



massachusetts institute of technology — computer science and artificial intelligence laboratory

Selectivity of Local Field Potentials in Macaque Inferior Temporal Cortex

Gabriel Kreiman, Chou Hung, Tomaso Poggio and James DiCarlo

AI Memo 2004-020
CBCL Memo 240

September 2004

Abstract

While single neurons in inferior temporal (IT) cortex show differential responses to distinct complex stimuli, little is known about the responses of populations of neurons in IT. We recorded single electrode data, including multi-unit activity (MUA) and local field potentials (LFP), from 618 sites in the inferior temporal cortex of macaque monkeys while the animals passively viewed 78 different pictures of complex stimuli. The LFPs were obtained by low-pass filtering the extracellular electrophysiological signal with a corner frequency of 300 Hz. As reported previously, we observed that spike counts from MUA showed selectivity for some of the pictures. Strikingly, the LFP data, which is thought to constitute an average over large numbers of neurons, also showed significantly selective responses. The LFP responses were less selective than the MUA responses both in terms of the proportion of selective sites as well as in the selectivity of each site. We observed that there was only little overlap between the selectivity of MUA and LFP recordings from the same electrode. To assess the spatial organization of selective responses, we compared the selectivity of nearby sites recorded along the same penetration and sites recorded from different penetrations. We observed that MUA selectivity was correlated on spatial scales up to 800 μm while the LFP selectivity was correlated over a larger spatial extent, with significant correlations between sites separated by several mm. Our data support the idea that there is some topographical arrangement to the organization of selectivity in inferior temporal cortex and that this organization may be relevant for the representation of object identity in IT.

Note: Gabriel Kreiman and Chou Hung contributed equally to this work

This report describes research done within the Center for Biological and Computational Learning in the Department of Brain and Cognitive Sciences and at the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. This research is sponsored by DARPA, ONR and a Whiteman fellowship to G.K.

Introduction

Evidence from recordings of spiking activity in macaque inferior temporal (IT) cortex suggests that neurons change their firing rates in response to complex visual stimuli which include faces but also objects and arbitrary shapes (Schwartz et al., 1983), reviewed in (Logothetis and Sheinberg, 1996; Tanaka, 1993). The fundamental role of IT in object recognition is supported by evidence that IT lesions impair the monkey's performance in object recognition tasks (Holmes and Gross, 1984). Selective responses in IT may arise from non-linear interactions of input from a hierarchical cortical ventral visual system (Riesenhuber and Poggio, 1999). Little is known about the types of features encoded by IT neurons and the neuronal coding mechanisms in IT.

One potential clue to understanding feature selectivity in IT may come from knowledge of the preferences of nearby neurons. Topographical maps are prevalent throughout the brain in multiple cortical areas including visual areas V1 and MT, auditory area A1, somatosensory cortex and motor cortex (Kandel et al., 2000). In IT, Fujita and colleagues made the interesting observation that single cells recorded along a penetration presumably normal to IT cortex tend to respond best to similar, but not identical objects (Fujita et al., 1992; Tanaka, 2003). Furthermore, data simultaneously recorded from multiple neurons also shows that nearby neurons in IT may show similar selectivity (Gochin et al., 1991). Based on these measurements, it has been suggested that there is a topographical columnar organization for feature selectivity in IT, with column size on the order of 200 to 500 μm (Fujita et al., 1992).

Given the proposed columnar organization of responses in IT, it is possible that other coarser measures of neural activity involving larger ensembles of neurons could show selective responses to complex visual stimuli. Interestingly, several investigators have reported selectivity using other tools that show lower spatial and/or temporal resolution including optical imaging in macaque IT cortex (Wang et al., 1996), functional magnetic resonance imaging (fMRI) data in the human fusiform gyrus (see for example (Kanwisher et al., 1997), and evoked response potential observations in human epileptic patients (Allison et al., 1994). One possible interpretation of these observations is that a topographical organization underlies the selectivity observed at the level of neuronal ensembles.

Local field potentials (LFP) constitute another electrophysiological measurement that is thought to arise from the combined activity of large numbers of neurons (Mitzdorf, 1985) and date back at least to the 1930s (Berger, 1930, as quoted in Hobson, 1995). The main component of the LFP signal in cortical structures seems to be the excitatory post-synaptic potentials at the dendritic level (Mitzdorf, 1985). The LFP signal shows higher correlation than spiking activity with the BOLD measurements used in fMRI experiments (Logothetis, 2003).

Given the spatial and temporal smoothing of the LFP signal, it is unclear whether these potentials would show any image selective responses in IT. Furthermore, if the selectivity of spike responses relies on complex temporal dynamics, the averaged response in the form of LFPs may fail to show selective responses. To examine this issue, we recorded MUA and LFP responses to visual images at many sites within IT cortex. We observed that the selective spiking responses from nearby recording sites were correlated as previously reported (Fujita et al., 1992; Gochin et al., 1991). We observed that at least one fourth of the recorded sites showed LFP responses that were significantly selective to images within our object set. The LFP responses were only weakly correlated to the MUA responses. Data from spiking neurons as well as data from LFP recordings could be used to build a classifier to decode neuronal responses.

These results provide evidence that LFPs in IT are selective to complex stimuli. Furthermore, our observations support the notion that there may be a topographical representation of visual information in IT, thus suggesting that it is of interest to study in further detail the properties of responses from ensembles of neurons. Finally, our data has important implications for the interpretation of the large body of functional imaging data.

1 Results

We recorded extracellularly from a total of 618 sites in 133 penetrations in the inferior temporal cortex of two macaque monkeys (512 sites in 107 penetrations and 106 sites in 26 penetrations respectively) while the animals passively viewed 78 images of complex shapes (the images are shown in Figure 4). The images, including human and monkey faces, toys, foodstuffs, vehicles, cats and white

boxes, were presented at the center of gaze for 100 ms and there was a 100 ms interval between pictures. We recorded multi-unit activity (MUA) by high-pass filtering the signal with a corner frequency of 400 Hz and thresholding. We also recorded local field potentials¹ (LFPs) by low-pass filtering the extracellular electrophysiological signal at 300 Hz and considering the [1;300] Hz interval (see Methods).

1.1 **Selectivity of spikes and local field potentials**

We observed that spike counts in windows of 100 or 200 ms from MUA recordings showed selectivity for some of the pictures that were presented to the monkey. This is in accordance with previous reports that date back to IT recordings from Gross and colleagues (Gross et al., 1972; Schwartz et al., 1983). An example of a site where the MUA showed a distinct response to different pictures is shown in Figure 1. We observe that there was a strong enhancement in the number of spikes approximately 110 ms after presentation of a picture of a hand (in contrast to the presentation of the other two images). Another example of a site where the MUA showed selective responses is shown in Figure 4 (and further samples of MUA and LFP responses are shown in 6-8). To address whether the responses to different stimuli were significantly different from each other we performed a one-way ANOVA with object identity as the main factor (we define r_v as the ratio of the across-picture variance to the within-picture variance, see Appendix 1). Spikes were counted within the *stimulus* interval defined as the period from 100 ms to 300 ms after stimulus onset².

For the example shown in Figure 4, the value of r_v in the stimulus interval was 5.1 ($p < 10^{-10}$), indicating that the neuron responded with different numbers of spikes to different pictures. To assess which images the site preferred, we performed a post-hoc analysis by comparing the responses to each

¹ To avoid artifacts due to line noise, we also applied a digital 60 Hz notch filter to the LFP signal (see Methods). We performed several preprocessing steps to assess the quality of the MUA and LFP data; these steps are described in the Methods Section. We computed the power spectral density for every site (see Figure 1 and Appendix 2). Most of the power in the spectrum was concentrated below 40 Hz. Sites with unusual peaks in the spectrum or with unstable recordings were removed from further analysis (see below).

² In order to determine an optimum spike count interval for the statistical analysis, we computed the proportion of sites that yielded a p value < 0.001 in the ANOVA test in successive 200 ms spike count windows shifted by 10 ms from 200 ms before stimulus onset until 200 ms after stimulus onset. The maximum of this plot corresponded to the interval [80;280) ms and the curve was rather flat in the top (similar results were observed for other p value thresholds).

image to the overall distribution of responses using a t test. This indicated that the site in Figure 4 responded to several images including many cars.

We asked whether the LFP in IT would also show selectivity for complex visual stimuli. The example in Figure 3 shows an enhancement in the LFP response to the photograph of a masked face over the LFP responses to the other two images. Figure 5 shows another example of a site where the LFP showed a selective response to different stimuli. For the LFP data, the response was defined as the total power in the [100;300) ms time interval (see Appendix 2 for the estimation of the of power spectral density). For the example in Figure 5, the value of r_v was 4.4 ($p < 10^{-8}$); the post-hoc analysis indicated that the power was enhanced for several pictures including many images of cats and dogs.

Of the 356 sites with MUA data³, we observed that 267 sites (75 %) showed p values < 0.001 in the ANOVA analysis and 238 (67 %) showed a p value < 0.001 for at least 1 image in the post-hoc t test. Of the 309 sites with LFP data, we observed that 147 sites (48 %) showed p values < 0.001 in the ANOVA analysis and 70 sites (23 %) showed a p value < 0.001 for at least 1 image in the post-hoc t test. Unlike spike waveforms, voltage signals recorded in the LFP frequency range are not well understood. The observation that the LFP modulation was object selective and showed a time course roughly similar to the MUA responses suggests that our observations of LFP selectivity were not recording artifacts. However, we performed several additional controls. First, we repeated the same analysis in the *background* interval, defined as the period from -200 to 0 ms with respect to stimulus onset. While this interval contains the previous presentation the images were presented in pseudo-random order and therefore this should only contain spurious selectivity⁴. None of the sites showed selectivity during the background

³ We recorded from a total of 618 sites, MUA data were recorded in all cases, LFP data were recorded in 524 sites. A total of 30 sites were outside IT (12 in STS and 18 control sites outside visual cortex). Some sites showed electrical noise and were not analyzed further (12 sites for the MUA data and 108 sites for the LFP data). Some recordings were unstable (26 sites for the MUA data and 25 sites for the LFP data). To be sure that the activity changes we studied were not due to non-stationary conditions of the recordings or other artifacts, we fit a line to the response signal versus repetition number (where response signal = spike count for MUA and power for the LFP, see Figure 2). We computed the stationarity ratio as the response in the last repetition divided by the response in the first repetition (after the linear fit) for each site. We removed from analysis those sites that showed more than a 30 % change in the signal (166 sites for the MUA data and 62 sites for the LFP data). The average ratio was 1.04 ± 0.13 for MUA and 0.98 ± 0.12 for LFPs. We also computed the short-term stationarity ratio by fitting a line to the response signal versus presentation time within a trial. The mean ratio was 1.01 ± 0.22 for MUA and 1.09 ± 0.36 for LFPs. The proportion of selective units did not change significantly upon removal of these last two constraints or for different values of the thresholds in these constraints.

⁴ A poor randomization of the stimuli would tend to increase the proportion of selective sites in the analysis of the background interval.

interval either for the MUA or for the LFP data at the 0.001 level⁵. Second, we also recorded MUA and LFP activity in 18 sites outside visual cortex in response to the same images. None of these sites showed any selective responses⁶. Third, some of the images were white contours without any information in the image center. Since we presumed that these images may potentially lead to lower neuronal activation, we repeated the analysis after excluding these images. The proportion of selective sites was only slightly lower (76% for MUA and 38% for LFP, the number of sites with at least 1 selective image was also only slightly decreased to 67 % for MUA and 21 % for LFP). Finally, several of the responses seemed to start after stimulus offset (see the examples in Figures 1, 2 and 3). To distinguish between a neuronal latency greater than 100 ms (typical of IT MUA) or a very short-latency, selective offset response; we recorded from 11 sites while presenting the stimuli for 200 ms (instead of 100 ms). These sites also showed selective responses with latencies around 100 ms.

The number of pictures that passed the post-hoc analysis was 2.9 ± 2.1 (range 0 to 11) for the MUA selective sites and 1.2 ± 1.8 (range 0 to 14) for the LFP selective sites. The ratio of the across-picture variance to the within-picture variance was 3.6 ± 2.1 for the MUA selective sites and 3.1 ± 1.4 for the LFP selective sites. Therefore, more sites showed selectivity according to the MUA data (ANOVA analysis), the MUA data responded selectively to more images (post-hoc t test analysis) and the selectivity was higher for MUA recordings (larger across picture to within picture variance ratio). Different thresholds for the p values also lead to the same conclusions.

Different frequency components of the LFP signal may have distinct neuronal origins and thus show different selectivity levels (Buzsaki, 1998; Kahana et al., 1999; Klausberger et al., 2003; Mitzdorf, 1985). We therefore divided the power of the LFP signal into different 10 Hz frequency bands⁷. We also compared the results with those obtained from the maximum amplitude of the LFP signal, a simple variable that still seems to contain much of the information in the LFP signal (Mehring et al., 2003). The number of sites that showed selectivity was lowest in the 20 to 40 Hz frequency band. The 0 to 20 Hz

⁵ At the 0.01 level, 1.4% of the sites showed selectivity according to the MUA responses and 0.7% of the sites showed selectivity according to the LFP data.

⁶ Sites recorded in the superior temporal sulcus (STS) also showed selective responses. Since they were very few in comparison with the IT data, they were not included in the results reported here.

⁷ Stimuli were presented for 100 ms and there was a 100 ms blank interval between presentations. The sampling frequency of 1 kHz makes the bandwidth resolution $1/(200 \times 1 \text{ ms}) = 5 \text{ Hz}$.

frequency band, the 100 to 300 Hz frequency band and the two measures of overall LFP power yielded similar results of approximately 40% selective sites.

Most of the LFP responses consisted of transient biphasic or triphasic potential changes. Given the short presentation times, it is difficult to make claims about the duration of the LFP response. In the case of MUA, the responses typically included transient increases in the number of spikes although selective decreases in the spike rate were also apparent. The selective responses were non-uniformly distributed among the 78 pictures that we presented. For example, some of the food photos elicited selective responses in many MUA recordings, one of the masked human faces elicited selective responses in many MUA and LFP recordings, one of the hand images was preferred by many LFP sites, and most of the white boxes generally failed to induce selective MUA or LFP responses.

The latencies of the LFP responses were shorter than the latencies of the MUA responses (MUA: 137 ± 25 ms, LFP: 99 ± 29 ms, t test $p < 10^{-4}$). The distribution of all the latencies is shown in 9A. Furthermore, for those sites that were selective both according to MUA and LFPs and where we could estimate the latency, we found no correlation between the latency of the responses from the two types of signals ($r=0.0047$, Figure 9B). A time course plot of the average MUA and LFP signals also shows that, on average, the LFP response starts to increase slightly before the MUA response (Figure 9C).

1.2 Comparison of MUA and LFP data

How does the selectivity of action potentials compare to that of local field potentials from the same recording site? As the example in Figure 3 illustrates, we often observed cases where there was an enhancement in the spike rate without concomitant changes in the LFP response and vice versa. Figure 10 shows a comparison of the selective responses for all sites that had appropriate MUA and LFP data. Overall, we observed that spikes and LFPs were generally not selective to the same images. From a total of 249 sites, 90 sites (36%) showed selectivity according to the MUA responses and not with the LFP data. There were also 14 sites (5.6%) where the LFPs were selective and the MUA was not. Of the 87 sites that showed selectivity according to both signals (35%), 39 sites (45% of the 87 sites, 16% of the 249 sites) had at least 1 image that passed the post-hoc test for both MUA and LFP. The mean proportion of overlap in the selective images was 5%. The lack of overlap in the selectivity of the two

signals can be seen in Figure 10 where the preferred images of the MUA and LFP data in each site are shown. The poor overlap in the selectivity of MUA and LFP was also observed for other threshold parameters and for other frequency bands.

The above analysis depends on several arbitrary statistical thresholds. We therefore directly compared the MUA and LFP responses by computing the correlation coefficient between the spike counts and the LFP power. We show three examples of sites where there were different levels of correlation between the LFP and MUA signals (Figure 11). Overall, we observed a unimodal distribution (Figure 12

A) with only a very small correlation between the two signals ($r=0.09\pm 0.20$ for the stimulus interval and 0.04 ± 0.15 for the background interval)⁸. This shows that the lack of overlap of the LFP and MUA responses is not due to the arbitrary thresholds used in the statistical analysis. The only frequency band that yielded a higher correlation with the spike responses was the [100;300) Hz portion of the LFP spectrum. However, it should be noted that these higher frequencies contain some components of the spikes themselves. It is quite remarkable to obtain such a low correlation coefficient between the two signals given that they were recorded from the same electrode.

1.3 Selectivity of nearby sites

Earlier work suggests that there is a topographical structure to the selectivity in IT cortex (Fujita et al., 1992; Gochin et al., 1991; Tanaka, 2003). We therefore explored whether there was any correlation between the patterns of selectivity at spatially separate recording sites. We examined cases in which the recording electrode sampled different sites along a single penetration (during a single recording session) and cases in which different sites were sampled on different penetrations. We observed several cases where the preferred stimuli in one site (from the spike counts in the MUA responses) were very similar to those in other sites from the same penetration. As an example, in one penetration, one recording position showed enhanced firing rates upon presentation of all the images of cats and dogs. At least two sites 200 μm apart from this one, showed enhanced firing rates for some images of cats and dogs. A site recorded several hundred μm away from the first one showed enhanced activity for one dog and one vehicle while another site even farther away showed a weak selective response to the same vehicle (Figure 13).

This description may be dependent on the thresholds used in the p values which are somewhat arbitrary. We therefore, computed the correlation coefficient between the spike counts between all possible pairs of sites including those within the same penetration and those in different penetrations (see

⁸ Considering only those sites that showed selectivity for the MUA and LFP responses, the correlation coefficient increased only slightly to 0.14 ± 0.25 (Figure 12B). Spearman correlation coefficients yielded similar results. We also computed the correlation coefficient between the odd and even repetitions of each picture for the same type of signal (MUA-MUA and LFP-LFP correlations). This constitutes a way of measuring the variability of the responses and provides a comparison basis for the MUA-LFP correlation coefficients. The average MUA-MUA correlation coefficient was 0.52 ± 0.15 and the average LFP-LFP correlation coefficient was 0.47 ± 0.16 .

Figure 14A). We observed that there were clear cases where sites within the same penetration showed much higher correlation coefficients than what was observed for a random pair of sites (boxes with colors closer to red near the diagonal in

Figure 14A). Since nearby sites were also recorded in the same day and the same recording session, we asked whether this observation could be attributed to some overall stimulus-independent similarity of the recording conditions. We computed the same distribution but counting spikes within the background interval ($[-200;0]$ ms). The correlation matrix did not show any clusters of high correlation coefficients (

Figure 14B). The Pearson correlation coefficient shown in

Figure 14A may be quite sensitive to outliers in the data. We therefore performed the same analysis using the non-parametric Spearman correlation coefficient (Lehmann, 1975) and reached the same conclusions. The analogous correlation matrix for LFPs looked rather different (

Figure 14C). We observed that correlations in the LFP power between sites extended to other recording sites in different penetrations as well. This effect was stimulus-dependent since we did not observe these strong correlations in the background period (

Figure 14D).

The clusters of high correlation coefficients were not observed in every single penetration. To quantify the degree of similarity between sites within the same penetration and to assess its statistical significance we plotted the distribution of correlation coefficients within a penetration and the distribution of correlation coefficients for all pairs of sites (Figure 15A). The average correlation coefficient for all site pairs was 0.02 and the average correlation coefficient for sites within the same penetration was 0.28 (the difference was highly significant, t test, $p < 10^{-10}$). For the background interval, the difference between sites within a penetration and all sites was not significant (Figure 15B, $p=0.65$). Strikingly, we observed correlation coefficients of about 0.2, above chance levels, even for distances of up to 800 μm between sites. The dependence on the correlation coefficient between sites on the distance between the two sites is shown in Figure 16A. For the LFP recordings, correlations were observed at much larger distances, spanning several mm (Figure 16B).

2 Discussion

2.1 About LFPs

What do local field potentials actually measure? The frequency component up to several hundred Hz recorded extracellularly is composed of multiple signals over relatively large areas of cortex including glia activity, axonal spikes, somatic spikes, dendritic potentials and capacitive components. The contribution of these different signals to the LFP measurements has been studied experimentally as well

as theoretically and some general observations can be made. In general, the contribution from glial activity and capacitive components is small and can be neglected (Mitzdorf, 1985; Mitzdorf, 1987; Speckmann and Elger, 1993). Under most circumstances, the contribution of spikes seems to be small compared to that from dendritic potentials⁹. Among dendritic potentials, inhibitory potentials are negligible and excitatory post-synaptic potentials seem to account for most of the LFP signal. Thus, roughly, the LFP signal could be thought of as a spatially and temporally low-pass filtered version of the input to a given area. The connection between the LFP signal and the input to an area is further supported by the shorter latency of the LFPs compared to spiking activity (Mehring et al., 2003).

2.2 Selectivity of LFPs

We observed that a large fraction of sites showed a selective response to some of the stimuli that were presented both in the MUA responses as well as in the LFP responses. The exact proportion of selective cells depends on several experimental conditions including the number of images presented as well as the exact definition of selectivity. Still, several other studies have reported similar observations showing that many neurons in IT respond differentially to distinct complex shapes¹⁰. The observation that fewer sites showed selectivity according to the LFP measurements could be dependent on the type of stimuli (although *a priori* the stimuli were not particularly chosen to enhance MUA responses versus LFP responses). Alternatively, it could represent differences between the input to and the output of (pyramidal cells in) inferior temporal cortex. On a more speculative note, perhaps the LFPs in IT show properties that resemble more the spiking output of earlier areas (like V4) after the non-linear transformations imposed by neurotransmitter release and dendritic processing.

⁹ Care should be taken in studying the upper frequency band (~ 100 to 300 Hz) where there can be significant leakage of the action potentials into the LFP recordings. This is particularly problematic when analyzing LFPs and spikes from the same electrode as in the current case but, given the spatial spread of the LFP signal, spikes several hundred microns away from the electrode can also permeate the higher frequencies of the LFP when recorded from multiple electrodes.

¹⁰ For example, the original studies of Gross and colleagues reported that more than 30% of the neurons tested preferred some shapes over others (Gross et al., 1972). Gochin and colleagues report that 52% of the neurons were selective to one or more of the 6 images presented to the monkey (Gochin et al., 1991). Tamara and Tanaka presented 100 visual stimuli and reported that 79% of 213 cells showed selectivity (similar criteria to the ones used here, Tamara and Tanaka, 2001).

Selective LFP responses have been observed in other brain areas. For example, Mehring and colleagues report selective LFP responses related to the direction of hand movements in motor cortex (Mehring et al., 2003). Mitzdorf and colleagues have shown that LFP recordings in cat visual cortex also show selectivity to different visual stimuli (Mitzdorf, 1985; Mitzdorf, 1987). Pesaran et al showed selective LFP activity in the macaque parietal cortex during an arm reach task (Pesaran et al., 2002). Douchin and colleagues report that up to 80% of their LFP recordings in the macaque SMA or M1 cortex showed selective responses (Donchin et al., 2001). The relationship between spiking activity and local field potentials may be quite complex. In some cases, the LFPs have been shown to be correlated to the spiking responses but there are several reports that show that the LFPs show different properties than the spikes (Donchin et al., 2001; Eggermont and Mossop, 1998; Murthy and Fetz, 1996; Pesaran et al., 2002). In our observations, LFPs show selectivity that can be correlated with the spiking response in some cases but in most sites, the correlation was only very weak¹¹.

2.3 Correlations between sites

We observed strong correlations in the selective responses of sites within the same penetration for MUA (Figure 14). Fujita and colleagues also made similar observations in their recordings although with a slightly smaller spatial scale (Fujita et al., 1992). Here we observed in some cases correlated responses between sites separated by more than 800 μm . Studies from cross-correlation of the activity of presumably two nearby neurons have also shown similar observations (Gochin et al., 1991). It seems highly unlikely that these results could be due to studying recordings from the same neurons by chance. Spike recordings typically register the somatic spikes (in contrast to axonal spikes). Simulation studies show that the magnitude of the extracellular recordings from spikes located in a soma farther than 200 μm away from the electrode are very small (Holt and Koch, 1999). Therefore, we could estimate a limit of a separation of at most 400 μm for two electrodes to be registering the activity of the same neuron. But

¹¹ There are multiple interesting reports that show that spikes can occur at particular phases of the LFP response (Fries et al., 2001; Kay and Laurent, 1999; Klausberger et al., 2003; Laurent and Davidowitz, 1994; O'Keefe and Speakman, 1987). We have not yet investigated this possibility; but it is likely that more rigorous claims can be made using data from multiple electrodes in this regard.

we have observed that correlations extend beyond 400 μm ¹². Therefore, our observations support the notion that there is a topographical representation of visual information in IT. Tanaka and colleagues have suggested that this topographical representation is rooted in the columnar structure of IT cortex (Fujita et al., 1992; Tanaka, 2003).

In the case of LFP recordings, we observed correlations that extended to sites several mm away from the recording electrode. This agrees with other estimates of the spatial extent of the LFP recordings (Mitzdorf, 1985; Mitzdorf, 1987; Speckmann and Elger, 1993). Murthy and Fetz made simultaneous recordings with multiple electrodes and they reported that sites separated by up to 14 mm in precentral cortex showed significant correlation in the LFP activity in the 20 to 40 Hz band (Murthy and Fetz, 1996).

2.4 Implications of LFP selectivity, some speculation

The observation that the average activity from large numbers of neurons measured in the form of local field potentials shows selectivity for complex shapes seems to be quite remarkable. One possible explanation for the selectivity of LFP in IT could be the topographical arrangement of selective neurons which is so prevalent in several areas of cortex such as V1 and MT. Selective LFP responses in the MTL have also been observed in humans (McCarthy et al., 1999; Nobre et al., 1994). A topographical map in IT has been suggested by Tanaka and colleagues (Fujita et al., 1992). The observations in

Figure 14 also suggest that nearby sites show similar selectivity. It seems that it would be hard to explain selectivity at the LFP level with a completely random arrangement of neurons (where the responses of one neuron are completely uncorrelated with those of its near or long distance neighbors). Thus, the LFP selective responses, together with the observations of Tanaka and those in

Figure 14 support the suggestion of topography to IT. Unfortunately, what exactly the topographical map represents and how these features may change from one location to another still remains mysterious.

As we briefly discussed in a previous Section, from a biophysical viewpoint, the origin of the LFP signal may be more correlated with EPSPs and therefore with the input to IT than with its output. This may explain why there does not seem to be a strong correlation between the LFP selectivity (input) and

¹² Furthermore, given the neuronal density in IT, it would seem highly unlikely for two electrodes separated by 400 μm to be monitoring the same neuron. With single electrode data, it is hard to establish that the recordings separated by 200 μm or less actually correspond to separate neurons. We visually inspected the waveforms in the two recording sites in several cases. While in many cases the waveforms do appear to be different, a more rigorous analysis after spike sorting is necessary to better assess this issue (and even then, the waveforms might change depending on the relative orientation of the electrode with respect to the neurons Holt and Koch, 1999). Stronger evidence comes from tetrode recordings or recordings from multiple electrodes; if most or all spikes from single units in two electrodes occur at the same time (or with a jitter less than the conduction velocity), then it is possible to state that the recordings come from the same neuron. The work of Gochin and colleagues suggest that this is generally not the case in IT (Gochin et al., 1991).

the MUA selectivity (output). The differences are given by the processing of the IT neurons and the difference between a single neuron (or a few neurons) and the average over large numbers of neurons.

Logothetis and colleagues have simultaneously measured SUA, MUA, LFP and BOLD changes in the hemodynamic response in fMRI experiments. Interestingly, the blood oxygen level dependent changes in the fMRI data correlate better with the LFP responses than with the spikes (Logothetis, 2002; Logothetis, 2003)¹³. Selective responses in fMRI have been reported by numerous investigators (see for example Kanwisher et al., 1997). Interestingly, fMRI responses seem to be much less selective than responses from spikes which directly parallels our observations about the selectivity of LFPs. A rough comparison would seem to indicate that LFP data is more selective than hemodynamic changes but a more detailed study would be required to further evaluate this. The spatial scale of LFP recordings (Figure 15) also seems to approximately correlate with the resolution of fMRI measurements.

2.5 Future directions

In spite of several decades of electrophysiological research in inferior temporal cortex, what kinds of features IT neurons are selective to remain quite mysterious. Our current report does not solve the problem, nor do we answer what the LFP measures are actually selective to. Several types of invariances have been observed in the responses of individual neurons in IT, including scale and position invariance (Gross et al., 1972; Logothetis et al., 1995, see however DiCarlo and Maunsell, 2004). It remains to be investigated whether LFPs also show invariance and how robust the LFP signal is to different transformations of the input images. It is possible to separate the neuronal components of the MUA by using spike sorting algorithms (Quiroga et al., 2004; Yu and Kreiman, 1999). It will be interesting to compare the SUA and MUA in terms of their selectivity, spatial extent and classification performance. Finally, it will also be of interest to understand how smaller ensembles of neurons (intermediate between the level of 1 or a few neurons and the level of several tens of thousands of neurons) encode visual information in IT. This will most likely require the use of multiple electrode recordings in IT. As the LFP

¹³ There are several filtering steps in the processing of the BOLD responses, these are described in further detail by Logothetis and colleagues (Logothetis, 2002; 2003).

recordings show in our study, ensembles can show properties that are hard to predict by recording exclusively spiking activity from a single electrode.

3 Methods

3.1 *Animals and Surgery*

Recordings were made from two monkeys (*Macaca mulatta*) weighing 6.0 kg (monkey K, female) and 5.0 kg (monkey N, male). Prior to behavioral training, aseptic surgery was performed to attach a head post to the skull and to implant a scleral search coil in one eye (monkey K, left eye; monkey N, right eye). After 2-3 months of behavioral training, a second surgery was performed to place a recording chamber (centered at Horsley-Clark coordinates A 16, L 12) to reach the anterior portion of right inferotemporal cortex (AIT). A second chamber was implanted in the left hemisphere of monkey K after one year of recording. All procedures were approved by the MIT Committee on Animal Care.

3.2 *Stimuli*

We used a set of 78 complex grayscale stimulus objects. Stimuli were arbitrarily divided prior to the experiment into 8 sets: toys, foodstuffs, human faces, hand/body parts (from monkey K), monkey faces (monkey K and animals from adjacent cages), vehicles, white boxes, and synthetic images of cats and dogs (Freedman et al., 2001).

Visual stimuli were presented on a video monitor (47.5 x 30.5 cm, 52 x 34 deg, 75 Hz, 1600 x 1200 pixels) positioned 52 cm from the monkey. Stimulus objects were presented on a constant gray background and had a size of 3.4 deg (maximum of width or height). Animals performed a passive fixation task during which each trial consisted of 20 stimuli sequentially presented at the fixation point (0.2 deg, red) at 5 Hz over 4 seconds. Each stimulus was on the screen for 100 msec, interleaved by a 100 msec blank matching the background luminance (28 Cd/m²). Stimuli ranged in luminance from 0.5 to 56 Cd/m². To preserve the approximate physical appearance of the objects, stimuli were not normalized for mean luminance level (average 23 Cd/m², range 6.1 to 36.5 Cd/m²) or contrast (mean 76%, range 16%

to 98%). To minimize response variability, the presentation position of each stimulus object was matched to the animal's center of gaze for each video frame, while the eyes were required to remain within a fixation window of ± 2 deg. Up to 20 stimuli were presented for each trial, and responses to the first 3 stimuli of each trial were ignored. Stimuli were presented in pseudorandom order, randomly presenting one entire set of 78 stimuli before beginning a second randomized set, until each stimulus had been shown 10 times.

3.3 Neural recording

Recordings were made from both hemispheres of monkey K and the right hemisphere of monkey N. The majority of data in monkey K was from the right hemisphere. A 26G guide tube was used to reach AIT from a dorsal approach. The superior temporal sulcus (STS) and the ventral surface were identified by comparing gray and white matter transitions and the depth of the skull base with MRI structural images. Penetrations were made over a $\sim 10 \times 10$ area of the ventral surface (Horsley-Clark AP: 10-20 mm, ML: 14-24 mm). Multi-unit recordings were made using glass-coated Pt/Ir electrodes (0.5-1.5 Mohm at 1 kHz). Spiking activity (400 Hz-6 kHz) and LFPs (1-300 Hz) were amplified, filtered, and stored using conventional equipment.

3.4 Data analysis

3.4.1 Data pre-processing

We performed several pre-processing steps before proceeding to further analysis to avoid potential artifacts and to ensure, as far as possible, that we were dealing with high-quality neuronal data. The first and the last 3 presentations in each trial were not considered for analysis. Also, trials where the monkey broke fixation were discarded from further analysis. Recordings without a minimum of 7 repetitions per image were also discarded.

For the MUA data, we first eliminated every spike that could arise from double detection by removing two spikes if the interspike interval was less than 2 ms (this constraint also removes any tightly

coincident spikes from separate neurons). We then plotted the interspike interval distribution and removed from further analysis those cases that showed sharp peaks at harmonics of 60 Hz (12 sites out of 618 recording sites). For each site, we then computed the distribution of firing rates and removed points that were beyond 4 standard deviations of the mean. We then plotted the average firing rate in each repetition as a function of repetition number to assess whether there were any non-stationary changes in the recordings. We compared each presentation against the overall distribution of firing rates using a t test; a presentation was skipped if the p value was less than 0.01. Finally, we considered the slope of the change in firing rate versus repetition number (longer-term stationarity) as well as the change in firing rate within a trial (short-term stationarity).

A similar procedure was applied to the LFP data. We first applied a digital notch filter to the data at 60 Hz and harmonics (4th order elliptic filter, 0.1 db peak-to-peak ripple, 40 db stopband attenuation). We computed the power spectral density using a 256-point Hanning window, and then using an overlap of 10. Sites that showed 60 Hz noise or high-frequency noise in the power spectral density were removed from further analysis (108 sites out of 560 LFP recording sites). For each site, we computed the distribution of the spectral power ($P_{LFP} = \sum_i |s_i|^2$, see Appendix 2). For each site, we removed those points that were beyond 4 standard deviations of the mean. We subsequently plotted the power as a function of repetition number. Each repetition was compared against the overall distribution using a t test; a repetition was skipped if the p value was less than 0.01. Finally, we also computed the slope of the change in spectral power as a function of repetition number (longer-term stationarity) and the LFP change within a trial (short-term stationarity).

3.4.2 Selectivity

In order to assess whether a site was selective or not, we used a non-parametric analysis of variance (np-ANOVA). np-ANOVA measures the variance across-stimuli compared to the variance within-stimuli. The larger the ratio of these two variances, the more significant is the result. Instead of making

assumptions about the distribution of this ratio, we estimated its empirical distribution under the null hypothesis of equal means among all the pictures by shuffling the picture labels (see Appendix 1).

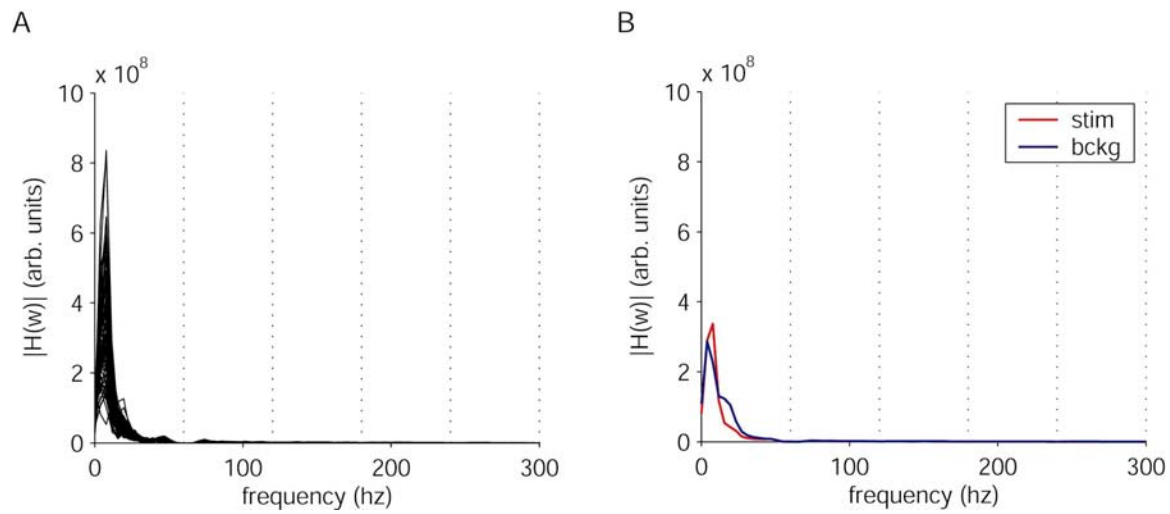
For the selectivity analysis, the response r was defined in the following ways: for the MUA, we used the spike count in the window [100,300) ms after stimulus onset, for the LFPs, we used the power in the window [100,300) ms after stimulus onset. We also studied different analysis windows as well as for other measures of LFP activity including different frequency bands (Pesaran et al., 2002) and the maximum amplitude (Mehring et al., 2003).

If a site was selective according to the np-ANOVA measure, we used the following ad-hoc definition to determine which pictures the site was selective to. The response was selective to picture j ($1 \leq j \leq 78$) if the response to j was significantly different from the baseline and from the response to the other pictures using a t test, $p < 0.01$. The baseline was defined as the corresponding signal in the interval [-200,0) ms before stimulus onset.

We compared the responses across different sites and between MUA and LFP by computing the correlation coefficients (both Pearson and Spearman) between 78-dimensional vectors containing the mean response for each image. We also compared the degree of overlap between MUA and LFP for the selective images or groups of images.

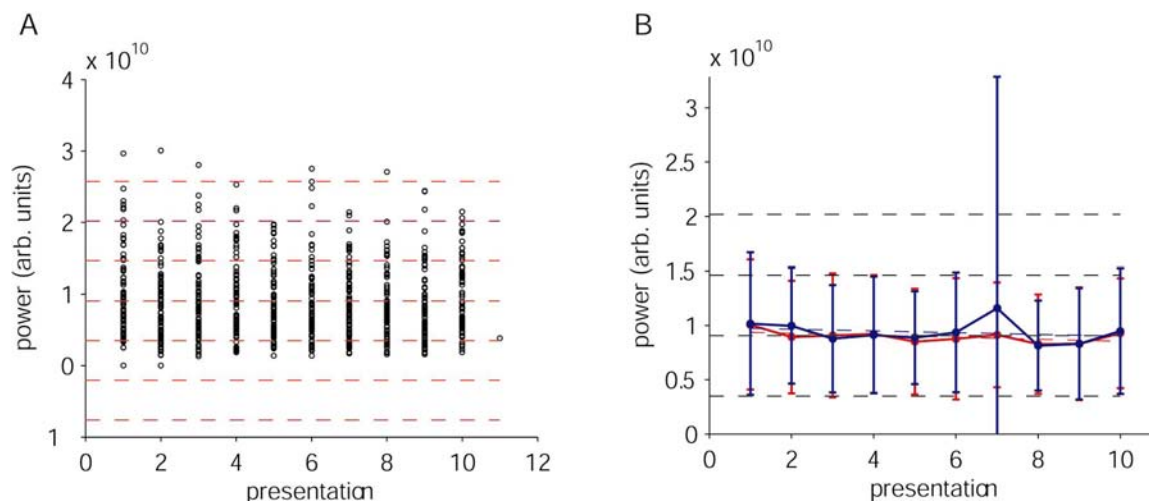
4 Figure and Table Legends

Figure 1: Example of power spectral density of local field potentials




Power spectral density (PSD) of the LFP responses (site = ckbea10). A. PSD for each of the 78 pictures (average of all repetitions for each picture). B. Average PSD over all pictures and all repetitions during the stimulus interval (red) or the background interval (blue).

Figure 2: Stability of LFP recordings



Stability of LFP recordings across repetitions. A. Power of the LFP response for each of the 78 pictures during each of the 10 repetitions of each stimulus. B. Average power for each repetition during the stimulus interval (red) or the background interval (blue). Site = ckea10.

Figure 3: Example of selective responses

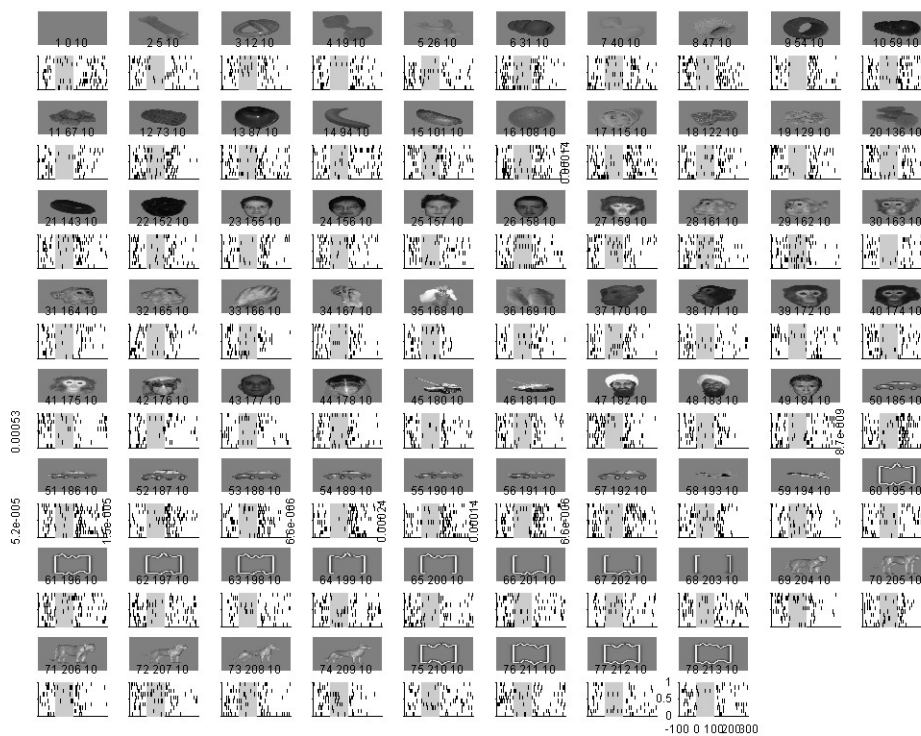


Example of the MUA and LFP responses of one site (ckbfa12) to repeated presentations of 3 different images (food, a hand and a masked human face). The MUA responses (middle part) are shown as raster plots where each row shows a separate repetition and each tick indicates a spike. The gray box denotes the stimulus presentation period (100 ms). There is a clear enhancement in the spike rate approximately 110 ms after appearance of the hand picture. The LFP responses (arbitrary units) are shown in the bottom part, the gray traces show the individual repetitions and the red trace shows the average LFP for

each stimulus. The dashed lines indicate the stimulus presentation time. There is a change in the LFP signal starting approximately 100 ms after onset of the masked human face. The MUA response to the hand was statistically significant (post-hoc t test $p < 10^{-10}$) and the LFP response to the masked face was statistically significant (post-hoc t test $p < 10^{-4}$), see Methods for data analysis.

Figure 4: Selective MUA example

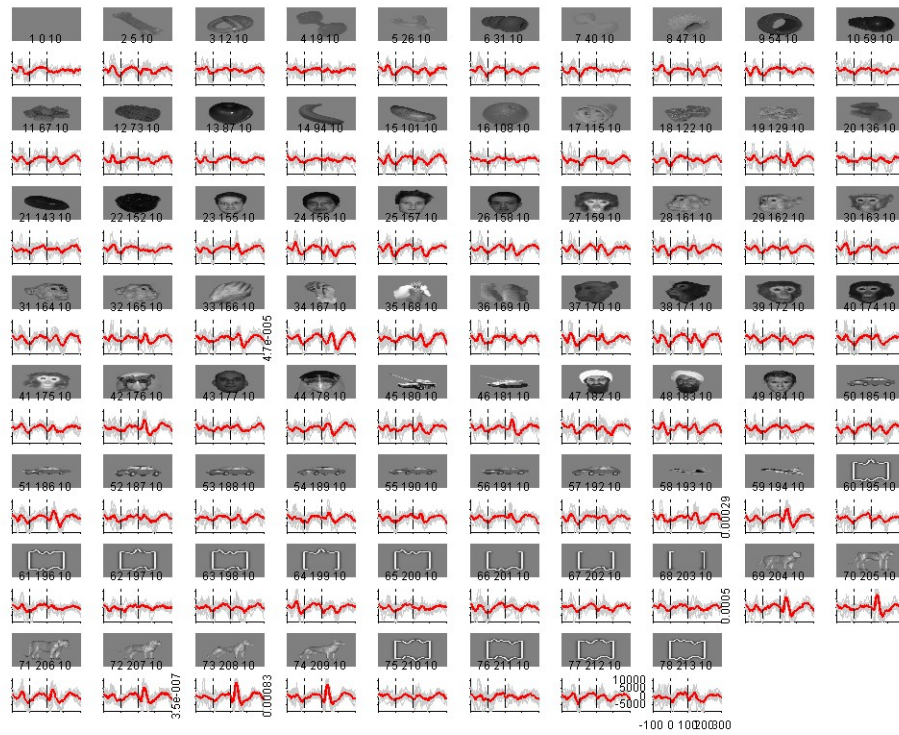
CKAFA10-c.mat



Example of MUA response to the 78 different pictures presented to the monkey. The raster plot shows one mark per spike and each row shows a separate presentation; spikes are aligned to stimulus onset time. Stimulus presentation (100 ms) is denoted by the shaded box. Below each picture we indicate the two internal image indices and the number of repetitions. When the p value for the post-hoc t test was $< 10^{-3}$, the p value is indicated to the left of the raster plot. Site = CKAFA10.

Figure 5: Selective LFP example

CKBEA01-c.mat



Example of LFP response to the 78 different pictures presented to the monkey. The gray lines show the LFP data for each repetition, the thick red line shows the average LFP signal (aligned to stimulus onset time). Stimulus presentation (100 ms) is denoted by the vertical dashed lines. Below each picture we indicate the two internal image indices and the number of repetitions. When the p value for the post-hoc t test was $< 10^{-3}$, the p value is indicated to the left of the raster plot. Site = CKBEA01.

Figure 6: Example III of MUA and LFP responses

Another example of MUA and LFP responses (the format follows the one in Figures 4 and 5). Here we show both the MUA and LFP data.

ckaya05-c



Figure 7: Example IV of MUA and LFP responses

Another example of MUA and LFP responses (the format follows the one in Figures 4 and 5). Here we show both the MUA and LFP data.

ckbca15-c



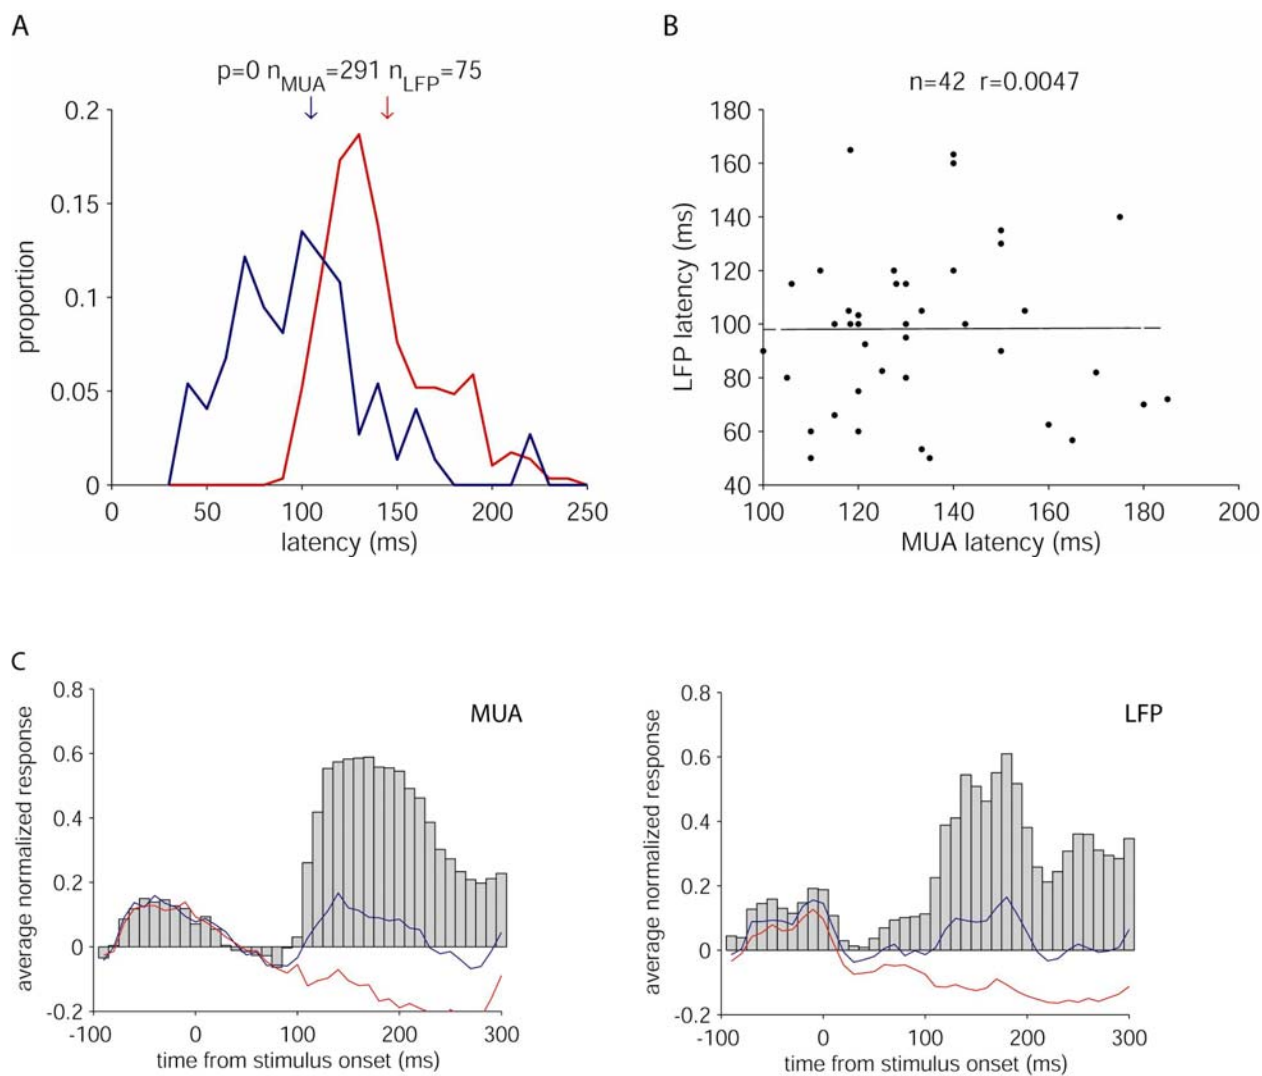
Figure 8: Example V of MUA and LFP responses

Another example of MUA and LFP responses (the format follows the one in Figures 4 and 5). Here we show both the MUA and LFP data.

ckboa31-c



Figure 9: Latency of MUA and LFP responses



A. Distribution of the latency of MUA (red) and LFP (blue) responses. The latency was computed only for those sites that passed the ANOVA selectivity criterion and for those images that passed the post-hoc t test criterion (see Methods). The arrows show the mean values. Bin size = 10 ms. **B.** Comparison of the latency of the MUA responses versus the latency of LFP responses. The line shows a fit to the data ($r < 0.01$). **C.** Time course of the MUA responses (left) and LFP responses (right). For each site, we took the activity for the best picture (gray bars), random picture (blue trace) and worst picture (red trace). The activity was binned (bin size = 10 ms, spike count for MUA and signal power for LFP), the average activity in the [0,100) ms interval was subtracted and the resulting signal was normalized by the 95th percentile. The resulting time course was then averaged over all the sites.

Figure 10: Comparison of the selectivity of MUA and LFP responses

Comparison of the MUA and LFP selective responses for each of the 78 different images. The x-axis goes from 1 through 78 and corresponds to each of the images presented to the monkey (the images are shown in **Figure 4**). Each row corresponds to a separate recording site. Sites recorded within the same penetration are grouped within horizontal dashed lines (separate penetrations are presented in the order in which they were recorded). A site was considered to be selective to an image if the ANOVA analysis and the post-hoc t test yielded $p < 10^{-3}$. Red boxes correspond to MUA selectivity, blue boxes show LFP selectivity and green boxes show overlap between the two. The MUA response is defined by the spike count and the LFP response is defined by the total power (both within the [100;300) ms interval). Vertical lines (and background gray tone) denote separate images from the 8 arbitrarily pre-defined stimulus groups (toys, foodstuffs, human faces, monkey faces, hands and body parts, vehicles, white boxes and cats or dogs). Site names and the number of sites per penetration are marked on the right and left of the panel. Overall, the overlap between the MUA and LFP responses (green boxes) was 0.05.




Figure 11: Examples of correlation between MUA and LFP responses

Three examples of the local field potential power (y-axis) as a function of the MUA spike count (x-axis) for each of the 78 pictures (averaged over all repetitions). A, C, E Correlations during the stimulus interval. B, D, F Correlations during the background interval. The overall distribution is given in Figure 12; here we show examples from the extremes of the distribution and a more average case. The line shows a linear fit to the data. Sites = ckbca15 (A,B), ckbda05 (C,D) and cnada02 (E,F).

Selectivity of local field potentials in macaque inferior temporal cortex

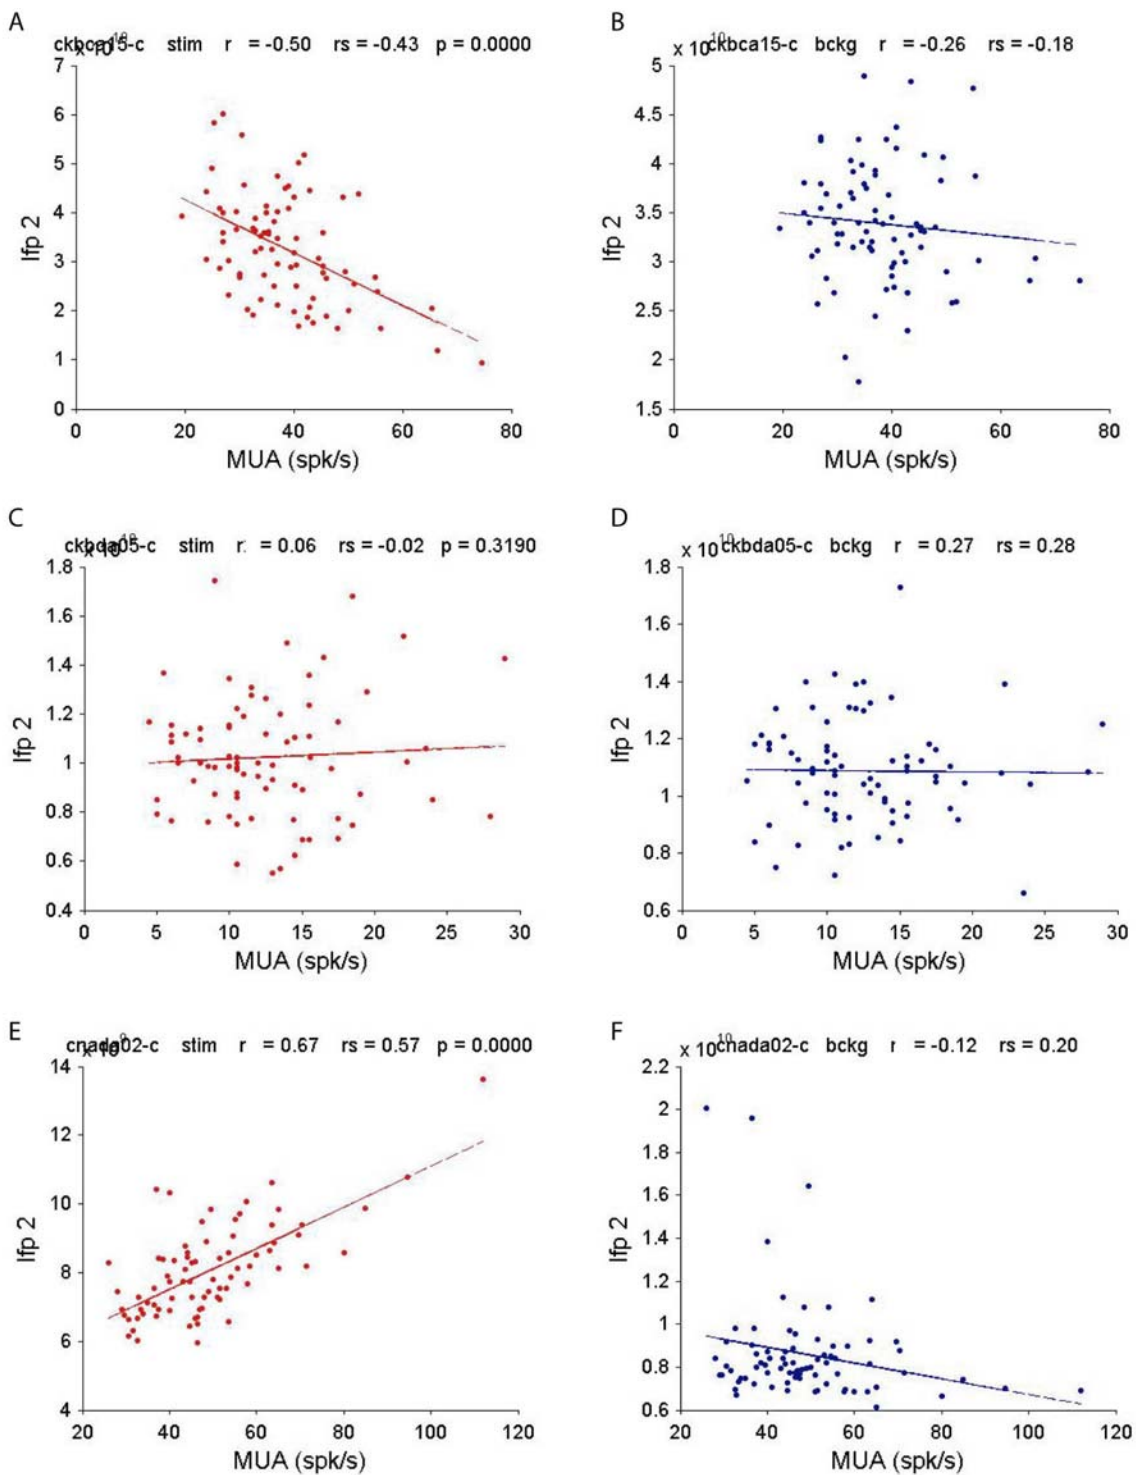


Figure 12: Correlation between MUA and LFP data

Distribution of Pearson correlation coefficients between MUA and LFP data from the same sites during the stimulus interval (red) or the background interval (blue). The arrows indicate the mean values. The correlation was computed from the 78 values indicating the average response (spike count or spectral power for MUA and LFP respectively) to each image and therefore does not depend on any of the statistical thresholds used in the previous figure. A Distribution for all sites with MUA and LFP data. B. Distribution showing only those sites that showed a selective response according to both the MUA and the LFP signal.

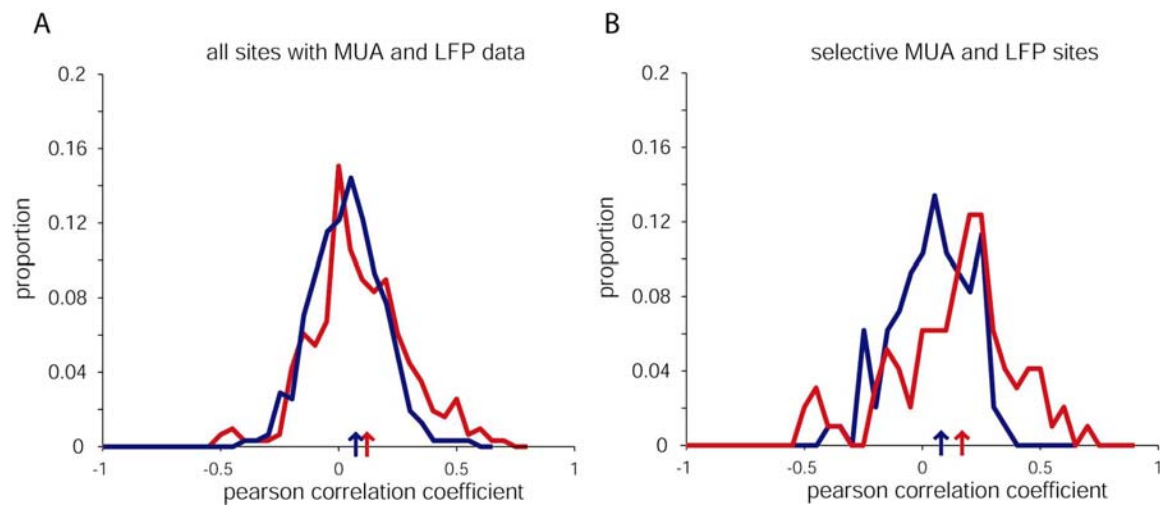


Figure 13: Example of MUA responses in multiple sites during one penetration

Mean firing rates in response to each of the 78 pictures presented to the animal for 8 sites separated by 200 μm steps recorded during the same penetration. Electrode positions are indicated to the left of the plots in microns relative to the first position. The y axis shows the spike counts in the [100;300) ms interval after subtracting the spike count in the [-200;0) ms baseline interval. The scale is shown in the bottom subplot. Site = ckaua.




Figure 14: Correlation in responses between different recording sites

Matrix of Pearson correlation coefficients between the MUA spike counts (A, B) or between the LFP power (C, D) for all pairs of sites. The matrix is symmetrical and the diagonal is trivially 1. The correlation coefficients are color coded (see scale at the bottom of the matrix). The correlation coefficients were computed between the 78-dimensional vectors composed of the average response to each of the images

Selectivity of local field potentials in macaque inferior temporal cortex

where the response was: (A) MUA during the stimulus interval, (B) MUA during the background interval, (C) LFP during the stimulus interval, (D) LFP during the background interval. Sites between dashed lines were recorded in the same penetration. Different penetrations here are shown in the order in which they were recorded.

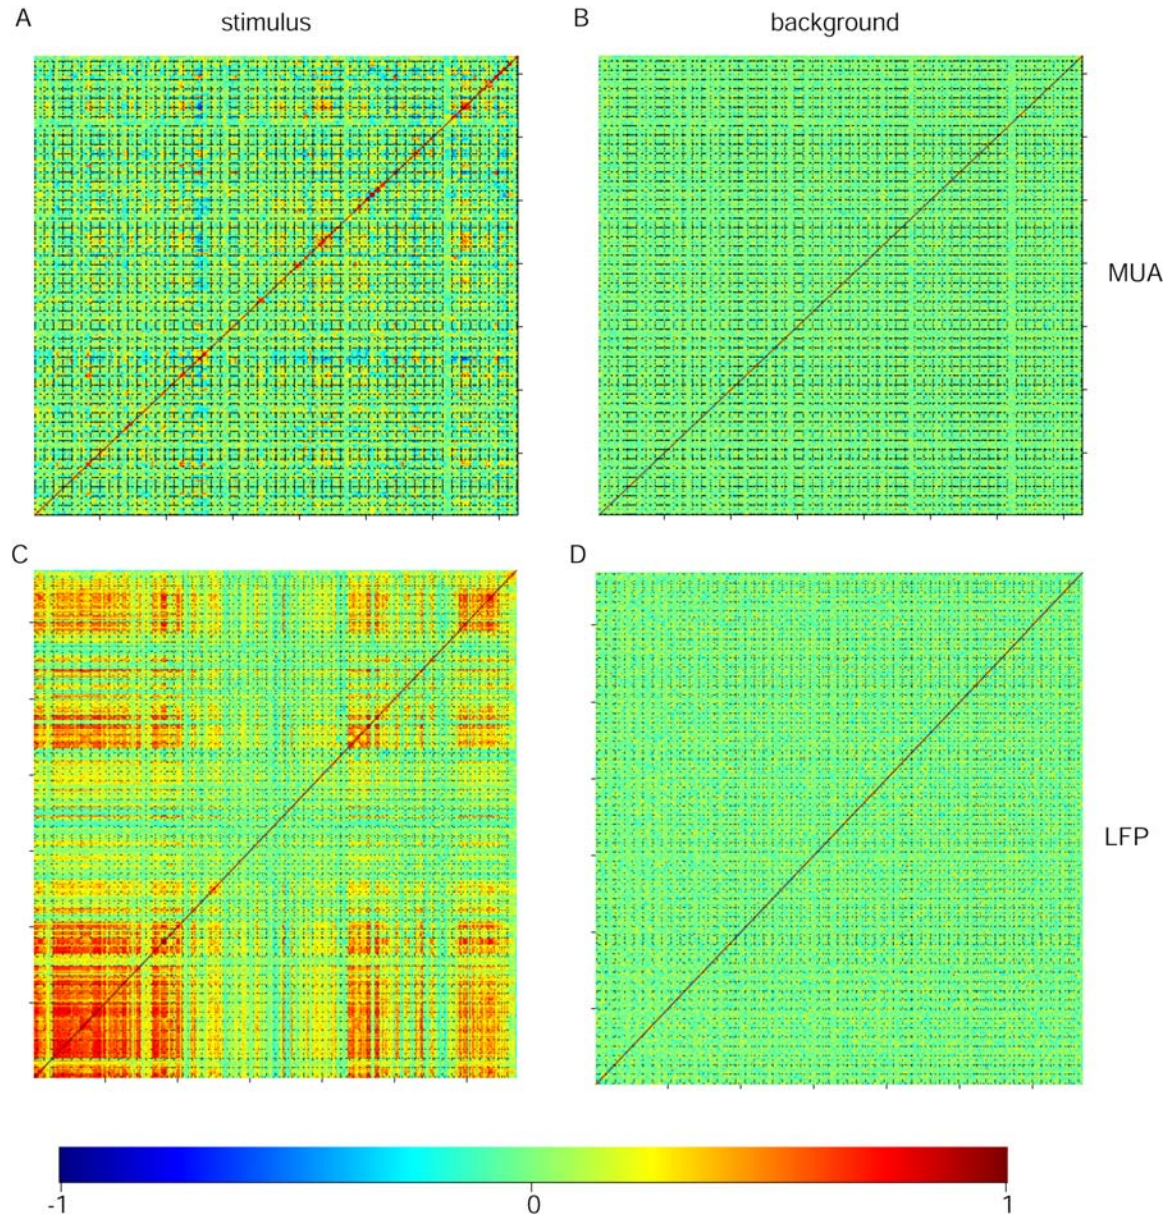


Figure 15: Distribution of correlation in the responses between different recording sites

Distribution of correlation coefficients for all pairs of sites (blue) or for pairs of sites recorded within the same penetration (red). The data for this plot is taken from the previous Figure. (A) Distributions for MUA, stimulus interval ([100;300) ms with respect to stimulus onset). (B) Distributions for MUA, background interval ([-200;0) ms with respect to stimulus onset). (C) Distributions for LFPs, stimulus interval. (D) Distributions for LFPs, background interval.

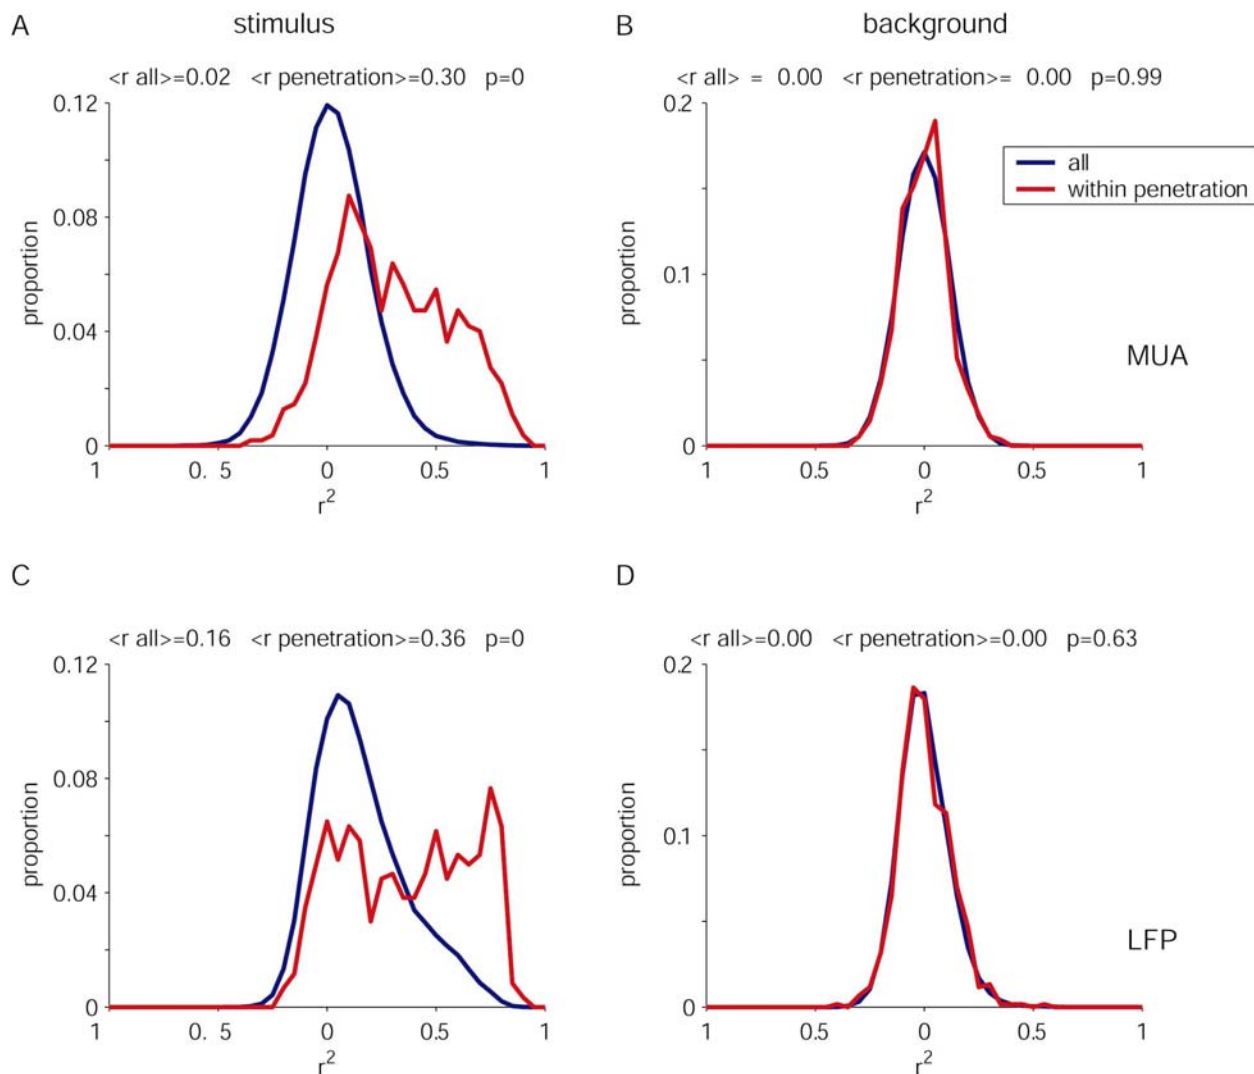


Figure 16: Correlation coefficients as a function of distance between sites

Pearson correlation coefficients between responses in different sites as a function of the separation between sites. For recording sites within the same penetration, the distance is given by the advancement of the microdrive; for recording sites in different penetrations the distance is computed from the estimated anterior/posterior and lateral/medial location of the electrodes. (A) MUA spike count responses. (B) LFP power. The error bars correspond to S.E.M. The distance and the number of pairs used to compute each value are shown above each bar. The red points show the mean correlation coefficient computed using the responses in the background time interval.

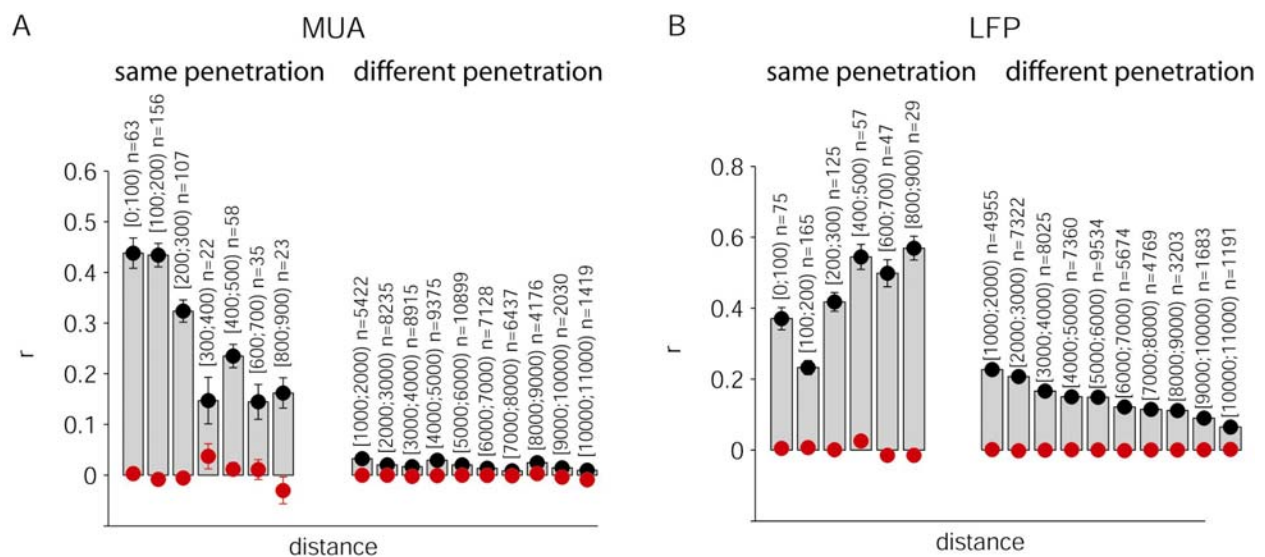


Table 1

Summary of the number of sites that showed selectivity for MUA and LFP responses as well as the degree of selectivity of each site (see Methods for selectivity criteria). The threshold p value for the numbers reported in this table was 0.001 (both for the ANOVA and the t test; similar conclusions were obtained for other p value thresholds).

	Number of selective sites, stimulus period	Number of selective sites, background period	Number of sites with at least 1 selective image	Number of images that passed the post-hoc t test	Across picture variance / within picture variance

Selectivity of local field potentials in macaque inferior temporal cortex

MUA	267 (75%)	0 (0%)	238 (67%)	2.9 ± 2.1	3.6 ± 2.1
LFP power	130 (42%)	0 (0%)	70 (22%)	1.2 ± 1.8	3.1 ± 1.4

5 Appendix 1: np-ANOVA

For a given site, let r_{ij} denote the response during presentation i of stimulus j ($i=1, \dots, n_{rep}$; $n_{rep} = 10$ in most cases and $j=1, \dots, 78$). We define the sum of squares *within-groups* S_w : $S_w =$

$$\sum_{j=1}^{78} \sum_{i=1}^{n_{rep_j}} (r_{ij} - \bar{r}_j)^2$$

where \bar{r}_j is the mean response for stimulus j . Let S_b denote the sum-of-squares

between groups: $S_b = \sum_{j=1}^{78} n_j (\bar{c}_j - \bar{c})^2$ where \bar{c} is the overall mean ($\frac{1}{N} \sum_{i=1}^{n_{rep}} \sum_{j=1}^{78} c_{ij}$) with $N = \sum_j n_j$. Note

that S_w measures the variability in repeated presentations of the same stimulus while S_b measures the variability in the responses to different stimuli.

Our statistic parameter will be the ratio between the two estimators of the standard deviation:

$$r_v = \frac{S_b / (k - 1)}{S_w / (N - k)}$$

Under the assumption of a normal distribution of firing rates in each group, r_v follows

an F distribution with $k-1$ and $N-k$ degrees of freedom. This is the classical ANOVA (Keeping, 1995; Scheffe, 1953). We estimated the actual distribution of t^* by using a bootstrap procedure (Efron and Tibshirani, 1993). For each iteration with a different shuffled label of the presentation, we computed the sum-of-squares values and the statistic r_v^* (the * denotes the fact that this was computed after shuffling the stimuli indices). Under the null hypothesis, H_0 , there is no difference among the responses to distinct stimuli within the selective category. If the neuron does indeed respond more strongly to some stimuli

than to others, then the observed value of the statistic \hat{r}_v computed from the actual order should be separated from most of the random shuffles. The probability that this happens under the null distribution can be estimated from the number of iterations in which the statistic r_v^* is larger than the observed value

$$\hat{r}_v :$$

$$p = \Pr_{H_0} (r_v^* \geq \hat{r}_v) \approx \#\{r_v^* \geq \hat{r}_v\} / n_{\Pi}$$

where n_{II} represents the number of different permutations that were analyzed and # denotes the cardinal or number of elements in the set. We used $n_{II} = 10,000$. This gives an accurate estimation of p (Efron and Tibshirani, 1993).

6 Appendix 2. Power spectral density

Let the local field potential signal be denoted by $c(t)$. We sample the signal at a frequency f_s (1 kHz) to produce the values c_j , $j = 1, \dots, N$ where N is the total number of samples. We define the mean squared amplitude of the signal as $\frac{1}{N} \sum_{j=1}^{N-1} |c_j|^2$. By Parseval's theorem, this is equal to the sum of the

power spectral density over all frequencies and therefore this is referred to throughout the text as the spectral power of the signal. The discrete Fourier transform of the signal is $C_k = \sum_{j=0}^{N-1} c_j e^{2\pi i j k / N}$ for $k=0, \dots, N-1$. When the data is multiplied by a window function to avoid leakage we obtain

$D_k = \sum_{j=0}^{N-1} c_j w_j e^{2\pi i j k / N}$ and the periodogram estimator of the power spectral density is:

$$P(f_0) = \frac{|D_0|^2}{W_{ss}}, P(f_k) = \frac{(|D_k|^2 + |D_{N-k}|^2)}{W_{ss}}, P(f_c) = \frac{|D_{N/2}|^2}{W_{ss}}$$

where f_c is the Nyquist frequency ($f_s/2$), $W_{ss} = N \sum_{j=0}^{N-1} w_j^2$ and $f_k = k f_s / N$ ($k=0, 1, \dots, N/2$). For further

detail, see (Oppenheim et al., 1997; Press et al., 1996).

7 References

- Allison, T., Ginter, H., McCarthy, G., Nobre, A. C., Puce, A., Luby, M., and Spencer, D. D. (1994). Face recognition in human extrastriate cortex. *Journal of Neurophysiology* 71, 821-825.
- Berger, H. (1930). Ueber das Elektroenkephalogramm des Menschen. *Journal of Psychology and Neurology* 40, 160-179.
- Buzsaki, G. (1998). Memory consolidation during sleep: a neurophysiological perspective. *J Sleep Res* 7, 17-23.
- DiCarlo, J., and Maunsell, J. (2004). Anterior Inferotemporal Neurons of Monkeys Engaged in Object Recognition Can be Highly Sensitive to Object Retinal Position. *Journal of Neurophysiology* 89, 3264-3278.
- Donchin, O., Gribova, A., Steinberg, O., Bergman, H., Cardoso de Oliveira, S., and Vaadia, E. (2001). Local field potentials related to bimanual movements in the primary and supplementary motor cortices. *Exp Brain Res* 140, 46-55.
- Eggermont, J. J., and Mossop, J. E. (1998). Azimuth coding in primary auditory cortex of the cat. I. Spike synchrony versus spike count representations. *J Neurophysiol* 80, 2133-50.
- Freedman, D., Riesenhuber, M., Poggio, T., and Miller, E. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291, 312-316.
- Fries, P., Reynolds, J., Rorie, A., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 23, 1560-1563.
- Fujita, I., Tanaka, K., Ito, M., and Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. *Nature* 360, 343-346.
- Gochin, P., Miller, E., Gross, C., and Gerstein, G. (1991). Functional interactions among neurons in inferior temporal cortex of the awake macaque. *Experimental Brain Research* 84, 505-516.
- Gross, C., Rocha-Miranda, C., and Brender, D. (1972). Visual properties of neurons in inferotemporal cortex of the Macaque. *Journal of Neurophysiology* 35, 96-111.
- Hobson, J. (1995). *Sleep* (New York: Scientific American Library).
- Holmes, E., and Gross, C. (1984). Stimulus equivalence after inferior temporal lesions in monkeys. *Behavioral Neuroscience* 98, 898-901.
- Holt, G. R., and Koch, C. (1999). Electrical interactions via the extracellular potential near cell bodies. *Journal of Computational Neuroscience* 6, 169-184.
- Kahana, M., Sekuler, R., Caplan, J., Kirschen, M., and Madsen, J. R. (1999). Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature* 399, 781-784.
- Kandel, E., Schwartz, J., and Jessell, T. (2000). *Principles of Neural Science*, 4th Edition (New York: McGraw-Hill).
- Kanwisher, N., McDermott, J., and Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience* 17, 4302-4311.

Kay, L. M., and Laurent, G. (1999). Odor- and context-dependent modulation of mitral cell activity in behaving rats. *Nature* 2, 1003-1009.

Keeping, E. S. (1995). *Introduction to Statistical Inference* (New York: Dover).

Klausberger, T., Magill, P. J., Marton, L. F., Roberts, J. D. B., Cobden, P. M., Buzsaki, G., and Somogyi, P. (2003). Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* 421, 844-848.

Laurent, G., and Davidowitz, H. (1994). Encoding of olfactory information with oscillating neural assemblies. *Science* 265, 1872-1875.

Lehmann, E. L. (1975). *Nonparametrics: statistical methods based on ranks* (New York: McGraw-Hill).

Logothetis, N. K. (2002). The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. *Philos Trans R Soc Lond B Biol Sci* 357, 1003-37.

Logothetis, N. K. (2003). The underpinnings of the BOLD functional magnetic resonance imaging signal. *J Neurosci* 23, 3963-71.

Logothetis, N. K., Pauls, J., and Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biology* 5, 552-563.

Logothetis, N. K., and Sheinberg, D. L. (1996). Visual object recognition. *Annual Review of Neuroscience* 19, 577-621.

McCarthy, G., Puce, A., Belger, A., and Allison, T. (1999). Electrophysiological studies of human face perception. II: Response properties of face-specific potentials generated in occipitotemporal cortex. *Cerebral Cortex* 9, 431-444.

Mehring, C., Rickert, J., Vaadia, E., Cardoso de Oliveira, S., Aertsen, A., and Rotter, S. (2003). Inference of hand movements from local field potentials in monkey motor cortex. *Nature Neuroscience* 6, 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, 37-99.

Mitzdorf, U. (1987). Properties of the evoked-potential generators: current source-density analysis of visually evoked-potentials in the cat cortex. *International Journal of Neuroscience* 33, 33-59.

Murthy, V. N., and Fetz, E. E. (1996). Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. *J Neurophysiol* 76, 3949-67.

Nobre, A., Allison, T., and Mc Carthy, G. (1994). Word recognition in the human inferior temporal lobe. *Nature* 372, 260-263.

O'Keefe, J., and Speakman, A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Exp Brain Res* 68, 1-27.

Pesaran, B., Pezaris, J., Sahani, M., Mitra, P., and Andersen, R. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature Neuroscience* 5, 805-811.

Quiroga, R., Nadasdy, N., and Ben-Shaul, Y. (2004). Unsupervised spike sorting with wavelets and superparamagnetic clustering. *Neural Computation* 16, 16161-1687.

Riesenhuber, M., and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience* 2, 1019-1025.

Scheffe, H. (1953). *The Analysis of Variance* (New York: Wiley).

Schwartz, E., Desimone, R., Albright, T., and Gross, C. (1983). Shape-recognition and inferior temporal neurons. *PNAS* 80, 5776-5778.

Speckmann, E., and Elger, C. (1993). Introduction to the Neurophysiological Basis of the EEG and DC potentials. In *Electroencephalography*, E. Niedermeyer and F. Da Silva, eds. (New York: Lippincott Williams and Wilkins).

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

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, 90-9.

Tanaka, K. (1993). Neuronal mechanism of object recognition. *Science* 262, 685-688.

Wang, G., Tanaka, K., and Tanifuji, M. (1996). Optical imaging of functional organization in the monkey inferior temporal cortex. *Science* 272, 1665-1668.

Yu, A., and Kreiman, G. (1999). Classification of extracellular microelectrode recordings from the human brain (Pasadena: Caltech).