Current Biology 16, 2066–2072, October 24, 2006 ©2006 Elsevier Ltd All rights reserved DOI 10.1016/j.cub.2006.08.064

Report

A Single-Neuron Correlate of Change Detection and Change Blindness in the Human Medial Temporal Lobe

Leila Reddy,[1,4,*] Rodrigo Quian Quiroga,[1,5] Patrick Wilken,[1,6] Christof Koch,[1] and Itzhak Fried[2,3]
[1] Computation and Neural Systems
California Institute of Technology
Pasadena, California 91125
[2] Division of Neurosurgery and
Semel Institute for Neuroscience and Human Behavior
University of California, Los Angeles
Los Angeles, California 90095
[3] Functional Neurosurgery Unit
Tel-Aviv Medical Center and
Sackler Faculty of Medicine
Tel-Aviv University
Tel-Aviv 69978
Israel

Summary

Observers are often unaware of changes in their visual environment when attention is not focused at the location of the change [1–4]. Because of its rather intriguing nature, this phenomenon, known as change blindness, has been extensively studied with psychophysics [5–7] as well as with fMRI [8–11]. However, whether change blindness can be tracked in the activity of single cells is not clear. To explore the neural correlates of change detection and change blindness, we recorded from single neurons in the human medial temporal lobe (MTL) during a change-detection paradigm. The preferred pictures of the visually responsive units elicited significantly higher firing rates on the attended trials when subjects correctly identified a change (change detection) compared to the unattended trials when they missed it (change blindness). On correct trials, the firing activity of individual units allowed us to predict the occurrence of a change, on a trial-by-trial basis, with 67% accuracy. In contrast, this prediction was at chance for incorrect, unattended trials. The firing rates of visually selective MTL cells thus constitute a neural correlate of change detection.

Results

Participants were patients with pharmacologically intractable epilepsy who were implanted with depth electrodes to localize the focus of seizure onset for potentially curative resection. For all patients, the locations of the implanted electrodes were determined solely by clinical criteria [12, 13]. In a total of nine patients, we recorded from 534 units located in the amygdala, entorhinal cortex, parahippocampal gyrus, and hippocampus. During screening sessions, patients were first presented with a large number of images, displayed in isolation at fixation, to ascertain which preferred stimuli elicited a significant response in at least one unit [14]. On the basis of these sessions, in which we found 110 units visually responsive to at least one preferred image, the stimulus set was chosen for the change-detection sessions. Given that a typical change-detection session lasted only 20–30 min because of clinical constraints, we had to choose a subset of 25 preferred images, which elicited responses in 43 units during the screening sessions. Twenty-nine of these units maintained significant visual responses to their preferred stimuli during the change-detection sessions that involved the simultaneous presentation of four images (see the Supplemental Data available online). On the basis of magnetic resonance imaging (MRI) confirmation, six of these units were located in the amygdala, three in the entorhinal cortex, five in the parahippocampal gyrus, and 15 in the hippocampus. Because of the low number of cells recorded in each area, we cannot make conclusive claims about differences in the latencies and firing patterns among the different areas. The remaining 505 units were also analyzed for specific changes in firing activity during the different conditions (correct versus incorrect trials) of the change-detection experiment. For these units, no consistent modulations were observed.

Participants were presented with two stimulus displays for 1 s each, separated by a blank interval of 1.5 s. Each display consisted of four peripherally presented images of celebrities, animals, objects, or landmark buildings. On half the trials, one of the images was replaced by a new image in the second display. Participants were asked to detect changes that occurred between these two displays (Figure 1A and Supplemental Data). Over the nine patients, the average hit rate was 0.78 ± 0.02 (mean ± standard error of the mean [SEM]), the average false-alarm rate was 0.19 ± 0.02, and the average response time in reporting whether a change had occurred was 1.4 ± 0.1 s following the offset of the second image. The difference in reaction times between correct (i.e., hits and correct rejections) and incorrect (i.e., false alarms and misses) trials was not significant ($p > 0.05$).

On any trial in the experiment, only one of the four images shown was a preferred stimulus for the targeted cells. Compatible with monkey electrophysiology reports [15–19], the firing activity to preferred stimuli was significantly reduced during the change-detection sessions (stimuli presented with three nonpreferred stimuli) compared to the screening sessions (stimuli appeared in isolation) (Figure 1C). This decrease in firing activity could also have been due to the more peripheral stimulus locations in the change-detection sessions compared to the screening sessions. On average, the firing

*Correspondence: lreddy@klab.caltech.edu
[4] Present address: McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, Massachusetts.
[5] Present address: Department of Engineering, University of Leicester, United Kingdom.
[6] Present address: Otto-von-Guericke-Universität, Magdeburg, Germany.

activity was three times lower during change-detection sessions (Figure S1).

In the experiment, there were two different trial types—change and no-change trials. However, for a particular visually responsive neuron, four different trial types could be determined depending on when its preferred stimulus was presented (Figure 1B). In "disappear" trials, the preferred picture was one of the four images in the first display period and was replaced by a nonpreferred picture in the second period. In "appear" trials, the preferred picture was absent in the first and appeared in the second interval. On no-change trials, the preferred picture of the cell was present in both display periods ("both" trials) or not in either ("none" trials).

The firing activity of the population of visually responsive cells in these four trial types was compared for correct and incorrect trials (Figure 2A). In all conditions, during the intervals when the preferred stimuli were present, neuronal responses were significantly higher on correct trials (black trace) compared to incorrect trials (white trace). For example, during "disappear" trials, firing rates during the first display presentation, when the preferred stimulus was present, were significantly higher (two-tailed t test, $p < 0.0005$) on correct versus incorrect trials. Similar results were obtained for the "appear" ($p < 0.005$) and "both" trials ($p < 0.05$) during display intervals when the preferred stimuli were present. These results indicate that on correct trials, the preferred stimuli were represented at the neuronal level more strongly than on the incorrect trials, regardless of the display interval they appeared in. On the other hand, in display periods when the preferred stimuli were absent, no significant difference was observed between correct and incorrect trials ($p > 0.05$), indicating that the observed results are specific to the selectivity of the neurons and are not due to an increased level of arousal on correct trials.

Interestingly, decreased activity on incorrect trials was also observed during the 1.5-s-long blank interval between the two display periods on "disappear" and "both" trials ($p < 0.05$). In both, a preferred stimulus was presented during the first display period, and this delay-period activity might reflect a memory of the stimulus, as has been observed in monkey prefrontal and inferotemporal cortex [20–22]. Although previous psychophysical studies have suggested that change blindness results from an overwriting of an iconic buffer by the "change" stimulus—rather than a failure of sensory memory [23]—our data indicate that at least on some incorrect trials, the encoding of the preferred stimulus was weaker during both stimulation and memory intervals.

The left panel in Figure 2B presents (in a summarized form) the average firing activity in all those intervals of the four trial types when the preferred stimuli were present. These results again indicate that firing rates were significantly higher on correct trials (black trace, mean firing rate = 4.6 ± 0.6 Hz) compared to incorrect trials (white trace, 3.4 ± 0.5 Hz) when the preferred stimulus was present on the screen ($p < 0.0001$). These data thus demonstrate neural correlates of change blindness and change detection in the human MTL. Note that although we use the terms "change detection" and "change blindness" broadly to refer to correct (hits and correct rejections) and incorrect (misses and false alarms) trials, respectively, the same pattern of results was obtained within the stricter definition of change blindness and change detection on change-present trials only (i.e., hits versus misses, $p < 0.0001$).

The right panel in Figure 2B summarizes the neuronal activity in those intervals when the preferred stimuli were absent. No significant differences in firing rates were observed between correct and incorrect trials ($p = 0.2$, mean firing rates = 2.3 ± 0.4 Hz and 2.2 ± 0.4 Hz, respectively), which is expected because only images that drove these cells weakly were presented to subjects.

Similar results were also obtained when each cell was required to contribute equal numbers of correct and incorrect trials to the population averages in each condition. As in Figure 2A, the mean firing activity was significantly higher on correct trials in the "disappear" ($p = 0.0007$), "appear" ($p = 0.003$), and "both" ($p = 0.01$) conditions during the intervals when the preferred stimuli were present. Consistent with the results in Figure 2B, the firing activity averaged over all intervals in which the preferred stimuli were displayed was also significantly greater on correct versus incorrect trials ($p = 0.0003$), but such an effect was not observed in those time windows when the preferred stimuli were absent ($p > 0.05$).

How reliably can the activity of a neuron on an individual trial tell us whether or not a change occurred in the stimulus? We addressed this issue quantitatively via a receiver operating characteristic (ROC) analysis [24] that is used to link stimulus information or a subject's behavior with neuronal activity [25–28]. For each neuron, we determined how well an ideal observer could use the number of spikes fired on individual trials to predict whether or not a change had occurred in the stimulus (see Experimental Procedures). The ROC curves for each cell are shown in Figure 3A. The area under the curves is a measure of the ability of each neuron to estimate whether a change occurred in the stimulus on a trial-by-trial basis (independent of the subject's behavioral report). A value of 0.5 would correspond to chance performance, and a value of 1 would reflect perfect accuracy in signaling a change. The distribution of the areas under the curves across all cells (Figure 3B) was centered at 0.64 ± 0.02 and was significantly greater than 0.5 ($p < 0.0001$). The outcome of our ROC analysis is compatible with previous experiments linking stimulus information with firing activity of neurons in the middle temporal area (MT) of the monkey brain [25], where ROC performance varied between 0.51 and 0.99 depending on the coherence of the motion stimuli.

We next calculated the ROC analysis separately over correct and incorrect trials. The distribution of the area under the ROC curves for the correct trials (Figures 3C and 3D) was centered at 0.67 ± 0.02 and was significantly shifted to the right of 0.5 ($p < 0.0001$). In contrast, the average ROC area computed over incorrect trials (Figures 3E and 3F) was significantly lower (0.49 ± 0.04; $p < 0.005$) and was not different from chance ($p = 0.95$). These results indicate that on trials when subjects correctly reported the presence or absence of a change, the firing of the recorded MTL cells allowed an ideal observer to make this judgment better than chance. On the other hand, when subjects made an error, MTL units did not carry the relevant information. We also studied

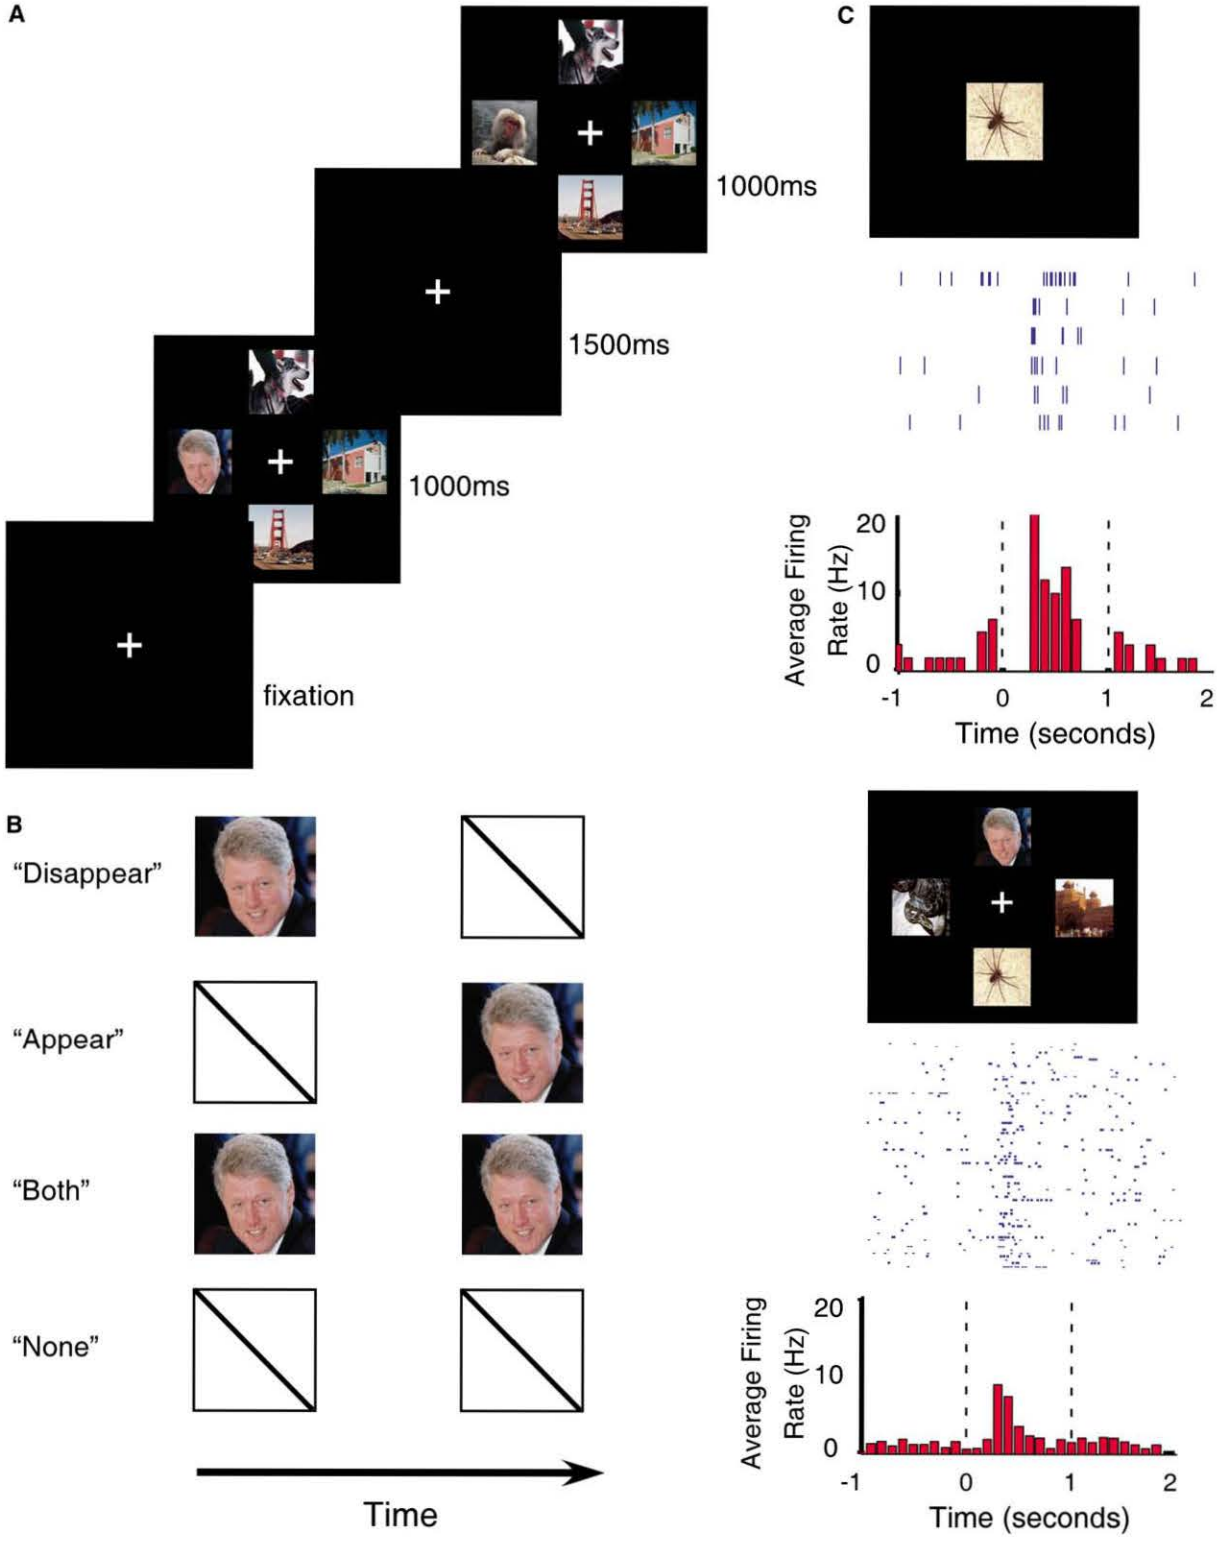

Figure 1. Experimental Timeline and Examples of Neuronal Responses

(A) Timeline for one trial in the change-detection experiment. Each trial began with a fixation cross presented for a random interval between 1000 and 1200 ms. Four images then appeared for 1000 ms, on a circle with a 6° radius. A delay interval of 1.5 s followed, after which four pictures were again presented at the same locations as previously. On roughly half of the trials, a change occurred and participants had to report whether they had noticed the change after the offset of the second display.

(B) From the point of view of a cell selective to a particular stimulus (e.g., Bill Clinton), four different trial types could occur: In the change trials, the picture of Clinton is present in the first display period and absent in the next ("Disappear"), or it is absent in the first and appears in the second ("Appear"). In the no-change trials, the picture is present in both display periods ("Both") or is absent all together ("None").

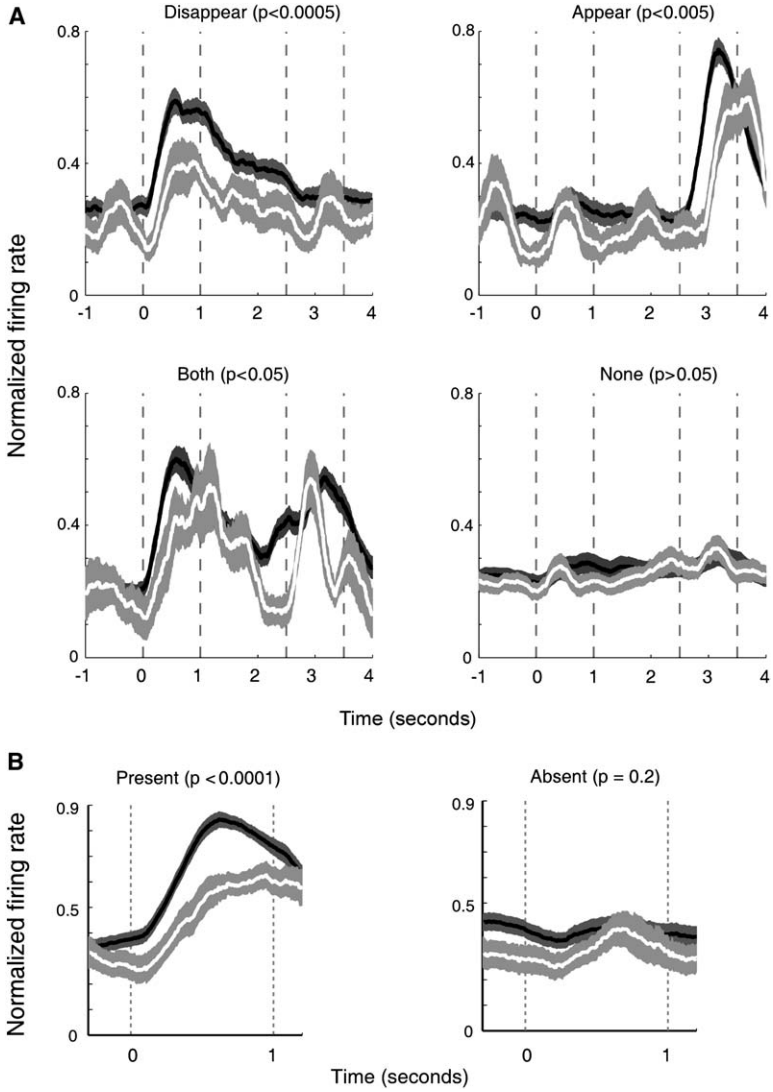

Figure 2. Responses during Change-Blindness and Change-Detection Trials

(A) The mean normalized spike-density function was calculated over the 29 cells in the four trial conditions for correct (hits and correct rejections, black curve) and incorrect (false alarms and misses, white curve) trials. The shaded areas represent the SEM. Note the significantly higher levels of activity on correct trials during all intervals when the preferred stimuli were present (corresponding significance values are presented at the top of each plot). During intervals when the preferred stimuli were not presented, there was no significant difference between correct and incorrect trials ($p > 0.05$).

(B) Population responses averaged over all stimulus intervals when the preferred stimuli were present (left panel) and absent (right panel). The stimulus appeared on the screen at 0 ms and stayed on for 1 s (the length of each display period). The firing rates during the correct trials (black curve) are significantly higher than firing rates during incorrect trials (white curve). On the right, there is no significant difference in firing rates between correct and incorrect trials, which is expected because the preferred stimuli were absent during these intervals.

whether the firing activity on individual trials could predict the behavioral report of a change (independent of whether a change was present in the stimulus or not). This analysis is sometimes referred to as "choice probability" [26], and we found that the firing activity of individual cells could register subjects' behavioral report significantly better than chance (58%, $p < 0.001$; Supplemental Data).

Discussion

Our results demonstrate that for the appearance or disappearance of one out of four images in a change-detection paradigm, the firing activity of MTL single neurons selective to this image was larger during stimulus presentation on correct trials compared to incorrect trials. These data are compatible with recent functional MRI (fMRI) studies reporting that visual areas in the ventral stream are strongly activated during change detection and to a lesser extent during change blindness [8, 9]. These studies have not, however, reported similar effects in the MTL. It is possible that fMRI currently does not have the necessary spatial resolution to examine these effects in the MTL. No spatial topography has been found in the MTL in animal studies. Our single-unit recordings in the human MTL have shown that this region has a sparse representation of familiar faces, scenes, etc. [14], but in such low numbers that it is possible that fMRI is not sensitive enough to pick these up. These factors conspire against finding significant differences in MTL blood-oxygen-level-dependent (BOLD) activity in a change-detection paradigm.

Previous studies have demonstrated correlates for change detection in early visual areas in monkeys [29],

(C) Comparing visual responses between screening (top panel) and change-detection (bottom panel) sessions. In the screening session, this unit showed a significant increase in firing to its preferred image that was presented foveally and in isolation ($p < 0.001$). The image appeared at $t = 0$ and was presented for 1 s. The response of the unit to the same image during the change-detection session is still significantly above baseline ($p < 0.0001$) but is about half as strong as the response in the screening session (mean firing rate = 3.3 ± 0.4 Hz versus 7.62 ± 2.0 Hz). During the change-detection experiment, the preferred image was presented peripherally along with three others known to not drive the cell.

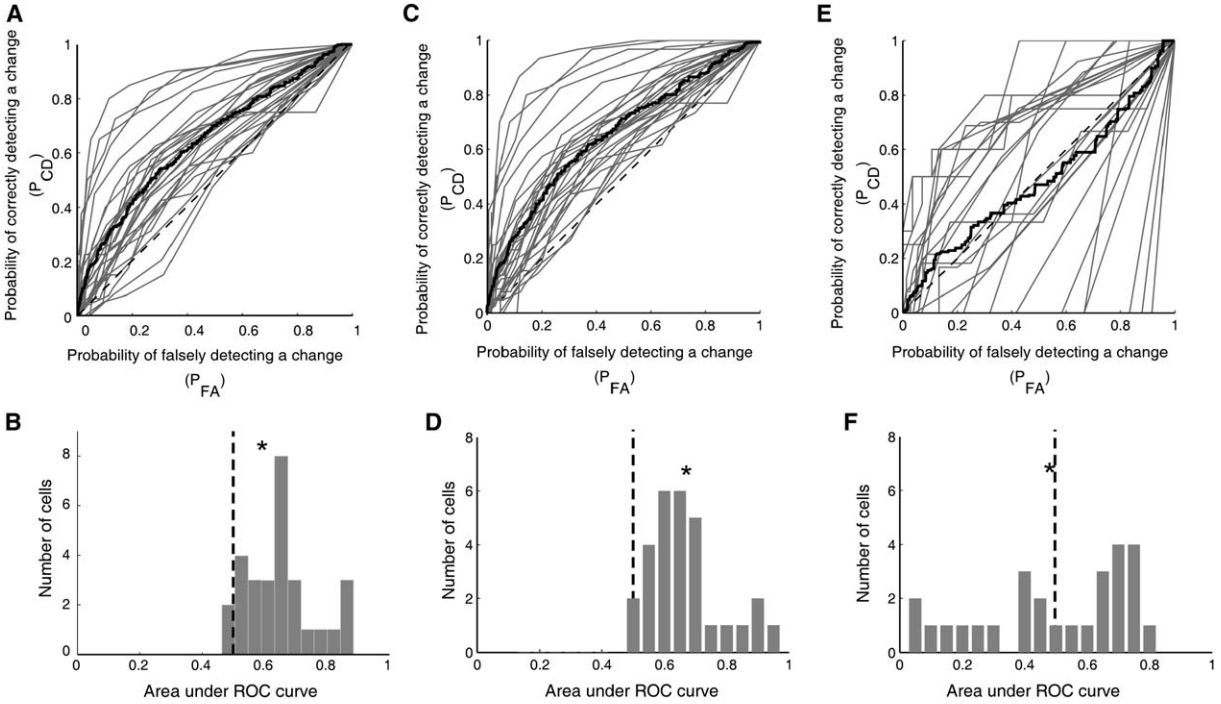

Figure 3. Predicting a Change in the Input—ROC Analysis

(A) The probability of classifying a trial as a change trial, P_{CD} ("correct detection"), is plotted against the probability of "false alarms," P_{FA} (falsely detecting a change), for all trials. The dashed line indicates chance performance ($P_{CD} = P_{FA}$). The different lines show the result of this calculation for each cell. The solid black line is the average ROC curve.

(B) The distribution of the area under the curves for each cell is significantly shifted to the right of 0.5, indicating that the 29 cells can signal a change above chance on a trial-by-trial basis ($p < 0.0001$). The mean area under the ROC curves is marked by a * and equals 0.64 ± 0.02.

(C) ROC curves for all cells calculated over the correct trials only.

(D) The distribution of the area under the curves is significantly shifted to the right of 0.5 ($p < 0.0001$). The mean area under the ROC curves is marked by a * and equals 0.67 ± 0.02.

(E) ROC curves for all cells calculated over incorrect trials only.

(F) The mean area under the ROC curves for incorrect trials (*) is 0.5 ± 0.04. In contrast to data in (D), on incorrect trials, the population of cells were at chance at signaling a change ($p = 0.95$). The two distributions in (D) and (F) are significantly different from each other ($p < 0.005$).

but the present findings constitute the first report of single-neuron correlates of this phenomenon in the human brain. Remarkably, the activity of single neurons allowed us to predict behavioral choices of participants above chance on a trial-by-trial basis (choice-probability analysis, Supplemental Data). It should be noted, however, that the MTL cells are unlikely to be explicitly signaling change detection or blindness. The ROC analyses certainly suggest that these cells carry the relevant information; however, this information is obtained by comparing the two stimulation periods. Neurons other than those reported here would thus be required to compare the activity in these intervals and to directly signal the occurrence of a change.

Attention is believed to play a key role during change detection—in the absence of attention, observers are found to be blind to changes made to objects [7]. With this hypothesis in mind, our findings show that the effect of attention during change detection is to modulate the firing rates of neurons encoding the preferred stimuli, such that activity is higher on correct versus incorrect trials. Note that because we did not measure neuronal activity during passive fixation, it is not possible for us to determine whether the modulations we observe are due to detections of changes or blindness to them—in other words, whether the presence of attention strengthened the representation of objects relative to some baseline level, or whether its absence resulted in a weaker encoding. Nevertheless, during change detection and blindness, attention is seen to significantly influence the activity of neurons so much so that the information carried by individual neurons on a trial-by-trial basis can be qualitatively very different. On attended trials, an ideal observer could reliably predict the occurrence of a change on individual trials, whereas on unattended trials both MTL neurons and subjects were "blind" to changes.

fMRI studies of the human MTL have suggested that these structures are more strongly activated for novel than for familiar stimuli [30, 31]. A recent electrophysiological study likewise showed modulation of MTL single-neuron activity on the basis of the familiarity or novelty of stimuli [32]. Could the effects we observe therefore be related to changes in familiarity or novelty? Given that subjects had equal exposure to all images during the screening sessions, which were conducted prior to the change-detection testing sessions (see Experimental Procedures), our findings cannot easily be attributed to modulation by either novelty or familiarity. Stimulus familiarity was thus comparable for the selected

preferred and nonpreferred stimuli, as well as for change-detection and change-blindness trials. Additionally, familiarity and novelty influence neuronal activity only gradually, over the course of several trials in the monkey brain [33–35], and even when single-trial modulations are observed in the human brain, their effects are only apparent after approximately 30 min [32]. In our data, we observed no significant effect of repetition on the firing activity of our cells over the course of each experimental session (each session's duration was split into four quarters, and no effect of quarter number on cells' responses was observed; one-way ANOVA, $p = 0.22$). Thus, it is unlikely that the effects we observe on each trial, within the space of a few seconds, are significantly influenced by changes in familiarity or novelty.

Ever since the initial reports of amnesic patient H.M. [36], MTL structures have been thought to lie at the core of the declarative-memory system and to function independently of other perceptual and cognitive functions [37, 38]. Recently, however, this view has been challenged by arguments that MTL structures might also participate in visual perception, at least of complex feature conjunctions [39–41], or short-term memory [42, 43]. Our results do not directly address this debate: Although the significantly higher delay-period activity on correct trials is suggestive of short-term memory, the responsiveness of our cells can be attributed to either the perception of their preferred stimuli or the memories associated with them. Yet the modulations we observe in the change-blindness paradigm arise from a comparison of correct or incorrect perceptual processing of these stimuli on each trial. Whether these modulations are the cause or the consequence of change blindness or detection cannot be addressed by this paradigm, and this issue remains open to further investigation.

Experimental Procedures

Recordings

The data were obtained from 17 sessions in nine patients with pharmacologically intractable epilepsy (eight right-handed, three male, 17 to 47 years old). Extensive noninvasive monitoring did not yield concordant data corresponding to a single resectable epileptogenic focus. Therefore, patients were implanted with chronic depth electrodes for 7–10 days to determine the seizure focus for possible surgical resection [12]. Here we report data from sites in the hippocampus, amygdala, entorhinal cortex, and parahippocampal gyrus. All studies conformed to the guidelines of the Medical Institutional Review Board at UCLA. The electrode locations were based exclusively on clinical criteria and were verified by MRI or by computed tomography (CT) coregistered to preoperative MRI. Each electrode probe had a total of nine microwires at its end, eight active recording channels, and one reference. The differential signal from the microwires was amplified with a 64-channel Neuralynx system (Tucson, Arizona) and filtered between 1 and 9000 Hz. Signals were sampled at 28 kHz. Each recording session lasted about 30 min. Spike detection and sorting was applied to the continuous data with a previously implemented clustering algorithm [14, 44].

Statistical Analysis

In the comparison between firing activity on correct versus incorrect trials, significance values were computed by using a two-tailed t test. Neuronal activity was integrated over a 1 s interval for the periods when the preferred stimuli were present (starting at 300 ms to take into account the latency in responses). For the analysis involving the blank interstimulus interval (ISI), a one-tailed t test was computed over activity in the delay interval.

ROC Analysis

For each neuron that we recorded from, we determined how well an ideal observer could use the firing rate on individual trials to predict whether or not a change had occurred by computing a ROC analysis [24]. On each trial, we computed the absolute difference between the firing activity in the two display intervals and used a sliding threshold to determine the probability of false alarms (P_{FA}) and the probability of correct detections (P_{CD}) [45]. Correct detections (the "hits" for the ROC analysis) correspond to the probability P (cell indicating change occurred | change occurred) and false alarms correspond to P (cell indicating change occurred | no change occurred). Several time intervals were considered for this analysis: 300–700 ms, 300–1000 ms, 300–1500 ms, 300–2000 ms, and 300–2500 ms (aligned to the onset of each display period). For all calculations, the best ROC performance was observed in the 300–1500 ms window (Figure S3). For each integer value of the threshold T, a difference greater than T in the spike number between the two display periods means that the neuron signaled that a change had occurred on that trial. Depending on whether a change had actually occurred or not, this signal counts either as a "correct detection" or a "false alarm." These values were computed for all trials together, as well as separately for correct and incorrect trials.

We estimated statistical significance of the ROC analysis by comparing the distribution of areas under the ROC curves with chance performance (0.5) by using a two-tailed t test ($p < 0.05$).

Supplemental Data

Supplemental Data include Results, Experimental Procedures, and four figures and are available with this article online at: http://www.current-biology.com/cgi/content/full/16/20/2066/DC1/.

Acknowledgments

We thank all the patients for their participation; I. Wainwright and B. Salaz for administrative assistance; E. Behnke, T. Fields, E. Ho, E. Isham, A. Kraskov, P. Steinmetz, and C. Wilson for technical assistance; and R. VanRullen for comments on the manuscript. This work was supported by grants from the National Institute of Neurological Disorders and Stroke (NINDS), National Institute of Mental Health (NIMH), National Science Foundation (NSF), Defense Advanced Research Projects Agency (DARPA), the Office of Naval Research, the W.M. Keck Fund for Discovery in Basic Medical Research, the Gordon Moore Foundation, the Sloan Foundation, and the Swartz Foundation for Computational Neuroscience.

Received: March 31, 2006
Revised: August 29, 2006
Accepted: August 29, 2006
Published: October 23, 2006

References

1. Rensink, R., and Simons, D. (2000). Change detection, attention, and the contents of awareness. Conscious. Cogn. *9*, S18–S19.
2. Rensink, R.A. (2002). Change detection. Annu. Rev. Psychol. *53*, 245–277.
3. Simons, D.J. (2000). Current approaches to change blindness. Vis. Cogn. *7*, 1–15.
4. Simons, D.J., and Levin, D.T. (1997). Change blindness. Trends Cogn. Sci. *1*, 261–267.
5. Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In Perception, *Volume 5*, K. Akins, ed. (New York: Oxford University Press).
6. Pashler, H. (1988). Familiarity and visual change detection. Percept. Psychophys. *44*, 369–378.
7. Rensink, R.A., O'Regan, J.K., and Clark, J.J. (1997). To see or not to see: The need for attention to perceive changes in scenes. Psychol. Sci. *8*, 368–373.
8. Beck, D.M., Rees, G., Frith, C.D., and Lavie, N. (2001). Neural correlates of change detection and change blindness. Nat. Neurosci. *4*, 645–650.
9. Pessoa, L., and Ungerleider, L.G. (2004). Neural correlates of change detection and change blindness in a working memory task. Cereb. Cortex *14*, 511–520.

10. Huettel, S.A., Guzeldere, G., and McCarthy, G. (2001). Dissociating the neural mechanisms of visual attention in change detection using functional MRI. J. Cogn. Neurosci. *13*, 1006–1018.
11. Ress, D., and Heeger, D.J. (2003). Neuronal correlates of perception in early visual cortex. Nat. Neurosci. *6*, 414–420.
12. Fried, I., MacDonald, K.A., and Wilson, C.L. (1997). Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. Neuron *18*, 753–765.
13. Fried, I., Wilson, C.L., Maidment, N.T., Engel, J., Behnke, E., Fields, T.A., MacDonald, K.A., Morrow, J.W., and Ackerson, L. (1999). Cerebral microdialysis combined with single-neuron and electroencephalographic recording in neurosurgical patients - Technical note. J. Neurosurg. *91*, 697–705.
14. Quiroga, R.Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single neurons in the human brain. Nature *435*, 1102–1107.
15. Luck, S.J., Chelazzi, L., Hillyard, S.A., and Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. J. Neurophysiol. *77*, 24–42.
16. Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. Science *229*, 782–784.
17. Reynolds, J.H., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. J. Neurosci. *19*, 1736–1753.
18. Treue, S., and Maunsell, J.H. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. Nature *382*, 539–541.
19. Treue, S., and Maunsell, J.H. (1999). Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. J. Neurosci. *19*, 7591–7602.
20. Fuster, J.M., and Jervey, J.P. (1981). Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli. Science *212*, 952–955.
21. Fuster, J.M., and Jervey, J.P. (1982). Neuronal firing in the inferotemporal cortex of the monkey in a visual memory task. J. Neurosci. *2*, 361–375.
22. Miller, E.K., Erickson, C.A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. J. Neurosci. *16*, 5154–5167.
23. Landman, R., Spekreijse, H., and Lamme, V.A. (2003). Large capacity storage of integrated objects before change blindness. Vision Res. *43*, 149–164.
24. Green, D., and Swets, J. (1966). Signal Detection Theory and Psychophysics (New York: Wiley).
25. Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. J. Neurosci. *12*, 4745–4765.
26. Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis. Neurosci. *13*, 87–100.
27. Cook, E.P., and Maunsell, J.H. (2002). Dynamics of neuronal responses in macaque MT and VIP during motion detection. Nat. Neurosci. *5*, 985–994.
28. Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. Nature *341*, 52–54.
29. Landman, R., Spekreijse, H., and Lamme, V.A. (2004). Relationship between change detection and pre-change [corrected] activity in visual area V1. Neuroreport *15*, 2211–2214.
30. Stern, C.E., Sherman, S.J., Kirchhoff, B.A., and Hasselmo, M.E. (2001). Medial temporal and prefrontal contributions to working memory tasks with novel and familiar stimuli. Hippocampus *11*, 337–346.
31. Stern, C.E., Corkin, S., Gonzalez, R.G., Guimaraes, A.R., Baker, J.R., Jennings, P.J., Carr, C.A., Sugiura, R.M., Vedantham, V., and Rosen, B.R. (1996). The hippocampal formation participates in novel picture encoding: Evidence from functional magnetic resonance imaging. Proc. Natl. Acad. Sci. USA *93*, 8660–8665.
32. Rutishauser, U., Mamelak, A.N., and Schuman, E.M. (2006). Single-trial learning of novel stimuli by individual neurons of the human hippocampus-amygdala complex. Neuron *49*, 805–813.
33. Asaad, W.F., Rainer, G., and Miller, E.K. (1998). Neural activity in the primate prefrontal cortex during associative learning. Neuron *21*, 1399–1407.
34. Li, L., Miller, E.K., and Desimone, R. (1993). The representation of stimulus-familiarity in anterior inferior temporal cortex. J. Neurophysiol. *69*, 1918–1929.
35. Rolls, E.T., Cahusac, P.M., Feigenbaum, J.D., and Miyashita, Y. (1993). Responses of single neurons in the hippocampus of the macaque related to recognition memory. Exp. Brain Res. *93*, 299–306.
36. Scoville, W.B., and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. J. Neurochem. *20*, 11–21.
37. Milner, B., Corkin, S., and Teuber, H.-L. (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of H.M. Neuropsychologia *6*, 215–234.
38. Squire, L.R., and Zola-Morgan, S. (1991). The medial temporal lobe memory system. Science *253*, 1380–1386.
39. Buckley, M.J., Booth, M.C., Rolls, E.T., and Gaffan, D. (2001). Selective perceptual impairments after perirhinal cortex ablation. J. Neurosci. *21*, 9824–9836.
40. Buckley, M.J., and Gaffan, D. (2006). Perirhinal cortical contributions to object perception. Trends Cogn. Sci. *10*, 100–107.
41. Lee, A.C., Buckley, M.J., Pegman, S.J., Spiers, H., Scahill, V.L., Gaffan, D., Bussey, T.J., Davies, R.R., Kapur, N., Hodges, J.R., et al. (2005). Specialization in the medial temporal lobe for processing of objects and scenes. Hippocampus *15*, 782–797.
42. Lee, A.C., Buckley, M.J., Gaffan, D., Emery, T., Hodges, J.R., and Graham, K.S. (2006). Differentiating the roles of the hippocampus and perirhinal cortex in processes beyond long-term declarative memory: A double dissociation in dementia. J. Neurosci. *26*, 5198–5203.
43. Olson, I.R., Moore, K.S., Stark, M., and Chatterjee, A. (2006). Visual working memory is impaired when the medial temporal lobe is damaged. J. Cogn. Neurosci. *18*, 1087–1097.
44. Quiroga, R.Q., Nadasdy, Z., and Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. Neural Comput. *16*, 1661–1687.
45. Wilken, P., and Ma, W.J. (2004). A detection theory account of change detection. J. Vis. *4*, 1120–1135.