Chapter 6

Change Blindness in the Human MTL

6.1 Introduction

In the earlier chapters of this thesis we have addressed the degree of visual processing that can occur outside the focus of attention. However, as has already been pointed out, under some circumstances when attention is not directly involved, we can be unaware of certain aspects of our surroundings. On a regular basis we surprise ourselves by how frequently we fail to notice something that was "right in front of our eyes," even when we were actively looking for it. This phenomenon, as annoying as it may be, can however serve as an interesting tool to probe visual perception and attention. Consider a familiar situation—you are desperately looking for your car keys on a kitchen table that is long overdue for a good tidying up. All through your frantic search, the keys are before you, but for some interval of time, you simply cannot catch sight of them. Then, all of a sudden your search ends as you spot them, lying in a pile of clutter. Throughout this process, the visual input to your eyes has been constant but your perceptual state changes at some point. That is, somewhere in your brain some group of neurons has changed its response in order to signal the detection of the keys. As we discuss below, visual attention is a key factor that brings about this change. Thus, if one has access to the neurons that encode keys, then the effects of attention on neuronal responses can be directly investigated by comparing responses between the two perceptual states.

In this chapter, we use a paradigm that demonstrates this limitation in perception, outside the focus of attention, to study the effects of attention on neuronal responses in our patient population. We record from single neurons in the human MTL that respond selectively to certain stimuli, and we compare their responses in two conditions: when patients noticed changes made to these stimuli and when they did not.

6.1.1 Change Blindness—another attention paradigm

Subjective experience suggests that we are fully aware of our environment. However, as is often the case with subjective opinions, this is a misleading notion—under some circumstances, we can fail to notice changes that occur in a scene. Although it is difficult to demonstrate this phenomenon in real life where unexpected changes are rare, and almost always attract attention, it can be simply verified in the laboratory under controlled conditions. Movie makers are also well acquainted with this phenomenon, and to some extent it can even be argued that this failure in vision makes their job a little easier. For during moviemaking, different sequences in a scene are not always filmed in order, and consequently a sequence filmed today might be continued tomorrow. Although filmmakers go to extreme lengths to make sure that the details of a scene are as identical as possible across camera cuts, failures do sometimes occur (even in such classics as 'The Lord of the Rings' where, amongst other things, in the scene when Aragorn is talking to the dying Boromir, in some shots Boromir's right hand is gripping Aragorn's shoulder, while in others it is not; see http://www.movie-mistakes.com). If vision were so finely tuned as to always catch these inconsistencies, we would probably not enjoy our classics as much as we do. However, fortunately for the entertainment industry we can often be blind, even while seeing!

This failure of the visual system to notice changes made to scenes was discussed in the introduction to this thesis, and is known to vision scientists as “change-blindness” (Simons & Levin, 1997; Rensink & Simons, 2000; Simons, 2000; Rensink, 2002) (Figure 1.1). Striking demonstrations of this rather troubling aspect of vision show that the appearance of bizarre figures in a scene (like a gorilla) (Simons & Chabris, 1999) or the disappearance of entire objects (Rensink, O'Regan et al., 1997; O'Regan, Rensink, & Clark, 1999) can go entirely unnoticed. Because of its rather intriguing nature, change-blindness has been extensively studied over the years, under various guises. I will discuss a few of these below.

One version of the experimental paradigms used to study change-blindness is contingent on eye-movements made by observers. For instance, in one experiment, Grimes showed observers various pictures, which they were supposed to study for a later memory test (Grimes, 1996). While observers were looking at the picture, freely moving their eyes from one detail to the next, Grimes made a change to the picture during the saccades, taking advantage of saccadic suppression. He found that under these conditions, observers were surprisingly poor at noticing these changes, even when they were significant (such as two people exchanging heads). Other experimenters have used different types of stimuli to study such saccade-contingent changes (Bridgeman, Hendry, & Stark, 1975; McConkie & Zola, 1979) and have come to the same conclusion: observers fail to detect changes made during saccades, except when the change is made at the targeted location of the eye movement (Currie, McConkie, Carlson-Radvansky, & Irwin, 1995).

Saccades are not necessary for change-blindness to occur. Observers can fail to detect changes made to a scene even when they are fixating. In one set of experiments, O’Regan and colleagues showed that changes go undetected when they are made concurrently with some other global disturbance on the screen (O'Regan, Rensink et al.,

1999). In their “mud splash” paradigm, they made changes to a picture and simultaneously presented a few, small, high-contrast patterns in different parts of the picture. These patterns, which resembled mud splashes on a car windshield, never covered the location of the change or obscured it in anyway. Subjects were instructed to fixate on each picture and to press a button as soon as they noticed a change, which could be as large as an object or region moving, changing color, or appearing or disappearing. The results showed that when changes occurred at the same moment as the “mud splashes,” subjects often failed to detect changes.

The situations described in these paradigms differ from everyday life in one crucial aspect—under normal circumstances, when changes occur suddenly in a particular place they cause disturbances or transients at the location of change. These transients are readily detected because they draw attention (Posner, 1980; Klein, Kingstone, & Pontefract, 1992) and thereby signal the location of the change. Thus normally, because the transient motion signals associated with changes can attract our attention, we are fairly efficient at detecting changes. Under the experimental conditions described above, however, transients associated with changes are prevented from grabbing attention.

In the saccade-contingent paradigm, this happens because of a phenomenon that occurs during saccades. Rapid eye movements such as saccades blur the visual image on the retina (Burr, Morrone, & Ross, 1994; Castet & Masson, 2000), thereby also blurring the change-related transient signal. This blurring effectively decreases the strength of the transient signal as a result of which it fails to grab attention. This results in change-blindness. In the “mud splash” paradigm, the transient signal itself is not blurred. Instead, in this case, the “mud splashes” themselves generate a number of other local transients all over the scene. Consequently the transient signal corresponding to the true change is just one in a number of such signals that arise and

therefore becomes less salient. Attention is not drawn automatically to the true transient signal, and if it focuses on the location of a “mud splash,” change-blindness occurs.

Another simpler way of diminishing the effective strength of the transient signal to attract attention is demonstrated in the flicker paradigm (Phillips, 1974; Pashler, 1988; Rensink, O'Regan et al., 1997). In this paradigm, the original and changed images are presented to observers, but with a brief blank screen shown between them (Figure 1.1). The images are presented to observers repeatedly until they detect the changing object, which can often take more than 1 minute. In each frame, the flicker introduces transients at every location in the scene, and just as with the previous paradigms, this ensures that the transient motion signals at the actual location of change are not as salient, thus hindering change detection. Additionally, other interesting observations have been made with the flicker paradigm, that further illustrate the role of attention in change detection. Firstly, in their experiment with the flicker paradigm, Rensink and colleagues tried to ascertain the influence of attention on change detection by comparing differences in responses to objects that attracted attention with those that did not. To determine which objects attracted attention, they asked 5 naïve observers to briefly describe a number of scenes verbally. The objects in the scenes that attracted attention were those mentioned by more than 3 observers and were defined as objects of central interest (CI). Objects of marginal interest (MI) were those that were mentioned by none of the participants. 10 different subjects were then tested in the change-blindness experiment using these scenes. As expected, Rensink et al., found that these participants demonstrated change-blindness. However, they also noticed that observers took much longer to detect changes made to the MI objects compared to the CI objects. They concluded that “visual perception of change in an object occurs only when that object is given focused attention …” (Rensink, O'Regan et al., 1997).

To provide further support for this claim, in another experiment, Rensink et al. also studied the influence of cueing attention to different locations in the scene. On some trials, subjects were presented with a valid cue that correctly predicted the location of the change. On other trials, the cue was misleading and directed attention to some other location in the picture. Their results showed that on the valid cue trials reaction times to both the MI and the CI objects were greatly sped up, and in fact there was no difference in these times between the two sets of objects. The invalid cues, on the other hand, had no significant effect on the time it took observers to notice the change. Thus when subjects' attention was directed to an object, they were faster in noticing changes to that object, confirming the conclusion that attention is a critical factor in change detection.

These observations made in the different change-blindness paradigms demonstrate the role of attention in detecting changes. In brief, it appears that on those trials when subjects miss noticing changes, they were less likely to be paying attention to the changing object. Conversely, on change detection trials, attention was more likely to be engaged by the changing object. Thus, along with the visual search and dual-task paradigms, change-blindness can also serve in the study of attention.

6.1.2. Attention in the human MTL

In this chapter, rather than using the dual-task paradigm, we study the effects of attention at the single-cell level using the change-blindness paradigm in our patient population. The main reason behind this switch in paradigms is the difficulty involved in running the dual task paradigm in the hospital setting. In the first place, the dual-task paradigm is very demanding on subjects given the short presentation times of the stimuli in the two tasks. For this reason, we run the risk of tiring patients out in a very short period of time. A logistical difficulty also comes about because of the great amount of

time required for training subjects and determining their SOAs in the dual-task. On average, subjects take about 10 hours for training, and since we cannot spend 10 hours before surgery with the patient for the sole purpose of training, running the dual-task paradigm in the hospital setting is not practical. For these reasons it was necessary to engage another simpler attentional paradigm, such as the change-blindness paradigm.

As has been discussed in the previous section, a significant amount of work has gone into the study of change-blindness at the psychophysical level. Its neuronal correlates on the other hand have only recently come under investigation. Much of this work has used fMRI to localize the different brain areas involved in the detection of changes (Ress, Backus, & Heeger, 2000; Beck, Rees, Frith, & Lavie, 2001; Ress & Heeger, 2003; Pessoa & Ungerleider, 2004). However, to the best of my knowledge, the single neuron correlates of this processing have not yet been studied. By using the change-blindness paradigm in the patient population, we thus not only examine the effects of attention on neuronal responses, but also investigate the neuronal correlates of change-blindness at the single cell level.

6.2 Methods

6.2.1 Behavioral task

The experimental timeline for 1 trial in the change-blindness paradigm is shown in Figure 6.1a. Each trial began with a fixation cross that was presented for a random interval between 1000 ms and 1200 ms. Following the fixation cross, 4 images appeared on the screen for 1000 ms as shown in the figure. The pictures were presented at 4 locations on a circle with radius of 6 degrees of visual angle. The midpoint of each

picture was located on the circle; each picture subtended approximately 1.5 degrees of visual angle. A blank interval of 1500 ms (black screen) followed the first set of pictures, and then a second display of 4 pictures appeared for 1000 ms. In the second display period, the pictures occupied the same location as the previous set of pictures. Additionally, in roughly half the trials, one of the 4 pictures was changed between the two display periods. Patients were instructed to report at the end of each trial if they noticed a change or not. The presentation times for the two display periods as well as the inter-stimulus interval (ISI) were determined beforehand following extensive psychophysics on a non-clinical population. Six subjects were tested on this paradigm, with ISI intervals of 100 ms and 1500 ms. Additionally, different display sizes (2, 4, or 6 images) were also tested. The performance of the non-clinical population was determined with each of these parameters, and the current set of parameters was chosen because it resulted in a hit rate of about 70% on average, thus ensuring some proportion of incorrect trials (Figure 6.2 a, b) .

The stimuli were presented to patients on a Macintosh G3 laptop computer. The laptop was placed on the patient's lap or a tray table in front of the patient depending on the patient's preference. Although we could not precisely constrain the viewing distance for each patient because of constraints in the hospital ward, in all cases, the laptop was roughly 60 cm away from the patient. Patients pressed two keys ('Y' and 'N') on the keyboard to report their responses on each trial.

During the experiment, patients were always instructed to fixate on the central fixation cross and to only covertly explore the pictures in each display period. Although we did not explicitly control for eye movements because of the difficulty involved with introducing the equipment into the clinical ward, there is evidence that suggests that eye movements do not influence neuronal responses in higher-level visual areas. Recordings in area IT in monkeys have shown that eye movements have virtually no

effect on the neuronal representation underlying object recognition (DiCarlo & Maunsell, 2000), particularly in plain backgrounds (as opposed to complex natural scenes) (Rolls, Aggelopoulos, & Zheng, 2003). Ringo et al. (Ringo, Sobotka, Diltz, & Bunce, 1994) have shown some modulation in firing rates due to eye movements in the anterior and medial temporal lobes but found no correlation between cells that were modulated by saccades and cells that were visually responsive. Perhaps most importantly, the strongest evidence for our set-up comes from control experiments carried out by Kreiman for his thesis (Kreiman, 2002). In these experiments, conducted in the same hospital setting as our experiments, Kreiman demonstrated that there was no modulation in firing rates of visually responsive cells as a result of eye movements. Based on his experiments and the literature discussed above, he concluded that "these data strongly support the idea that eye movements do not significantly alter the visual representation in MTL" (p.110 in (Kreiman, 2002)).

6.2.2 Stimulus presentations and trial types

In most paradigms that study attention in monkeys (e.g. (Moran & Desimone, 1985; Reynolds, Chelazzi et al., 1999)), for each recorded cell, stimuli that drive the cell well and weakly (preferred and non-preferred stimuli, respectively) are first selected, and then used to investigate the effects of attention on the cell's response. Similarly, in this experiment, it was necessary to select a set of preferred (usually, only one) and non-preferred (4 to 8 on average) stimuli for each targeted cell. This was greatly facilitated by the screening/testing experiment format described in Chapter 5. Given the data from the screening session, we usually had a good idea of the visually responsive neurons we were recording from and the stimuli that drove these neurons strongly or weakly. Based on this information, we could select the stimuli to use in the change-blindness

experiment. As a result of our time constraints, we were limited in the number of preferred pictures we could select for each session, since we wanted to ensure that we had enough repetitions of the change and no-change trials for each preferred picture that we presented. Thus, given the 30-minute recording session on average, and taking the length of each trial into account, this limited us to 4 preferred pictures in each session.

The limitation on the number of preferred pictures also constrained the number of neurons we could target for our study during a single experimental session. In some sessions, the 4 stimuli we chose drove 4 independent neurons that would count towards the total number of neurons in the change-blindness experiment. In most cases though, in the screening session only two or three neurons would have the strongest responses and would respond significantly to several pictures. Accordingly, in these cases, a smaller number of neurons would contribute to the total count. All in all, we recorded 17 change detection sessions in 9 patients. In these sessions, although we had recorded approximately 110 visually responsive neurons (for a definition of responsiveness, see Section 6.3.4) from a total of approximately 740 cells in the corresponding screening sessions, because of the limitations explained previously, only 49 of these neurons were targeted (i.e. had been pre-selected on the basis of the screening session, and preferred and non-preferred stimuli had been determined for these cells). Only these 49 targeted cells were included in our analysis.

The general design of the experiment also mimicked monkey experiments in that, on each trial, only one of the stimuli was selected from one cell's set of preferred pictures. The other stimuli in each display period were chosen from the group of non-preferred pictures (common to all targeted cells). Based on task demands, from the subject's point of view, there were only two different trial types—change and no-change trials. However, as Figure 6.1b shows, from the point of view of a particular neuron, 4

different trial types could be determined with respect to its preferred stimulus. For example, consider a neuron that responds selectively to a particular picture of Bill Clinton. In the first type of trial, the picture of Clinton could have been one of the four images in the first display period, but could have been replaced by another picture in the second display period. We will call these types of trials "disappear trials." In the "appear trials", the preferred picture of the cell, could have been absent in the first set of four pictures but then have appeared in the second display period. On trials when no change occurred, the preferred picture of the cell would have been present in both display periods—the "both trials". And finally, the preferred picture could have been absent all together on the "none trials". In order to study the effects of attention on the cell encoding this picture of Clinton, we will have to consider its activity during all the relevant display periods (where the preferred stimulus was present), that is, the first display period of all the "disappear" trials, the second period of all the "appear" trials, and both display periods of the "both" trials. Over these intervals, to get at the effects of attention on this cell, we will then look at the difference in firing rates between correct and incorrect trials since, as discussed before, the probability of attention being focused on this picture is higher on the correct trials than the incorrect ones. For every preferred stimulus, there were 20 "disappear", 20 "appear," and 14 "both" trials over the recording session. Overall for all preferred stimuli, there were at least 50 "none" trials.

6.2.3 Recordings

The procedures for electrode implantation, and data acquisition and pre-processing have been discussed previously in Chapter 5. The current chapter describes data obtained in 9 patients (3 female, all right-handed, 17-47 years old) who participated in a total of 17 sessions of the change-blindness experiment.

6.2.4 Data analysis

Visual responsiveness

As was discussed in Section 6.3.2, for each cell, preferred and non-preferred stimuli were determined based on the results of the screening experiment. In the testing experiment, we verified if the preferred stimulus from the screening session still elicited significant increases in firing rates by computing a paired t-test. This test was uncorrected for multiple comparisons because only one hypothesis was being tested. The t-test compared the distribution of firing activity during the [-1000, -300] ms baseline interval and the [300, 1000] ms interval during which the stimulus was presented over all trials. (0 ms represents the time of stimulus onset.) A cell was considered to have remained visually responsive if the p value of the t-test was < 0.05.

Population responses

Population responses were computed using the normalized spike density function (sdf) (Richmond, Optican, & Spitzer, 1990; Sheinberg & Logothetis, 1997). For each neuron, the sdf was obtained by convolving the spike train on each trial with a 200 ms fixed width Gaussian and then averaging over all trials. The spike trains were binned in 5 ms bins before convolution. For each unit, the sdf was normalized by dividing by its peak activity. The normalized sdf was then averaged over the population of cells. Although, for each neuron the maximum of this normalized sdf was equal to 1, over the population of cells, the maximum value is less than 1 because the activity in different cells peaked at different times. The average responses were computed separately for correct and incorrect trials in all trial types. Cells were included in this analysis in each

condition only if there was a minimum of 4 trials in that condition. This method is similar to the one used by Kreiman (Kreiman, 2002).

6.3 Attention in the MTL

6.3.1. Visual responses in a crowded environment

During the change-blindness experiments, we targeted 49 neurons that had been visually responsive in the screening sessions in 9 patients. Of the 49 targeted neurons, 29 were still responsive during the testing sessions. In this section, we will first describe the responsiveness of these 29 cells, and then speculate on why the other 20 neurons appear to have "lost" their visual selectivity.

6.3.1.1 Reduction in the strength of responses

Of the 49 visually responsive neurons from the screening sessions, 29 were still responsive when presented with 4 other images and were further analyzed (see Table 6.1 for the location of these neurons). An example of such a neuron, located in the right amygdala is shown in Figure 6.3. The response of this neuron in the screening session is shown in Figure 6.3a while the response in the change-blindness experiment is shown in Figure 6.3b. In comparison to the response in the screening session, in the testing session the response is much weaker when the same stimulus is placed together with 3 other stimuli known to drive this cell poorly. While in the screening session, the neuron had an average firing rate of about 14 Hz; the response in the testing session was about 7 Hz on average. The response of another unit located in the right hippocampus is

shown in Figure 6.4. Over all 29 cells that were visually responsive in the testing sessions, the relationship between firing activity in the screening and testing sessions is shown in Figure 6.5. As the graph shows, there is a substantial drop in the firing activity over the population of neurons in the change-blindness sessions compared to their response in the screening experiments. The slope of the graph is 0.2 signifying that the responses in the testing session were 5 times lower on average.

This significant drop in firing rates observed in the testing session can be accounted for by the competitive interactions that would arise amongst our stimuli. As is well known from monkey electrophysiology, the response of a neuron to a preferred stimulus in a crowded environment is much lower compared to its response to the same stimulus presented in isolation. In these experiments (see Fig 1.5b), a drop in firing rate is observed if even one additional stimulus is placed in competition with the preferred stimulus of the cell (Moran & Desimone, 1985; Sato, 1989; Treue & Maunsell, 1996; Luck, Chelazzi et al., 1997; Missal, Vogels, Li, & Orban, 1999; Reynolds, Chelazzi et al., 1999; Treue & Maunsell, 1999; Chelazzi, Miller, Duncan, & Desimone, 2001; Rolls, Aggelopoulos et al., 2003). In comparison, in our experiment, 3 additional stimuli known to generate a poor response in the targeted neuron were presented along with the preferred stimulus of the cell. Thus, in effect, the environment in the change-blindness experiment was more cluttered compared to these classic electrophysiology experiments, increasing the competitive interactions among stimuli. These interactions between stimuli would result in reduced firing activities.

6.3.1.2 Reduction in the number of responsive cells

Of the 49 cells we targeted from the screening session, only 29 maintained statistically significant responses (i.e. $p < .05$) in the testing session, while the other 20

lost their responsiveness. This loss in the number of responsive cells could simply be a consequence of the decrease in response amplitudes observed in the cluttered environment (as discussed above), which would make the detection of any signal more difficult. Furthermore, as we have discussed in Chapter 1, unattended stimuli in cluttered scenes typically get filtered out. The absence of focal attention on certain stimuli in our cluttered environment could similarly reduce the selectivity of the relevant neurons.[1] Another reason could be the time delay between the screening and testing sessions. In general, we were fairly successful in running the testing session shortly after the screening session. Typically, the delay was on the order of 3-4 hours, the time it took us to analyze the data acquired in the screening session. However, on a few occasions, we had to postpone the testing experiment till much later in the day, or even the next day, generally because the patient did not want to test right away. Of the 17 change-blindness sessions we ran, 7 were tested within 4 hours of the screening experiments, 6 were tested 6 hours later, while the other 4 were tested 12 or more hours later. These delays between screening and testing sessions could also account for the differences in neuronal responses we observed. It is possible that during these delays, an electrode could have moved relative to the skull, and that in the testing session we were no longer recording from the same neuron (for an in-depth discussion on the stability of recordings in our clinical setting, see (Kreiman, 2002)). Additionally, it is not necessarily true that a neuron would always maintain the same selective responses. One would assume *a priori* that a response would be stable over days, but even in our screening sessions we observed neurons that had more plastic responses. For example, there were instances when after a few days of recording, cells developed responses to the experimenters whom the patient had only ever seen in the previous few days.

[1] As we have already discussed, on incorrect trials, attention is less likely to be focused on the relevant stimuli.

6.3.2. Behavioral responses

Over the population of subjects during the change-blindness experiments, the average hit rate (i.e. the probability of reporting a change when it occurred) was 77.8 ± 2.0% (mean ± s.e.m.). The average false alarm rate (falsely reporting a change when there was none) was 18.8 ± 2.0% (Figure 6.2c). The average hit rate was close to that obtained for our non-clinical population (see Section 6.2.1), suggesting that despite being in the hospital and under medication, the clinical population performed fairly well on this task. The patients' percent correct was 78.0 ± 1.7% on average.[2]

6.3.3 The neural correlates of change blindness and change detection.

As we discussed in the introduction to this chapter, the question we are interested in here is what happens at the level of individual neurons in the MTL when observers detect a change, compared to when they do not. The response of one neuron in the anterior hippocampus is shown in Figure 6.6a, and another unit in the right amygdala is shown in Figure 6.6b. The responses in each of the 4 conditions are shown in these figures where the histograms in red correspond to the correct trials in each condition and the black histograms reflect the incorrect trials. The dashed vertical lines represent the time intervals of the first display period, the ISI interval, and the second display period, respectively. A similar analysis was conducted for the entire population of responsive cells. The average response over the population in the 4 separate trial types is shown in Figure 6.7. As discussed in Section 6.3.3, this figure shows the mean

[2] Note that chance performance was not always at 50% because the design of the experiment included a slightly higher number of change trials compared to no-change trials. On average, chance performance was 56.6 ± 0.9 (mean ± sem).

normalized spike density function averaged over all cells and all trial types, and computed separately for the correct and incorrect trials. The red trace shows the average activity in the correct trials, and the black trace corresponds to incorrect trials. The shaded regions reflect the standard error of the mean. It is evident from this figure that for all trial types, neuronal responses over our population of 29 cells were higher in correct trials compared to the incorrect trials, during the intervals that contained the preferred stimulus. For example, during the "disappear" trials, cell firing rates were significantly higher ($p < 0.001$)[3] during the first, 1s long display period (when the preferred stimulus was presented to subjects) for the correct trials compared to the incorrect ones. Similarly during the "appear" trials, firing rates were significantly higher ($p < 0.01$) in the second display period for the correct trials, and in the "both" trials, firing rates for correct trials were higher during both display periods ($p < 0.05$). These results thus 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, during those display periods when the preferred stimuli were not present, no significant difference was observed between correct and incorrect trials ($p > .05$), 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, increased activity on correct trials was also observed during the 1.5s long blank interval between the two display periods on the "disappear" and "both" trials that approached significance ($p = 0.06$).[4] In both these types of trials, a preferred stimulus was presented during the first display period, and this increased activity after the stimulus was removed might reflect a memory of the stimulus encoded in the neuronal trace.

[3] Significance values were computed using a two-tailed t-test to compare activity on correct and incorrect trials. Neuronal activity was integrated over a 1s interval for the periods when the preferred stimuli were present (starting at 300ms, to take into account the latency in responses).

[4] Significance values were computed as before using a two-tailed t-test to compare activity during the [1000 2500ms] time interval.

Figure 6.8a presents (in a summarized form) the average activity over our population of cells in all intervals when the preferred stimuli were present. In other words, the data shown here is averaged over the first display period of the “disappear” trials, the second display period of the “appear” trials, and both periods of the “both” trials. The red and black traces again correspond to correct and incorrect trials. As mentioned earlier, the results show that firing rates were significantly higher on correct trials compared to incorrect trials when the preferred stimulus was present on the screen ($p < 0.0002$). These data thus demonstrate the correlates of change-blindness and change detection in the human MTL. When subjects correctly detected changes, firing rates in the population of cells encoding the relevant stimuli were significantly higher during the time intervals when these stimuli were presented to the subjects. On the other hand, neuronal activity during these time intervals was reduced when subjects were blind to changes. In contrast to Figure 6.8a, the results in Figure 6.8b summarize the average neuronal activity in those time intervals when the preferred stimuli of the cell were absent. In this condition, no significant difference in firing rates are observed between correct and incorrect trials ($p > 0.4$), which is expected since only those stimuli that drove these cells weakly were presented to subjects.

6.3.4 Predicting change

The data presented thus far is averaged over many trials and many cells, and shows that on average the activity of individual neurons differs between correct and incorrect trials. However, on each trial, an individual cell encoding a particular stimulus does not have the liberty to compute averages. Instead, if it participates in change detection, it must contribute to the decision based on information it receives on a trial-by-trial basis. How could this neuron tell us on each trial, whether a change occurred or

not? A very simple model of change detection could assume that on each trial, a decision-making system would monitor the activity of neurons encoding the presented stimuli over the length of the trial. It would then compare the firing rates between the two display periods, and based on this comparison it could signal the presence or absence of a change. A change in activity between the first and second display intervals would signal that a change had occurred in one of the stimuli, while similar firing rates in the two intervals would indicate no change.

How reliably can the activity of a neuron on individual trials tell us whether or not a change occurred? We addressed this issue quantitatively using a receiver operating characteristic (ROC) analysis that is used classically in psychophysics (Green & Swets, 1966). 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. On each trial, we computed the difference between the firing activity in the two display periods[5] and used a sliding threshold to determine the probability of false alarms and the probability of correct detections. More precisely, on each trial, if we denoted the firing activity in the two display periods as r_1 and r_2 respectively, and the event of a change as X, we were interested in determining the quantity $P(X \mid \Delta r)$ for different values of a threshold T, (where Δr is the difference between r_1 and r_2). For each integer value of the threshold T, a difference in spike number between the two display periods greater than T, (i.e. $\Delta r > T$), would mean that the neuron signaled that a change had occurred on that trial. Depending on if a change had actually occurred or not, this signal would then count either as a “correct detection” or a “false alarm”. We computed these values for all trials together, as well as separately for correct and incorrect trials.

[5] We computed the ROC over different periods of time in both display periods (Figure 6.13). To be precise, the number of spikes in the interval [300 X] ms (aligned to the onset of each display period) was computed separately for X = [700, 1000, 1500, 2000, 2500]. For all the data presented in this chapter, ROCs were evaluated using spike counts in the [300 1500] ms interval.

The ROC curves we thus obtained for each of our cells are shown in Figure 6.9a and 6.9c. In these plots, the probability of correct detection is plotted against the probability of false alarms for different values of the threshold T. The area under each curve is a measure of the ability of each neuron to estimate whether a change occurred. A value of 0.5 would correspond to chance performance, while a value of 1 would reflect perfect accuracy in signaling a change. Over the population of cells, when computed over correct trials, the average value was 0.67 ± 0.02. In contrast, the average ROC area computed over incorrect trials was significantly lower (0.49 ± 0.04; $p<.005$) and not significantly different from chance ($p=0.95$).

The individual values for each cell in our population over correct and incorrect trials are plotted in Figure 6.9b and 6.9d, respectively. Over the correct trials, the distribution of the area under the ROC curves (Figure 6.9b) is centered at 0.67 and is significantly shifted to the right of 0.5 ($p < 0.0001$). On incorrect trials, however, the mean of this distribution is no different from chance (Figure 6.9d). These figures thus indicate that the ability to predict whether or not a change occurred on individual trials depends on whether the patient made correct or incorrect behavioral reports on the corresponding trials. In other words, on trials when the subject correctly reported the presence or absence of a change, the MTL cells also appeared to be able to correctly make this judgment. On the other hand, when subjects were in error, the MTL units did not appear to carry the relevant information either. These results are therefore compatible with the notion that MTL neurons could contribute to subjects' perception during change-blindness and change detection: on trials when subjects were "blind" to changes, the neurons appear to be "blind" as well, whereas on correct trials, both subjects and MTL neurons were able to detect changes.

Additionally, this data suggests that at the neuronal level, attention plays a key role in a neuron's ability to correctly signal a change. Under the assumption that on

correct trials attention was more likely to be focused on the relevant stimuli, the ROC values indicate that with the influence of attention, we can generally guess above chance whether or not a change occurred simply by observing the firing rates of individual neurons on individual trials. On the other hand, when attention is presumably not focused on the relevant stimuli, we are as likely to correctly predict a change as not. The significant, beneficial effect of attention in predicting a change is thus demonstrated by the difference of the two distributions shown in Figure 6.9b and d ($p<.005$).

This ability of individual cells to signal the occurrence of a change can be compared to the behavioral performance of our subjects. On average over subjects, the percentage correct in change detection was 0.78 ± 0.01. This value is higher than the performance of individual neurons (0.65 ± 0.02%),[6] but this is not unreasonable, given that the behavioral performance of subjects presumably reflects information contributed by more than one neuron. Although it has been observed that single neurons can sometimes be as sensitive as subjects in some tasks (Newsome, Britten, & Movshon, 1989; Britten, Shadlen, Newsome, & Movshon, 1992), in most cases it is reasonable to suppose that performance of single neurons is bounded by overall behavioral performance.

The performance achieved by the ROC analysis in predicting the occurrence of a change can be further understood by comparing the firing activity of each unit in the two display intervals. Theoretically, on change trials, the larger the differences in firing activity between the intervals when the preferred stimulus was present versus absent, the higher the probability of predicting a change. Conversely, on no-change trials when the preferred stimulus was present in both intervals, for high prediction values, firing activity should be almost identical in the two display periods.

---

[6] The ROC value presented here is computed over both correct and incorrect trials together. This value over all trials is closer to the ROC value for correct trials as a result of the fact that there were many more correct than incorrect trials.

For our set of cells, the relationship between firing activity in the two display periods of the change trials is shown in Figure 6.10a. It is evident from the figure that activity in the interval when the preferred stimulus was absent was significantly lower ($p < .0001$) than the corresponding activity in preferred stimulus "present" intervals. On average, firing activity during the former period was 60% of the activity in the latter interval. The relationship between firing activity in the two display intervals of the no-change trials when the preferred stimuli were present (i.e. "both" trials) is shown in Figure 6.10b. This figure demonstrates that firing activity during the second interval was significantly ($p < .00001$) lower in the second display period compared to the first. This difference in firing activity might appear to be surprising since the physical stimulus was identical in the two display periods. However, previous reports of decreases in firing activity with prolonged exposure to a particular stimulus (e.g. due to adaptation or habituation) are consistent with this observation (Miller, Li, & Desimone, 1991; Lueschow, Miller, & Desimone, 1994; Miller & Desimone, 1994; Desimone, 1996).

Another way of examining the data that is summarized in the ROC plots is to compare the distributions of the difference of firing activity between the two display intervals in the change and no-change trials. In principle, the further apart these distributions are, the easier it is to correctly predict the occurrence of a change. The distribution of firing activity in the change and no-change trials is shown in Figure 6.10c, d for the correct and incorrect trials, respectively. The black distributions in these plots, centered at 0, correspond to the "null" distribution—the distribution of differences in firing activity between the two display periods of the no-change trials. Because of our earlier observation that in the ("both") no-change trials, firing activities are not identical in the two display periods, to obtain the null distributions, we have averaged the two distributions obtained by subtracting r_1 from r_2, and vice versa. By thus averaging, we remove the bias introduced by choosing either one of the two subtractions. Indeed, in

the change trials, with which we wish to compare these null distributions, the target picture occurs with equal probability in the first or second periods.

To estimate the ability of cells to be able to predict changes, the distributions of firing activity during change trials can be compared to these "null" distributions. The distributions for the change trials are shown in red. These distributions correspond to differences in firing activity between the stimulus present and stimulus absent intervals of the change trials. The mean of this distribution over correct trials is centered at 2.6, indicating that on average, for each cell, the difference in firing activity between the two display intervals is 2.6 spikes on change trials. Although this value appears to be fairly low at first, it is interesting to recall reports that suggest that subjects are able to distinguish between stimuli on the basis of 1-2 spikes in the somatosensory system (Vallbo & Johansson, 1976; Johansson & Vallbo, 1979; Parker & Newsome, 1998). On incorrect trials, however, the mean of the distribution is much smaller (0.8 spikes), again demonstrating that on trials when subjects made errors, the difference in firing activity between the two display periods on change trials was not readily distinguishable from the corresponding difference in activity on no-change trials.

6.3.5 Predicting behavior

The firing activities of individual neurons in the MTL are thus able to give us information about the visual stimulation presented to subjects—they are able to signal above chance whether or not a change occurred. But to what extent do these cells participate in the subjects' decisions? Does their activity on individual trials allow us to guess above chance the behavioral report of the subject? In other words, using the terminology mentioned in the previous section, and denoting the behavioral report of subjects by Y, we are interested in estimating $P(Y \mid \Delta r)$.

To understand this issue, we conducted an ROC analysis to determine the correlation between neuronal activity and the behavioral report of subjects on a trial-by-trial basis. For this analysis we again estimated the probability of whether or not a change had occurred by comparing spike counts in the first and second display periods, as described above. We then compared this estimate not with the actual stimulus, but with the behavioral report of the subjects, for each value of the sliding threshold T. A "correct detection" or "false alarm" would be counted depending on whether the prediction from the spike counts matched the subjects' behavioral reports or not.

The ROC curves for this analysis are shown in Figure 6.11a. As before, the probability of correct detection is plotted against the probability of false alarms for each value of the threshold. The average area under the curve in this case is 0.58 ± 0.01. The distribution of these values for all cells is shown in Figure 6.9b and as before is significantly shifted to the right of 0.5 ($p < 0.001$). This indicates that on a trial-by-trial basis, the behavior of subjects, (their perception of the change, or the corresponding change-blindness) can be predicted above chance by individual neurons in this population.

The relationship between predicting changes in the stimulus and behavioral decisions reflecting the percept of the subject is shown in Figure 6.12 for each cell. On average, the population of cells appears to be better at predicting stimulus related information than the behavioral percept of the subjects (slope = 0.4). Thus, in a continuum of processing from the retina (which would be perfect at predicting changes in the stimulus) to the output motor neurons (which would presumably perfectly reflect the behavioral decision of the subject), MTL cells appear to be closer to the sites of stimulus rather than perceptual/decision-making processing.

This conclusion appears to be at odds with a previous report that demonstrated that the activity of MTL neurons correlates strongly with the behavioral percept of

subjects during flash suppression (Kreiman, Fried, & Koch, 2002). However, the high correlation between neuronal activity and subjective percept observed in that study could be related to the relatively unnatural manner in which stimuli are presented during flash suppression. Flashing an image to one eye could completely suppress neural activities in response to the other stimulus, thereby giving rise to the observation that suppressed (i.e. unperceived) stimuli elicit no neuronal response in the MTL. Under more natural viewing conditions, however, MTL neurons could maintain weak (and non-zero) responses to stimuli that are not perceived, thus suggesting that these neurons are not strongly correlated with perception.

6.4 Discussion

In this experiment, we set out to investigate the effects of attention on neuronal responses in the human MTL. Our results demonstrate that in a change-blindness paradigm, when subjects correctly detect the presence or absence of a change made to particular stimuli, firing activity of neurons encoding these stimuli are high during stimulus presentation. On the other hand, neuronal activity during presentation of the relevant stimuli is significantly reduced when change-blindness occurs. Thus, we demonstrate for the first time at the single-cell level the neuronal correlates of change-blindness and change detection. Furthermore, under the reasonable assumption that on correct trials attention was more likely to be focused on the relevant stimuli than on incorrect trials, these results also demonstrate the effect of attention on the population of neurons. From our data it appears that attention enhances firing rates of neurons encoding the attended stimuli. This attentional effect, which is consistent with observations made in monkeys (Moran & Desimone, 1985; Treue & Maunsell, 1996;

Reynolds, Chelazzi et al., 1999) is demonstrated here for the first time in single neurons of the human brain.

Our results also showed a tendency towards higher firing rate on correct trials during the blank interval (ISI) between the two stimulus periods. This activity was observed on those trials when the preferred stimulus of the cell was presented in the display period preceding the ISI (i.e. the “disappear” and “both” trials). This enhanced activity which persists after the stimulus is removed might reflect a memory for the stimulus, that is more strongly encoded on correct rather than incorrect trials. It is not unreasonable to suppose that a stronger memory of stimuli from the preceding interval could contribute to better performance on a trial. Indeed some theories of change detection assume a stable representation of a limited number of objects in a scene, which is used for online comparison with incoming information (Rensink, 2000, 2002). The memory trace observed after the offset of the stimulus could serve as a neuronal correlate of this representation. This interpretation is also in line with the role of the hippocampus and other medial temporal lobe structures in the formation of short-term memories (Squire & Zola-Morgan, 1991; Kandel & Hawkins, 1992; Zola-Morgan & Squire, 1993; Kandel, Schwartz, & Jessell, 2000; Buckner & Wheeler, 2001). Damage to these structures results in learning deficits and impairs formation of recent memories, as is strikingly demonstrated in the case of patient H.M. (Penfield & Milner, 1958; Milner, 1972). Thus, MTL structures are involved in the temporary storage of information during visual processing, and the trace we observe during the ISI interval could represent this storage.

Based on the neurons they record from, monkey electrophysiologists have often tried to determine how the firing of individual neurons relates to the animal’s behavior (Celebrini & Newsome, 1994; Britten, Newsome, Shadlen, Celebrini, & Movshon, 1996; Shadlen, Britten, Newsome, & Movshon, 1996). Similarly, given the neural correlates of

change detection in the MTL cells, it is interesting to ask how close these MTL neurons are to the site where the decision is made regarding change perception. On the one hand, if the decisions were made in the MTL itself, one could reasonably expect that the firing rate of MTL cells should be highly correlated with the behavioral report of the subjects on a trial-by-trial basis. However, as our ROC analysis demonstrates, while individual neurons can predict the subjects' behaviors significantly above chance, this performance is correct only 58% of the time.[7] Furthermore, if MTL cells perfectly reflected the percept of subjects, on correct trials one would expect high firing rates during the presentation of the relevant stimulus, but on those trials when subjects made a mistake and presumably did not perceive the relevant stimuli, the neurons representing these stimuli should be silent. However, our data does not indicate that neurons are silent on incorrect trials. Instead, the data shows that even on these trials, neuronal activity in response to the preferred stimuli was significantly higher when compared to baseline activity. It is as though in these situations, although the subject does not see the particular stimulus (resulting in change-blindness), the recorded neuron does to a certain extent. Thus, at the neuronal level, the subjective percept and neuronal activity do not match up exactly, resulting in lower levels of performance in predicting behavioral reports.

Another factor to explain the moderate performance value from the ROC analysis could be that this analysis reflects a behavioral prediction that is based on the activity of a single neuron. In contrast, the decision systems of the brain probably monitor large populations of neurons simultaneously. The limited information provided by individual

[7] Note however, that this value is comparable with several earlier reports of the relationship between neural responses and perceptual judgments in monkey area MT (Britten, Newsome et al., 1996; Cook & Maunsell, 2002). In these studies, the authors report that MT neurons could predict the animal's behavioral choice roughly 56% to 60% of the time.

cells might not be sufficient to always reflect behavior (Britten, Newsome et al., 1996; Cook & Maunsell, 2002).

In addition, if one assumes that a neuron carries visual information that contributes to the behavioral decision, then the information it can carry about the change is likely to be bounded by what the neuron "sees" of the change. In other words, its ability to accurately signal the presence or absence of a change might be constrained by its ability to reliably encode the presence or absence of its preferred stimulus. As we saw in the ROC analysis, based on the firing rate of individual neurons, changes in the presence or absence of preferred stimuli can be guessed 65% of the time on average. Given our assumption that the MTL neurons contribute to the decision of change perception, their performance at predicting behavior would then have an upper bound of 65%. Under this constraint, it turns out that the MTL neurons would actually be fairly efficient in using the stimulus-related information available to them—their performance in estimating behavior being 77% of their ability to signal a change.[8] (Note again that this upper bound is only valid under the assumption that we are dealing with neurons that are involved in the perceptual and cognitive operations underlying change detection. For example, a motor neuron encoding subjects' button presses would be accurate 100% of the time in predicting the behavior of subjects, without carrying any independent stimulus-related information.)

Alternatively, the seat of the behavioral decision could be located in other brain regions. Evidence for this hypothesis comes from recent fMRI work investigating the neuronal correlates of change-blindness and change detection. These studies have reported that visual areas in the ventral stream are activated strongly during change

[8] This calculation scales performance in behavioral prediction relative to performance in signaling a change based on stimulus-related information, while keeping chance at 50%. The exact formula used is: 0.5+0.5*(B-0.5)/(S-0.5) where B is performance on behavioral prediction (for us 0.58), and S is performance on signaling change (0.65).

detection and to a lesser extent (but still significantly) during change-blindness (as we show here for the medial temporal lobe cells). However, activity in dorsal areas was more clearly differentiated depending on whether subjects perceived the change or not. In fronto-parietal sites for instance, enhanced activity was observed during change detection, but no activity was detected during change-blindness (Beck, Rees et al., 2001; Pessoa & Ungerleider, 2004). Given this pattern of activation, it is thus possible that individual neurons in dorsal areas would be better correlated with subjects' behavioral reports than the MTL neurons we record from. Additionally, these dorsal networks are less likely to carry stimulus-related information since they are mostly silent during change-blindness trials, even though the visual stimulation is identical to change detection trials.

Taken together these results suggest that the mere activation of ventral stream and MTL neurons might not directly predict change detection (since correlation with the stimulus is higher than with the behavioral outcome). Rather, the ventral and MTL neurons could carry information related to the visual input (such as what we observe during both change detection and change-blindness trials), but the contents of awareness would arise through interactions with areas in the dorsal pathway, as suggested by Rees and colleagues (Rees, Kreiman, & Koch, 2002).

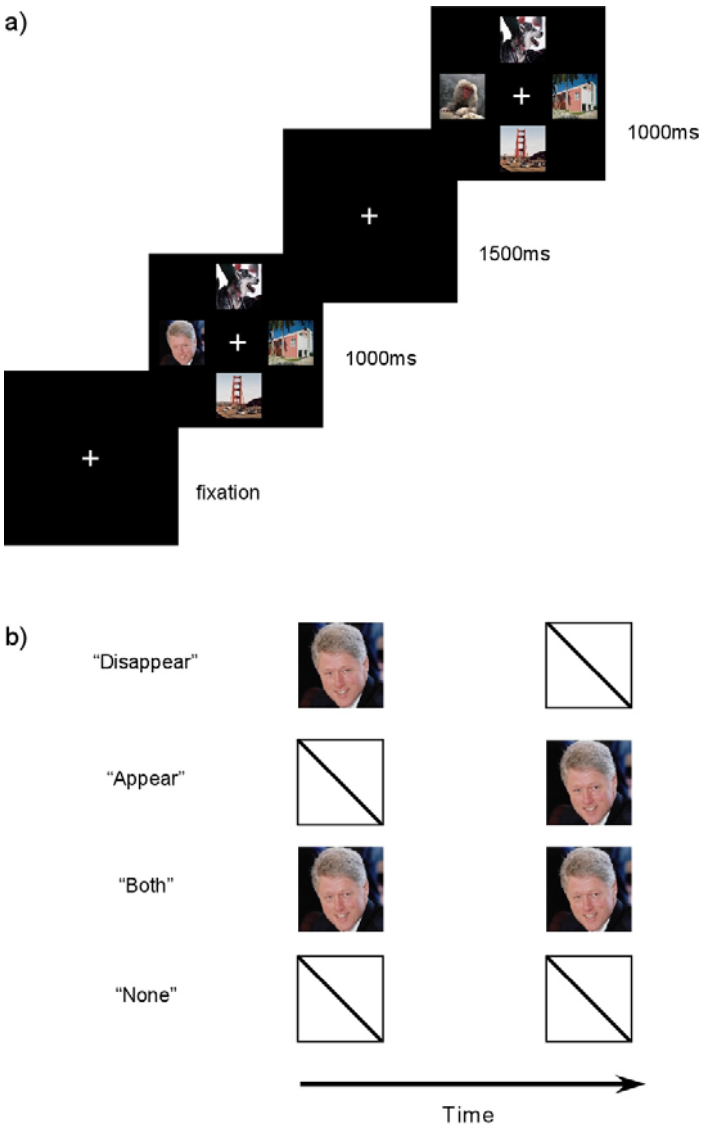

Figure 6.1 Experimental timeline for one trial in the change-blindness experiment. a) Each trial begins with a fixation cross, followed by four pictures that are presented for 1000ms. A blank interval for 1500ms follows and is replaced by another set of 4 pictures, again presented for 1000ms. On roughly half of the trials, one of the four pictures changes from the first display to the second (in this case the picture of Bill Clinton changes). At the end of the trial, subjects have to report whether or not they noticed a change. b) From the point of view of a cell selective to Bill Clinton, 4 different trial types could occur: In the change trials, the picture of Clinton could be present in the first display period and then be absent in the next ("disappear trials"). Or it could be absent in the first and appear in the second ("appear trials"). In the no-change trials, the picture could be present in both display periods ("both periods") or could be absent all together ("none trials"). By comparing activity between correct and incorrect trials, during the intervals when the picture was present, the effects of attention on the responses of this neuron can be examined.

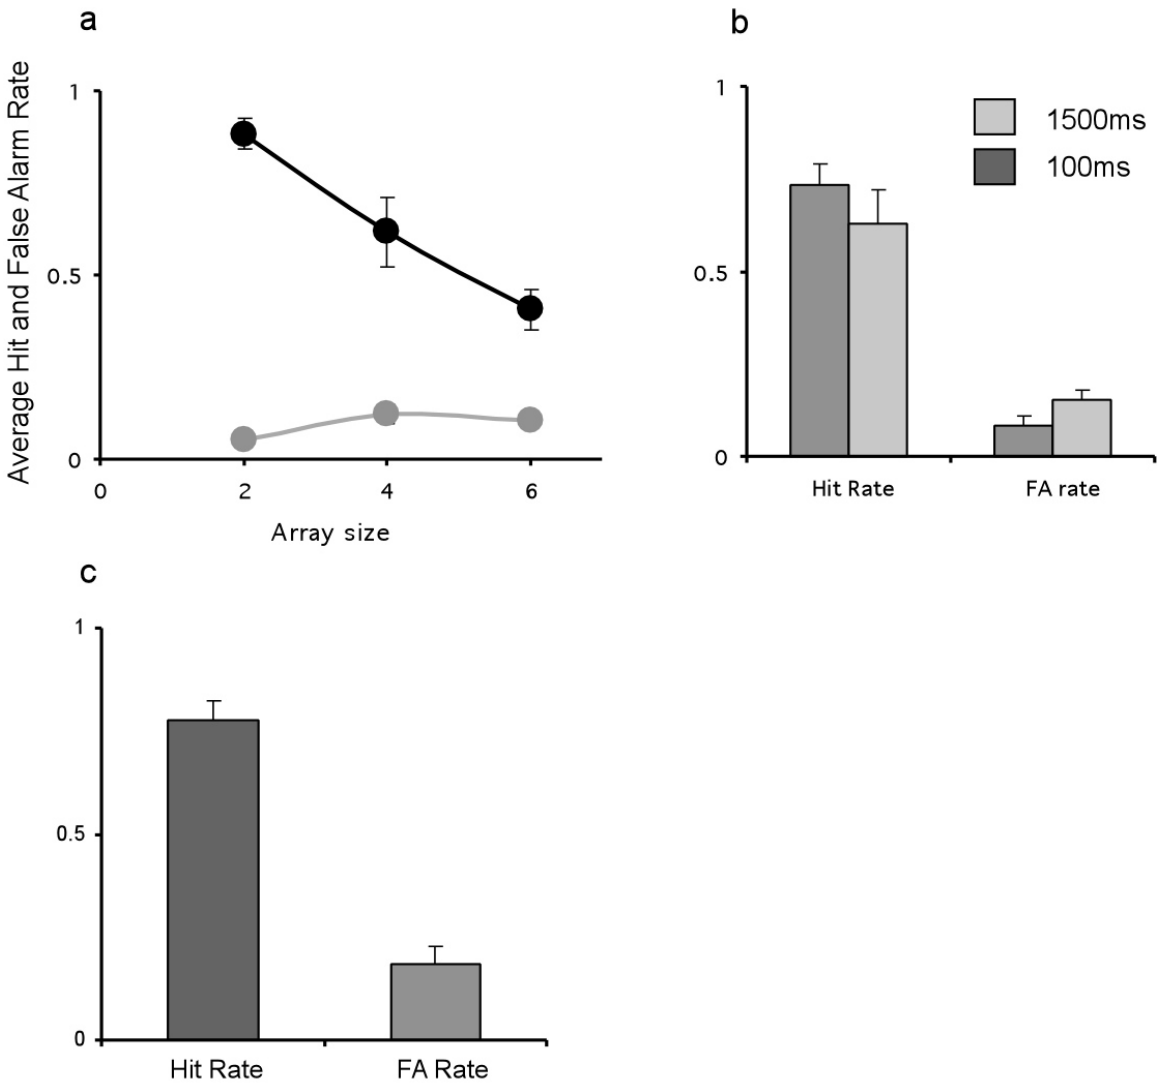

Figure 6.2 Behavioral responses of the clinical and non-clinical populations. a) Average hit rate (black) and false alarm rate (gray) for the non-clinical population (n=6 subjects) for different sizes of the stimulus display (2, 4, or 6 images in each display period; data for an ISI interval =1.5s is shown here). b) The hit rates and false alarm rates for the 6 subjects, when tested with two different ISI intervals (100ms and 1500ms; data with a 4-image display in each display period is shown here). c) Average hit rate and false alarm rate observed for the clinical population (n=9 subjects).

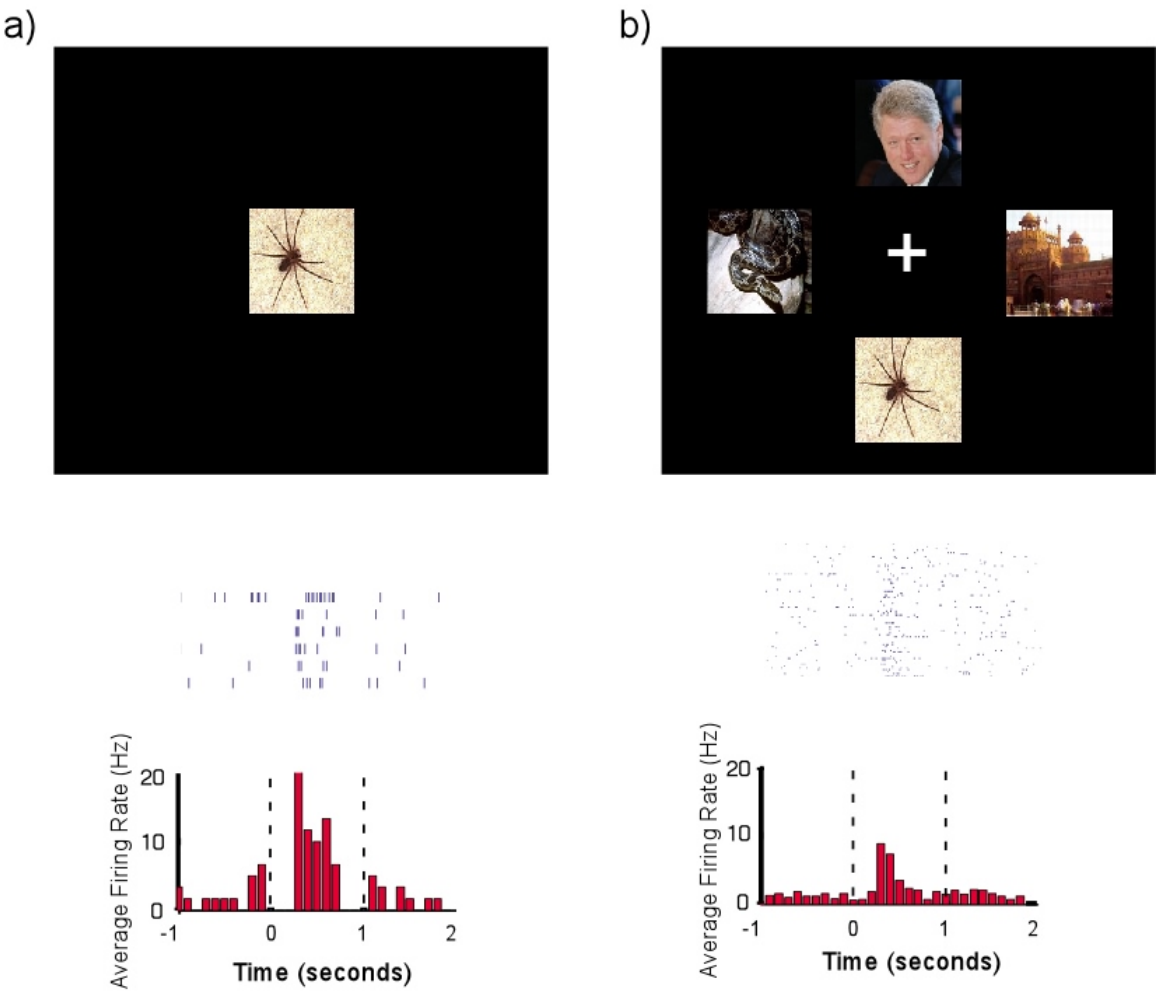

Figure 6.3 Comparing visual responses between screening and testing sessions. This unit is located in the right amygdala. a) In the screening session, this unit showed a significant increase in firing activity to an image of a spider that was presented foveally and in isolation ($p<.001$). The image appeared at t=0 and was presented for 1 second. b) The response of the same unit to the spider during the change-blindness set up is still significantly above baseline ($p<.0001$) but is about half as strong as the response in a). During the change-blindness experiment, the preferred image was presented peripherally along with 3 others known to drive the cell weakly. Visual responsiveness was determined by comparing firing activity during the pre-stimulus baseline interval with firing activity during the intervals when the stimulus was presented (see Section 6.3.4).

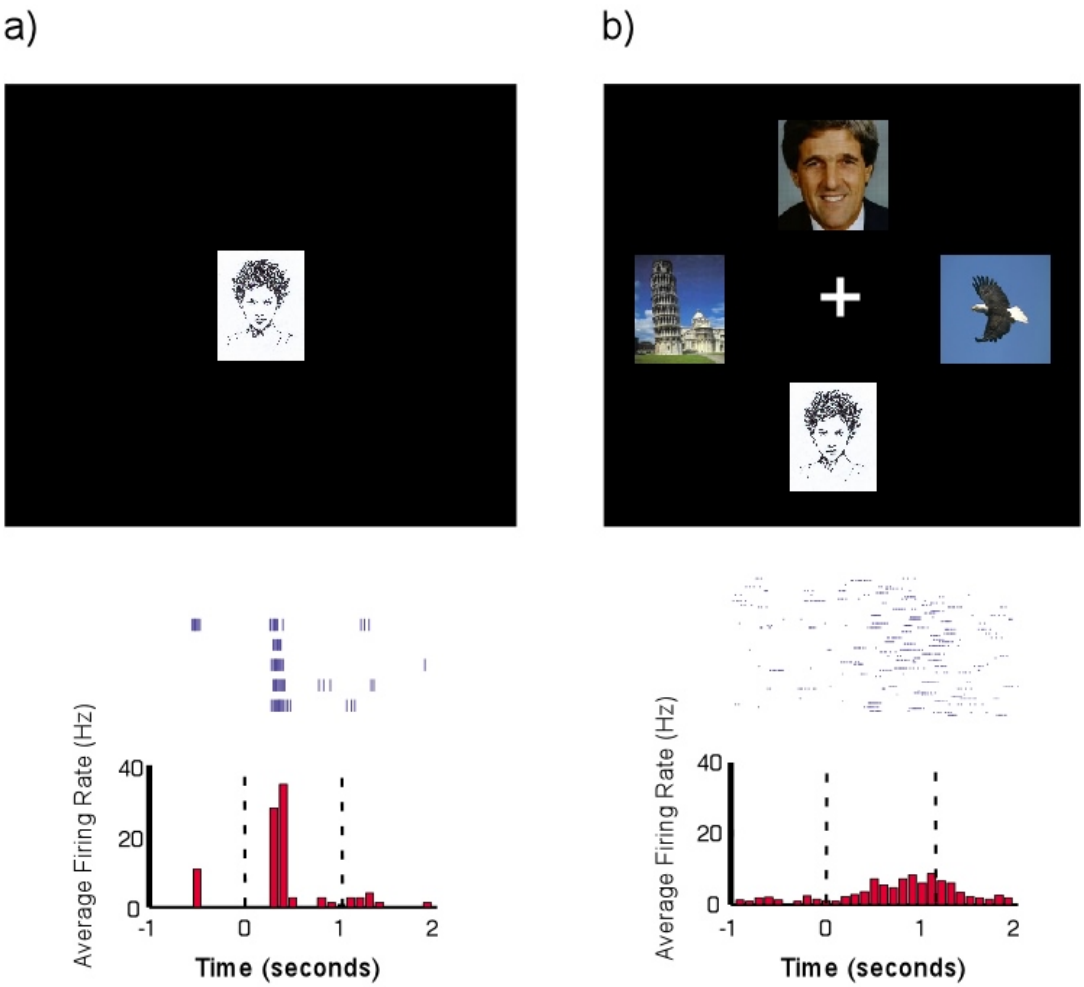

Figure 6.4 Comparing visual responses between screening and testing sessions. This unit is located in the right hippocampus. a) In the screening session, this unit showed a significant increase in firing activity to an image of actress Halle Berry that was presented foveally and in isolation ($p<.005$). The image appeared at t=0 and was presented for 1 second. b) The response of the same unit to Halle Berry during the change-blindness set up is still significantly above baseline ($p<.001$) but is much weaker than the response in a). Note that the actual image of the actress that was presented to the patients is different from the one shown here; we cannot show the original image because of image-copyright issues.

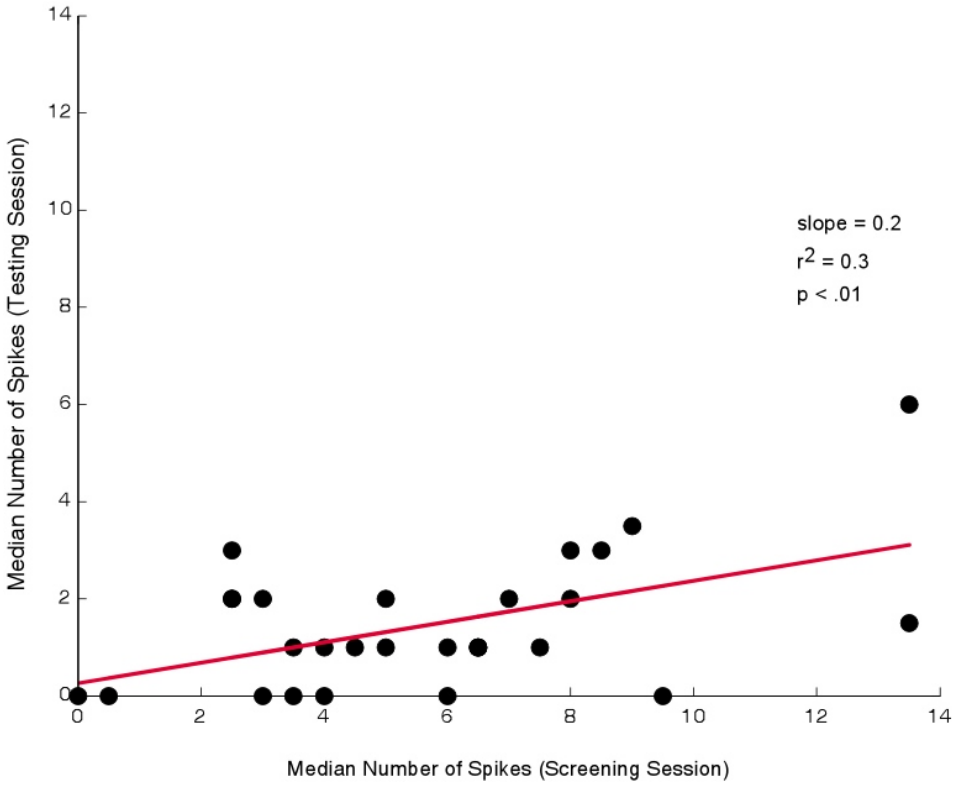

Figure 6.5 A comparison of the firing activity of all 29 neurons during the testing and screening sessions. The neuronal activity in the testing session is plotted against activity during the screening session. Over the population of cells, there is a significant decrease in activity in the testing session. This can be explained by the fact that in the testing session, each preferred stimulus was presented peripherally along with 3 other stimuli known to drive the cell weakly, while in the screening session, a single image was foveally presented (see Figures 6.3 and 6.4). On average the response to the preferred stimulus during the testing session was 20% of the response during the screening session.

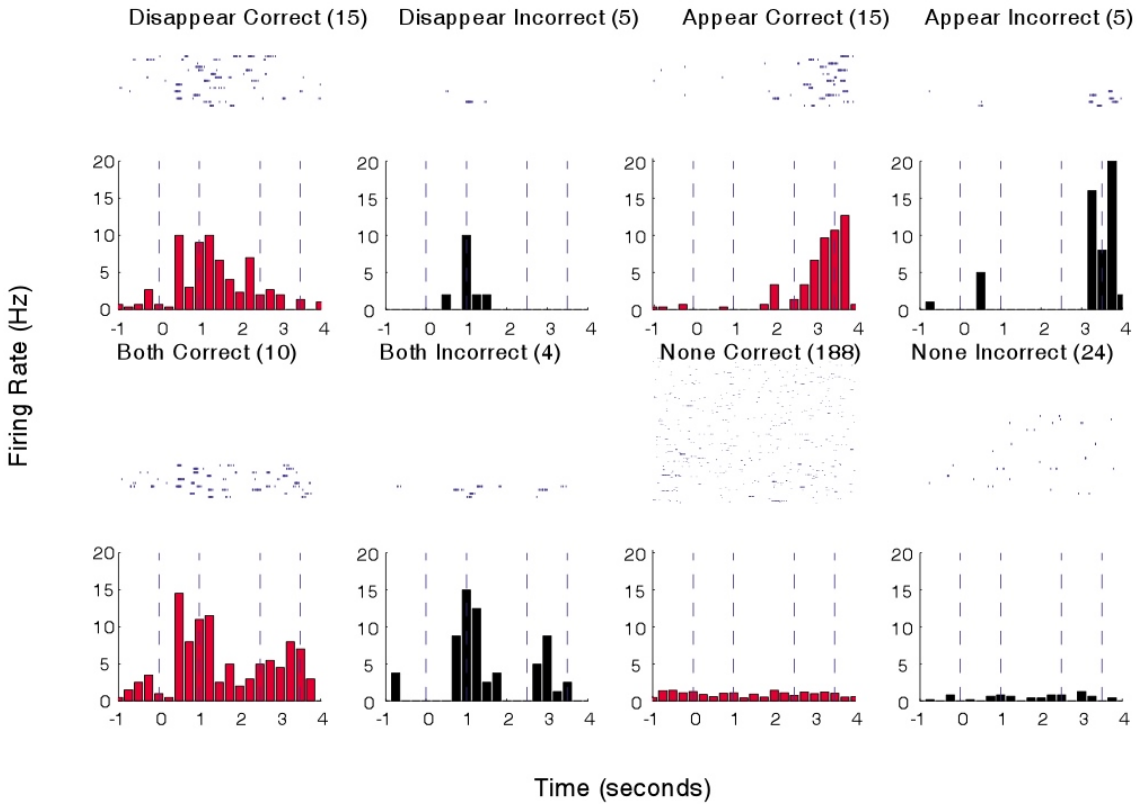

Figure 6.6 Responses during the 4 different trial types of the change-blindness experiment. a) Response of 1 unit located in the right amygdala. The red histograms correspond to correct trials in each trial type. The black histograms correspond to incorrect trials. Firing rate in Hz is plotted against time (ms). The first dashed vertical line represents the onset of the first set of 4 pictures, which stay on the screen for 1s (second vertical line). Following an ISI interval of 1.5s the second set of 4 pictures appears for 1s (last two vertical lines). A pre-stimulus baseline of 1s is also shown. The numbers in parenthesis represent the number of different trials in each condition. Although the effect is not significant at the single-cell level, the average neuronal response during the interval when the preferred stimulus is present tends to be higher on correct compared to incorrect trials. The average response time in reporting whether a change had occurred was 4.3 ±0.1s over the trials presented in this figure, aligned relative to the onset of the first set of 4 stimuli.

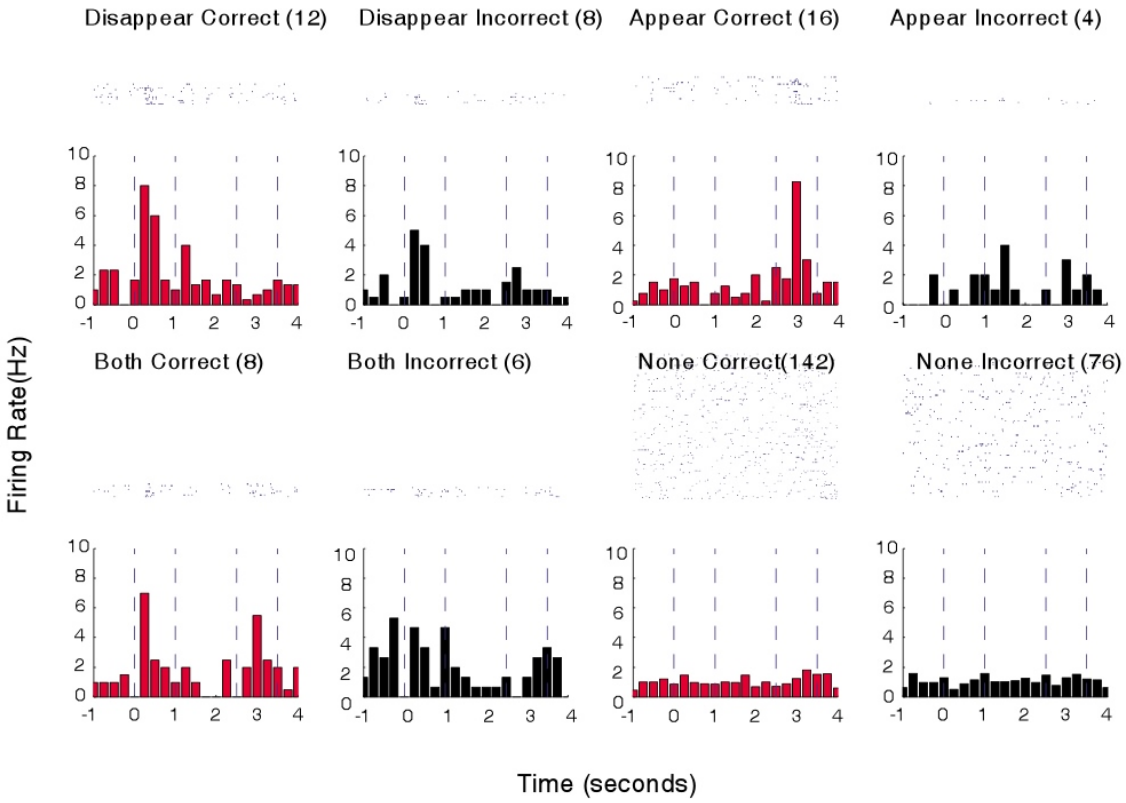

Figure 6.6b) Responses of another unit during the 4 different trial types of the change-blindness experiment. The data is presented in the same format as the previous figure. The average response time in reporting whether a change had occurred was 5.1 ± 0.1s over the trials presented in this figure, aligned relative to the onset of the first set of four stimuli. Again, activity on correct trials tends to be higher than on incorrect trials during those display periods when the preferred stimulus was present.

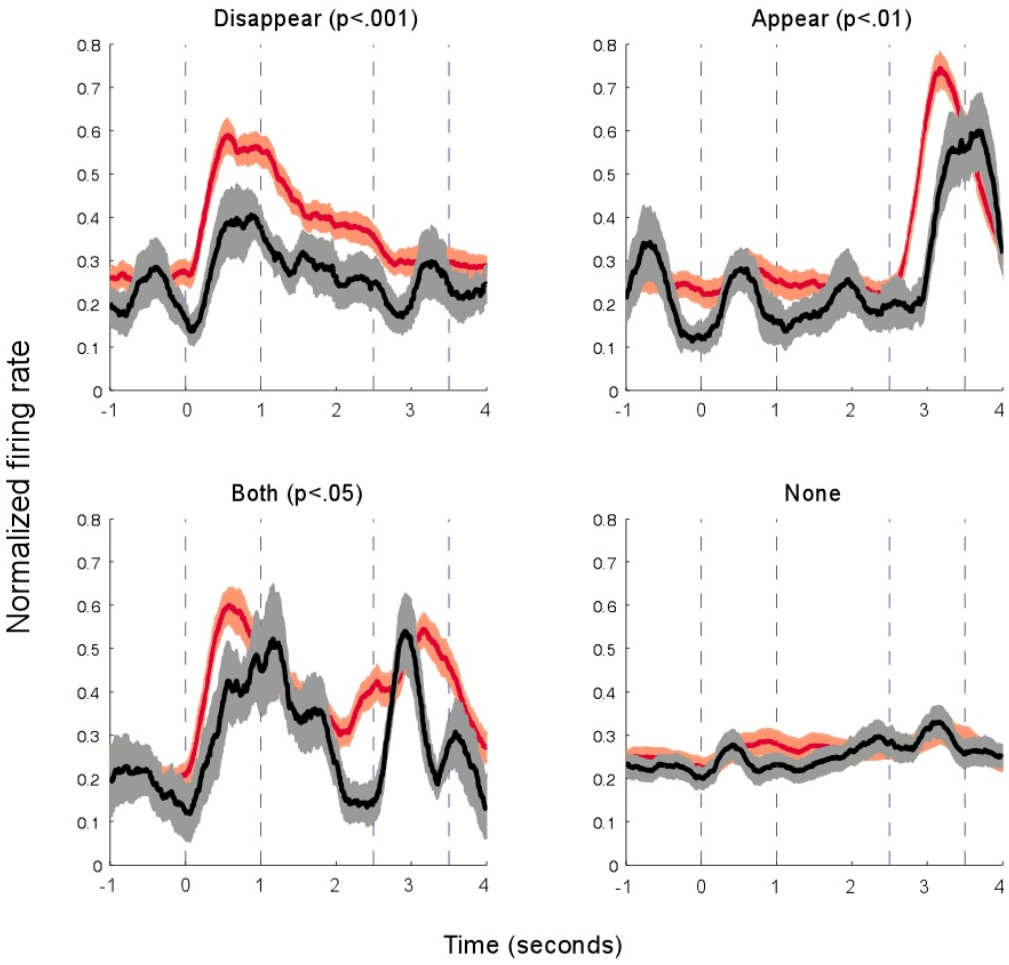

Figure 6.7 Responses during change detection and blindness for all cells. The mean normalized spike density function was calculated over the 29 cells in the 4 trial conditions for the correct (red curve) and incorrect (black curve) trials. Population responses were computed using the normalized spike density function (sdf). For each neuron, the sdf was obtained by convolving the spike train on each trial with a 200 ms fixed width Gaussian and then averaging over all trials. The spike trains were binned in 5 ms bins before convolution. For each unit, the sdf was normalized by dividing by its peak activity. The normalized sdf was then averaged over the population of cells. The average responses were computed separately for correct and incorrect trials in all trial types. The shaded areas represent the s.e.m. 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 ($p>.05$). The average response time in reporting whether a change had occurred was 4.9 ± 0.1s over all patients. The difference in reaction times between correct and incorrect trials is not significant ($p>.05$).

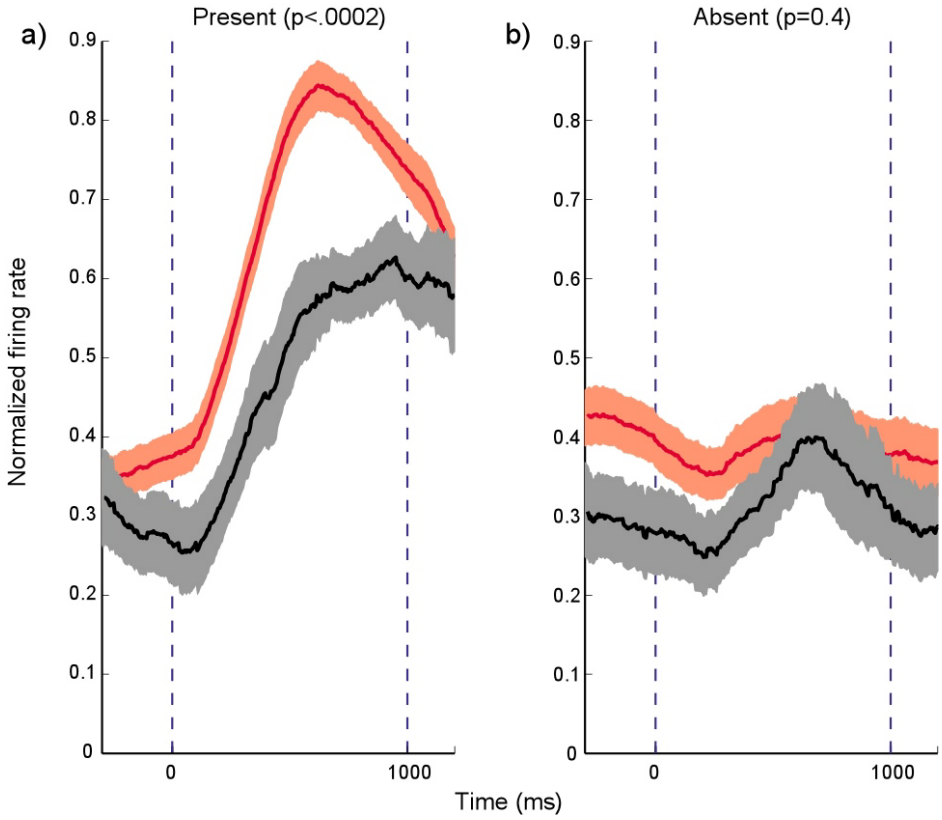

Figure 6.8 Population responses averaged over a) all stimulus intervals when the preferred stimuli were present and b) all stimulus intervals when the preferred stimuli were absent. In a), the data is averaged over the first display period of the “disappear” trials, the second period of the “appear” trials, and both periods of the “both” trials. In b), the data is averaged over the second interval of the disappear trials, the first interval of the appear trials, and both intervals of the none trials. The x-axis represents time in milliseconds. The stimulus appeared on the screen at 0ms and stayed on for 1000ms (the length of each display period). The y-axis represents normalized firing rates. The firing rates during the correct trials are significantly higher than firing rates during incorrect trials. In b) there is no significant difference in firing rates between correct and incorrect trials, which is expected since the preferred stimuli were absent during these intervals.

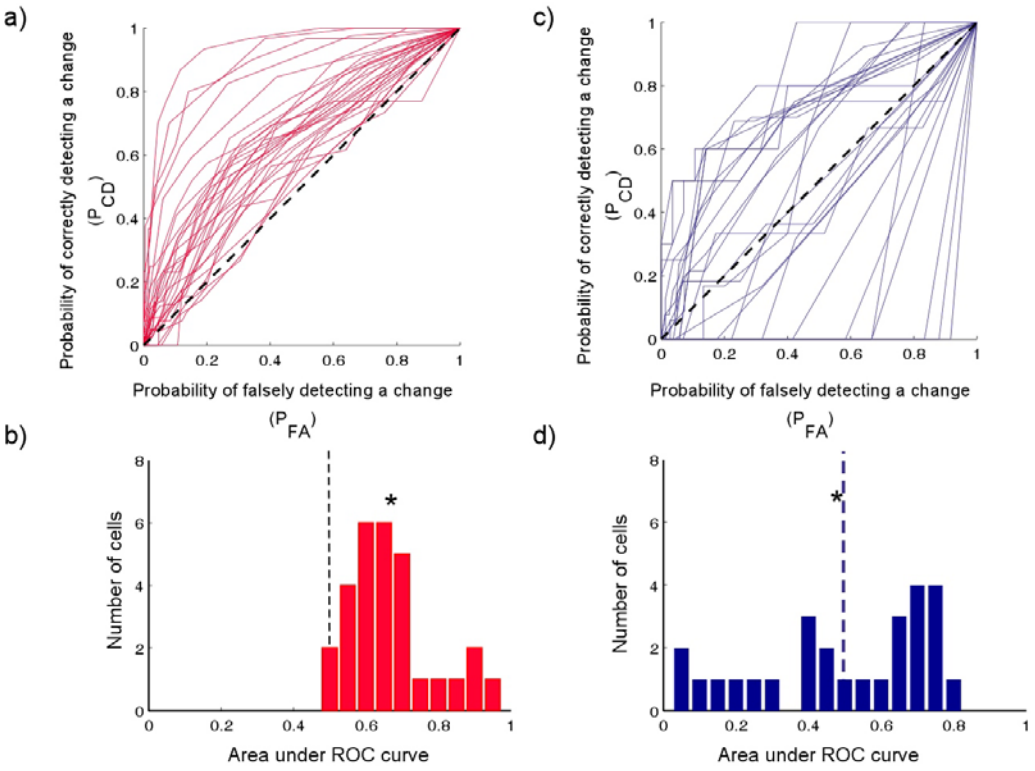

Figure 6.9 Predicting a change–ROC analysis over correct and incorrect trials. a) On each correct trial, the difference in the number of spikes occurring in the two display intervals was computed. If the absolute value of this difference exceeded a certain threshold T (for successive integer values of T), a change was predicted. This prediction was then compared to the actual visual stimulation—if a change had actually occurred, the signal would count as a "correct detection"; otherwise it would count as a "false alarm." 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). The dashed line indicates chance performance ($P_{CD} = P_{FA}$). The different lines show the result of this calculation for each cell. b) The distribution of the area under the curves for each cell over correct trials is significantly shifted to the right of 0.5 indicating that the population of cells can signal a change above chance on a trial-by-trial basis ($p<.0001$). The mean area under the ROC curves is marked by a * and equals 0.67 ± 0.02. c) ROC curves for all cells calculated over incorrect trials. d) The mean area under the ROC curves (*) is 0.5 ± 0.04. In contrast to data in (b), on incorrect trials, the population of cells were at chance at signaling a change ($p=.95$). Under the assumption that attention was more likely to be focused on the relevant stimuli on correct trials, the significant difference ($p<.005$) between ROC areas on correct versus incorrect trials reflects the advantageous effect of attention at predicting a change. The data presented demonstrates that only on those trials when patients detected changes made to the relevant stimuli (i.e. the correct trials) could cells signal a change significantly above chance on a trial-by-trial basis.

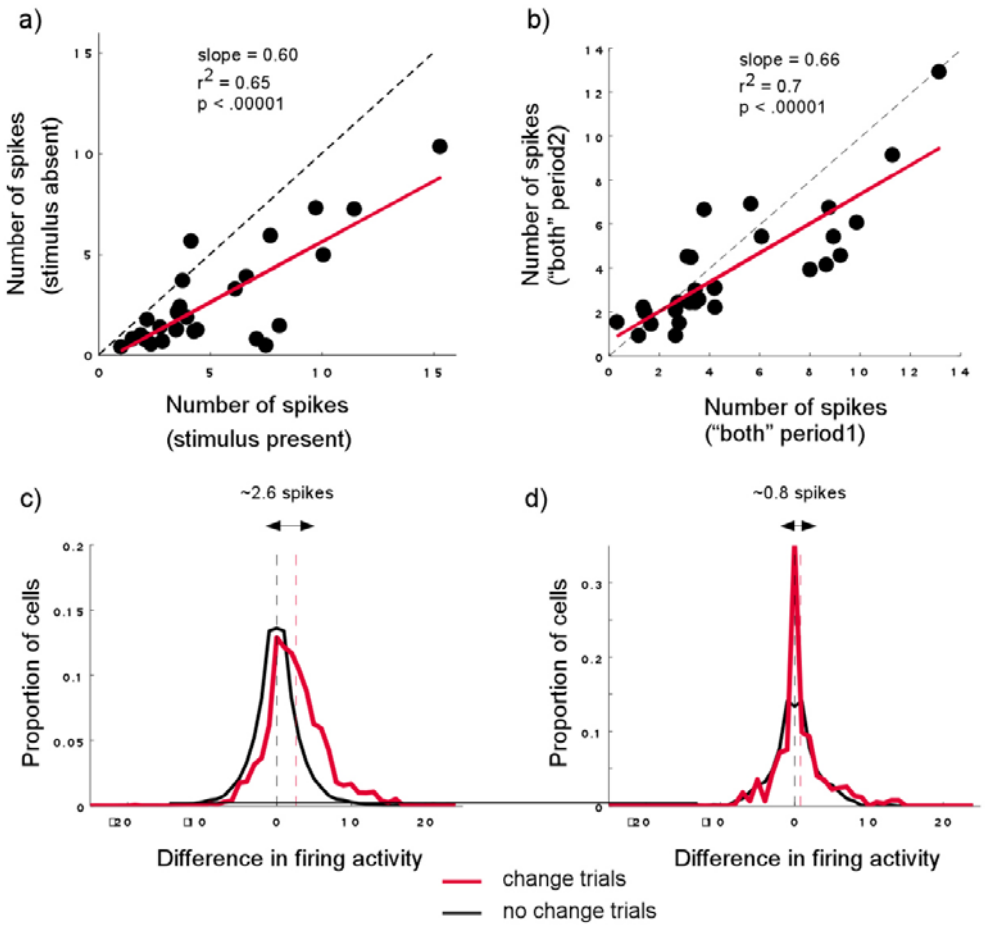

Figure 6.10 Firing activity during change and no-change trials. a) Activity for each cell during change trials. In this plot, the number of spikes fired in the absence of the preferred stimuli is plotted against activity in the presence of the preferred stimuli. The slope of 0.6 reflects that on average, firing activity in the absence of the preferred stimuli was six-tenths of activity when the stimuli were present. b) Activity in the two display periods of the no-change trials when the preferred stimuli were present (i.e. "both" trials). On average, the activity of each cell is lower during the second display period compared to the first, even though the visual stimulation was identical in both periods. c, d) The distribution of the difference in firing activity between the two display intervals for the no-change (black) and change (red) trials. These distributions are shown in c) for correct trials and in d) for incorrect trials. The distribution on correct change trials (i.e. activity in "present" —activity in "absent" intervals) is shifted to the right of distribution for no-change trials by about 2.6 spikes. Thus, on average, on correct change trials, the difference in activity between the two display periods is ~ 2.6 spikes. On incorrect change trials, however, the mean difference in activity between the two display intervals is only ~0.8 spikes indicating again that on incorrect trials, firing activity between change and no-change trials is not easily distinguishable.

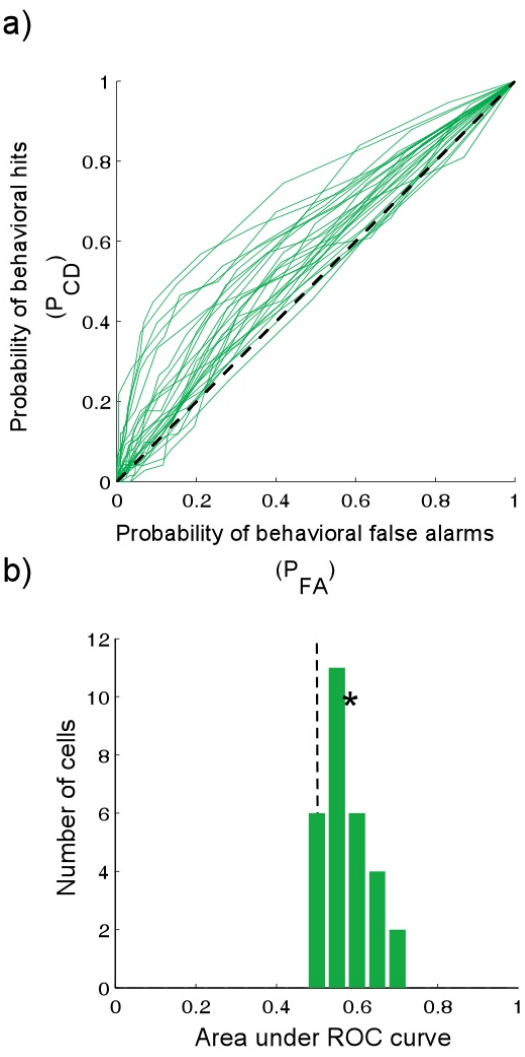

Figure 6.11 Predicting Behavior–ROC analysis. a) As before (Figure 6.8) on each trial, the difference in the number of spikes occurring in the two display intervals were computed, and the probability of a change occurring was predicted. This prediction was then compared with the behavior of the patient—if the neuron and patient reported that a change had occurred, the signal would count as a "correct detection"; otherwise if the neuron predicted a change but the patient reported no-change, the signal would count as a "false alarm." The probability of predicting behavior correctly, P_{CD} ("correct detection") is plotted against the probability of incorrect predictions ("false alarms," P_{FA}). The dashed line indicates chance performance ($P_{CD} = P_{FA}$). The different lines show the result of this calculation for each cell. b) The distribution of the area under the curve for each cell. The histogram is significantly shifted to the right of 0.5 indicating that the population of cells can predict behavior above chance on a trial-by-trial basis ($p<.001$). The mean area is marked by a * and equals 0.58 ± 0.01.

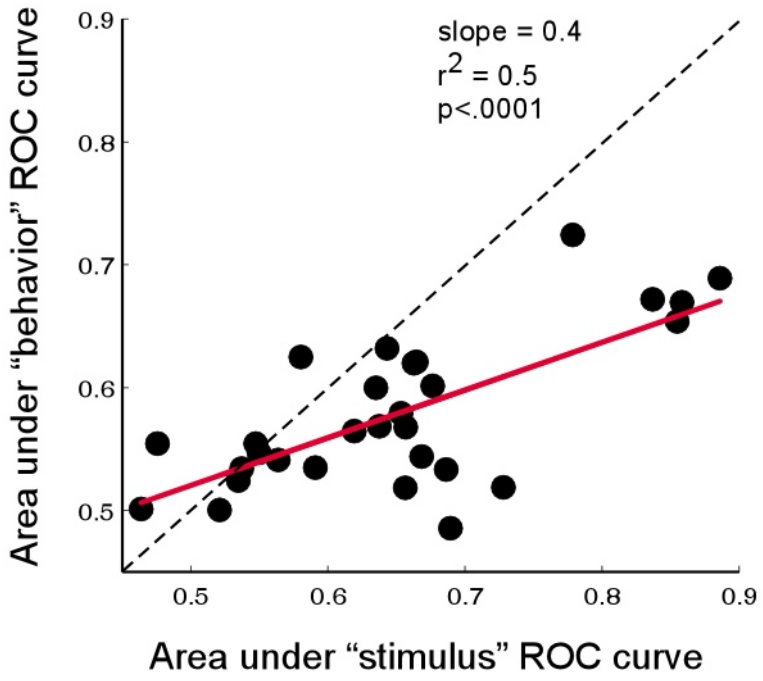

Figure 6.12 Comparison of the ability to predict change and the ability to predict behavior on a trial-by-trial basis, for each cell. On average, each cell is better able to predict changes in the stimuli rather than the behavioral choice of the subject on a trial-by-trial basis.

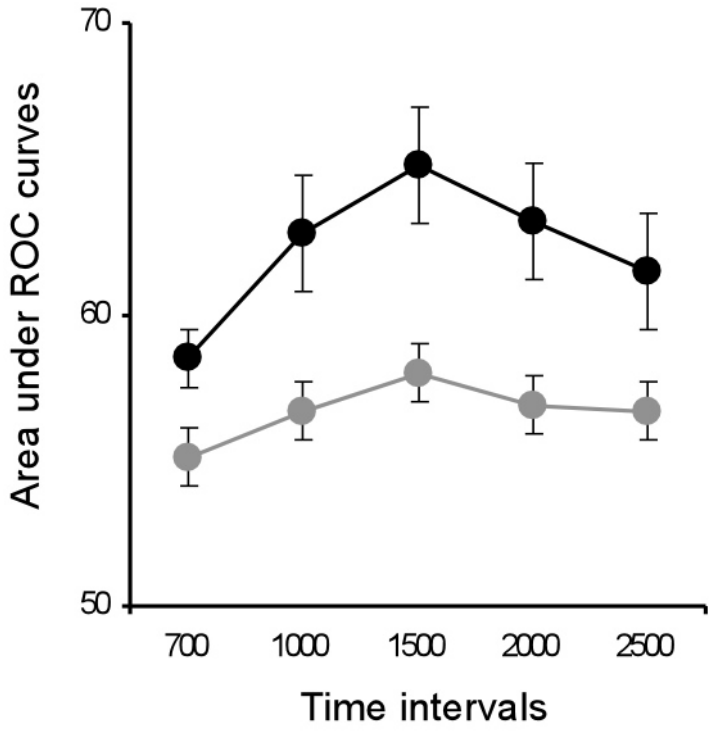

Figure 6.13 Comparison of the area under the stimulus and behavior ROC curves. ROC curves for predicting the stimulus (black line, corresponding to data in Figure 6.9a, b), and behavior (gray line, corresponding to data in Figure 6.11a) were computed over different time intervals. Each point on the graph corresponds to the ROC area when computed between [300 X] ms (aligned to the onset of the relevant display period) for 5 different values of X (shown on the x-axis). The ROC data reported in this chapter was computed over [300 1500] ms, which is optimal for both stimulus and behavior predictions.

Table 6. 1 Distribution of responsive cells during the change-blindness paradigm

Brain Area	Number of cells
RA	5
REC	2
LPG	4
LEC	1
RAH	11
RMH	1
RPG	1
LMH	2

The brain areas are as follows: RA: Right Amygdala, REC: Right Entorhinal Cortex, LPG: Left Parahippocampal Gyrus, LEC: Left Entorhinal Cortex, RAH: Right Anterior Hippocampus, RMH: Right Medial Hippocampus, RPG: Right Parahippocampal Gyrus, LMH: Left Medial Hippocampus.