Intracranial Recordings Reveal Transient Response Dynamics During Information Maintenance in Human Cerebral Cortex

Niv Noy,^{1,2} Stephan Bickel,³ Elana Zion-Golumbic,^{4,5} Michal Harel,²
Tal Golan,⁶ Ido Davidesco,⁶ Catherine A. Schevon,⁷ Guy M. McKhann,⁷
Robert R. Goodman,⁷ Charles E. Schroeder,^{4,5}
Ashesh D. Mehta,³ and Rafael Malach^{2*}

¹The Gonda Multidisciplinary Brain Research Center, Bar-Ilan University, Ramat Gan, Israel

²Department of Neurobiology, Weizmann Institute of Science, Rehovot, Israel

³Department of Neurosurgery, Hofstra North Shore LIJ School of Medicine, Hempstead,
New York

⁴Department of Psychiatry, Columbia University College of Physicians and Surgeons,
New York, New York

⁵Cognitive Neuroscience and Schizophrenia Program, Nathan Kline Institute, Orangeburg,
New York

⁶The Edmond and Lily Safra Center for Brain Sciences, Interdisciplinary Center for Neural
Computation, Hebrew University, Jerusalem, Israel

⁷Department of Neurology, Columbia University College of Physicians and Surgeons,
New York, New York

Abstract: Despite an extensive body of work, it is still not clear how short term maintenance of information is implemented in the human brain. Most prior research has focused on “working memory”—typically involving the storage of a number of items, requiring the use of a phonological loop and focused attention during the delay period between encoding and retrieval. These studies largely support a model of enhanced activity in the delay interval as the central mechanism underlying working memory. However, multi-item working memory constitutes only a subset of storage phenomena that may occur during daily life. A common task in naturalistic situations is short term memory of a single item—for example, blindly reaching to a previously placed cup of coffee. Little is known about such single-item, effortless, storage in the human brain. Here, we examined the dynamics of brain responses during a single-item maintenance task, using intracranial recordings implanted for clinical purpose in patients (ECoG). Our results reveal that active electrodes were dominated by transient short latency visual and motor responses, reflected in broadband high frequency power increases in occipito-

Additional Supporting Information may be found in the online version of this article.

Correction added after online publication 3 July 2015. The corresponding author's last name was corrected.

Contract grant sponsor: EU FP7 VERE; Contract grant number: 257695; Contract grant sponsor: EU FP7 Flagship HBP; Contract grant number: 604102; Contract grant sponsor: ICORE program (ISF); Contract grant number: 51/11; Contract grant sponsor: Helen and Martin Kimmel (to R.M.); Contract grant sponsor: The Page and Otto Marx Jr. Foundation (to S.B., A.D.M.); Contract

grant sponsor: National Institute of Mental Health; Contract grant numbers: MH093061 (to E.Z.) and MH086385 (to C.E.S.)

*Correspondence to: Rafael Malach, Weizmann Institute of Science, Rehovot 76100, Israel. E-mail: rafi.malach@gmail.com

Received for publication 22 September 2014; Revised 17 June 2015; Accepted 18 June 2015.

DOI: 10.1002/hbm.22892

Published online 3 July 2015 in Wiley Online Library (wileyonlinelibrary.com).

temporal, frontal, and parietal cortex. Only a very small set of electrodes showed activity during the early part of the delay period. Interestingly, no cortical site displayed a significant activation lasting to the response time. These results suggest that single item encoding is characterized by transient high frequency ECoG responses, while the maintenance of information during the delay period may be mediated by mechanisms necessitating only low-levels of neuronal activations. *Hum Brain Mapp* 36:3988–4003, 2015. © 2015 Wiley Periodicals, Inc.

Key words: ECoG; iEEG; gamma; short term memory; working memory; delayed activity; maintenance

INTRODUCTION

A central characteristic of human cognition is the ability to maintain necessary information on a short time basis and use it for task performance after a delay. This short term information maintenance has been studied extensively in behaving monkeys [Miller et al., 1996; Rainer et al., 1998]. In the human cortex, a large body of research examined visual working memory using noninvasive functional magnetic resonance imaging (fMRI) and scalp electroencephalography (EEG) methods [Courtney et al., 1997; Schluppeck et al., 2006; Tong, 2013]. A series of studies have used ECoG in examining human working memory processes and their modulations under memory load. Using the well-known Sternberg paradigm involving memory of letters and digits, a parametric increase in gamma (~30–90 Hz) power during the delay period, was reported [Axmacher et al., 2007; Howard et al., 2003; Meltzer et al., 2008].

These studies, however, focused on a specific kind of maintenance, in which subjects are typically asked to maintain information about a number of items simultaneously. Such multi-item maintenance necessitates the activation of specific strategies—for example, rehearsing the different items' names in phonological loops [Bor et al., 2003; Conway et al., 2005; Curtis and D'Esposito, 2003; Curtis and Lee, 2010; Miller, 1956]. Furthermore, strategies associated with this kind of tasks are often mentally demanding and require allocation of attentional resources. The attentional and effort requirements may activate general-purpose cortical areas that are not necessarily specialized for working memory per se. For example, recent evidence has suggested the existence of a general-

purpose frontal system [Fedorenko et al., 2013] that is likely to be engaged during difficult working-memory tasks.

In contrast to the literature mentioned earlier, during naturalistic conditions there are numerous effortless situations that require the maintenance of single item information (e.g., blindly picking up a cup of coffee). Little is known about the neuronal dynamics that underlie such natural and simple behaviors.

In this study, we aimed to examine the cortical dynamics during an easy (one of two), single item maintenance task. The experiment was conducted in six epileptic patients who performed a delayed response (DR) task—overall number of examined ECoG recording sites was 617. Our results demonstrate a striking dominance of short latency transient visual responses evident throughout the cerebral cortex.

In contrast, delayed maintenance activity in the period between stimulation and response was remarkably weak. Particularly, no cortical site showed a significant and sustained delay period activation in the broad high frequency range that lasted up to the time of motor execution. This suggests that during the mental maintenance of a single and easy to remember item, cortical responses are dominated by transient activations at short latency, while the delay period is characterized by minimal (if any) levels of activity.

METHODS

Participants

Recordings of electrical activity were obtained from six neurosurgical patients (four female, aged 27.5 ± 9.7), with pharmacoresistant epilepsy, monitored for potential

TABLE I. Patients' characteristics

Patient ID	Gender	Age	Handedness	Implanted hemisphere	# of implanted electrodes	Seizure location
1	F	22	Right	L	116	Left O
2	F	35	Right	Bilateral	128	Left F-T
3	M	21	Right	R	128	Right F-T
4	F	21	Right	R	104	Right O
5	F	22	Left	L	110	Left T
6	M	44	Right	L	92	Left F-P

F, female; M, male; L, left; R, right; O, occipital; F-T, fronto-temporal; T, temporal; F-P, fronto-parietal.

surgical treatment (see Table I). Four patients were hospitalized at the North Shore University Hospital (NSUH) and two at Columbia University Medical Center. All provided fully informed consent according to the National Institutes of Health guidelines, as monitored by the local institutional review board, in accordance with the ethical standards of the Declaration of Helsinki.

ECoG recordings were made over the course of clinical monitoring for spontaneous seizures. The decision to implant, the electrode location, and the duration of implantation were determined entirely on clinical grounds without reference to this investigation. The patients were informed that participation in this study would not jeopardize their clinical care in any way and that they could withdraw from it at any time. In fact, it should be noted that participating in these studies may have actually benefited the patients by closer monitoring, multiple observers, and scientific personal confirmation of electrode identities and localization.

Stimuli and Task

The participants were asked to perform a DR task in which images were presented on a standard laptop display (60 Hz refresh rate) in a distance of about 60 cm. Participants' responses were recorded via a mouse button press.

The pictures were color images of faces [Minear and Park, 2004], $\sim 15^\circ$ (640 pixels) in width, superimposed with a small red fixation cross. Each trial in the experiment consisted of a target image presented for 250 ms followed by a 1,500–4,000 ms blank interval and then a 250 ms long auditory cue (300–540 Hz, pure tone). The task was to indicate the gender of the face by a button press (male left button, female right button) only after the auditory cue was given (Fig. 1).

Each participant underwent 88 trials, divided to two 5 min runs consisting of 44 trials each. The intertrial interval was 1,500–4,000 ms.

Data Acquisition

Each participant was implanted with subdural electrode arrays containing 92–128 contact electrodes (see Table I). In total, 678 electrodes were implanted. Electrodes were arranged in one-dimensional strips or in two dimensional grids placed directly on the cortical surface.

The signal was sampled at a rate of 2,000 Hz (NSUH) or 1,000 Hz (Columbia). Stimulus-triggered electrical pulses were recorded along with the ECoG data for precise synchronization of the stimuli with the recorded electrical responses.

Electrode Localization

Computed tomography (CT) scans following electrode implantation were coregistered to the postoperative mag-

Figure 1.

Experimental paradigm. On each trial, a face image was presented for 250 ms, followed by a fixation only interval of 1,500–4,000 ms and a 250 ms auditory cue. On hearing the cue the participant pressed a button to indicate whether the presented image was of a male or a female. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

netic resonance imaging (MRI). Preoperative and postoperative MRIs were both skull-stripped using the BET algorithm from the Oxford Centre for Functional MRI of the Brain (FMRIB) software library (FSL; www.fmrib.ox.ac.uk/fsl/) followed by coregistration to account for possible brain shift caused by electrode implantation and surgery. Electrodes were identified in the CT using BioImagesuite (www.bioimagesuite.org). The coordinates of the electrodes were then normalized to Talairach space [Talairach and Tournoux, 1988] and rendered in BrainVoyager software in two dimensions as a surface mesh, enabling precise localization of the electrodes both with relation to the participant's anatomical MRI scan and in standard coordinate space. For joint presentation of all participants' electrodes, the locations were projected onto a cortical reconstruction of healthy subject from a previous study—in this process, we left out 20 electrodes which were undetectable in the CT and their coordinates could not be extracted.

Following the definition of each electrode Talairach coordinates, we divided them into seven functional groups, corresponding to rough anatomical segregation: low order visual (LOV), high order visual (HOV), parietal

(Prt), auditory (Aud), postcentral (PoC), precentral (PrC), and prefrontal (Frt) electrodes. This division left out seven electrodes that could not be well defined according to their Talairach location. These electrodes, together with the 20 electrodes mentioned above, 8 that recorded no signal at all, and 26 with highly corrupted signal (61 altogether) were left out of analysis leaving us with 617 electrodes.

Precise information about the location of the epileptic center and the corresponding electrodes was not available to us. As detailed in Table I, only a very rough definition, at the level of hemisphere and lobe was present. We therefore did not exclude any electrode based on clinical criteria (although analytical rejection of trials was performed—see below). Rather, we defined all electrodes in the reported areas as “suspected electrodes”—a set of electrodes to be examined more carefully in the group analysis.

Data Preprocessing

Potential 60 Hz electrical interference and its harmonics (multiplications of 60 Hz) were removed from the raw electrical signals using a linear-phase notch FIR filter.

Then, to discard non-neuronal contributions from the extra cranial reference electrode, a de-referencing procedure was adopted. For each patient separately, the ECoG signal from all (not excluded) electrodes was averaged, producing a “common” signal time course. This “common” component was subtracted from each electrode raw time course.

The recordings provided by NSUH, which were sampled at 2,000 Hz, were down sampled to 1,000 Hz to match the sampling frequency of the data from the Columbia University Medical Center. The z-score of this signal was used to create averaged event related potentials (ERP) while for induced responses we used the broadband high frequency (BHF, 50–150) and beta (15–25) power modulations.

For the calculation of the BHF power modulations, the signal was band passed in the frequency range 50–150 Hz using linear-phase FIR filters (window size ranging from 20 to 60 ms according to the subrange—see below). The power modulations were extracted by taking the absolute value of the filtered signal’s Hilbert transform [Davidesco et al., 2013].

ECoG power spectrum typically has a roughly $1/f^2$ profile, which results in a more dominant contribution from the lower end of the frequency spectrum. To compensate for this, we calculated a “flattened” form of the BHF by dividing the 50–150 frequency range into 10 Hz subranges, filtering each separately, dividing the power modulation by its mean value, and finally averaging across subranges [Davidesco et al., 2013; Fisch et al., 2009; Lachaux et al., 2005; Vidal et al., 2010]. A similar process was used in order to extract beta power modulations in the frequency range 15–25 (divided to 2.5 Hz subranges).

Data processing was performed using MATLAB. For filtering, we used original and adapted EEGLAB [Delorme and Makeig, 2004] code.

Noisy trials rejection was performed by automatically detecting extreme peaks (>3 or <-3 standard deviations) in the time course that appear in more than 20 electrodes simultaneously (separately for BHF (50–150), beta (15–25), and raw ECoG signal). These trials were excluded from analysis altogether for the sake of consistency between electrodes and brain regions.

To evaluate possible unwanted effects of the de-referencing procedure, the entire preprocessing was conducted twice: once as described above and once without the “common” signal removal. The second dataset was used for comparison with the results reported below.

Quantitative Definitions of Electrode Responses

Three time periods, representing different stages during each trial, were of interest in our analysis. The first is the stage of picture presentation, termed here “visual period” (V); the second is the delay or maintenance period, which was divided to two time intervals as explained in the next paragraph; and the third is the cue and report, termed here “auditory\motor period” (AM).

Because of the jitter between the presentation of the picture and auditory cue, there was no single delay period, lasting from picture onset to the onset of the auditory cue, which was present in all trials. To address this issue, we analyzed two time intervals representing the delay period. The first interval was defined relative to the picture onset using the longest possible delay available in all trials—this interval is referred to as “transient delay period” (TD); conversely, the second delay interval was chosen relative to the auditory cue to evaluate maintenance activity up to the time of the cue—termed “sustained delay period” (SD).

The exact time intervals that we used for the definition of these four periods are specified in Table II. Slightly different definitions were used for the BHF (50–150) and beta (15–25) signals; this is because different temporal resolution and smoothing is associated with each. An illustration of the BHF (50–150) intervals can be found at the bottom of Figure 2.

TABLE II. Baseline and response intervals

	Interval	BHF (50–150)	Beta (15–25)
Relative to picture	Baseline	–250 to 0	–400 to –100
	V	0 to 750	100 to 650
	TD	1,000 to 1,750	1,100 to 1,650
Relative to cue	SD	–750 to 0	–650 to –100
	AM	0 to 750	100 to 650

The specific time intervals that were analyzed are presented for the two frequency ranges. The power modulations in these intervals were used for statistical testing and for AUC calculation. BHF, broadband high frequency; V, visual; TD, transient delay; SD, sustained delay; AM, auditory\motor.

Figure 2.

Trial matrices of selected electrodes. Trial matrices of representative selected electrodes from several brain regions are presented. In each panel, rows are trials ordered by their delay length (picture to auditory cue interstimulus interval—from 1.5 s at the top to 4 s at the bottom). The trial matrices are divided to two parts. In the left part of the panels, trials are locked to the onset of the visual stimulus (red vertical line) while in the right side they are locked to the onset of the auditory cue (green vertical line). The x-axes represent time in seconds relative to the stimuli. Diagonal

Activity was defined as the area under the curve in a certain time interval (the mean of the power in this interval—AUC). Electrodes were considered responsive if a significant (corrected, see below) increase in activity was detected in one of the time intervals (V, TD, SD, or AM) compared with a baseline interval [Davidesco et al., 2013]. We used a time interval prior to the beginning of each trial as the baseline—that is, the time before picture presentation onset, when no (experiment related) task should have taken place.

Responsive electrodes were divided into three groups according to the timing of their responses: visual electro-

des (V electrodes)—those that showed an increase in AUC in the short latency time interval relative to picture presentation; transient delay electrodes (TD)—those that showed an increase at the longer latency; and auditory\motor responsive electrodes (AM), that showed an increase after the presentation of the auditory cue. It should be noted that no electrode was significantly activated during the sustained delay time interval preceding the auditory cue (SD).

The statistical significance of the responses was determined by means of one sample one tailed Student's *t* test performed on the correct trials' AUC distribution. *P*-values

The statistical significance of the responses was determined by means of one sample one tailed Student's *t* test performed on the correct trials' AUC distribution. *P*-values

were calculated for each of the 617 electrodes examined in this study separately and then corrected for multiple comparisons to a 0.05 false discovery rate [FDR, Benjamini and Hochberg, 1995]. This process was done separately for each of the time interval categories mentioned above.

The latency of visually induced BHF (50–150) responses was evaluated as well. It was defined as the first time point in which a signal, averaged across trials, reached a quarter of its maximum activity in the V time window.

Eventually, we calculated three average responses for each brain region (see Electrode Localization section) by averaging the induced BHF (50–150) responses of the significant electrodes in each set—this was done once locked to the picture presentation and once locked to the auditory cue. The same procedure was done for the ERP signal as well. In this case, the absolute values of the single electrodes ERPs were used to reduce sensitivity to local dipole inversions.

Finally, we have also checked for the existence of negative responses in both frequency ranges. This was done by the exact same process described above only this time with a search for activity decrease instead of increase.

Bayesian analysis: In addition to the standard statistical approach, we have conducted a Bayesian analysis of the BHF (50–150) responses. Following the approach of Bayesian inference to the multiple comparison problem [Gelman et al., 2012], we integrated data across electrodes into a single model. Single trial responses were modeled by a mixed-effects multilevel linear model, with electrode, trial and patient ID as random factors and the sampling window (activation or baseline) as a fixed factor. This was done separately for each of the time intervals of interest.

We compared three model variations: the full model (H2), allowing for different effects of the activation versus baseline in different electrodes, main effect model (H1), allowing only for a uniform effect of activation versus baseline, and the null model, with no activation versus baseline effect at all (H0). Computations were conducted using the R-project BayesFactor package [Morey et al., 2015]. The JSZ priors with the default parameters set by the “lmBF” function were used. Each hypothesis model produced a likelihood score and a Bayesian factor (BF) was extracted by comparing pairs of models.

Behavioral Task Performance

Valid trials were defined as those in which participants managed to perform the DR task, that is, only trials in which no button press was made during the delay period. The percent of correct responses out of the valid trials, together with the mean reaction time (RT) in correct response trials were calculated for each participant. Both were calculated once for the whole session and once, to evaluate possible practice effects, for the first and last quarter of trials separately. At all stages of analysis only correct trials were considered.

We used the nonparametric Wilcoxon signed rank test to examine the difference between the first and the last quarter of trials. The use of a nonparametric statistic was preferred here as the small sample size ($n = 6$) did not allow relying on the central limit theorem to assume a normal distribution of the means.

RESULTS

Our study was based on ECoG recordings in six participants (see Table I), each implanted with 92–128 electrodes (678 in total, of which 617 were analyzed—see Methods). Figure 1 depicts the experimental paradigm. The study was aimed at examining a simple single item maintenance task (one of two options). In the experiment, termed here DR, pictures of faces of both genders were presented for 250 ms followed by an interval of a 1.5–4 s terminated by an auditory cue (beep sound). Participants were instructed to press one of two buttons on hearing the cue, to indicate whether the preceding picture was of a male or a female.

Note that no attempt was made to constrain the possible single item maintenance strategy of the patients, who could maintain the information either as pictorial information, linguistic information (male/female), or even as the final motor output decision (left/right button press). The critical point is that to correctly perform the task, the patients had to maintain decision-relevant information during the delay period and use this information to execute the correct motor response. Indeed, each maintenance strategy was expected to produce an activation signal in different parts of the cortex (e.g., somato-motor for anticipatory button press, visual cortex for maintained visual images, and language cortex for maintenance of the gender name). However, given the substantial electrode coverage, it was reasonable to expect at least some indication to any of the possibilities mentioned above.

Behavioral Responses

As expected from the simple nature of the single item DR task, the patients performed very well right from the beginning. Examining the difference between the first and the last quarter of trials revealed no significant practice effect, both in RT ($P = 0.39$, two sided Wilcoxon signed rank test, see methods) and in the percent of correct responses ($P = 0.79$, two sided Wilcoxon signed rank test). Overall accuracy in valid trials and the RT for correct responses were $93 \pm 4.5\%$ and 621 ± 155.2 ms, respectively.

Omitted, nonvalid, trials constitute $6.4 \pm 3.2\%$ (mean $\pm$ SD) of trials. As the reason for failure in these trials could not be specifically attributed to failure in maintenance, they could not serve as an estimate of the difficulty level of the maintenance task.

Indeed, considering the simplicity of the task, it was reasonable to expect even higher performance levels. However, it is important to remember that an ECoG

experimental setting is far from ideal and that patients' task performance is usually poor in comparison to that of healthy subjects. Patients' performance is mainly effected by their postoperative overall condition as well as by the noisy and distracting environment of the hospital room.

In this light, when evaluating the difficulty level of the task, it is impossible to separate the actual difficulty of maintaining information throughout the delay period, from other confounds that impair performance. Using only valid trials leads to a more accurate representation of the actual difficulty level, but may still result in an overestimate as there is no way to determine the reason for failure in a valid trial—wheatear, it is a maintenance failure or some other irrelevant reason (e.g., confusion in the button press itself or lack of attention to the picture presentation).

Neuronal Responses

As we and others have reported previously [Henrie and Shapley, 2005; Lachaux et al., 2012; Mukamel et al., 2005; Nir et al., 2007; Ray and Maunsell, 2011], an important index of overall neuronal firing rate is the power in the BHF range (50–150 Hz, BHF) of the local field potential, which was shown to play a role in working memory as well [Howard et al., 2003]. We, therefore, focused our analysis on this frequency band, as well as the more conventional stimulus locked evoked potential (ERP, see Methods).

In addition, we examined power modulation in the beta (15–25) frequency range, which was also demonstrated to be linked to maintenance activity [Tallon-Baudry et al., 2001]. Lower frequencies analysis was not included in the current report as the temporal resolution associated with it is too low, and therefore, with regards to the DR experimental design, do not allow a clear distinction between perception, encoding, and maintenance processes.

To examine the dynamics of the neuronal responses during the various stages of the DR task, we averaged the BHF (50–150) power modulations as well as the ECoG signals (ERP) and beta (15–25) power modulations across individual trials, time locked to the visual stimulus onset and time locked to the auditory cue. In this article, we mainly discuss positive BHF (50–150) effects. Hence, “responsive” or “active” electrodes, unless otherwise specified, refers to electrodes that showed a significant increase in power at this frequency range.

In the BHF (50–150) analysis we searched, separately, for four types of neuronal responses: transient visual responses (V responses—significant increases in BHF (50–150) power only during the time windows of 0–750 ms from picture onset), transient DR (TD—increased BHF (50–150) power during the time window of 1,000–1,750), sustained DR (SD—increased BHF (50–150) power during the time window –750 to 0 prior to the auditory cue) and auditory/motor responses (AM—increased BHF (50–150) power during the time window 0–750 postauditory cue).

All responses were compared with a baseline of –250 to 0 before picture presentation (see Methods for details and Fig. 2 for an illustration of the time intervals).

This analysis was conducted for beta (15–25) power modulations as well, with slightly different time windows (see Methods and Table II). In addition, a search for negative responses (decrease in activity), in both frequency bands, was performed.

Remarkably, of the 617 electrodes examined in the study *none* showed a sustained increase in delay-period activity at both frequency ranges—that is significant activity ($P < 0.05$ corrected, see Methods) that lasted up to the auditory cue/motor response (SD interval). None showed *decrease* in BHF (50–150) SD activity as well, while in the beta frequency range (15–25) three electrodes demonstrated a weak negative SD effect.

An obvious reservation concerns the statistical power of the tests we used. It is inherently problematic to *prove* a null hypothesis, hence, the fact that no evidence for sustained neuronal activity during the delay period was found does not necessarily mean it does not exist.

To further investigate whether a few cortical sites did show a sustained neuronal activation that could underlie maintenance processes, we specifically examined the electrodes showing the highest level of such activity (still non-significant although). Figure 2 shows the responses of the two electrodes that manifested the highest sustained activity in our entire sample. It depicts trial matrices sorted by the trials' interstimulus intervals (ISI). The top two panels are of electrodes with a low P -value for SD interval activity, these electrodes are the ones that provided the strongest evidence for activity in this interval. Still, by inspecting their trial matrices, it can be readily appreciated that there was no discernible sustained activation effect. The other five panels present representative electrodes from several regions and activity profiles.

In addition, we adopted a meta-analysis-like approach by treating the whole set of SD P -values (without correcting for multiple comparison) as a result of multiple independent tests aimed at refuting the same zero hypothesis. This was done by combining all of the P -values using Fisher's method, which accumulate the evidence from small values and can be very sensitive even to small effects “hidden” within the dataset. This test came out not significant as well with a $P = 0.89$.

Finally, we examined the results using the more sensitive Bayesian statistical approach [Gelman et al., 2012]. For the SD condition, there was a strong evidence in favor of the *null* hypothesis (H_0) compared with the full (H_2) and main (H_1) models ($BF_{20} < 10^{-10}$ and $BF_{10} = 0.09 \pm 2.98\%$, respectively).

In the other time intervals we have inspected, the picture was quite different. At the BHF (50–150) range, 162 electrodes (~26.3%) showed significant V activation, 25 (~4.1%) electrodes showed TD activity and 88 (~14.3%) showed a significant AM response. Negative responses

Figure 3.

Location of ECoG cortical recording sites in all participants. Sites were obtained from CT and MRI scans, superimposed on a cortical reconstruction of one Subject from a previous study. (A) The electrodes are shown on a lateral view of the cortical hemispheres. Color coding represents the electrode category: red is for visual electrodes, green for TD electrodes, and blue is for auditory/motor electrodes. Gray represents nonsignificant electrodes. Electrodes that responded to more than one category are marked with a pie chart representing the relative activity in each category (defined by the AUCs). (B) Same

electrodes, with the same color coding as in panel A are shown on a flattened map of the brain. Color contours mark the division of the cortex to regions. Filled arrowheads indicate occipital and temporal cluster of electrodes and unfilled arrowheads indicate frontal and parietal clusters. LOV, low order visual; HOV, high order visual; Prt, parietal; Aud, auditory; Pos, post-central; Pre, precentral; Frt, prefrontal; P, posterior; A, anterior; LH, left hemisphere; RH, right hemisphere. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

were evident at 40 (~6.5%) electrodes in the V interval, 24 (~3.9%) in the TD interval, and 13 (~2.1%) in the AM interval. At the beta frequency range (15–25), the distribution of positive significant electrodes at the V, TD, and AM intervals was 49 (~7.9%), 9 (~1.5%), and 6 (~1%), while the distribution of negative significant electrodes was 94 (~15.2%), 54 (~8.8%), and 116 (~18.8%), respectively.

The Bayesian analysis results were in line with the standard statistics results, showing a very strong evidence in favor of the full model (H2) in each of the three time interval ($BF_{20} < 10^{-10}$ and $BF_{21} < 10^{-10}$ for all).

Figure 3 depicts the distribution of the entire set of positive responsive BHF (50–150) electrodes presented on folded (A) and unfolded cortical formats (B). Color coding indicates which of the three types of responses (V, TD, and AM) was significant—electrodes that belonged to more than one of the categories were marked with a pie chart representing the relative strength of activity (AUC) in each category. Nonresponsive electrodes are depicted in

gray. Note that our coverage was widespread across all cortical lobes.

As can be seen, and in agreement with previous studies, many of the visually responsive electrodes overlapped the well-known visual areas—both early visual areas in the occipital lobe as well as high order temporal lobe regions (filled arrow-heads).

Importantly, consistent visually induced responses could be discerned also outside visual areas, with prominent clusters in the parietal and frontal lobes (unfilled arrow-heads). A few electrodes were active at more than one interval: both V and AM (26, ~4.2%), both V and TD (6, ~1%), both TD and AM (4, ~0.6%), or V, TD and AM (8, ~1.3%).

To simplify the presentation of ECoG signal responses, the electrodes were subdivided into seven groups based on anatomical considerations (see color contours in Fig. 3B). It is important to note that these are rough estimates based only on anatomical criteria, as fMRI mapping was unavailable in these patients.

TABLE III. Distribution of electrodes across patients

Patient ID	LOV	HOV	Prt	Aud	PoC	PrC	Frt	All
1	11 (6)	36 (6)	7 (0)	15 (2)	17 (8)	9 (5)	13 (5)	108 (32)
2	10 (3)	28 (5)	8 (0)	17 (3)	9 (4)	12 (5)	40 (1)	124 (21)
3	11 (3)	20 (1)	7 (0)	19 (2)	10 (3)	24 (9)	32 (3)	123 (21)
4	24 (24)	16 (9)	15 (8)	10 (7)	15 (11)	9 (5)	1 (1)	90 (65)
5	21 (15)	28 (11)	12 (5)	10 (4)	14 (6)	3 (3)	0 (0)	88 (44)
6	4 (4)	16 (9)	0 (0)	16 (6)	19 (12)	17 (7)	12 (2)	84 (40)
All	81 (55)	144 (41)	49 (13)	87 (24)	84 (44)	74 (34)	98 (12)	617 (223)

The table presents the number of electrodes that were analyzed in each patient (different from the number of implanted electrodes in Table I), divided by brain region. In brackets, the number of BHF (50–150) responsive electrodes—V, TD, or AM. LOV, low order visual; HOV, high order visual; Prt, parietal; Aud, auditory; PoC, postcentral; PrC, precentral; Frt, prefrontal.

Region averages were extracted across active electrodes. The distribution of electrodes and significance across patients and brain regions is presented in Table III. It should be noted that with the exception of prefrontal cortex, where the majority of electrodes was derived from four patients—the coverage of other cortical areas was fairly balanced across patients.

Figure 4 depicts the average (across electrodes) BHF (50–150) power in the DR experiment in each cortical area presented separately for the three electrode sets: visually responsive (V), transient delay responsive (TD), and audio/motor responsive (AM). Responses are presented both time locked to the visual image presentation (red, left), and time locked to the auditory cue (green, right) after the variable delay indicated by a dotted line.

Inspecting this rich set of data revealed several important aspects. First, with regard to V electrodes (left column) it can be seen that transient visual responses were prominent, as expected, in visual cortex, but could be found, albeit with reduced amplitude, throughout the entire cortical mantle. Of particular interest are V electrodes located in the auditory and frontal areas (postcentral, precentral, and prefrontal). Note that in these sets, the typical response consisted of a short-latency visual transient and a re-activation on the auditory cue onset. Surprisingly, a substantial number of electrodes in precentral cortex (23, ~31.1% of the electrodes in this region) belonged to the V electrode group—more than the number of AM electrodes in this region (19, ~24.4%). These electrodes' average V activity was stronger than their AM response.

Searching for TD electrodes—that is, electrodes responding during the delay (middle column) revealed a limited set (25, ~4.1% of all electrodes) that showed rather weak and variable response. These were located mainly in postcentral and HOV cortices and, as our statistical analysis revealed, their responses were not sustained throughout the delay period and declined to baseline prior to cue onset.

Finally, AM electrodes were, as expected, most prominent in auditory and frontal cortices. Interestingly, a lim-

ited number of the electrodes in HOV areas (6, ~6.8%) belonged to this group and produced higher average activation in response to the auditory cue compared with the response to the visual stimuli.

ERP responses measured in the same electrodes largely recapitulated the main effects observed in the BHF (50–150) range—that is, widespread activation in V electrodes on the one hand, a robust response to the auditory cue in AM electrodes in frontal regions on the other and a transient activation during the delay period for the subset of TD electrodes (see Fig. 5).

To allow a direct depiction of the latency of the BHF (50–150) visual responses outside visual cortex, the averaged responses from all electrodes in all of the groups were superimposed on each other. Figures 6 and 7 depicts such superposition for BHF (50–150) and ERP, respectively—showing the relative amplitudes of the visual and auditory\motor average responses in all significantly responsive electrodes (V, TD and AM sets together). As can be seen, while there were substantial reductions in visual response amplitudes at more frontal cortical regions, the onsets of the responses were quite similar—suggesting a rapid progression of signals from visual to frontal regions.

A quantitative analysis of the average latencies in each brain region supported the impression arising from Figure 6. Table IV presents the estimated latencies of neuronal responses for each patient and brain region separately. The important point here is that the latencies in the different regions are all up to 250 with great similarity between region averages. Inspecting the distribution of latencies between patients indicates that this effect was not driven by any specific one of them and is a general phenomenon, as is also indicated by comparing random effect (RFX, giving equal weights to each patient) with the fixed effect (FFX, pooling all electrodes together) averages.

To rule out the possibility that the similar latencies in the BHF (50–150) signal were an artifact of the temporal smoothing associated with its calculation (see Methods), we repeated the same analysis focusing only on the high end of this band (100–150 Hz). Here, as well, the onset of

Figure 4.

Induced BHF (50–150) averages in different cortical regions in the three electrode sets. Averaged BHF (50–150) activation for visual (red) and auditory/motor (green) stimuli. Data is shown separately for V electrodes (A), TD electrodes (B), and AM electrodes (C). Thick lines indicate average, and transparent

area indicates $\pm$ SEM. The event related averaging of the BHF (50–150) is locked to the face stimuli and to the auditory cues (left and right vertical lines, in each graph, respectively). Y axes are percent signal change. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

the visually induced response in precentral areas coincided with that of the responses in HOV areas.

Could the minimal activity in the delay period be a consequence of practice or repeated performance? To examine this possibility we divided the entire recording session to four parts and analyzed both delay periods' responses (TD and SD) separately for the first and last quarters of trials. Comparing the delay period responses in the beginning (first quarter of trials) versus the end (last quarter of trials) of the experiment failed to reveal any electrode that shows a significant difference, neither for TD nor for SD (all P values >0.05 , FDR corrected).

Finally, we set to examine the electrode set that may have been located in epileptic tissue (see Methods). The proportions of responsive electrodes out of the total number of electrodes in the two sets were practically the same, with 35.8% for "healthy" electrodes and 38.4% for potentially epileptic ones. There were not enough responsive electrodes in this subset to allow for a systematic comparison between regions and significance types, but, as can be appreciated from Supporting Information Figures 1 and 2 there did not seem to be any specific trend for these. In most regions, the activity average of the potentially epileptic electrodes corresponded to that of the other electrodes.

Figure 5.

Absolute ERP averages in different cortical regions in the three electrode sets. Same as Figure 4 except here we plot the absolute ERP average. Y axes are absolute z-scores. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

Most importantly, even when we excluded the set of potentially epileptic electrodes from the multiple comparison correction, we still failed to find a single electrode that passed the significance threshold for SD responses.

DISCUSSION

Dominance of Transient Responses During the Delayed Response Task

The most clear cut outcome of this study was the significant imbalance between delayed period activity and the transient visual and auditory/motor responses during the

DR experiment. In contrast to a substantial number of cortical sites showing significant V (~26.3%) or AM (~14.3%) responses—we failed to find even a single cortical site exhibiting SD period activation lasting up to the motor response. Furthermore, the Bayesian statistics approach also provided a strong support for the null hypothesis in the SD time interval while supporting the H2 model (condition $\times$ electrode interaction) for all other time intervals of interest, hence, providing a converging support for our findings.

Even when considering transient responses during the delay period as a DR—these constituted a small minority of sites (25, ~4.1%). Direct comparison of the activation levels (AUCs) of all of the electrodes during the visual

Figure 6.

Comparison of BHF (50–150) responses in all cortical regions. Averages of the BHF (50–150) response in all responsive electrodes (from all sets) in each cortical region, superimposed. Thick lines indicate average, and transparent areas indicate $\pm$ SEM. Red bar on the horizontal axis indicate the duration of the visual stimuli and green the duration of the auditory cue. The event related averaging of BHF (50–150) is locked to the onset of the visual stimuli and to the auditory cue (left and right vertical lines at time zero, respectively). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

and the two delay periods (separately) revealed extremely significant differences ($P < 0.001$ for both comparisons, two sample paired t -test).

These results appear to differ from previous ECoG studies that used the Sternberg letter memorizing test and showed a load-dependent change in delay period activity, mainly in fronto-lateral sites [Axmacher et al., 2007; Howard et al., 2003; Mainy et al., 2007; Meltzer et al., 2008]. However, it should be noted that a number of cognitive aspects other than the memory itself are expected to be modulated by load, for example, attentional demand and arousal levels, so it is not clear whether the reported association of delay activity and load were specific to working memory [Curtis and Lee, 2010].

Figure 7.

Comparison of ERP responses in all cortical regions. Same as Figure 6 for absolute ERP responses. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

In fact, the DR experiment was specifically designed to keep such processes at a minimal and constant level throughout. Even the minimal load condition in the studies mentioned above is more demanding than the DR task as remembering one out of 26 letters is more challenging than holding a binary choice. Additionally, the Sternberg memory paradigm is usually based on serial presentation of letters which means participants have no way of knowing what the difficulty level of a trial is until its end—this alone may promote different strategies and increase the load.

Furthermore, it should be noted that the Sternberg working memory task is based on linguistic (letters) material while in our experiment patients could perform the task also by memorizing nonverbal aspects such as face, gender, or motor plans which may have engaged different, nonlinguistic cortical processes.

Even if we disregard all of the differences, a closer look at the studies cited above reveals that most of them do not

TABLE IV. Response latency distribution

Patient ID	LOV	HOV	Prt	Aud	PoC	PrC	Frt	FFX mean
1	108 (6)	195 (5)	–	–	501 (1)	189 (3)	96 (2)	170 (17)
2	79 (3)	111 (5)	–	251 (2)	261 (2)	197 (1)	371 (1)	170 (14)
3	64 (3)	40 (1)	–	–	–	137 (8)	158 (1)	114 (13)
4	136 (24)	169 (8)	193 (8)	192 (5)	177 (4)	141 (3)	282 (1)	161 (53)
5	115 (14)	188 (9)	146 (5)	131 (2)	247 (3)	162 (3)	–	153 (36)
6	90 (4)	161 (5)	–	114 (3)	89 (10)	102 (5)	231 (2)	116 (29)
RFX mean	99 (6)	144 (6)	170 (2)	172 (4)	255 (5)	155 (6)	228 (5)	147 (6)
FFX mean	117 (54)	164 (33)	175 (13)	172 (12)	168 (20)	143 (23)	209 (7)	149 (162)

Average latency in each brain region for the six patients—only V responsive electrodes latencies were averaged. In brackets, the number of observations—corresponding, in most cases, to the number of V responsive electrodes or, in the RFX row, to the number of patients. The FFX average is of all electrodes pooled together, while the RFX average is produced by averaging across patients' averages.

deal separately with the minimal load condition. The only one to investigate this angle is Axmacher et al. study [2007]. In the paper, an evidence for gamma (52–98 Hz) power *suppression* in the delay period of the minimal load condition is provided.

Aside from the differences between our study and this one, which shade a doubt on the comparability of the two conditions, there are three additional issues regarding this minimal condition result which leave an opening to alternative interpretations. First, in the delay period a different visual stimulus than the one in the baseline period is presented—this, by itself, may be a cause for the effect; second, the length of the delay is only 2.5 s, while in our case the maximum interval is 4 s, so short term effect can still be present; and third, the delay interval is of a constant length, a fact that may promote anticipation effect.

A possible alternative explanation for our results is that the cognitive demands in the DR experiment were too low and failed to highlight the neural correlates of information maintenance. While this possibility cannot be completely ruled out, two aspects argue against it.

First, note that the memorized content on its own—regardless of whether it was the visual image or the motor plan—was sufficiently complex to elicit robust transient visuo/motor responses. It is difficult to see why memorizing these contents over an extended period should be so much easier compared with the encoding and execution stages. Second, an even easier cognitive task, that of a mere detection of a simple tone (i.e., the auditory cue), generated a robust activation (Fig. 4, right panels). Again, it is difficult to see why remembering to press the right or left button should be easier than tone detection. Finally, it is likely that the task became progressively easier as the patients became familiarized with their task—yet no indication of such training effects could be discerned in the delayed period activity.

It could be argued that the patients adopted a motor strategy to maintain the decision information over the delay period. For example, by starting to lightly press the targeted button in anticipation of the tone cue. However, we have seen no evidence of such “cheating” strategy in relevant electrode activity—that is, in our recordings from somato-sensory or premotor cortex.

More generally, the patients may have adopted a strategy in which a simple and isolated internal cue, that is the finger to be used, some linguistic cue or the face image, was maintained and these limited traces activated neuronal groups that were too small and isolated to be detected by the limited sampling of the ECoG recordings. However, it is important to note that precisely these isolated and limited stimuli produced a sufficiently robust activations that were consistently detected by our ECoG recordings when presented directly, but the same electrodes that showed a robust response to, say, a tone, or the button press, failed to show a similar activation during the delay period.

The strength-of-manipulation issue, discussed above, is strongly related to an important point regarding the baseline we used for comparison. Given that our study examined a single-item memory effect, a concern could be raised that the pretrial period, used as the baseline, may by itself reflect memory-related activations for example, general task instructions. It could be argued that such baseline activity may have masked the single-item working memory effect induced by our task.

However, it is reasonable to assume, and this was also substantiated by the maintained performance of the patients, that background memory processes that were present during the experiment continued throughout the trial—while the task relevant memory effect was added to these processes rather than replacing them. Hence, this additional memory activation should be detectable in the SD period. A more direct exploration of this interesting issue will necessitate mapping the selectivity of the memory processes to different contents. However, in our experimental paradigm, we failed to find a significant category-related effect even in the robust visual responses. Thus, additional studies using more distinguishable categories are needed to address this point.

Memory Maintenance Through a Hidden Dynamic State

Thus, it appears that our results emphasize the transient activation during visual stimuli and decision periods and demonstrate weak activity during the delay period in our task. This raises an obvious question—what neuronal mechanism could underlie the successful maintenance of information during the delay periods. A likely mechanism that has recently been proposed in modeling work, is a process by which memory is stored at the synaptic level using slow synaptic events that necessitating a minimum of “refresh” neuronal firing during the delay period [Mongillo et al., 2008].

In fact, both single unit results [Stokes et al., 2013] as well as noninvasive recordings [Lewis-Peacock et al., 2012] provide further support for the notion that some hidden dynamical state, not evident in neuronal spiking, may underlie the memory maintenance [Barak et al., 2010].

Short Latency Visually Induced Responses Outside Visual Cortex

Our ECoG recording results reveal that short latency responses, induced by the visual stimuli could be detected in fronto-parietal cortical regions, showing similar characteristics to the responses observed in HOV areas. Short latency responses were evident both for increases in BHF (50–150) power as well as the evoked responses.

The short latency responsive electrodes could be found both posteriorly and anteriorly to the central sulcus, as well as in the prefrontal cortex. This confirms previous ECoG

reports of short latency responses in frontal cortical sites [Blanke et al., 1999; Kirchner et al., 2009; Noy et al., 2015]. This study, as well as the latter reference, extends these findings by showing that similar short latency responses are present also in power increase at the BHF (50–150) range, which we and others have related to neuronal firing and perceptual awareness [Fisch et al., 2009; Henrie and Shapley, 2005; Mukamel et al., 2005; Nir et al., 2007; Ray and Maunsell, 2011; Tallon-Baudry and Bertrand, 1999].

Thus, our results are in line with suggestions that the cortical connectivity follows small-world architectures optimized for rapid signal transfer along distant cortical sites [Singer, 2013; Sporns, 2013; Van den Heuvel and Sporns, 2011, 2013].

Interestingly, the complement of the rapid spread of visual responses into fronto-parietal cortical regions was also observed—that is, spread of audio/motor responses following the auditory cue into HOV areas (see Figs. 4 and 5). Such rapid signals could be related to a number of top-down processes such as attentional modulation, eye-movement related signals, and so forth. Further research will be needed to clarify the role of these responses.

It could be argued that the observed fast spread of signals may be a trivial by-product of signal processing—that is, it could be that the process of subtracting the global “common” signal from all electrodes (see methods) have inadvertently mixed signals between regions. To rule out this possibility, we have repeated all steps of the analysis with a dataset that was not subjected to the de-referencing procedure.

The percent of responsive electrodes in this dataset was lower in each of the brain regions. This is not surprising as non-neuronal contributions from the reference electrode are expected to lower the signal to noise ratio. The important thing is that all of the effects found in the “clean” dataset were evident in this dataset as well, including a considerable amount of V electrodes in fronto-parietal region and of AM electrodes in the occipital lobe.

Furthermore, if the de-referencing procedure was indeed mixing signals between areas it should have introduced a uniform effect across all electrodes. Our data, on the contrary, show that the responsive electrodes are confined to specific regions (Fig. 3).

The Role of Fronto-Parietal Visually-Induced Signals in Single Item Maintenance

What could be the function of the short latency responses to visual stimuli found in fronto-parietal regions during the DR experiment (Table IV and Figs. 6 and 7)? Within the context of the present issue, it is tempting to hypothesize that at least part of the functions subserved by these signals is to encode the decision information that needs to be maintained during the delay period until motor execution. Thus, within the framework of synaptic models of working memory, the initial burst of activation initiated the sustained synaptic events that on the cue

onset will deliver the correct decision information to the motor cortex.

The Measurement Window

In interpreting the significance of these results, and particularly with reference to detecting delayed period activity, it is important to emphasize that the ECoG recordings were limited to a specific spatio-temporal “window” of the neuronal activity. ECoG recordings do not contact deep sulci, leaving a substantial expanse of fronto-parietal cortex inaccessible to our examination. In addition, sparsely distributed neurons—that is, active neurons that were not clustered in groups, may produce signals that are too weak to be discernible in the ECoG signals which reflect averaged mass activity.

Finally, our results are based mainly on analysis of BHF (50–150) power modulations. While we and others have previously argued that increases in BHF (50–150) power provide a reliable index of mass increases in firing rates [Nir et al., 2007], this has been demonstrated in a limited set of areas and tasks. It could be that during maintenance activity the coupling between BHF (50–150) and the average firing rates is reduced, limiting our ability to deduce neuronal firing from BHF (50–150) power.

With regards to lower frequencies, such as the beta or theta range, it is important to note that here the link to neuronal firing is even more tentative [Mukamel et al., 2005]. So, while previous research has suggested that changes in these lower frequencies’ power may be linked to memory maintenance [Raghavachari et al., 2001; Tallon-Baudry et al., 2001] the link of such changes to spiking activity should be viewed with caution at this stage.

One relevant point that should be considered within this context is that high levels of neuronal activity should lead to high levels of BOLD signals when conducting a multiple-item working memory tasks. Thus, when considering the long integration time of BOLD-fMRI, which effectively “amplifies” sustained periods of neuronal firing [Mukamel et al., 2005; Nir et al., 2008] one would expect particularly high BOLD signals during the delay period in a task demanding maintenance of information.

However, all BOLD-fMRI studies of delay period activity during maintenance report the opposite—weak signals that are typically far lower than the encoding sensory stimuli, although a load dependency has been demonstrated here as well [Courtney et al., 1997; D’Esposito et al., 2000; Jha and McCarthy, in press; Pessoa et al., 2002].

ACKNOWLEDGMENT

We are grateful to the six patients who participated in this study.

REFERENCES

Axmacher N, Mormann F, Fernández G, Cohen MX, Elger CE, Fell J (2007): Sustained neural activity patterns during working

- memory in the human medial temporal lobe. *J Neurosci* 27:7807–7816.
- Barak O, Tsodyks M, Romo R (2010): Neuronal population coding of parametric working memory. *J Neurosci* 30:9424–9430.
- Benjamini Y, Hochberg Y (1995): Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J R Stat Soc Series B Methodol* 57:289–300.
- Blanke O, Morand S, Thut G, Michel CM, Spinelli L, Landis T, Seeck M (1999): Visual activity in the human frontal eye field. *Neuroreport* 10:925.
- Bor D, Duncan J, Wiseman RJ, Owen AM (2003): Encoding strategies dissociate prefrontal activity from working memory demand. *Neuron* 37:361–367.
- Conway AR, Kane MJ, Bunting MF, Hambrick DZ, Wilhelm O, Engle RW (2005): Working memory span tasks: A methodological review and user's guide. *Psychon Bull Rev* 12:769–786.
- Courtney SM, Ungerleider LG, Keil K, Haxby JV (1997): Transient and sustained activity in a distributed neural system for human working memory. *Nature* 386:608–611.
- Curtis CE, D'Esposito M (2003): Persistent activity in the prefrontal cortex during working memory. *Trends Cogn Sci* 7:415–423.
- Curtis CE, Lee D (2010): Beyond working memory: The role of persistent activity in decision making. *Trends Cogn Sci* 14:216–222.
- Davidesco I, Zion-Golumbic E, Bickel S, Harel M, Groppe DM, Keller CJ, Schevon CA, McKhann GM, Goodman RR, Goelman G (2013): Exemplar selectivity reflects perceptual similarities in the human fusiform cortex. *Cereb Cortex* 24:1879–1893.
- Delorme A, Makeig S (2004): EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J Neurosci Methods* 134:9–21.
- D'Esposito M, Postle BR, Rypma B (2000): Prefrontal cortical contributions to working memory: Evidence from event-related fMRI studies. *Exp Brain Res* 133:3–11.
- Fedorenko, E, Duncan, J, Kanwisher, N. (2013): Broad domain generality in focal regions of frontal and parietal cortex. *Proc Natl Acad Sci USA* 110:16616–16621.
- Fisch L, Privman E, Ramot M, Harel M, Nir Y, Kipervasser S, Andelman F, Neufeld MY, Kramer U, Fried I (2009): Neural “ignition”: Enhanced activation linked to perceptual awareness in human ventral stream visual cortex. *Neuron* 64:562–574.
- Gelman A, Hill J, Yajima M (2012): Why we (usually) don't have to worry about multiple comparisons. *J Res Educ Eff* 5:189–211.
- Henrie JA, Shapley R (2005): LFP power spectra in V1 cortex: The graded effect of stimulus contrast. *J Neurophysiol* 94:479–490.
- Howard MW, Rizzuto DS, Caplan JB, Madsen JR, Lisman J, Aschenbrenner-Scheibe R, Schulze-Bonhage A, Kahana MJ (2003): Gamma oscillations correlate with working memory load in humans. *Cereb Cortex* 13:1369–1374.
- Jha AP, McCarthy G (2000): The influence of memory load upon delay-interval activity in a working-memory task: An event-related functional MRI study. *J Cogn Neurosci* 12:90–105.
- Kirchner H, Barbeau EJ, Thorpe SJ, Gis R, Li geois-Chauvel JC (2009): Ultra-rapid sensory responses in the human frontal eye field region. *J Neurosci* 29:7599.
- Lachaux JP, George N, Tallon-Baudry C, Martinerie J, Hugueville L, Minotti L, Kahane P, Renault B (2005): The many faces of the gamma band response to complex visual stimuli. *Neuroimage* 25:491–501.
- Lachaux JP, Axmacher N, Mormann F, Halgren E, Crone NE (2012): High-frequency neural activity and human cognition: Past, present and possible future of intracranial EEG research. *Prog Neurobiol* 98:279–301.
- Lewis-Peacock JA, Drysdale AT, Oberauer K, Postle BR (2012): Neural evidence for a distinction between short-term memory and the focus of attention. *J Cogn Neurosci* 24:61–79.
- Mainy N, Kahane P, Minotti L, Hoffmann D, Bertrand O, Lachaux JP (2007): Neural correlates of consolidation in working memory. *Hum Brain Mapp* 28:183–193.
- Meltzer JA, Zaveri HP, Goncharova II, Distasio MM, Papademetris X, Spencer SS, Spencer DD, Constable RT (2008): Effects of working memory load on oscillatory power in human intracranial EEG. *Cereb Cortex* 18:1843–1855.
- Miller EK, Erickson CA, Desimone R (1996): Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16:5154.
- Miller GA (1956): The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychol Rev* 63:81.
- Minear M, Park DC (2004): A lifespan database of adult facial stimuli. *Behav Res Methods* 36:630–633.
- Mongillo G, Barak O, Tsodyks M (2008): Synaptic theory of working memory. *Science* 319:1543.
- Morey R, Rouders J, Jamil T (2015): BayesFactor: Computation of Bayes factors for common designs [R package]. Version 0.9.11-1. Available at: <http://CRAN.R-project.org/package=BayesFactor>. Accessed on 1 June, 2015.
- 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.
- Nir Y, Fisch L, Mukamel R, Gelbard-Sagiv H, Arieli A, Fried I, Malach R (2007): Coupling between neuronal firing rate, gamma LFP, and BOLD fMRI is related to interneuronal correlations. *Curr Biol* 17:1275–1285.
- Nir Y, Dinstein I, Malach R, Heeger DJ (2008): BOLD and spiking activity. *Nat Neurosci* 11:523–524.
- Noy N, Bickel S, Zion-Golumbic E, Harel M, Golan T, Davidesco I, Schevon C, McKhann G, Goodman R, Schroeder C (2015): Ignition's glow: Ultra-fast spread of global cortical activity accompanying local “ignitions” in visual cortex during conscious visual perception. *Conscious Cogn* 35:206–224.
- Pessoa L, Gutierrez E, Bandettini PA, Ungerleider LG (2002): Neural correlates of visual working memory: fMRI amplitude predicts task performance. *Neuron* 35:975–987.
- Raghavachari S, Kahana MJ, Rizzuto DS, Caplan JB, Kirschen MP, Bourgeois B, Madsen JR, Lisman JE (2001): Gating of human theta oscillations by a working memory task. *J Neurosci* 21:3175–3183.
- Rainer G, Asaad WF, Miller EK (1998): Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature* 393:577–579.
- Ray S, Maunsell JHR (2011): Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biol* 9:e1000610.
- Schluppeck D, Curtis CE, Glimcher PW, Heeger DJ (2006): Sustained activity in topographic areas of human posterior parietal cortex during memory-guided saccades. *J Neurosci* 26:5098–5108.
- Singer W (2013): Cortical dynamics revisited. *Trends Cogn Sci* 17:616–626.
- Sporns O (2013): Network attributes for segregation and integration in the human brain. *Curr Opin Neurobiol* 23:162–171.

- Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J (2013): Dynamic coding for cognitive control in prefrontal cortex. *Neuron* 78:364–375.
- Talairach J, Tournoux P (1988): Co-planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System: An Approach to Cerebral Imaging. New York: Georg Thieme.
- Tallon-Baudry C, Bertrand O (1999): Oscillatory gamma activity in humans and its role in object representation. *Trends Cogn Sci* 3: 151–162.
- Tallon-Baudry C, Bertrand O, Fischer C (2001): Oscillatory synchrony between human extrastriate areas during visual short-term memory maintenance. *J Neurosci* 21:177.
- Tong F (2013): Imagery and visual working memory: One and the same? *Trends Cogn Sci* 17:489–490.
- van den Heuvel M, Sporns O (2013): Network hubs in the human brain. *Trends Cogn Sci* 17:683–696.
- van den Heuvel MP, Sporns O (2011): Rich-club organization of the human connectome. *J Neurosci* 31:15775–15786.
- Vidal JR, Ossandón T, Jerbi K, Dalal SS, Minotti L, Ryvlin P, Kahane P, Lachaux JP (2010): Category-specific visual responses: An intracranial study comparing gamma, beta, alpha, and ERP response selectivity. *Front Hum Neurosci* 4: 195.