-

ceived antibiotics and analgesics. All procedures were approved by local authorities (Regierungspraesidium Baden-Wuerttemberg) and conducted in accordance with applicable guidelines of the European Commission (EUVD 86/609/EEC) and the National Institutes of Health for the care and use of laboratory animals.

Recording Technique

Neural activity was recorded from two to eight tungsten microelectrodes (UEWLGDSMNN1E, FHC INC, Bowdoinham ME) with impedance of 1–2 MΩ at 1 kHz, which were moved up or down in pairs using custom-made mini-microdrives mounted on a plastic grid (Crist Instruments, Hagerstown MD) allowing a minimum interelectrode separation of 1 mm. The signal from each electrode was initially preamplified (×20, Thomas Recording, Giessen, D) using the recording chamber as reference and then split into two signals which were filtered and amplified separately (BAK electronics, Germantown, MD) to yield local field potential and single unit activity. The LFP was obtained by high-pass filtering the raw signal above 0.1 Hz and digitizing at 4464 Hz. The analyses in this study were performed on the LFP in units of the analog/digital (A/D) converter. Note that a 1 A/D sample corresponds to 4.9 μV. For the analyses presented here, we used a digitally low-pass filtered (4-pole Butterworth, cutoff 100 Hz) version of the LFP downsampled to a sampling rate of 200 Hz. SUA was obtained by first band-pass filtering the raw signal from 300 Hz to 4 kHz and digitizing the signal at 22.231 kHz and sorting action potentials into clusters corresponding to signals from single isolated neurons (Datawave, Longmont, CO). During each recording session, we advanced each pair of electrodes until the activity of one or more neurons was well isolated and began data collection after a wait period of about 1 hr.

Behavioral Task

Monkeys performed a modified version of delayed-matching-to-sample. Each trial, shown schematically in Figure 1A, began when the monkey grasped a lever. After fixating on a fixation point for a period of 1 s, a sample stimulus (S) was presented for 350 ms. The sample was $10^\circ \times 10^\circ$ in size, presented at the center of gaze, and could be one of eight different stimuli at one of six different contrasts (100%, 75%, 50%, 25%, 10%, and 5%). Contrast was modulated by subtracting—for each gun independently—the mean value over the entire image from each pixel, scaling down the pixel value to the appropriate percentage, i.e., multiplication by 0.05 for the 5% contrast condition, and then adding back the mean that was originally subtracted. We used colored natural images with normalized Fourier amplitude spectra (Rainer et al., 2001, 2004) with mean intensity of 24 cd/m^2 presented on an isoluminant gray background as stimuli. After a brief 1 s delay, a probe stimulus (P) was presented for 600 ms. Probe stimuli were always presented at full (100%) contrast. On 50% of the trials, the probe matched the sample and monkeys had to release the lever to obtain a juice reward. On the remaining trials, monkeys had to wait for an additional short (200 ms) delay, after which they were required to release the lever to the presentation of the correct matching stimulus. This was done to ensure that monkeys were actively engaged in every trial. We performed 37 electrode penetrations (monkey A, 19; monkey B, 18) in extrastriate area V4 collected during seven recording sessions (monkey A, 4 sessions; monkey B, 3 sessions). We thus acquired LFP data from a total of 37 recording sites and were able to isolate 72 single neurons from these channels (monkey A, 44; monkey B, 28). Monkeys completed on average 640 trials or about 13 repetitions for each stimulus at each contrast level during each session. Stimuli were presented on a graphics workstation (TDZ 2000, Intergraph Systems, Huntsville AL) at a resolution of 1280×1024 on a gamma-corrected 21 inch monitor placed 97 cm away from the animals, with refresh rate of 75 Hz running an OpenGL-based stimulation program under Windows NT. Behavior control was maintained by a network of PCs running the QNX real-time operating system (QSSL, Ontario, CA).

Data Analysis

Data Analysis was performed in MATLAB (MathWorks, Natick, MA). Trial-averaged spectrograms were computed for the LFP from the beginning of the fixation period to the end of the delay period averaged across all eight objects for each contrast separately, using an FFT with length 256, windows length 250 ms, and overlap of 125 ms.

For the time course analyses, the raw LFP was filtered in the theta range (4–8 Hz) using a zero-phase forward and reverse digital 4-pole Butterworth band-pass filter to yield the signal LFP_θ. A continuous estimate of the energy was then computed as $E(LFP_\theta) = real(H)^2 + image(H)^2$, where H is the Hilbert transform of the LFP_θ signal, with the real part of H being LFP_θ itself and the imaginary part of H a 90° phase-shifted version of LFP_θ. The Hilbert transform provides a convenient continuous estimate of signal energy without the need for low-pass filtering that is typically performed if only $(LFP_\theta)^2$ is used as an energy estimate.

We used a Rayleigh test (Fisher, 1993) to ask whether action potential timing varied systematically with the angle of the theta band-filtered LFP. We first computed the mean resultant length R as

$$R = \sqrt{C^2 + S^2}/n,$$

where

$$C = \sum_i \cos\theta_i$$

and

$$S = \sum_i \sin\theta_i$$

where n is the number of trials and θ_i is the angle at which a given action potential occurred. The significance probability for rejecting uniformity was $P = e^{-Z}[1 + (2Z - Z^2)/(4n) - (24Z - 132Z^2 + 76Z^3 - 9Z^4)/(288n^2)]$, with $Z = nR^2$. We computed the preferred angle for each neuron as the direction of the resultant vector:

$$\hat{\theta} = \begin{cases} \tan^{-1}(S/C) & S > 0,\ C > 0 \\ \tan^{-1}(S/C) + \pi & C < 0 \\ \tan^{-1}(S/C) + 2\pi & S < 0,\ C > 0 \end{cases}$$

Acknowledgments

This work was supported by the Max Planck Society. G.R. is a DFG Heisenberg Investigator (RA1025/1-1).

Received: July 8, 2004
Revised: September 28, 2004
Accepted: November 22, 2004
Published: January 5, 2005

References

Abeles, M. (1982). Local Cortical Circuits: An Electrophysiological Study (Berlin: Springer).

Asaad, W.F., Rainer, G., and Miller, E.K. (1998). Neural activity in the primate prefrontal cortex during associative learning. Neuron *21*, 1399–1407.

Asaad, W.F., Rainer, G., and Miller, E.K. (2000). Task-specific neural activity in the primate prefrontal cortex. J. Neurophysiol. *84*, 451–459.

Buzsaki, G. (2002). Theta oscillations in the hippocampus. Neuron *33*, 325–340.

Chelazzi, L., Miller, E.K., Duncan, J., and Desimone, R. (2001). Responses of neurons in macaque area V4 during memory-guided visual search. Cereb. Cortex *11*, 761–772.

Compte, A., Constantinidis, C., Tegner, J., Raghavachari, S., Chafee, M.V., Goldman-Rakic, P.S., and Wang, X.J. (2003). Temporally irregular mnemonic persistent activity in prefrontal neurons of monkeys during a delayed response task. J. Neurophysiol. *90*, 3441–3454.

Csicsvari, J., Hirase, H., Czurko, A., Mamiya, A., and Buzsaki, G. (1999). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving Rat. J. Neurosci. *19*, 274–287.

Diesmann, M., Gewaltig, M.O., and Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. Nature *402*, 529–533.

Fisher, N.I. (1993). Statistical Analysis of Circular Data (Cambridge: Cambridge University Press).

Fries, P., Reynolds, J.H., Rorie, A.E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. Science *291*, 1560–1563.

Funahashi, S., and Takeda, K. (2002). Information processes in the primate prefrontal cortex in relation to working memory processes. Rev. Neurosci. *13*, 313–345.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1993). Dorsolateral prefrontal lesions and oculomotor delayed-response performance: evidence for mnemonic scotomas. J. Neurosci. *13*, 1479–1497.

Fuster, J.M., and Alexander, G.E. (1971). Neuron activity related to short-term memory. Science *173*, 652–654.

Fuster, J.M., and Jervey, J.P. (1981). Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli. Science *212*, 952–955.

Fuster, J.M., Bodner, M., and Kroger, J.K. (2000). Cross-modal and cross-temporal association in neurons of frontal cortex. Nature *405*, 347–351.

Gail, A., Brinksmeyer, H.J., and Eckhorn, R. (2000). Contour decouples gamma activity across texture representation in monkey striate cortex. Cereb. Cortex *10*, 840–850.

Gevins, A., Smith, M.E., McEvoy, L., and Yu, D. (1997). High-resolution EEG mapping of cortical activation related to working memory: effects of task difficulty, type of processing, and practice. Cereb. Cortex *7*, 374–385.

Goldman-Rakic, P.S. (1996). Regional and cellular fractionation of working memory. Proc. Natl. Acad. Sci. USA *93*, 13473–13480.

Gundel, A., and Wilson, G.F. (1992). Topographical changes in the ongoing EEG related to the difficulty of mental tasks. Brain Topogr. *5*, 17–25.

Hopfield, J.J. (1995). Pattern recognition computation using action potential timing for stimulus representation. Nature *376*, 33–36.

Howard, M.W., Rizzuto, D.S., Caplan, J.B., Madsen, J.R., Lisman, J., Aschenbrenner-Scheibe, R., Schulze-Bonhage, A., and Kahana, M.J. (2003). Gamma oscillations correlate with working memory load in humans. Cereb. Cortex *13*, 1369–1374.

Jensen, O., Gelfand, J., Kounios, J., and Lisman, J.E. (2002). Oscillations in the alpha band (9–12 Hz) increase with memory load during retention in a short-term memory task. Cereb. Cortex *12*, 877–882.

Kahana, M.J., Seelig, D., and Madsen, J.R. (2001). Theta returns. Curr. Opin. Neurobiol. *11*, 739–744.

Kaiser, J., Ripper, B., Birbaumer, N., and Lutzenberger, W. (2003). Dynamics of gamma-band activity in human magnetoencephalogram during auditory pattern working memory. Neuroimage *20*, 816–827.

Klimesch, W. (1996). Memory processes, brain oscillations and EEG synchronization. Int. J. Psychophysiol. *24*, 61–100.

Kubota, K., and Niki, H. (1971). Prefrontal cortical unit activity and delayed alternation performance in monkeys. J. Neurophysiol. *34*, 337–347.

Lisman, J.E., and Idiart, M.A. (1995). Storage of 7 +/− 2 short-term memories in oscillatory subcycles. Science *267*, 1512–1515.

Lutzenberger, W., Ripper, B., Busse, L., Birbaumer, N., and Kaiser, J. (2002). Dynamics of gamma-band activity during an audiospatial working memory task in humans. J. Neurosci. *22*, 5630–5638.

Makeig, S., Delorme, A., Westerfield, M., Jung, T.-P., Townsend, J., Courchesne, E., and Sejnowski, T.J. (2004). Electroencephalographic brain dynamics following manually responded visual targets. PLoS Biol *2*(6), e176 DOI: 10.1371/journal.pbio.0020176.

Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. *24*, 167–202.

Murthy, V.N., and Fetz, E.E. (1996). Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. J. Neurophysiol. *76*, 3968–3982.

O'Keefe, J., and Recce, M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus *3*, 317–330.

Pesaran, B., Pezaris, J.S., Sahani, M., Mitra, P.P., and Andersen, R.A. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. Nat. Neurosci. *5*, 805–811.

Raghavachari, S., Kahana, M.J., Rizzuto, D.S., Caplan, J.B., Kirschen, M.P., Bourgeois, B., Madsen, J.R., and Lisman, J.E. (2001). Gating of human theta oscillations by a working memory task. J. Neurosci. *21*, 3175–3183.

Rainer, G., and Ranganath, C. (2002). Coding of objects in the prefrontal cortex in monkeys and humans. Neuroscientist *8*, 6–11.

Rainer, G., Asaad, W.F., and Miller, E.K. (1998). Selective representation of relevant information by neurons in the primate prefrontal cortex. Nature *393*, 577–579.

Rainer, G., Rao, S.C., and Miller, E.K. (1999). Prospective coding for objects in primate prefrontal cortex. J. Neurosci. *19*, 5493–5505.

Rainer, G., Augath, M., Trinath, T., and Logothetis, N.K. (2001). Nonmonotonic noise tuning of BOLD fMRI signal to natural images in the visual cortex of the anesthetized monkey. Curr. Biol. *11*, 846–854.

Rainer, G., Lee, H., and Logothetis, N.K. (2004). The effect of learning on the function of monkey extrastriate visual cortex. PLoS Biol *2*(2), E44 DOI: 10.1371/journal.pbio.0020044.

Riehle, A., Grun, S., Diesmann, M., and Aertsen, A. (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. Science *278*, 1950–1953.

Robinson, D.A. (1963). A method for measuring eye movements using a scleral search coil in a magnetic field. IEEE Trans. Biomed. Eng. *10*, 137–145.

Rols, G., Tallon-Baudry, C., Girard, P., Bertrand, O., and Bullier, J. (2001). Cortical mapping of gamma oscillations in areas V1 and V4 of the macaque monkey. Vis. Neurosci. *18*, 527–540.

Romo, R., Brody, C.D., Hernandez, A., and Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. Nature *399*, 470–473.

Salinas, E., and Sejnowski, T.J. (2001). Correlated neuronal activity and the flow of neural information. Nat. Rev. Neurosci. *2*, 539–550.

Sarnthein, J., Petsche, H., Rappelsberger, P., Shaw, G.L., and von Stein, A. (1998). Synchronization between prefrontal and posterior association cortex during human working memory. Proc. Natl. Acad. Sci. USA *95*, 7092–7096.

Seung, H.S., Lee, D.D., Reis, B.Y., and Tank, D.W. (2000). Stability of the memory of eye position in a recurrent network of conductance-based model neurons. Neuron *26*, 259–271.

Singer, W., and Gray, C.M. (1995). Visual feature integration and the temporal correlation hypothesis. Annu. Rev. Neurosci. *18*, 555–586.

Stam, C.J., van Cappellen van Walsum, A.M., and Micheloyannis, S. (2002). Variability of EEG synchronization during a working memory task in healthy subjects. Int. J. Psychophysiol. *46*, 53–66.

Tallon-Baudry, C., Bertrand, O., Peronnet, F., and Pernier, J. (1998). Induced gamma-band activity during the delay of a visual short-term memory task in humans. J. Neurosci. *18*, 4244–4254.

Tallon-Baudry, C., Kreiter, A., and Bertrand, O. (1999). Sustained and transient oscillatory responses in the gamma and beta bands in a visual short-term memory task in humans. Vis. Neurosci. *16*, 449–459.

Tallon-Baudry, C., Bertrand, O., and Fischer, C. (2001). Oscillatory synchrony between human extrastriate areas during visual short-term memory maintenance. J. Neurosci. *21*, RC177.

Tallon-Baudry, C., Mandon, S., Freiwald, W.A., and Kreiter, A.K. (2004). Oscillatory synchrony in the monkey temporal lobe correlates with performance in a visual short-term memory task. Cereb. Cortex *14*, 713–720.

Tesche, C.D., and Karhu, J. (2000). Theta oscillations index human hippocampal activation during a working memory task. Proc. Natl. Acad. Sci. USA *97*, 919–924.

Wang, X.J. (2001). Synaptic reverberation underlying mnemonic persistent activity. Trends Neurosci. *24*, 455–463.

Watanabe, M. (1996). Reward expectancy in primate prefrontal neurons. Nature *382*, 629–632.

Worden, M.S., Foxe, J.J., Wang, N., and Simpson, G.V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. J. Neurosci. *20*, RC63.

Behavioral/Systems/Cognitive

Brain Oscillations Control Timing of Single-Neuron Activity in Humans

Joshua Jacobs,[1] Michael J. Kahana,[2] Arne D. Ekstrom,[3] and Itzhak Fried[4,5]

[1]Neuroscience Graduate Group and [2]Department of Psychology, University of Pennsylvania, Philadelphia, Pennsylvania 19104, [3]Center for Cognitive Neurosciences, Semel Institute, Department of Psychiatry and Biobehavioral Sciences, University of California School of Medicine, Los Angeles, California 90095, [4]Division of Neurosurgery and Semel Institute for Neuroscience and Human Behavior, University of California, Los Angeles, California 90095, and [5]Functional Neurosurgery Unit, Tel-Aviv Medical Center and Sackler Faculty of Medicine, Tel-Aviv University, Tel-Aviv 69978, Israel

A growing body of animal research suggests that neurons represent information not only in terms of their firing rates but also by varying the timing of spikes relative to neuronal oscillations. Although researchers have argued that this temporal coding is critical in human memory and perception, no supporting data from humans have been reported. This study provides the first analysis of the temporal relationship between brain oscillations and single-neuron activity in humans. Recording from 1924 neurons, we find that neuronal activity in various brain regions increases at specific phases of brain oscillations. Neurons in widespread brain regions were phase locked to oscillations in the theta- (4–8 Hz) and gamma- (30–90 Hz) frequency bands. In hippocampus, phase locking was prevalent in the delta- (1–4 Hz) and gamma-frequency bands. Individual neurons were phase locked to various phases of theta and delta oscillations, but they only were active at the trough of gamma oscillations. These findings provide support for the temporal-coding hypothesis in humans. Specifically, they indicate that theta and delta oscillations facilitate phase coding and that gamma oscillations help to decode combinations of simultaneously active neurons.

Key words: **phase locking; theta; gamma; intracranial EEG; navigation; local field potential**

Introduction

Many recent neurobiological theories of human memory and perception rely critically on brain oscillations to coordinate the timing of neuronal spiking. For example, Buzsáki (2005) proposed a theory of episodic memory in which theta oscillations (4–8 Hz) act as a timing signal to ensure that neurons representing related stimuli spike nearby each other in time. In addition, other researchers have suggested memory models that rely on theta for different functions such as separating the memory-encoding and the memory-retrieval processes (Hasselmo et al., 2002) or maintaining stimuli representations in working memory (Lisman and Idiart, 1995). In addition to theta, gamma oscillations (30–90 Hz) play a role in a different set of models in which they bind together sets of neurons to represent complex stimuli (for review, see Fries, 2005).

The experimental data supporting these theories come from simultaneous recordings of neuronal spiking and local-field potentials (LFPs) in animals. For example, the discovery that some rodent hippocampal neurons spike at different phases of the theta oscillation depending on the animal's behavioral state or spatial location suggested that theta may be a general mechanism for phase coding (Fox et al., 1986; O'Keefe and Recce, 1993). There is also considerable evidence for the role of gamma oscillations in animal cognition and behavior. For example, in olfactory and visual cortices, the amplitude and synchrony of gamma oscillations indicates properties of external stimuli (Freeman and Schneider, 1982; Gray et al., 1989). Furthermore, recent research indicates that gamma oscillations play a role in cognitive behaviors such as maintaining stimuli representations in working memory (Pesaran et al., 2002) and selectively attending to one visual stimulus in a crowded scene (Fries et al., 2001). Although excitatory and inhibitory neurons spike at different phases of the gamma oscillation (Eeckman and Freeman, 1990; Csicsvari et al., 2003), there are no reports of gamma-band phase coding in which a single neuron would spike at different phases of gamma according to behavior.

Consistent with these striking observations of brain oscillations in animals, there is now a large body of work demonstrating functional correlates of theta and gamma oscillations in the human brain (for review, see Kahana, 2006). These studies have used a variety of methods including scalp electroencephalography and magnetoencephalography, as well as LFPs recorded intracranially in patients undergoing neurosurgical treatments. Although brain oscillations, particularly in the theta band, are prominent both in humans and in animals, it has sometimes proven difficult to link human and animal research findings. The theta oscillation most often studied in humans is in the neocortex

Received Oct. 25, 2006; revised Jan. 26, 2007; accepted Feb. 20, 2007.

This work was supported by National Institutes of Health Research Grants MH61975 and MH062196, National Science Foundation Grant SBE0354378, and by the Swartz Foundation. We gratefully acknowledge help from Matthew Mollison and Igor Korolev for cluster cutting and Sean Polyn, Marieke van Vugt, and Per Sederberg for helpful discussions.

Correspondence should be addressed to either of the following: Dr. Itzhak Fried, Division of Neurosurgery, University of California, Los Angeles, Los Angeles, CA 90095, E-mail: ifried@mednet.ucla.edu; or Dr. Michael J. Kahana, Department of Psychology, University of Pennsylvania Philadelphia, PA 19104, E-mail: kahana@psych.upenn.edu.

DOI:10.1523/JNEUROSCI.4636-06.2007

 0270-6474/07/273839-06$15.00/0

and has many distinct local generators (Raghavachari et al., 2006), whereas the theta oscillation typically studied in animals is in the hippocampus and is primarily driven by one of a small number of sources (Buzsáki, 2002). To date, no human studies have compared brain oscillations to the timing of neuronal spiking.

Our goal in the present study was to characterize the temporal relationship between single-neuron activity and LFP oscillations in the human brain. Going beyond previous animal work, we sought to characterize this relationship across a wide range of frequencies and across a number of disparate brain regions. To accomplish this, we recorded simultaneous spiking activity and LFPs from intracranial microelectrodes in 20 patients undergoing treatment for drug-resistant epilepsy (see Materials and Methods). In these recordings, we identified a total of 1924 neurons, and compared the activity of each neuron with LFP oscillations recorded at the same microelectrode. During recording, subjects played "Yellow Cab," a virtual taxi driver video game. We selected this task because it induces task-related brain oscillations at a range of frequencies, including the theta band, in cortical and limbic structures (Ekstrom et al., 2005).

Materials and Methods

We examined data from intracranial microelectrodes in patients undergoing surgical treatment for drug-resistant epilepsy. Electrodes were positioned by clinical teams to isolate the epileptic seizure focus for subsequent potential surgical resection (surgeries performed by I.F.). Microelectrode coverage included hippocampus (377 neurons), parahippocampal region (Witter and Wouterlood, 2002) (421 neurons), amygdala (404 neurons), frontal regions (438 neurons), and occasionally temporal and parietal cortices (284 neurons). This study conformed to the guidelines of the Medical Institutional Review Board at University of California, Los Angeles. We examined data from a total of 46 recording sessions from 20 different subjects (individual subjects participated in 1–4 sessions) (for details, see supplemental Table 1, available at www.jneurosci.org as supplemental material). These recording sessions took place in patients' free time between medical treatments. During each recording session, subjects played "Yellow Cab" (Ekstrom et al., 2005) on a bedside laptop computer for 25–60 min. During the game, ≈75% of the participants' time was spent actively moving to particular landmarks in the environment. The remainder of the time was spent reading instruction screens and planning routes.

Each patient was implanted with 6–12 clinical intracranial depth electrodes. Each clinical electrode terminated with a set of nine 40 μm platinum–iridium microwires (Fried et al., 1999). The first eight microwires were insulated except for their tip and were used to record single-unit action potentials and LFPs. The ninth microwire had its insulation stripped for ≈1 cm and served as the recording reference for the other eight microwires on that depth probe. We recorded from each microwire at sampling rates of 28–32 kHz using a Cheetah recording system (Neuralynx, Tucson, AZ). Action potentials were manually isolated using spike shape, clustering of wavelet coefficients, and interspike intervals (Quiroga et al., 2004). Typically we isolated zero or one distinct neuron from each microwire, but in rare cases we observed up to three distinct neurons from a single microwire.

For analyses of LFP oscillations, we downsampled recordings to 2 kHz and then applied 60 and 120 Hz second-order Butterworth notch filters. To minimize the contribution of low-frequency components of spikes toward spectral calculations, we removed the samples from 2 ms before to 8 ms after each spike and replaced them with a linear interpolation of the underlying LFP signal. Then, we computed oscillatory phase and power of the LFP using Morlet wavelets (wave number, 4) at frequencies between 1 and 100 Hz ($2^{x/8}$ Hz for $x \in 0, \ldots, 53$). We considered a neuron phase locked at frequency f if the hypothesis of uniformity for its LFP f hertz phase distribution could be rejected at $p < 0.001$ using a Bonferroni-corrected Rayleigh test (Fisher, 1993). This Bonferroni correction compensated for our application of the Rayleigh test across each of 54 frequencies. In summary statistics, if a neuron fulfilled this phase-locking criterion at multiple frequencies, we considered it phase locked at the frequency at which its phase distribution was most nonuniform according to the Rayleigh statistic.

We studied neuronal spiking activity during different levels of LFP oscillatory power by first calculating the phase and power of the LFP of each phase-locked neuron throughout the entire recording session, at the frequency at which it was phase locked. We then labeled each point of the recording session according to whether the LFP exhibited high, medium, or low power. To perform this labeling, first we removed any potential artifacts caused by movement or electrical noise by discarding the 10% of the recording session corresponding to the lowest fifth percentile and the highest 95th percentile of LFP power throughout the recording session. Then, we labeled the remaining time points into one of three groups (high, medium, or low) according to the instantaneous log-transformed power (at the phase-locked frequency). The boundaries for these groups were equally spaced across log-transformed power values. Finally, separately for the time points in each group, we calculated the instantaneous firing rate of each neuron as a function of oscillatory phase. (The phase was divided into 16 groups equally spaced between 0 and 2π.)

Results

We identified 1215 neurons whose spiking was phase locked to LFP oscillations. This phenomenon was prominent in the theta band (4–8 Hz). Figure 1*A* illustrates the behavior of a neuron in the right superior temporal gyrus of subject 2 that was phase locked to 7.3 Hz theta oscillations. This firing rate of the neuron varied by more than twofold according to the instantaneous theta phase: it had a firing rate of 7.9 Hz at its preferred phase (5.7 rad) just before the peak of the oscillation and a firing rate of 3.7 Hz near the trough of the oscillation (Fig. 1*A*, right). Figure 1*B* shows the activity of a neuron from subject 1's right temporal cortex that was phase locked to the trough of 6.2 Hz theta oscillations.

In addition to the theta band, we also identified neurons that exhibited phase-locked spiking to oscillations in the delta- (1–4 Hz), alpha- (10–16 Hz), beta- (16–30 Hz), and gamma- (30–90 Hz) frequency ranges. Thus, human neuronal phase locking is a general phenomenon that is not isolated to a narrow frequency range. For example, Figure 1, *C* and *D*, depicts the activities of neurons in entorhinal and orbitofrontal cortices that were phase locked to oscillations at 3.4 and 3.7 Hz, respectively. Figure 1, *E* and *F*, shows the behavior of hippocampal neurons that were phase locked to gamma oscillations.

An additional pattern we observed was that neurons could be phase locked to oscillations at multiple frequencies. For example, Figure 1*E* shows the activity of a neuron from the right posterior hippocampus of subject 4 that was primarily phase locked to 70 Hz gamma oscillations but also exhibited moderate phase locking to delta oscillations. Figure 1*H* presents the activity of a neuron from the left amygdala of subject 3 that exhibited significant phase locking to both 9.5 and 1.8 Hz oscillations. Figure 1*C* describes a neuron from the left entorhinal cortex that was phase locked to oscillations at 3.4 and 12 Hz.

In the rodent hippocampus, the prominent oscillatory pattern is the 4–8 Hz theta rhythm (Buzsáki, 2002). However, in our recordings, we found that hippocampal phase locking was more prevalent in the 1–4 Hz range than in the 4–8 Hz range. For example, Figure 1*G* depicts the activity of a hippocampal neuron that was phase locked to oscillations at 1.1 Hz. In addition, the hippocampal neurons highlighted in Figure 1, *E* and *F*, also exhibit significant phase locking at 1–4 Hz.

Animal studies have identified neurons that were phase locked to local oscillations in hippocampus (Fox et al., 1986; Skaggs et al., 1996; Csicsvari et al., 2003) and entorhinal (Chrobak and

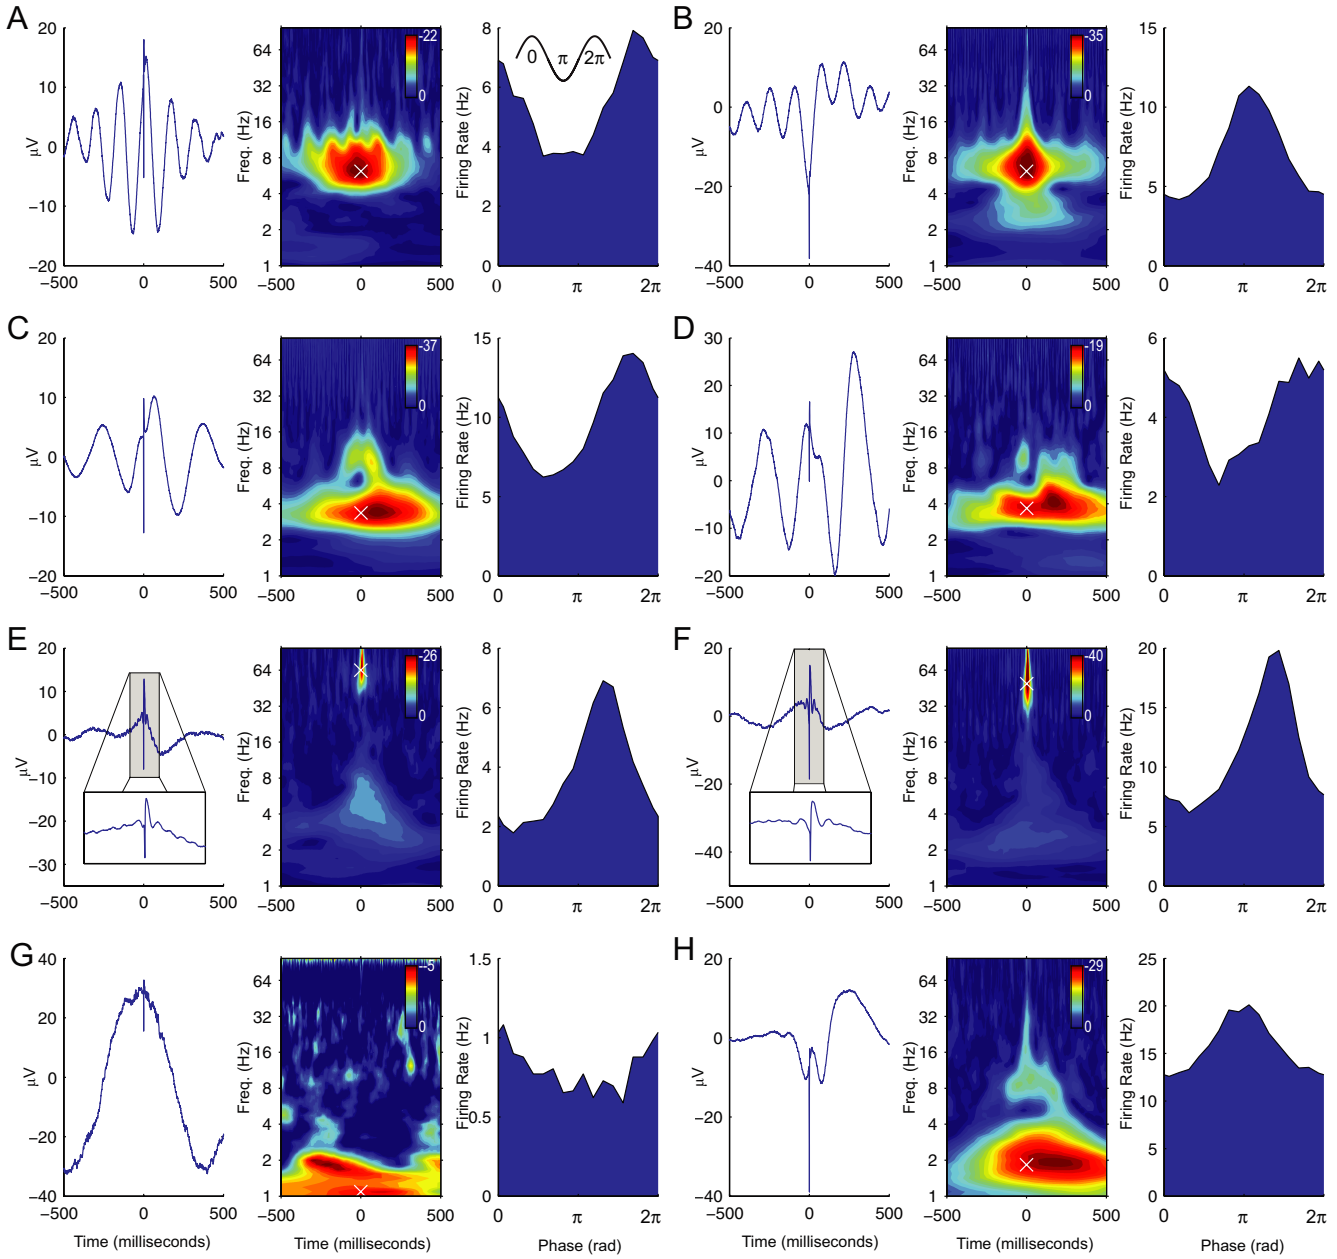

Figure 1. ***A***, Activity of a neuron from subject 2's right superior temporal gyrus, which was phase locked to the peak of 7.3 Hz theta oscillations. Left, Spike-triggered LFP average. Spike onset occurred at 0 ms. Middle, *Z* score from Rayleigh test evaluating the hypothesis of a uniform LFP phase distribution at the moment of spike onset. Shifts along the horizontal axis indicate the relation between spiking and phase of time-shifted LFP oscillations. The white × symbol indicates the frequency at which phase locking at 0 ms is most statistically significant. Right, Firing rate as a function of phase at 7.3 Hz (the frequency marked by the × in middle panel). The circular mean resultant vector length $\bar{R}$ (Fisher, 1993) of this 7.3 Hz LFP phase distribution is 0.19. The inset example waveform is a reminder that the phases of the peak and trough of an oscillation are 0 (or 2π) radians and π radians, respectively. ***B***, Activity of a neuron from subject 1's right temporal gyrus that was phase locked to the trough of 6.2 Hz theta oscillations ($\bar{R} = 0.24$). ***C***, Behavior of a neuron from subject 20's left entorhinal cortex, which was primarily phase locked to 3.4 Hz oscillations ($\bar{R} = 0.2$), but also exhibited some phase locking to oscillations at 9.5 Hz. ***D***, A neuron from subject 18's right orbitofrontal cortex, which was phase locked to 3.7 Hz oscillations ($\bar{R} = 0.17$). ***E***, Behavior of a neuron from subject 4's left posterior hippocampus that was phase locked to the trough of 70 Hz gamma oscillations ($\bar{R} = 0.17$). Inset in left panel depicts a zoomed plot of the spike-triggered average of this neuron. ***F***, Behavior of a neuron from subject 12's left anterior hippocampus that exhibited phase-locked spiking near the trough of 49 Hz gamma oscillations ($\bar{R} = 0.2$). ***G***, The activity of a neuron from subject 1's left anterior hippocampus, which was phase locked to the peak of 1.1 Hz oscillations ($\bar{R} = 0.1$). ***H***, Activity of a neuron from the left amygdala of subject 3 that was phase locked to oscillations primarily at 1.8 Hz ($\bar{R} = 0.11$), but also at 9.5 Hz.

Buzsáki, 1998), parietal (Pesaran et al., 2002), and visual cortices (Fries et al., 2001; Lee et al., 2005): this phenomenon was notably absent from frontal cortex (Siapas et al., 2005). Because our dataset included recordings from widespread brain regions, we had the opportunity to determine whether the prevalence of phase locking in particular frequency ranges localized to particular brain regions. We found that the prevalence of phase locking significantly varied across brain regions in the delta- [$\chi 2(4) = 89$; $p < 0.001$], theta- [$\chi 2(4) = 78$; $p < 0.001$], and gamma- [$\chi 2(4) = 71$; $p < 0.001$] frequency bands. (For details, see supplemental Table 1, available at www.jneurosci.org as supplemental material.) Figure 2*A* shows that phase locking was especially prevalent at theta frequencies in temporal and parietal cortices [$\chi 2(1) = 64$, $p < 0.001$, Bonferroni-corrected] and at delta frequencies in hippocampus [$\chi 2(1) = 47$, $p < 0.001$, Bonferroni-corrected]. In hippocampus, significantly more neurons were phase locked to

delta oscillations than to theta oscillations [$\chi 2(1) = 51$; $p < 0.001$]. Note that unlike some recent studies in animals (Siapas et al., 2005), we report some phase locking to theta oscillations in orbitofrontal cortices (20 of 438 cells) (see also Fig. 1*D*). Figure 2*C* depicts the strength of phase locking across all phase-locked neurons by plotting the circular mean resultant vector length $\bar{R}$, which is a measure of relative increase in firing rate of each neuron at its preferred phase (Fisher, 1993). The median $\bar{R}$ across the population of phase-locked neurons is 0.1.

The phase-coding hypothesis predicts that neurons encode information via the oscillatory phase at which they spike (Huxter et al., 2003). This theory thus predicts that different neurons would spike at varied phases of ongoing oscillations. To test the phase-coding hypothesis in humans, we labeled the preferred phase of each neuron as the oscillatory phase at which its firing rate was highest. (The peak and trough of each oscillation correspond to 0 [or 2π] radians and π radians, respectively.) Figure 2*B* shows the preferred phase of each phase-locked neuron. Neurons that were phase locked to delta or theta oscillations had varied preferred phases (Fig. 2*D*; see also Fig. 1*A–D*,*G*,*H*). However, neurons phase locked to gamma oscillations had preferred phases at or just after the trough of the oscillation (Fig. 2*E*; see also Fig. 1*E*,*F*). These dissimilar preferred-phase patterns ($p < 0.001$, Mardia–Watson–Wheeler test) indicate that neurons phase locked to gamma oscillations behave differently from those phase locked to delta or theta oscillations.

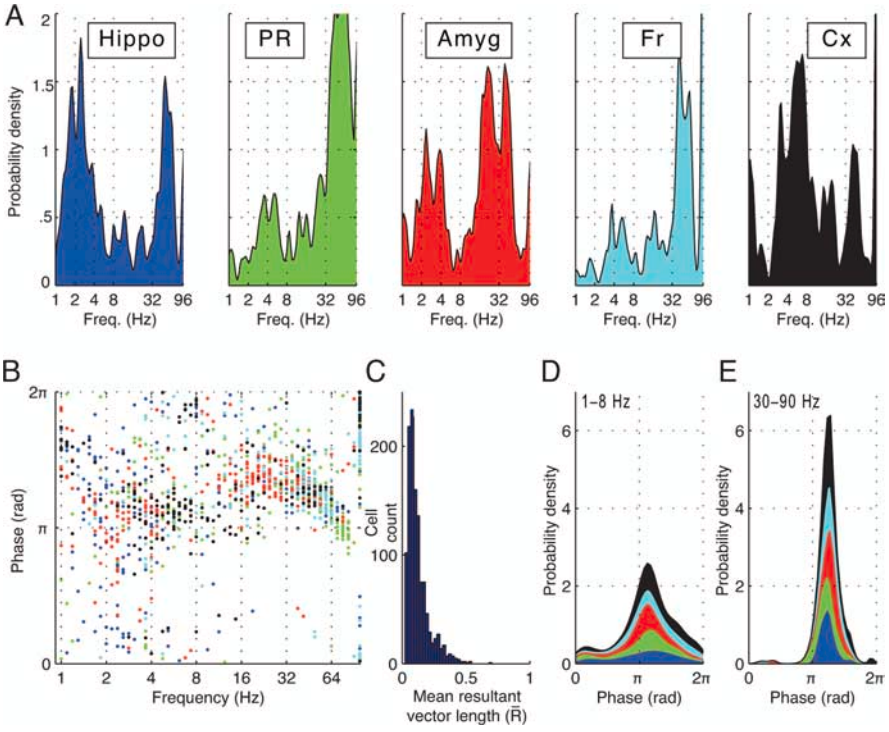

Figure 2. ***A***, Probability of a neuron exhibiting significant phase locking as a function of frequency (Freq.), grouped by region. Hippo, Hippocampus; PR, parahippocampal region; Amyg, amygdala; Fr, frontal region; Cx, temporal and parietal cortices. See Materials and Methods for phase-locking criteria. Total area under each curve indicates the fraction of neurons in each region that were phase locked (63, 70, 67, 49, and 72%, in hippocampus, parahippocampal region, amygdala, frontal region, and temporal and parietal cortices, respectively). Each curve is smoothed with a Gaussian kernel. ***B***, Distribution of preferred phases (i.e., the phase in which firing rate was highest) of all 1215 phase-locked neurons. The region of each neuron is indicated using color scheme from ***A***. ***C***, Histogram of the circular mean resultant vector length ($\bar{R}$) (Fisher, 1993) of each phase-locked neuron. ***D***, Preferred-phase probability density for all neurons that were phase locked at delta or theta frequencies (1–8 Hz). Coloring within each curve indicates the preferred-phase distributions of neurons in different regions. ***E***, Preferred-phase probability density for all neurons that were phase locked to gamma oscillations (30–90 Hz).

Although recent work indicates that the power of brain oscillations related to mean neuronal firing rates (Logothetis et al., 2001; Mukamel et al., 2005; Niessing et al., 2005), it is unknown how this phenomenon relates to the activity of phase-locked neurons. To study this relationship, we analyzed the firing rate of each phase-locked neuron during high-, medium-, and low-power oscillations (see Materials and Methods). We identified a number of neurons that exhibited excitation (i.e., increased firing) at their preferred phase during high-power oscillations, compared with their activity during low oscillatory power (Fig. 3*A*,*B*). Some neurons displayed the opposite pattern and were primarily inhibited (i.e., decreased firing) at nonpreferred phases during high-power oscillations (Fig. 3*C*). Finally, other neurons responded to high-power oscillations with substantial firing-rate changes at all phases (Fig. 3*D*).

This led us to ask whether, across all neurons, phase locking was primarily related to excitation at preferred phases or inhibition at nonpreferred phases. Our analysis showed that neurons exhibited significantly greater firing rates during high-power oscillations than during periods of low oscillatory power (mean increase, 0.7 Hz; $p < 0.001$, t test). This effect was especially robust at the preferred phase of high-power oscillations at which the mean firing rate was elevated by 3.8 Hz compared with periods of low oscillatory power ($p < 0.001$, *t* test). This indicates that the phase-locking phenomenon is primarily related to increased firing during high-power oscillations. This is consistent with the in vitro finding that some neurons act as bandpass filters by spiking in response to experimentally induced electrical oscillations at particular frequencies (Pike et al., 2000).

Discussion

Our results indicate that the neuronal phase-locking phenomenon is present in various brain regions of humans engaged in cognitive tasks. This phenomenon, in conjunction with the well, studied and diverse relationship between brain oscillations and human and animal learning (Berry and Thompson, 1978; Rizzuto et al., 2003; Jacobs et al., 2006), indicates that oscillation-modulated temporal coding plays a role in human cognition. Furthermore, our data suggest that gamma oscillations facilitate a different type of information coding compared with delta or theta oscillations. Delta and theta oscillations may facilitate phase-based temporal coding because we found that individual phase-locked neurons had widely varied preferred phases. This is consistent with the observations that different rodent hippocampal neurons were phase locked to varied theta phases (Fox et al., 1986; O'Keefe and Recce, 1993) and that the amplitude of gamma oscillations varies according to the instantaneous theta phase (Mormann et al., 2005). In contrast to delta and theta oscillations, neurons that phase locked to gamma oscillations consistently had preferred phases near the trough of the oscillation; thus, it seems unlikely that these oscillations facilitate phase cod-

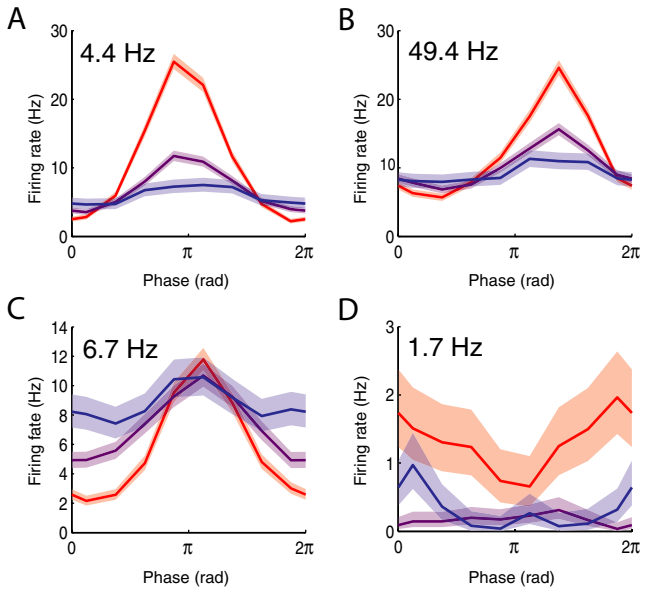

Figure 3. Relationship between oscillatory power and firing rates of phase-locked neurons. ***A***, Activity of a neuron from subject 5's left parietal cortex that exhibited maximal phase locking at 4.4 Hz. Red, purple, and blue lines indicate firing rates during high, medium, and low oscillatory power, respectively (for details, see Materials and Methods). The shaded area indicates 95% confidence intervals based on the binomial distribution. ***B***, Activity of a neuron from subject 12's left anterior hippocampus (same neuron as Fig. 1*F*) that was phase locked to LFP oscillations at 49 Hz. ***C***, Activity of a neuron from subject 1's right superior temporal gyrus (same neuron as Fig. 1*A*). ***D***, Activity of a neuron from subject 18's left anterior hippocampus.

ing. This trough-locked spiking is similar to data from animal studies demonstrating that neurons in both neocortex and hippocampus spike at or just after the trough of gamma oscillations (Eeckman and Freeman, 1990; Csicsvari et al., 2003). In these systems, it is believed that information is encoded by the combination of neurons that spike at least once in each oscillation cycle, rather than the phase of each spike. Thus, our findings are consistent with the idea that human neuronal networks oscillating at gamma frequencies encode information by the combination of active neurons in each 10–50 ms oscillation cycle (Harris et al., 2003; Fries, 2005).

One important difference between our human results and animal electrophysiological literature is our observation of prominent hippocampal oscillations at 1–4 Hz (see Figs. 1*E–G*, 2*A*) rather than the 4–8 Hz theta oscillation prominent in rodent hippocampus (Buzsáki, 2002). Although this human hippocampal oscillation falls into the delta range (1–4 Hz), it shares several characteristics with the 4–8 Hz rodent theta oscillation. For example, we found that individual hippocampal neurons were phase locked to widespread phases of this 1–4 Hz oscillation (Fig. 2*D*) and that some hippocampal neurons simultaneously phase locked both to these 1–4 Hz oscillations and to gamma oscillations (Fig. 1*E*,*F*). Analogous trends have been observed with rodent hippocampal theta (Fox et al., 1986; Bragin et al., 1995). An additional link between these two oscillatory patterns comes from Bódizs et al. (2001), who report that during rapid-eye-movement sleep, the human hippocampus exhibits robust 1.5–3 Hz delta-band oscillations, whereas animals in this state display 4–8 Hz theta oscillations. Based on this evidence, we suggest that the frequency range of human hippocampal theta may extend to ≈1 Hz.

The existence of neuronal phase locking helps to interpret behavior-related changes in brain recordings. For example, after humans view visual stimuli, theta oscillations reset to particular phases in response to task demands (Rizzuto et al., 2003). This phase reset simultaneously occurs throughout multiple brain regions; thus, it may cause phase-locked neurons throughout the brain to spike in precise temporal patterns. Thus, our findings suggest that oscillatory phase resetting temporally synchronizes phase-locked neurons in widespread brain regions. This is consistent with reports of behavior-related changes in the coherence of oscillations in different brain regions (Jones and Wilson, 2005; Hyman et al., 2005). These patterns of inter-region oscillatory synchrony may be elaborate because our findings suggest that neuronal phase locking spans a broad range of frequencies (Fig. 2*A*), and that some brain regions exhibit oscillations at multiple frequencies (Fig. 1*C*,*E*,*F*,*H*) (Chrobak and Buzsáki, 1998). Finally, recent imaging studies showed that hemodynamic activity closely relates to both power of gamma oscillations and mean neuronal firing rates (Logothetis et al., 2001; Mukamel et al., 2005; Niessing et al., 2005). Our findings substantially add to these results by indicating that during high-power gamma oscillations, neuronal activity is phase locked to these oscillations. Thus, we suggest that reports of increased hemodynamic activity, in regions in which they are correlated with the amplitude of gamma oscillations, should be interpreted both in terms of increased neuronal firing and increased neuronal gamma-band synchrony.

References

Berry SD, Thompson RF (1978) Prediction of learning rate from the hippocampal electroencephalogram. Science 200:1298–1300.

Bódizs R, Kántor S, Szabó G, Szűcs A, Erőss L, Halász P (2001) Rhythmic hippocampal slow oscillation characterizes REM sleep in humans. Hippocampus 747–753.

Bragin A, Jando G, Nadasdy Z, Hetke J, Wise K, Buzsaki G (1995) Gamma (40–100 Hz) oscillation in the hippocampus of the behaving rat. J Neurosci 15:47–60.

Buzsáki G (2002) Theta oscillations in the hippocampus. Neuron 33:325–340.

Buzsáki G (2005) Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. Hippocampus 15:827–840.

Chrobak J, Buzsáki G (1998) Gamma oscillations in the entorhinal cortex of the freely behaving rat. J Neurosci 18:388–298.

Csicsvari J, Jameison B, Wise KD, Buzsáki G (2003) Mechanisms of gamma oscillations in the hippocampus of the behaving rat. Neuron 37:311–322.

Eeckman F, Freeman W (1990) Correlations between unit firing and EEG in the rat olfactory system. Brain Res 528:238–244.

Ekstrom AD, Caplan J, Ho E, Shattuck K, Fried I, Kahana M (2005) Human hippocampal theta activity during virtual navigation. Hippocampus 15:881–889.

Fisher NI (1993) Statistical analysis of circular data. Cambridge, UK: Cambridge UP.

Fox SE, Wolfson S, Ranck JBJ (1986) Hippocampal theta rhythm and the firing of neurons in walking and urethane anesthetized rats. Exp Brain Res 62:495–508.

Freeman W, Schneider W (1982) Changes in spatial patterns of rabbit olfactory EEG with conditioning to odors. Psychophysiology 19:44–56.

Fried I, Wilson CL, Maidment NT, Engel Jr J, Behnke E, Fields TA, MacDonald KA, Morrow JW, Ackerson L (1999) Cerebral microdialysis combined with single-neuron and electroencephalographic recording in neurosurgical patients. J Neurosurg 91:697–705.

Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. Trends Cogn Sci 9:474–480.

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

Gray C, König P, Engel A, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. Nature 338:334–337.

Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsáki G (2003) Organization of cell assemblies in the hippocampus. Nature 424:552–556.

Hasselmo ME, Bodelon C, Wyble BP (2002) A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. Neural Comput 14:793–817.

Huxter J, Burgess N, O'Keefe J (2003) Independent rate and temporal coding in hippocampal pyramidal cells. Nature 425:828–832.

Hyman J, Zilli E, Paley A, Hasselmo M (2005) Medial prefrontal cortex cells show dynamic modulation with the hippocampal theta rhythm dependent on behavior. Hippocampus 15:739–749.

Jacobs J, Hwang G, Curran T, Kahana MJ (2006) EEG oscillations and recognition memory: theta correlates of memory retrieval and decision making. NeuroImage 15:978–987.

Jones MW, Wilson MA (2005, Dec) Theta rhythms coordinate hippocampalprefrontal interactions in a spatial memory task. PLoS Biol 3:e402.

Kahana MJ (2006) The cognitive correlates of human brain oscillations. J Neurosci 26:1669–1672.

Lee H, Simpson G, Logothetis N, Rainer G (2005) Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. Neuron 45:147–156.

Lisman J, Idiart MA (1995) Storage of 7 ± 2 short-term memories in oscillatory subcycles. Science 267:1512–1515.

Logothetis N, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysio-logical investigation of the basis of the fMRI signal. Nature 412:150–157.

Mormann F, Fell J, Axmacher N, Weber B, Lehnertz K, Elger C, Fernandez G (2005) Phase/amplitude reset and theta-gamma interaction in the human medial temporal lobe during a continuous word recognition memory task. Hippocampus 15:890–900.

Mukamel R, Gelbard H, Arieli A, Hasson U, Fried I, Malach R (2005) Coupling between neuronal firing, field potentials, and FMRI in human auditory cortex. Science 309:951–954.

Niessing J, Ebisch B, Schmidt KE, Niessing M, Singer W, Galuske RA (2005) Hemodynamic signals correlate tightly with synchronized gamma oscillations. Science 309:948–951.

O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3:317–330.

Pesaran B, Pezarais 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.

Pike F, Goddard RS, Suckling JM, Ganter P, Kasthuri N, Paulsen O (2000) Distinct frequency preferences of different types of rat hippocampal neurones in response to oscillatory input currents. J Physiol (Lond) 529:205–213.

Quiroga RQ, Nadasdy Z, Ben-Shaul Y (2004) Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. Neural Comput 16:1661–1687.

Raghavachari S, Lisman JE, Tully M, Madsen JR, Bromfield EB, Kahana MJ (2006) Theta oscillations in human cortex during a working memory task: evidence for local generators. J Neurophysiol 95:1630–1638.

Rizzuto DS, Madsen JR, Bromfield EB, Schulze-Bonhage A, Seelig D, Aschenbrenner-Scheibe R, Kahana MJ (2003) Reset of human neocortical oscillations during a working memory task. Proc Natl Acad Sci USA 100:7931–7936.

Siapas A, Lubenov E, Wilson M (2005) Prefrontal phase locking to hippocampal theta oscillations. Neuron 46:141–151.

Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus 6:149–172.

Witter M, Wouterlood F (2002) The parahippocampal region: past, present and future. In: The parahippocampal region: organization and role in cognitive functions (Witter M, Wouterlood F, eds), pp 3–19. Oxford: Oxford UP.

European Journal of Neuroscience, Vol. 15, pp. 1244–1254, 2002

Timecourse of object-related neural activity in the primate prefrontal cortex during a short-term memory task

Gregor Rainer and Earl K. Miller
Center for Learning and Memory, RIKEN-MIT Neuroscience Center, and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Keywords: delay activity, electrophysiology, *Macacca mulatta*, monkey, natural images

Abstract

We studied the timecourse of neural activity in the primate (*Macacca mulatta*) prefrontal (PF) cortex during an object delayed-matching-to-sample (DMS) task. To assess the effects of experience on this timecourse, we conducted the task using both novel and highly familiar objects. In addition, noise patterns containing no task-relevant information were used as samples on some trials. Comparison of average PF ensemble activity relative to baseline activity generated by objects and noise patterns revealed three distinct activity periods. (i) Sample onset elicited a transient sensory visual response. In this sensory period, novel objects elicited stronger average ensemble activity than both familiar objects and noise patterns. (ii) An intermediate period of elevated activity followed, which began before sample offset, and continued well into the delay period. In the intermediate period, activity was elevated for noise patterns and novel objects, but near baseline for familiar objects. (iii) Finally, after average ensemble activity reached baseline activity at the end of the intermediate period, a reactivation period occurred late in the delay. Experience had little effect during reactivation, where activity was elevated for both novel and familiar objects compared to noise patterns. We show that the ensemble average resembles the activity timecourse of many single prefrontal neurons. These results suggest that PF delay activity does not merely maintain recent sensory input, but is subject to more complex experience-dependent dynamics. This has implications for how delay activity is generated and maintained.

Introduction

The prefrontal (PF) cortex contains many neurons that show elevated activity during delay periods in cognitive tasks (Fuster, 1993; Goldman-Rakic, 1995; Miller & Cohen, 2001). Studies have implicated delay activity in short-term memory for spatial locations (Funahashi *et al*., 1989; Funahashi *et al*., 1993) and objects (Fuster & Alexander, 1971; Miller *et al*., 1996; Rainer *et al*., 1998a). Delay activity has also been associated with anticipatory or prospective coding for object stimuli (Rainer *et al*., 1999) and reward (Kubota & Niki, 1971; Watanabe, 1996), as well as target selection (Rainer *et al*., 1998b; Hasegawa *et al*. 2000), behavioural rules (White & Wise, 1999; Asaad *et al*. 2000) and motor preparation (Bruce & Goldberg, 1985; Watanabe, 1986; di Pellegrino & Wise, 1993). PF neurons exhibiting delay activity thus play an important role in bridging the gap between a sensory stimulus and a temporally delayed response during cognitive tasks. Lesion and electrophysiological studies have also shown that the PF cortex plays an important role in learning (Petrides, 1985; Parker *et al*., 1998), and that PF neural response properties can be modified strongly by experience (Bichot *et al*., 1996; Rainer & Miller, 2000).

In addition to this *in vivo* work correlating PF delay activity with various cognitive functions, several recent computational studies have investigated by which mechanisms delay activity might be generated and maintained. Models used to describe delay activity have focused typically on persistent activity, i.e. the prolongation of a neural response to a sensory stimulus after the removal of sensory stimulation (Durstewitz *et al*. 2000a; Wang, 2001). Approaches used to model persistent activity include discrete attractor models based on recurrent excitation and inhibition (Amit & Brunel, 1997; Amit *et al*., 1997) and detailed biophysical models (Compte *et al*. 2000; Durstewitz *et al*. 2000b). Other work has focused on networks with bistable solutions with a resting and an active state based on long NMDA-channel (Wang, 1999) or short AMPA-channel (Laing & Chow, 2001) activation timescales. While these models have contributed substantially to our understanding of delay activity and the mechanisms that may give rise to it, the electrophysiological studies suggest that activity in the monkey PF cortex is far more varied and complex than simple persistence of sensory information. Here, our aim is to provide a comprehensive account of the timecourse of neural ensemble activity during a simple cognitive task requiring short-term memory for objects.

We examined neural activity in a delayed-matching-to-sample (DMS) task (Fig. 1). When objects were used as samples, monkeys needed to retain this object information for a short delay to correctly perform the task. On some trials, noise patterns were used as samples. Although these noise patterns had similar image statistical properties as the objects, they did not provide any task-relevant information and did not need to be retained over the course of the delay. This allowed us to compare activity specific to processing and maintenance of task-relevant object information, with activity elicited by noise patterns that needed to be processed but not maintained in short-term memory.

Correspondence: Dr G. Rainer, *MPI, as below.
E-mail: gregor.rainer@tuebingen.mpg.de

**Present address*: Max-Planck-Institute for Biological Cybernetics, Spemannstrasse 38, D-72076 Tübingen, Germany.

Received 15 October 2001, revised 28 January 2002, accepted 19 February 2002

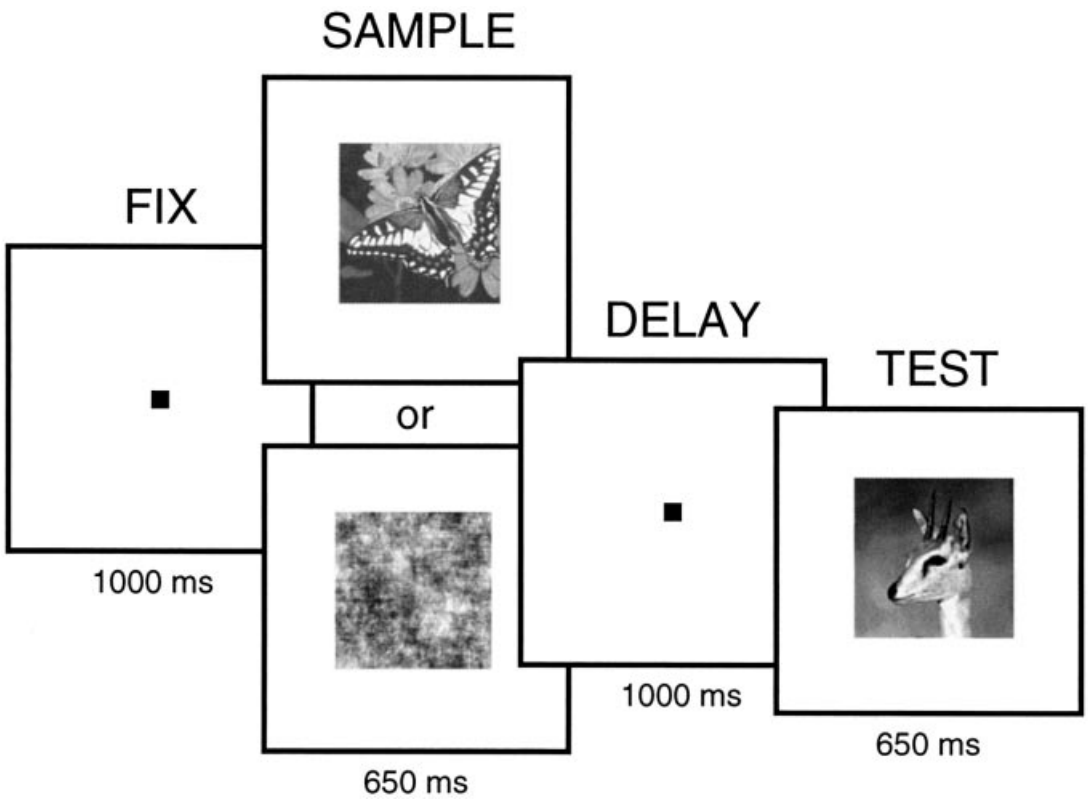

FIG. 1. Sequence of trial events. After a fixation period (FIX), a SAMPLE image (either one of five objects or one of four noise patterns) was presented. After a brief DELAY, a TEST image (one of the five objects) was presented, and monkeys had to release a lever if this TEST image matched the SAMPLE or in the case of a nonmatch hold the lever for the entire test duration and release to a subsequently presented correct match (not shown).

Furthermore, we addressed effects of experience on the timecourse of object processing by conducting the experiment with familiar as well as novel objects (see next section).

Materials and methods

Two adult rhesus monkeys, *Macacca mulatta* (monkey A, female, 8 kg; monkey B, male, 11 kg), participated in the experiments. All procedures were performed in accordance with National Institutes of Health guidelines and the recommendations of the Massachusetts Institute of Technology Animal Care and Use Committee.

Behavioural task

The behavioural paradigm was a modified version of delayed-matching-to-sample (see Fig. 1). Each trial began when the monkey grasped a metal lever. A fixation point ($0.3 \times 0.3°$) was then presented at the centre of a computer screen positioned in front of the animal. After attaining fixation, monkeys were required to maintain fixation within $\pm 1.25°$ of this fixation point throughout the rest of the trial. After 1000 ms of fixation, a sample object was presented for 650 ms. This sample object could be either one of five natural images (objects), one of four noise patterns (noise). During the experiment, intermediate images between objects and noise patterns generated by Fourier phase interpolation were also employed. The purpose of this was to study the ability of prefrontal neurons to communicate information about degraded images. Details about how these degraded images were generated as well as relevant behavioural and neural data have been described elsewhere (Rainer & Miller, 2000), and will not be further discussed here. After a brief delay period, one of the five natural images was presented as a test object. Monkeys had to release the lever if this test object matched the sample, or hold the lever for the entire test object duration in case of a nonmatch. In the nonmatch case, a brief second delay (200 ms) followed, which was always followed by a correct match object requiring a lever release. This second delay was included only to ensure that monkeys made a behavioural response on every trial and was not used in any of the analyses. Match and nonmatch trials occurred equally often. Monkeys received apple juice as a reward for correct performance on trials with object samples. On trials with noise patterns as the sample, half the trials were designated arbitrarily as match trials, the other half as nonmatch trials. Monkeys were thus rewarded randomly on half the trials, independent of whether they held or released the lever on trials with noise pattern samples. This reward protocol was chosen to ensure that monkeys were motivated to attempt identification of intermediate interpolated patterns. Both monkeys had extensive experience over several years with delayed matching tasks prior to participation in the present experiments.

Stimuli

Natural images were selected from a large database containing pictures of animals, faces, flowers and outdoor scenes. After adjusting the images to have equal mean intensity, we computed the Fourier amplitude spectrum for each image, and averaged the amplitude spectra to obtain a mean amplitude spectrum (MAS). The MAS had the spatial frequency ($f^{-\alpha}$) dependence characteristic of natural images (Field, 1987). The Fourier phase spectra of the images were then converted back into image space using the MAS. This ensured

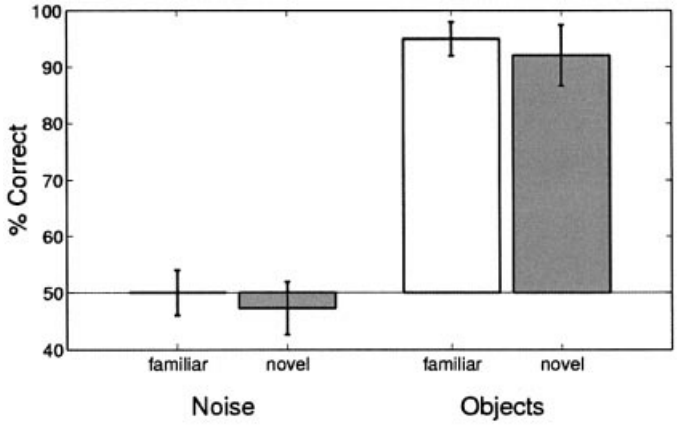

FIG. 2. Behavioural performance – behavioural performance pooled across a total of 25 sessions from two monkeys, with corresponding standard deviations.

that all five objects had identical power at all spatial frequencies. Noise patterns were created by first generating random phase spectra (i.e. each phase coefficient was randomly assigned within the range $-\pi$ to $+\pi$). These random phase spectra were then transformed to image space with the natural image MAS computed previously. Transformations were performed using custom-written software (MATLAB, Mathworks, Natick, USA). Noise patterns were thus matched to the natural images in terms of luminance and spatial frequency content. Stimulus size was 4° × 4°, and mean luminance of each entire image was 14 cd/m^2. Stimuli were always presented at the centre of gaze on a 17-inch computer monitor after appropriate gamma correction to ensure linearity of the display.

Novelty/familiarity

To assess effects of visual experience, monkeys performed the task using familiar and novel objects. During 'familiar object' sessions, we used objects that were highly familiar to the monkeys. They had

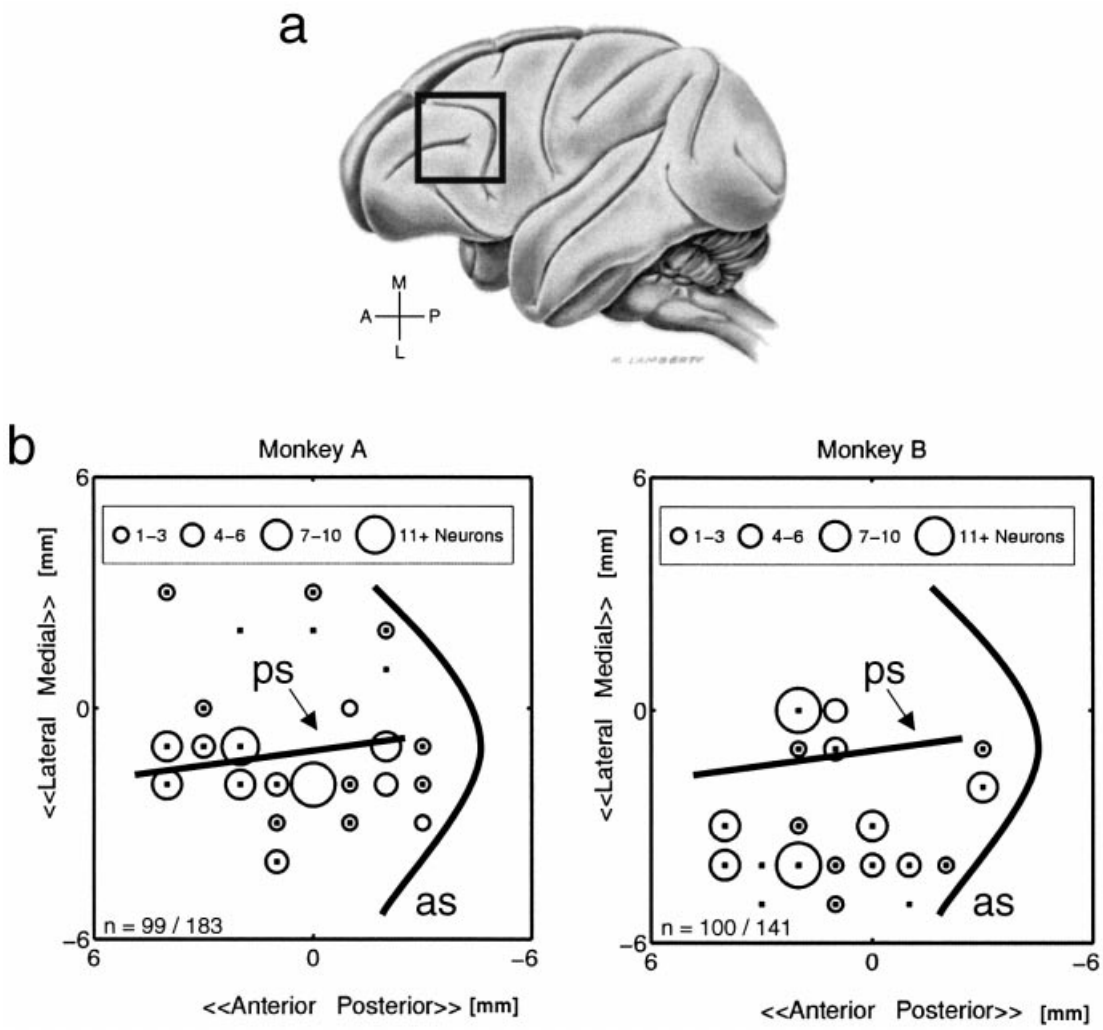

FIG. 3. Recording locations. (a) The box shows the general recording area on a lateral view of a generic *Macacca mulatta* brain. Abbreviations: M, medial; L, lateral; A, anterior; P, posterior. (b) Electrode penetration sites for each of the two monkeys are shown. Recordings were made from a region around and lateral to the principal sulcus (ps), anterior to the arcuate sulcus (as) of the prefrontal cortex. The size of the circles indicates the number of selective neurons recorded at that site; dots represent sites where nonselective neurons were isolated. A neuron was termed selective if it showed a significant difference in activity in response to objects vs. noise patterns in any of the three task periods (sensory, intermediate or reactivation). Significance was assessed using a Wilcoxon signed rank test (evaluated at $P < 0.01$, see Materials and methods). The number of selective neurons is shown at the bottom left for each of the monkeys, expressed as a fraction of the total number of neurons recorded in that animal with familiar and novel objects. (a) Drawing published with permission from The Max Planck Institute for Biological Cybernetics. Drawn by Mr. K. Lamberty for the Institute.

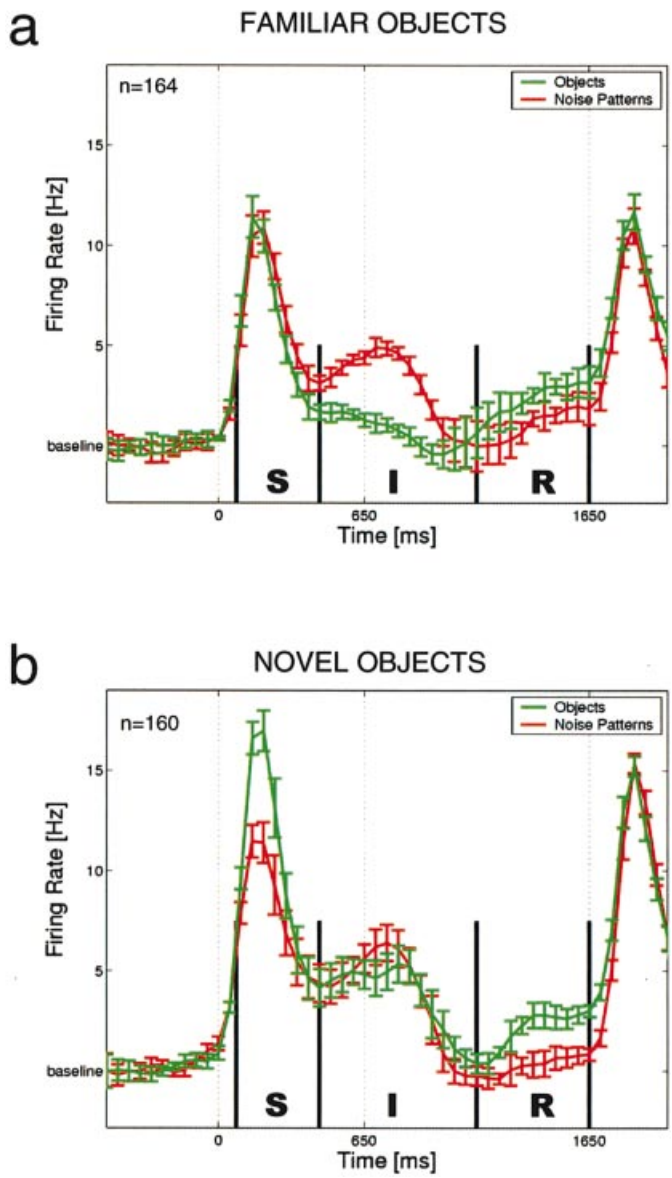

FIG. 4. Average population activity. Neural activity averaged across all recorded neurons in (a) the familiar and (b) the novel object experiments after subtraction of baseline activity. The vertical black bars represent the boundaries of the sensory (S, 80–450 ms), intermediate (I, 450–1150 ms) and reactivation (R, 1150–1650 ms) periods. Sample presentation occurred at 0 ms, the delay period started at 650 ms and ended at 1650 ms. Green curves represent activity to objects, and red curves represent activity to noise patterns. The number of neurons contributing to the graphs is shown on the upper left of each panel. Error bars represent standard deviations of mean activity to all five objects or all four noise patterns across the population, and are shown for illustrative purposes. Bin width, 50 ms.

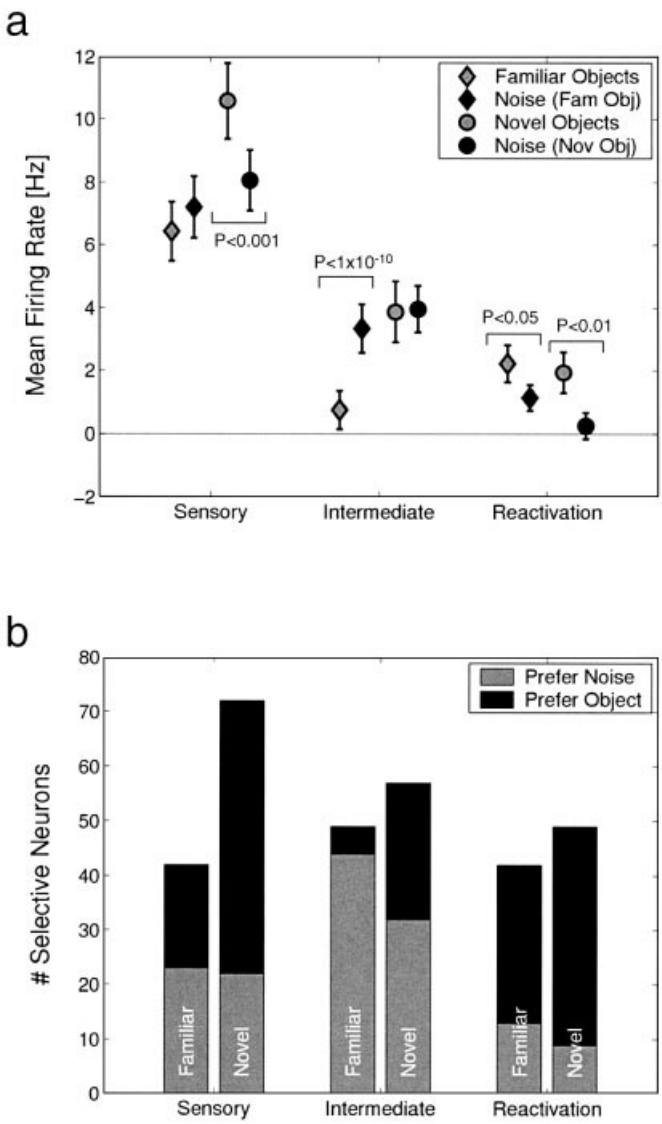

FIG. 5. Neural preference for objects or noise patterns. (a) Mean activity to objects and noise patterns is shown for each of the experiments in the three task periods for the entire population of neurons (familiar objects $n = 164$, novel objects $n = 160$) after subtraction of baseline activity. Standard errors (SEM) are shown for illustrative purposes (data are generally not distributed normally). A W-test ($P < 0.01$) revealed comparisons between objects and noise patterns that reached statistical significance. P-values are shown adjacent to the associated pair of datapoints. The abbreviations 'fam obj' and 'nov obj' refer to the familiar and the novel object experiments, respectively. (b) Each bar represents the number of neurons, a subset of the total population in each experiment that showed significant differences between objects and noise patterns during each of the three periods (W-tests, $P < 0.01$). The grey portion of each bar represents the fraction of neurons that preferred noise patterns (i.e. responded more on average to the four noise patterns than to the five objects), and the black portion represents the fraction of neurons that preferred objects.

extensive training with this particular set of objects for at least ten training sessions conducted prior to the experiments described here. A different set of familiar objects was used for each of the two monkeys. When familiar objects were used, noise patterns were always novel every session (i.e. we generated four new noise patterns for each session). During 'novel object' sessions, monkeys were presented with a new set of five objects, which they had never seen before. This set of novel objects did not change throughout the session. Thus, monkeys saw several repetitions of these novel objects during course of the session, but they did not have extensive prior experience with them, as was the case for familiar objects. For novel object sessions, noise patterns were kept constant across days such that the same four noise patterns were presented repeatedly while objects were changing across days.

Recording technique

A scleral search coil Robinson (1963), head restraint and recording chamber were implanted under aseptic conditions while the animals

were anaesthetized using isoflurane. Postoperatively, the animals received analgesics and antibiotics and were kept alive for participation in further experiments. During recording sessions, monkeys were seated in primate chairs within sound-attenuating enclosures (Crist Instruments, Damascus MD, USA). Their heads were restrained, and a juice spout was placed near their mouth for automated delivery of reward (apple juice). For extracellular recordings of action potentials, we employed a grid system (Crist Instruments, Damascus MD, USA) with custom-made modifications that allowed us to use eight tungsten electrodes (FHC instruments, Bowdoin ME, USA) simultaneously. Penetrations were made perpendicular to the surface of the skull, and the minimum separation between adjacent electrodes was 1 mm. Recording sites near the principal sulcus of the lateral prefrontal (PF) cortex were localized using magnetic resonance imaging (see Fig. 3). We did not screen neurons for involvement in the task, but instead advanced the electrodes until the activity of one or more neurons was well isolated. After a suitable wait period of 1–2 h, we then commenced recording. This was performed to ensure an unbiased estimate of PF neural activity. Due to the number of conditions required and the limitations on the number of trials a monkey can work on a given day, it was not possible to complete both the familiar and the novel object experiment during a single recording session. However, care was taken to record neurons at similar locations in the two experiments in each of the monkeys, and at similar recording depths. Monkeys completed an average of 865 trials during 25 recording sessions (familiar objects, 14 sessions; novel objects, 11 sessions), resulting in on average over 20 repetitions for each of the five objects and about 30 repetitions for each of the four noise patterns. Analyses were conducted on data from all attempted trials (both correct and incorrect), excluding only trials on which the monkey broke fixation or failed to respond at all. We did this because behavioural choice occurred only upon test object presentation, which happened after the trial period that we analysed. Using all attempted trials yielded about

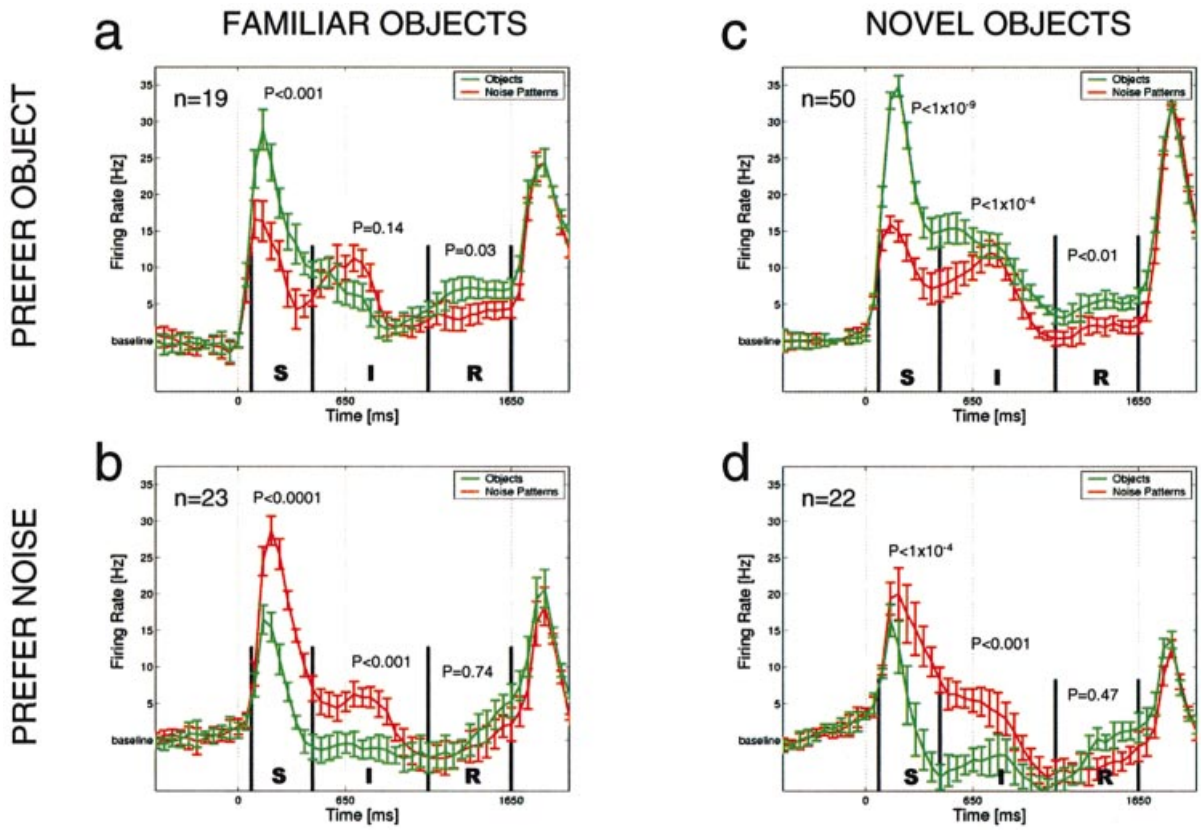

FIG. 6. Temporal dynamics of neural activity sorted by sample period preference. For each experiment, we selected neurons that during the sample period (S) showed statistically significant (*W*-test, $P < 0.01$) preference for (a and b) objects vs. noise patterns and (c and d) vice versa. Error bars represent standard deviations of the mean responses to objects or noise patterns across the population. The number of neurons contributing to each panel is shown in the upper left. Black vertical lines represent the boundaries of the sensory (S), intermediate (I) and reactivation (R) periods. The *P*-values represent significance of paired Wilcoxon tests comparing average activity to objects against average activity to noise patterns in each of these periods, for the neurons contributing to each graph.

FIG. 7. Single neuron responses to objects and noise patterns. Four single neuron examples from the familiar object experiment (a–d), and four single neuron examples from the novel object experiment (e–h) are shown. Responses to each of the objects and noise patterns are shown separately, with red curves representing the five objects and green curves representing the four noise patterns. Neural activity is shown relative to baseline firing rate. The symbols in the upper left part of each panel refer to the scatterplots showing activity for all neurons in Fig. 8. Black vertical lines represent the boundaries of the sensory (S), intermediate (I) and reactivation (R) periods. The top two rows represent single neuron examples from monkey A; the bottom two rows represent single neurons from monkey B. The stars drawn beside the period identity symbols (S, I or R) denote significant differences between average activity to objects relative to noise (*W*-tests, $P < 0.01$).

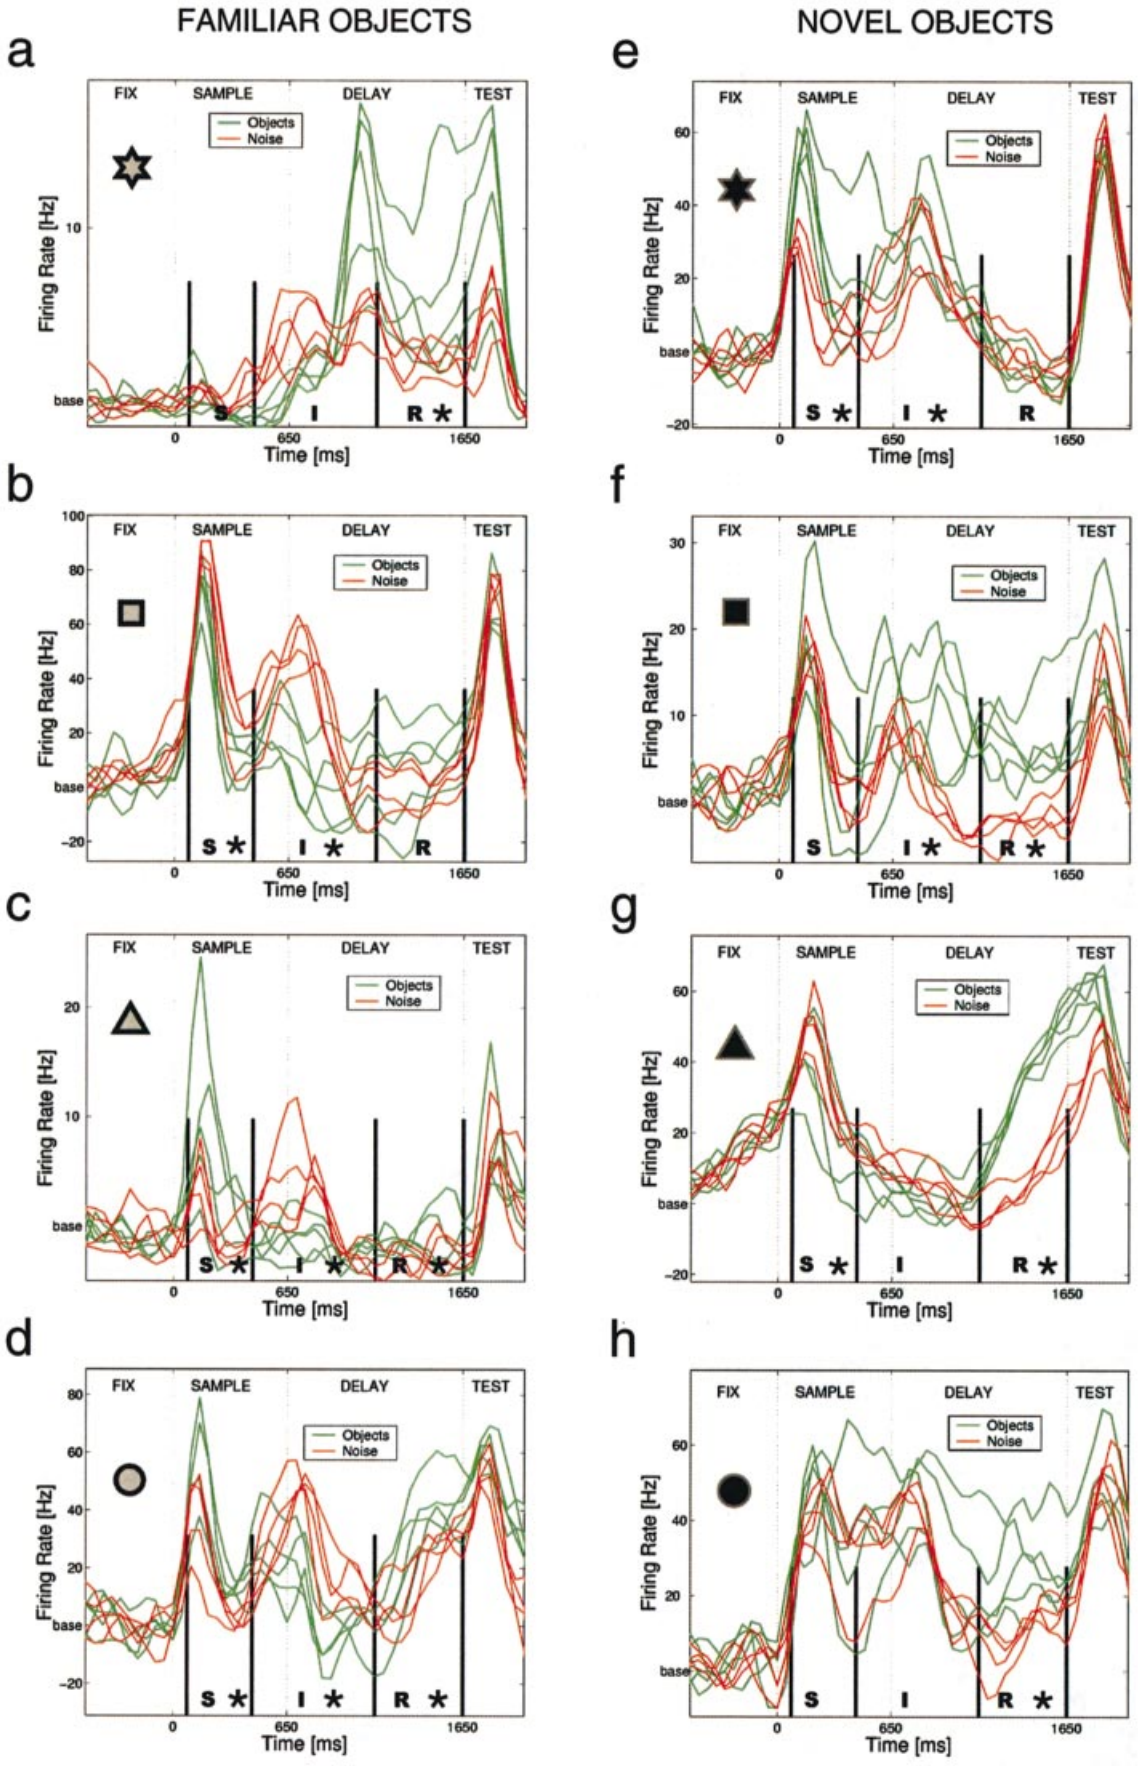
FAMILIAR OBJECTS
NOVEL OBJECTS
a
b
c
d
e
f
g
h
FIX
SAMPLE
DELAY
TEST
Objects
Noise
Firing Rate [Hz]
base
S
I
R
0
650
1650
Time [ms]

twice the data for the noise pattern condition, allowing more accurate quantification of neural activity. However, our results did not depend on this choice; we repeated the major analyses performed in the present study using correct trials only, which yielded very similar results to the findings presented here.

Data analysis

Neural activity was analysed relative to a baseline of activity during the fixation period prior to sample object onset. Baseline activity was assessed during a period lasting 800 ms, starting 200 ms after onset of fixation (to exclude visual transients related to acquisition of fixation) and ending at sample onset. This baseline activity was subtracted from task-related activity for all the analyses in this study. This was performed to facilitate averaging activity for neurons with different baseline firing rates. Baseline firing rates were similar in the two experiments (familiar objects: median, 10.9Hz; range, 0.1Hz-39.2Hz; novel objects: median, 12.9Hz; range, 0.5Hz-37.1Hz).

Statistical tests

To evaluate whether there were significant differences in firing rate, to objects and noise patterns in the different task periods, we used the Wilcoxon signed rank test for equality of medians. Unlike the *t*-test, the Wilcoxon test makes no assumptions about the underlying distributions. To evaluate whether there was a significant difference in average activity between objects and noise patterns, we first computed average firing rates across all objects and all noise patterns for each neuron. Then we performed a matched pairs Wilcoxon signed rank test, assessing significance of the differences between these two values across the population. When assessing whether single neurons showed significant differences between objects and noise patterns, we grouped trials into two groups – the first one comprised of trials with one of the five objects as sample, and the second one comprised of trials with one of the four noise patterns as sample. We then performed an unmatched pairs (as object and noise pattern data is generated on different trials) Wilcoxon test, examining whether there were significant systematic differences in neural activity between these two groups of trials. In this text, we refer to these statistical tests as *W*-tests. Tests were performed on average activity during the three task periods: sensory (S), from 80 to 450 ms after sample onset; intermediate (I), from 450 to 1150 ms after sample onset and reactivation (R), from 1150 to 1650 ms after sample onset. For the major statistical analyses, *P*-values are given for the entire population as well as for each of the monkeys separately to demonstrate consistency of the results across the two animals.

To examine whether an observed fraction of neurons preferring objects to noise patterns was significantly different from even proportions, we performed a χ^2-test comparing the observed ratio to an even split. For example, for a case where 50 neurons preferred objects and 22 preferred noise patterns, we tested whether this [50 : 22] was a significantly different distribution from 36 : 36. When comparing fractions of selective neurons between the two experiments, we used a standard χ^2-test examining equality of distributions.

Results

Behaviour

Behavioural performance is summarized in Fig. 2. We found that monkey's behavioural performance was near ceiling when objects were used as sample stimuli in both the familiar and the novel object experiment (95% and 92% correct, respectively; *t*-test, $P > 0.1$). Thus, they were able to perform the DMS task equally well with novel objects as with highly familiar objects. However, when noise patterns were used as sample stimuli, monkeys' performance was near chance level of 50% correct in both experiments, because noise patterns did not contain any task-relevant information.

Neural activity

We recorded neural activity from 164 neurons in the familiar object experiment (monkey A, $n = 79$; monkey B, $n = 85$), and 160 neurons in the novel object experiment (monkey A, $n = 104$; monkey B, $n = 56$). Sites where neurons showing significant differences in activity in response to objects and noise patterns relative to baseline firing rate are shown in Fig. 3 for the two monkeys separately. As can be seen, most neurons were found near, as well as ventral to the principal sulcus.

Mean population response

The mean response of the entire unscreened neural population for each of the two experiments relative to baseline firing rate is shown in Fig. 4. Taking the experiments together, three distinct periods were evident in these average response histograms. Stimulus presentation evoked a transient sensory visual response (S), lasting from 80 to 450 ms after sample onset. This was followed by an intermediate period (I), lasting from about 450–1150 ms after sample onset. Finally, after a brief return to baseline activity, there was a reactivation period (R) during the last 500 ms of the delay characterized by an increasing 'climbing' activity profile. We observed several trends in these population histograms, which are quantified (paired Wilcoxon test, see Materials and methods) in Fig. 5a. (i) There was great similarity between neural responses to noise patterns in the two experiments, despite the fact that noise patterns were actually novel each day in the familiar object experiment and kept constant across days in the novel object experiment (*W*-tests, $P > 0.1$). Thus, familiarity of the noise did not affect the neural response. Accordingly, we henceforth use the term noise patterns regardless of whether they were employed in the familiar or novel object experiment. (ii) While during the sensory (S) period familiar objects elicited similar average activity as noise (*W*-test, $P > 0.1$; monkey A, $P = 0.06$; monkey B, $P > 0.1$), activity was greater in response to novel objects compared to noise patterns (*W*-test, $P < 0.001$; monkey A, $P = 0.047$; monkey B, $P = 0.007$). (iii) During the intermediate (I) period there was no difference in activity between novel objects and noise (*W*-test, $P > 0.1$; monkeys A and B, $P > 0.1$), while there was significantly less activity on average to familiar objects than to noise patterns (*W*-test, $P < 1 \times 10^{-10}$; monkey A, $P < 1 \times 10^{-6}$; monkey B, $P < 1 \times 10^{-5}$). (iv) During the reactivation (R) period, both familiar (*W*-test, $P < 0.05$; monkey A, $P = 0.052$; monkey B, $P < 1 \times 10^{-5}$) and novel (*W*-test, $P < 0.01$; monkey A, $P > 0.1$; monkey B, $P < 1 \times 10^{-5}$) objects elicited higher average activity than noise patterns, although this effect reached significance in only one animal.

These trends were confirmed by examining the number of neurons showing significant differences in activity between objects and noise patterns. In Fig. 5b we summarize for how many neurons this difference was significant in each of the task periods for both experiments. Significance was assessed using an unpaired Wilcoxon test (see Materials and methods). Comparing the familiar and novel object experiments, similar numbers of neurons showed significant differences in activity to objects and noise patterns during the intermediate and reactivation periods (χ^2-test, $P > 0.1$), whereas, during the sensory period, more neurons showed such differences in

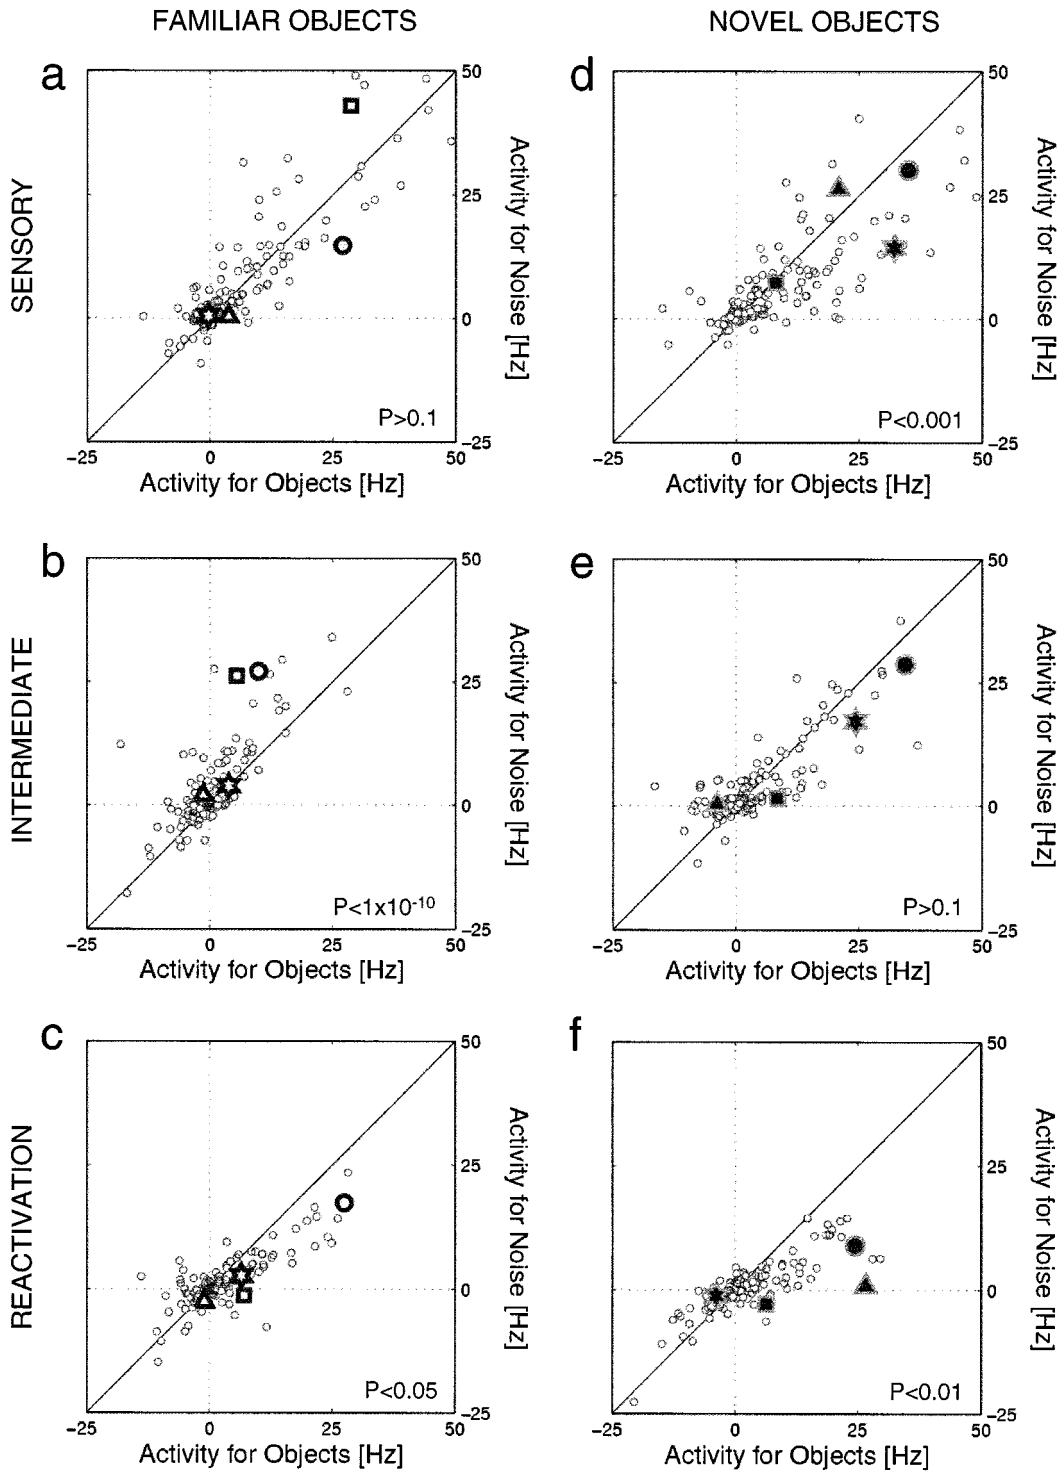

FIG. 8. Distribution of neural responses to objects and noise patterns. The columns display data from the familiar and novel object experiments. Each panel summarizes the response of the entire population of single neurons to objects and noise patterns during the sensory, intermediate and reactivation periods (familiar objects, $n = 164$; novel objects, $n = 160$). Each circle represents a single neuron, and large symbols represent the single neuron examples from Fig. 7. Neural activity is plotted relative to baseline response, such that positive values depict increases from the activity during the fixation period and negative values depict decreases. The significance level of a Wilcoxon test comparing mean activity between objects and noise patterns is shown at the bottom right of each panel for reference (same values as in Fig. 5a).

the novel object experiment (χ^2-test, $P < 0.001$). During the sensory period, preference for objects vs. noise patterns was evenly distributed in the familiar object experiment (χ^2-test, $P > 0.1$), while more neurons preferred novel objects to noise patterns (χ^2-test, $P < 0.05$). This indicates that although there was no systematic difference between activity to familiar objects and noise patterns during the

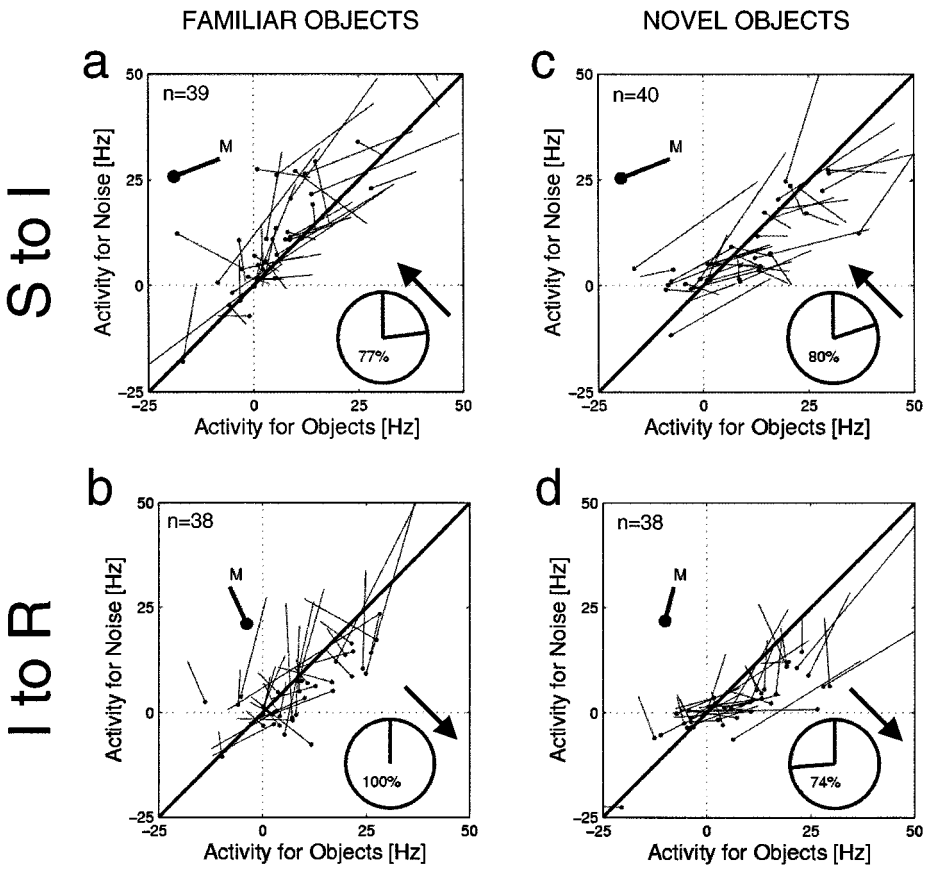

FIG. 9. Shifts in object preference. The vectors represent the shift in activity preference (a and c) from the sensory to the intermediate period (S→I) and (b and d) from the intermediate to the reactivation period (I→R), where the dots represent vector end points. Each vector represents a single neuron, and its' vertices represent that neuron's mean response to objects and noise patterns in the source and the target periods. For example, the vectors in Fig. 9a connect the vertices corresponding to each contributing neuron in Fig. 8a and Fig. 8b. Only neurons that showed shifts in preference between objects and noise patterns larger than 7 Hz are included in this analysis, i.e., neurons with a projection onto the bold vector beside the pie chart larger than 7Hz in length. The number of cases is shown in the upper left. The bold vector labelled 'M' represents the average shift for the population. The pie chart reports the fraction of neurons for which the vectors pointed in the same direction as the bold vector shown above the pie charts, which represents a systematic shift towards noise preference in (a and c), and a systematic shift towards object preference in (b and d).

sensory period, many single neurons actually did distinguish between familiar objects and noise. These subpopulations are analysed separately in the next section. During the intermediate period, a similar number of neurons preferred novel objects or noise patterns (χ^2-test, $P > 0.1$), while many more neurons showed greater activity to noise patterns than to familiar objects (χ^2-test, $P < 0.0001$). Finally, during reactivation, there was a trend for preference for objects over noise patterns in both experiments that reached significance only for novel objects (familiar objects, χ^2-test, $P = 0.12$; novel objects, χ^2-test, $P < 0.01$).

Mean response of neurons preferring objects or noise patterns

To examine the timecourse of the neural response in more detail, we plotted separately the temporal development of average activity for neurons that showed significant preference for objects over noise patterns and vice versa in each of the experiments. In Fig. 6a we show average activity for neurons that preferred familiar objects to noise patterns. Although these neurons were selected to show greater activity for objects during the sensory period, this activity preference was not maintained through the intermediate period but only reappeared during the reactivation period towards the end of the delay. During the intermediate period, neural activity was similar for objects and noise patterns (*W*-test, $P = 0.14$). For neurons that preferred novel objects to noise during the sensory period (Fig. 6b), this sensory preference was prolonged through the intermediate period and was also present during reactivation. Noise-preferring neurons showed similar trends in both experiments, as shown in Fig. 6c and d; preference for noise patterns during the sensory period was maintained during the intermediate period. During reactivation, these neurons did not show significant differences in response to objects and noise patterns (*W*-tests, $P > 0.1$).

Single neuron examples for familiar and novel objects

The above population results suggest that the temporal dynamics present in the average population do not appear to be a result of averaging together distinct neural populations, but rather might characterize response profiles of single neurons. Indeed, the development of neural activity over time seen in many single neurons resembled that of the average population. Single neuron examples for the familiar object experiment are provided in Fig. 7a–d, where each object and noise pattern is shown separately to allow assessment of the variability of neural activity among, as well as between, objects and noise patterns. In general, neurons were active during all three task periods, although Fig. 7a provides an example neuron that did not participate in sensory period processing. Consistent with the trends observed in the population, some single neurons preferred objects (e.g. Figure 7c), while others preferred noise patterns (Fig. 7b) during the sensory period. During the intermediate period, these neurons tended to prefer noise patterns, whereas, during reactivation, they tended to fire more vigorously to objects. Four examples for the novel object experiment are shown in Fig. 7e–h. During the sensory period, most neurons preferred objects (Fig. 7e and g) or showed no consistent preference (Fig. 7f and h). There was however, no consistent trend during the intermediate period, while there was systematic preference for objects during the reactivation period.

Distribution of neural preferences for objects or noise patterns

Scatterplots showing mean activity to objects and noise patterns for all neurons recorded in each experiment are shown in Fig. 8. The large symbols depict the single neuron examples shown in Fig. 7, demonstrating that this is a representative subset of the entire

population. These plots summarize the activity of all recorded neurons and demonstrate directly the main findings of this study. During the sensory period, many neurons showed transient visual responses in both experiments. While there was no consistent preference for familiar objects over noise patterns, neurons did tend to prefer novel objects relative to noise. During the intermediate period by contrast, there was no consistent trend for novel objects, but familiar objects led to markedly less activity than noise patterns. During reactivation, neurons in both experiments again showed preference for objects.

Using these population scatterplots, we assessed directly how relative preference between objects and noise changed between the three task periods. We first identified neurons that showed shifts in preference between objects and noise patterns [e.g. from the sensory (S) to the intermediate (I) period] greater than an arbitrary threshold of 7 Hz. We did this to eliminate neurons that showed little or no preference shift. The vector shifts for these neurons are shown in Fig. 9a for familiar objects. The mean vector points to the left and slightly downward, consistent with a general decrease in activity from S to I, and a shift towards a preference for noise patterns over objects. Indeed, the majority (77%) of neurons shifted towards noise preference. Interestingly, the same was true in the novel object experiment (Fig. 9c), both in terms of the mean vector shift and the proportion of neurons shifting towards preference for noise. Analysis of the shift between the intermediate (I) and the reactivation (R) periods revealed a strikingly different result. For familiar objects (Fig. 9b), all (100%) of the shifting neurons shifted towards object preference, and the direction of the mean vector suggests that there was little change in mean firing rate. The results were similar for novel objects (Fig. 9d), although shifts towards object preference were apparent for 74% of neurons and the mean shift vector revealed a modest decrease in mean activity from I to R.

Discussion

In this study, we examine the timecourse of the activity of an ensemble of PF neurons in a delayed-matching-to-sample (DMS) task. We compare population activity evoked by objects, which needed to be retained during a short delay, to activity evoked by noise patterns that contained no task-relevant information and did not need to be retained. To assess effects of experience on this timecourse, we conducted the experiment with novel and highly familiar objects.

Three distinct periods were evident in the average population timecourse. During an initial visual sensory period, novel objects elicited greater activity than noise patterns while there was no systematic difference between activity to familiar objects and noise patterns. This finding is not surprising, as familiarity tends to cause a decrease in neural activity in many primate brain regions including the inferior temporal (Li *et al*., 1993), the perirhinal (Fahy *et al*., 1993) and the prefrontal (Asaad *et al*., 1998) cortices, as well as the hippocampus (Cahusac *et al*., 1993) and amygdala (Wilson & Rolls, 1993). Indeed, results from a selective lesion disconnecting the frontal from the temporal lobe suggest that the preference for novel objects in PF neurons may be a consequence of feed-forward activity from temporal cortical areas (Parker & Gaffan, 1998; Parker *et al*., 1998).

During an intermediate period, there was elevated activity for both noise patterns and novel objects, but not for familiar objects. In fact, even PF neurons that preferred familiar objects vs. noise during the sensory period did not maintain this preference in the form of elevated activity during the intermediate period (Fig. 6a). Intermediate period activity could not be ascribed to the offset of the visual stimulus, as it began well before the stimulus was turned off (see Fig. 4a). Neither was it a result of spike-frequency dependent adaptation of neural activity, as novel objects elicited the largest transient visual response but also showed robust activity during the intermediate period (Fig. 4b). Object information needed to be maintained through the intermediate period for both familiar and novel objects, yet we found sustained elevated activity through this period only for novel objects and noise patterns that did not need to be maintained (Fig. 7h), but not for familiar objects. Thus, maintenance of sensory information as elevated activity does not characterize the intermediate period well. This suggests that some other type of processing might occur during the intermediate period; but what might it be? One possibility is that intermediate period activity provides a signal that is instrumental in learning and causes changes in prefrontal or related cortical networks. This signal may be absent for familiar objects because no further leaning is required, whereas it may be evoked both by novel objects, which have not been overlearned, as well as noise patterns. Perhaps experience leads to the formation of inhibitory circuits within the PF cortex, which silence learning-related processing when a highly familiar input pattern is detected. Another possibility is that a dopaminergic or cholinergic signal up-regulates the excitability of PF neurons and thus produces more spiking during the intermediate period for novel stimuli or noise patterns but not for highly familiar over-learned patterns. dopamine, by causing an elevation in NMDA-current, and acetylcholine by causing in increase in the voltage activated calcium current can both lead to transient increases in spike generation and could thus underlie the additional activity seen in the intermediate period for novel stimuli or for the ambiguous noise patterns. Consistent with this hypothesis, robust learning-related differences have been described in the dopamine neurons which project to wide cortical target regions including the prefrontal cortex (Schultz *et al*., 1993; Schultz *et al*., 1997).

A reactivation period occurred in the late delay during which activity was elevated for both novel and familiar objects relative to noise patterns. It is activity during this reactivation period that resembles what is classically known as delay or persistent activity (Goldman-Rakic, 1990). For example, information about the sample object was maintained as elevated activity, and in many neurons activity also showed an increasing trend as the end of the delay approached. Previous work has implicated such PF climbing activity in anticipatory coding for objects (Rainer *et al*., 1999), reward (Watanabe, 1996) as well as motor preparation (Bruce & Goldberg, 1985). As in the present study, over 90% of trials with object samples were rewarded, compared to about 50% for trials with noise pattern samples (Fig. 2), it is possible that differential anticipatory coding for reward played a role during the reactivation period. Note that reward asymmetry between objects and noise patterns cannot explain results during the intermediate period, however, as familiar and novel objects were rewarded equally. The example neuron shown in Fig. 7g exhibited reactivation period activity consistent with reward expectancy, in that it was systematically more active on object-sample trials than on noise-sample trials. More generally though, activity during the reactivation period was modulated differentially according to the identity of the sample object, consistent with a role in short-term memory for objects, as in the example neurons shown in Fig. 7a, b, f and h. Motor preparation is unlikely to have played much of a role in the reactivation period, as monkeys could decide whether to release the lever only when presented with the test object after the delay.

The shift analyses presented in Fig. 9 demonstrate most directly the dynamics in preference for objects vs. noise patterns during the

three task periods. From the sensory to the intermediate period, there was a similar shift for the two neural populations towards noise-preference (Fig. 9a and b). For familiar objects, this shift resulted in a marked and highly significant preference for noise during the intermediate period (Fig. 8b). For novel objects, however, preference for objects vs. noise patterns was distributed evenly during the intermediate period (Fig. 8e), despite this similar vector shift. The reason for this is the difference in the sensory period distributions in the two experiments. In the novel object experiment, the vector shift towards noise-preference abolishes the ensemble preference for objects evident during the sensory period. For familiar objects, evenly distributed ensemble preference during the sensory period is shifted towards noise-preference. The shifts from the intermediate to the reactivation period were, again, quite similar for the two experiments (Fig. 9b and d), yielding distributions that were biased significantly towards object preference. Thus, the population dynamics were actually quite similar between the two experiments, the major differences being a result for neural population preference for objects during the novel object experiment.

Taken together, the present results indicate that prefrontal delay activity – at least as assessed in the present study – is more complex than simple maintenance, and is subject to experience-dependent dynamics. Computational models of delay activity need to be extended to capture such dynamics, and provide more realistic accounts of activity profiles in different tasks to gain further insight into how the prefrontal cortex contributes to memory storage and manipulation.

Acknowledgements

We thank N. Brunel, D. Durstewitz, S. Fusi, B. Gutkin and X.J. Wang for comments on the manuscript. This work was supported by the RIKEN-MIT Neuroscience Research Center and by the National Institute of Neurological Disorders and Strokes.

Abbreviations

AMPA, α-amino-3-hydroxy-5-methyl-4-isoxazole propionate; PF, prefrontal; NMDA, *N*-methyl D-aspartate.

References

Amit, D.J. & Brunel, N. (1997) Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex*, **7**, 237–252.

Amit, D.J., Fusi, S. & Yakovlev, V. (1997) Paradigmatic working memory (attractor) cell in IT cortex. *Neural Comput.*, **9**, 1071–1092.

Asaad, W.F., Rainer, G. & Miller, E.K. (1998) Neural activity in the primate prefrontal cortex during associative learning. *Neuron*, **21**, 1399–1407.

Asaad, W.F., Rainer, G. & Miller, E.K. (2000) Task-specific neural activity in the primate prefrontal cortex. *J. Neurophysiol.*, **84**, 451–459.

Bichot, N.P., Schall, J.D. & Thompson, K.G. (1996) Visual feature selectivity in frontal eye fields induced by experience in mature macaques. *Nature*, **381**, 697–699.

Bruce, C.J. & Goldberg, M.E. (1985) Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.*, **53**, 603–635.

Cahusac, P.M., Rolls, E.T., Miyashita, Y. & Niki, H. (1993) Modification of the responses of hippocampal neurons in the monkey during the learning of a conditional spatial response task. *Hippocampus*, **3**, 29–42.

Compte, A., Brunel, N., Goldman-Rakic, P.S. & Wang, X.J. (2000) Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex*, **10**, 910–923.

Durstewitz, D., Seamans, J.K. & Sejnowski, T.J. (2000a) Neurocomputational models of working memory. *Nature Neurosci*, **3**, 1184–1191.

Durstewitz, D., Seamans, J.K. & Sejnowski, T.J. (2000b) Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. *J. Neurophysiol.*, **83**, 1733–1750.

Fahy, F.L., Riches, I.P. & Brown, M.W. (1993) Neuronal activity related to visual recognition memory: long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex. *Exp. Brain Res.*, **96**, 457–472.

Field, D.J. (1987) Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am.[a]*, **4**, 2379–2394.

Funahashi, S., Bruce, C.J. & Goldman-Rakic, P.S. (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.*, **61**, 331–349.

Funahashi, S., Chafee, M.V. & Goldman-Rakic, P.S. (1993) Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature*, **365**, 753–756.

Fuster, J.M. (1993) Frontal lobes. *Curr. Opin. Neurobiol.*, **3**, 160–165.

Fuster, J.M. & Alexander, G.E. (1971) Neuron activity related to short-term memory. *Science*, **173**, 652–654.

Goldman-Rakic, P.S. (1990) Cellular and circuit basis of working memory in prefrontal cortex of nonhuman primates. *Prog. Brain Res.*, **85**, 325–326.

Goldman-Rakic, P.S. (1995) Cellular basis of working memory. *Neuron*, **14**, 477–485.

Hasegawa, R.P., Matsumoto, M. & Mikami, A. (2000) Search target selection in monkey prefrontal cortex. *J. Neurophysiol.*, **84**, 1692–1696.

Kubota, K. & Niki, H. (1971) Prefrontal cortical unit activity and delayed alternation performance in monkeys. *J. Neurophysiol.*, **34**, 337–347.

Laing, C.R. & Chow, C.C. (2001) Stationary bumps in networks of spiking neurons. *Neural Comput.*, **13**, 1473–1494.

Li, L., Miller, E.K. & Desimone, R. (1993) The representation of stimulus familiarity in anterior inferior temporal cortex. *J. Neurophysiol.*, **69**, 1918–1929.

Miller, E.K. & Cohen, J.D. (2001) An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.*, **24**, 167–202.

Miller, E.K., Erickson, C.A. & Desimone, R. (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J. Neurosci.*, **16**, 5154–5167.

Parker, A. & Gaffan, D. (1998) Memory after frontal/temporal disconnection in monkeys: conditional and non-conditional tasks, unilateral and bilateral frontal lesions. *Neuropsychologia*, **36**, 259–271.

Parker, A., Wilding, E. & Akerman, C. (1998) The Von Restorff effect in visual object recognition memory in humans and monkeys. The role of frontal/perirhinal interaction. *J. Cogn. Neurosci.*, **10**, 691–703.

di Pellegrino, G. & Wise, S.P. (1993) Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. *J. Neurosci.*, **13**, 1227–1243.

Petrides, M. (1985) Deficits on conditional associative-learning tasks after frontal- and temporal-lobe lesions in man. *Neuropsychologia*, **23**, 601–614.

Rainer, G., Asaad, W.F. & Miller, E.K. (1998a) Memory fields of neurons in the primate prefrontal cortex. *Proc. Natl Acad. Sci. USA*, **95**, 15008–15013.

Rainer, G., Asaad, W.F. & Miller, E.K. (1998b) Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature*, **393**, 577–579.

Rainer, G. & Miller, E.K. (2000) Effects of visual experience on the representation of objects in the prefrontal cortex. *Neuron*, **27**, 179–189.

Rainer, G., Rao, S.C. & Miller, E.K. (1999) Prospective coding for objects in primate prefrontal cortex. *J. Neurosci.*, **19**, 5493–5505.

Schultz, W., Apicella, P., Ljungberg, T., Romo, R. & Scarnati, E. (1993) Reward-related activity in the monkey striatum and substantia nigra. *Prog. Brain Res.*, **99**, 227–235.

Schultz, W., Dayan, P. & Montague, P.R. (1997) A neural substrate of prediction and reward. *Science*, **275**, 1593–1599.

Wang, X.J. (1999) Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J. Neurosci.*, **19**, 9587–9603.

Wang, X.J. (2001) Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci.*, **24**, 455–463.

Watanabe, M. (1986) Prefrontal unit activity during delayed conditional Go/No-Go discrimination in the monkey. II. Relation to Go and No-Go responses. *Brain Res.*, **382**, 15–27.

Watanabe, M. (1996) Reward expectancy in primate prefrontal neurons. *Nature*, **382**, 629–632.

White, I.M. & Wise, S.P. (1999) Rule-dependent neuronal activity in the prefrontal cortex. *Exp. Brain Res.*, **126**, 315–335.

Wilson, F.A. & Rolls, E.T. (1993) The effects of stimulus novelty and familiarity on neuronal activity in the amygdala of monkeys performing recognition memory tasks. *Exp. Brain Res.*, **93**, 367–382.