

Learning Increases Stimulus Salience in Anterior Inferior Temporal Cortex of the Macaque

Bharathi Jagadeesh, Leonardo Chelazzi, Mortimer Mishkin and Robert Desimone
J Neurophysiol 86:290-303, 2001. ;

You might find this additional info useful...

This article cites 45 articles, 24 of which you can access for free at:
<http://jn.physiology.org/content/86/1/290.full#ref-list-1>

This article has been cited by 15 other HighWire-hosted articles:
<http://jn.physiology.org/content/86/1/290#cited-by>

Updated information and services including high resolution figures, can be found at:
<http://jn.physiology.org/content/86/1/290.full>

Additional material and information about *Journal of Neurophysiology* can be found at:
<http://www.the-aps.org/publications/jn>

This information is current as of August 12, 2013.

Learning Increases Stimulus Saliency in Anterior Inferior Temporal Cortex of the Macaque

BHARATHI JAGADEESH,^{1,2} LEONARDO CHELAZZI,³ MORTIMER MISHKIN,¹ AND ROBERT DESIMONE¹

¹Laboratory of Neuropsychology, National Institute of Mental Health, Bethesda, Maryland 20892; ²Department of Physiology and Biophysics, University of Washington, Seattle, Washington 98195; and ³Dipartimento di Scienze Neurologiche e della Visione, Sezione di Fisiologia, University of Verona, Verona 37134, Italy

Received 30 August 1999; accepted in final form 20 March 2001

Jagadeesh, Bharathi, Leonardo Chelazzi, Mortimer Mishkin, and Robert Desimone. Learning increases stimulus saliency in anterior inferior temporal cortex of the macaque. *J Neurophysiol* 86: 290–303, 2001. With experience, an object can become behaviorally relevant and thereby quickly attract our interest when presented in a visual scene. A likely site of these learning effects is anterior inferior temporal (aIT) cortex, where neurons are thought to participate in the filtering of irrelevant information out of complex visual displays. We trained monkeys to saccade consistently to one of two pictures in an array, in return for a reward. The array was constructed by pairing two stimuli, one of which elicited a good response from the cell when presented alone (“good” stimulus) and the other of which elicited a poor response (“poor” stimulus). The activity of aIT cells was recorded while monkeys learned to saccade to either the good or poor stimulus in the array. We found that neuronal responses to the array were greater (before the saccade occurred) when training reinforced a saccade to the good stimulus than when training reinforced a saccade to the poor stimulus. This difference was not present on incorrect trials, i.e., when saccades to the incorrect stimulus were made. Thus the difference in activity was correlated with performance. The response difference grew over the course of the recording session, in parallel with the improvement in performance. The response difference was not preceded by a difference in the baseline activity of the cells, unlike what was found in studies of cued visual search and working memory in aIT cortex. Furthermore, we found similar effects in a version of the task in which any of 10 possible pairs of stimuli, prelearned before the recording session, could appear on a given trial, thereby precluding a working memory strategy. The results suggest that increasing the behavioral significance of a stimulus through training alters the neural representation of that stimulus in aIT cortex. As a result, neurons responding to features of the relevant stimulus may suppress neurons responding to features of irrelevant stimuli.

INTRODUCTION

Certain types of objects are easy to pick out of visual displays. For example, a bright bar “pops out” of a display of dim bars. This phenomenon, normally referred to as stimulus saliency, can often be attributed to low level visual processing that precedes cognitive processing of the information. Following experience with a complex object, it is possible that the object may acquire saliency, even though it is not a “strong” stimulus based on brightness or other elementary image fea-

tures (Ahissar and Hochstein 1995; Ellison and Walsh 1998; Karni and Sagi 1991; Morris et al. 1997; Sireteanu and Rettenbach 1995). With sufficient acquired saliency, an object may pop out of a cluttered scene and be processed in preference to other objects, in much the same way that a brighter bar is processed before dimmer bars in a mixed display of bright and dim bars. The image of one’s spouse, for example, might pop out from a crowd containing many different people.

Physiological studies of the effects of behavioral relevance on the responses of visual neurons have generally been focused on top-down attentional effects rather than on learning effects. Selective attention modulates the visual responses of cells in several visual areas in the ventral processing stream (Chelazzi et al. 1993, 1998; Connor et al. 1996, 1997; Fuster 1990; Moran and Desimone 1985; Luck et al. 1997; Richmond and Sato 1987; Spitzer et al. 1988) including anterior inferior temporal (aIT) cortex, an area known to play an important role in visual recognition memory (Gaffan 1994; Meunier et al. 1993; Murray and Mishkin 1986; Suzuki et al. 1993; Zola-Morgan et al. 1989, 1993). Previous studies of visual search in aIT cortex have shown that, when monkeys are cued to search for a particular stimulus in a display of two or more items, neuronal responses, measured before the behavioral choice is made, are dominated by the properties of the relevant item and not by those of irrelevant distractors (Chelazzi et al. 1993, 1998). In these studies, at the start of the trial, animals were briefly presented with a cue stimulus at fixation, followed by a delay period and then an array of two or more stimuli at extrafoveal locations. The animal was rewarded for making a saccadic eye movement to the target stimulus in the array that matched the previous cue. The responses of aIT cells to the array were determined almost exclusively by the target stimulus; if the target was a good stimulus for the cell, the cell responded well, and if it was a poor stimulus for the cell, the cell responded poorly. The influence of distractor stimuli was filtered out, even though they were still present in the receptive field. In addition to this gating effect on the response to the distractors, aIT cells also often showed differential activity during the delay period, depending on which stimulus was shown as the cue. If the cue was a good stimulus for the cell, the cell tended to have higher activity during the delay than if

Address for reprint requests: B. Jagadeesh, Box 357330, Dept. of Physiology and Biophysics, University of Washington, Seattle, WA 98195 (E-mail: bjag@u.washington.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

the cue was a poor stimulus for the cell. Elevated delay activity may be evidence for a dynamic top-down “bias” in favor of cells coding the cue/target stimulus on a given trial (Chelazzi et al. 1993, 1998), originating from neural mechanisms for working memory outside of the temporal cortex.

In previous studies of visual search in aIT cortex, the stimuli used as target and distractor switched randomly from trial to trial (or between short blocks of trials). A target stimulus on one trial would be an irrelevant distractor on another trial. Therefore there was no opportunity for a particular stimulus to acquire intrinsic salience over the course of trials, thereby necessitating the use of working memory to solve the task. We hypothesized that if monkeys were trained on a similar task but with the same stimulus used consistently as a target on many trials, this would induce learning effects in aIT cortex that would not be dependent on working memory, leading to stable changes in target salience.

METHODS

Surgical procedures

Three adult male rhesus monkeys weighing 5.9–9.9 kg were used. The surgical procedures have been described previously (Miller et al. 1993). All procedures and animal care were conducted according to National Institutes of Health (NIH) guidelines and were approved by the animal care and use committee. The monkey was placed in an aluminum and plastic stereotaxic apparatus and scanned with magnetic resonance imaging (MRI). The MRI images were used to locate the recording chamber in stereotaxic coordinates over the aIT cortex. Under aseptic surgical conditions, a head post, recording chamber, and a scleral eye coil for monitoring eye position (Judge et al. 1980) were implanted while the monkeys were under isoflurane anesthesia. The centers of the chambers were located at 18–20 mm anterior to the interaural line and 15–17 mm lateral to the midline. The animals received antibiotics and analgesics after the surgery.

Neurophysiological recording

Single unit activity was recorded using lacquer-coated tungsten microelectrodes (Roboz Microprobes), which were advanced vertically through the brain to the cortex on the surface of the inferior temporal gyrus. Individual spike waveforms were discriminated using an on-line spike-sorting system (Signal Processing Systems, Prospect, Australia). Neurons were isolated on-line with the spike sorter, and then tested with a set of visual stimuli, presented at the fovea, while the monkey maintained fixation on a point on the visual display. When a cell appeared to be visually responsive and selective, i.e., responded above its spontaneous firing rate, and gave differential responses to the various stimuli in the set, the rest of the experiment was performed. If the cell appeared to be unresponsive or nonselective, the electrode was advanced and another cell isolated, until a cell that satisfied the criteria was found. Electrodes were placed stereotaxically at sites in the aIT on the basis of the MRI images, which were also in stereotaxic coordinates. The location of electrode penetrations in stereotaxic coordinates and the MRIs were used to estimate the location of the recording sites. The recording sites span a 7-mm-diam area centered approximately 1 mm medial to the anterior tip of the anterior middle temporal sulcus (AMTS), including parts of area TE and the perirhinal cortex.

Stimuli

The stimuli were colored pictures of complex objects (e.g., faces, random objects, and textures), typically 1.5–2.5° in size, presented on a computer monitor. In the discrimination tasks, two stimuli were

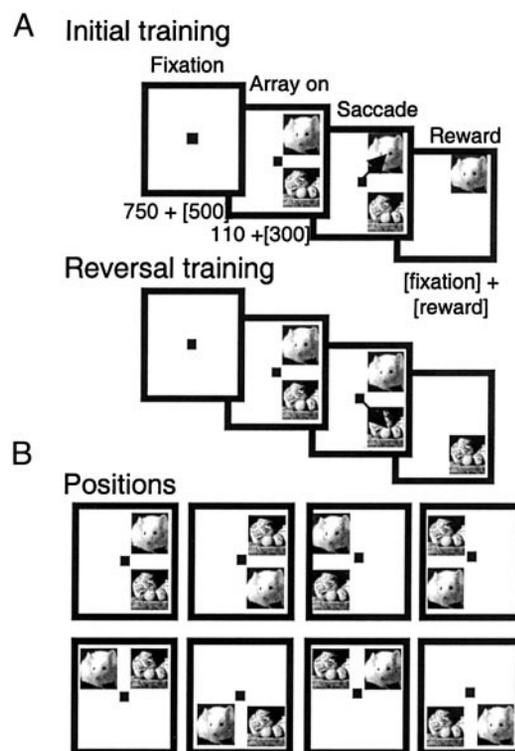


FIG. 1. *A*: discrimination learning and reversal (*task 1*). The task learned during neurophysiological recording. Monkeys maintained gaze on a fixation target (dark square); after a random delay, an array of 2 stimuli was presented at an unpredictable position in the peripheral visual field. The task was to learn, through trial and error, to saccade to the positive (target) stimulus in the display. Which stimulus would be positive was randomly assigned by the experimenter. After training with 1 reward contingency, the reward contingency was reversed, and the monkey re-learned the task. Numbers under panels refer to approximate times in milliseconds, with numbers in brackets referring to variable time intervals, ranging as high as the numbers in the brackets. The duration of required fixation was different for the 3 monkeys (100 ms, *monkey A*; 200 ms, *monkeys B* and *C*). Stimuli shown are representative of those used in the experiment. *B*: stimulus configurations. Arrays appeared in all the configurations shown above, in random order. Novel, previously unused stimuli were presented in each recording session. The task was always to saccade to the positive stimulus, wherever it appeared.

presented 3.5–4.5° from the fovea, each centered in any one of the four quadrants of the visual field. Thus the two stimuli could appear together in the contralateral or ipsilateral visual field, or across the vertical meridian from each other, in either the upper or lower visual fields (Fig. 1*B*).

Behavioral tasks

TASK 1: DISCRIMINATION LEARNING AND REVERSAL. Monkeys were trained in a discrimination task in which they learned to saccade to one picture in a display consisting of two pictures. A fixation spot was displayed on the screen, which the monkey was required to fixate. After a variable delay, an array of two images was displayed on the monitor at extrafoveal locations (see Fig. 1*B*). The monkey was rewarded for making an eye movement to the positive, or target, stimulus, and maintaining fixation on it for 100 ms (*monkey A*), 200 ms (*monkeys B* and *C*; Fig. 1*A*). At that time, the target stimulus remained on the screen for another 400 ms, while the reward was delivered. The other stimulus in the display was turned off during this period. If the monkey made an eye movement to the negative stimulus, i.e., an incorrect response, no reward was delivered, and a time-out period was imposed. The negative stimulus remained on during this time-out period for 800–1,600 ms, and the positive stimulus was

turned off. Responses that occurred at latencies shorter than 110 ms were scored as errors, regardless of the target fixated on that trial. The locations of the target and distractor varied randomly from trial to trial, and therefore the monkey had to use its extrafoveal vision to find the target based on its features (Fig. 1B). The stimulus in the pair that was initially to be used as the target was chosen by the experimenter randomly at the start of the session, and the monkey learned which of the two stimuli was the target through trial and error. Because the monkeys developed a learning set, they learned new pairs of stimuli quickly (see RESULTS).

New pairs of stimuli were introduced at the start of each recording session. After the monkey learned to discriminate the target from the distractor stimulus in the pair, and after sufficient neurophysiological data were collected in this condition (typically 80–160 total trials), the positive and negative stimuli were switched. That is, the previously relevant target stimulus became the irrelevant distractor, and the previous distractor became the target. The monkey learned the reversal task through trial and error, typically in just a small number of trials (see RESULTS). The reversal task was run until the monkey performed the task with high accuracy and until sufficient neurophysiological data were collected.

A passive fixation task was used to optimize the stimuli for each cell, outside the context of the discrimination task. A fixation spot appeared at the start of the trial, which the monkey was required to fixate. After a delay of 500 ms, a stimulus appeared, centered over fixation, for 300 ms. The monkey was rewarded for maintaining fixation throughout the stimulus period. While the monkey performed a passive fixation task, neuronal responses were recorded to 24 different randomly selected stimuli. Based on these responses, we selected both a “good” stimulus, i.e., the one that elicited the best or nearly the best response from the set, and a “poor” stimulus that elicited little or no response. The good and poor stimuli were then paired in the discrimination task. Because the cell only responded well to the good stimulus, one could compare the response to the good stimulus on trials when it was the target to the response to the same stimulus on trials when it was the distractor. The selection of the good or poor stimulus as target in the original learning at the start of each session was made randomly.

In some experiments, catch trials were added in an effort to decrease the reward value of guessing. Guessing consisted of responses at very short latencies and was accompanied by a decrease in the percent of correct responses. In catch trials, an array consisting of two identical stimuli (held constant over multiple recording sessions and different from both the good and poor stimulus) was displayed, and this was followed, after a delay, by the stimulus pair that the monkey was to discriminate. The stimulus pair was displayed at the same location as the catch trial stimuli. During the presentation of the array consisting of identical stimuli, the monkey was to withhold a saccade and remain focused on the fixation point. These catch trials were not included in data analysis.

TASK 2: CONCURRENT DISCRIMINATION PERFORMANCE. The monkey’s task in this experiment was identical to that in *task 1*, but a pretrained set of 10 stimulus pairs was used. One member of each pair was a target stimulus, and one member was a distractor. The set was learned by the monkey over a period of several weeks before recordings were initiated. During the recording sessions, any of the 10 stimulus pairs could appear on any given trial, so the stimulus pair to appear on any given trial was not predictable. In addition, stimuli were presented at the fovea (as in the fixation task) intermixed with the discrimination task (Fig. 2). The 10 stimulus pairs used are shown in Fig. 2.

Data analysis

Three different time intervals were used for the statistical analysis of both single-cell and population data: 1) baseline time interval, from



FIG. 2. Concurrent discrimination performance (*task 2*). Task in which any of 10 different pairs could appear on a given trial. Monkey was trained with this particular stimulus set before the recording sessions began. The 1st stimulus in each pair is the positive (target) stimulus in each pair. The stimuli, their pairing, and the reward contingencies remained consistent, throughout the training and recording periods. The timing and stimulus configurations were the same as in Fig. 1, *task 1*.

150 ms before stimulus onset to 25 ms after stimulus onset; 2) stimulus time interval, from 75 ms after stimulus onset to 250 ms after stimulus onset; and 3) saccade time interval, from 75 ms before saccade onset to 50 ms after saccade onset (since the latency of these cells is >50 ms, the response at 50 ms after the saccade should not reflect changes in neuronal response caused by the saccade). Data were analyzed using both ANOVA and *t*-tests, and evaluated at the $P < 0.05$ level of significance.

SINGLE-CELL SPIKE DENSITY HISTOGRAMS. Responses of cells were averaged across multiple trials, and the average response was convolved with a gaussian kernel with a SD of 20–40 ms. These histograms were constructed by time locking the events either to the onset of the visual stimulus or to the onset of the saccade. On average, 10–20 trials per stimulus were used for analysis of the stimulus selectivity of the cell, 20–80 trials per stimulus for analysis of the response during performance of the task, and 10–40 trials per stimulus for analysis of the response during performance of the task in the multiple pair version of the task.

POPULATION SPIKE DENSITY HISTOGRAMS. The population histograms shown in the figures were averaged from the data from individual cells. We chose not to normalize the responses before averaging, to illustrate the actual average firing rates in the population. Many of the same histograms were recalculated after normalizing the data from each cell to its maximum firing rate, and there was almost no difference in the shape or time course of the histograms before and after normalization.

RESULTS

Behavioral performance in task 1

In *task 1*, monkeys learned to saccade first to one stimulus in an array of two stimuli, and then, after learning the task, they were given reversal training with the other stimulus as the target. Behavioral results are reported only for those trials in which the monkey made a saccade to one of the two stimuli on the screen, excluding any catch trials (see METHODS). Trials were excluded from analyses if the monkey broke fixation before the onset of the array, if the saccade latency was <110 ms after the stimulus array onset, or if a saccade was made to

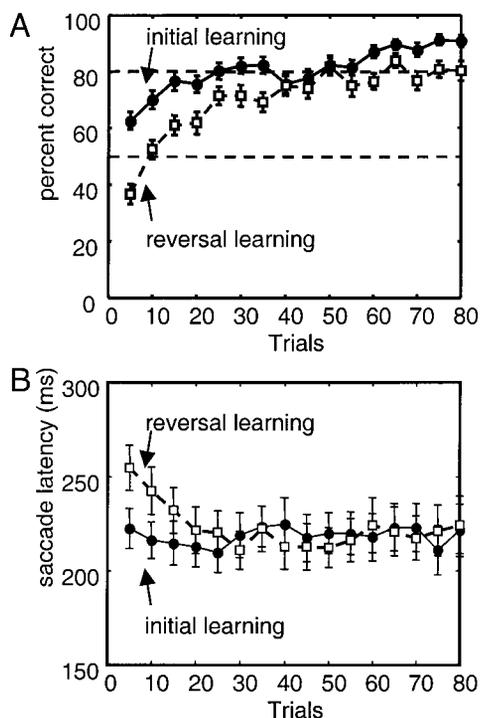


FIG. 3. Time course of learning. *A*: performance over trials during the training, averaged across the stimuli used during the neurophysiological recording. Each point connected to another by a line represents the average percent correct (\pm SE) over 5 trials of stimulus presentation, in any of the possible stimulus configurations. Circles correspond to the original training phase; squares correspond to the reversal training phase. Squares and circles unconnected by lines are the average percent correct in 2 trials for the 1st 4 trials. Chance was 50% correct, designated by the bottom dashed line. The top dashed line represents a behavioral performance of 80% correct (in 5 trials). *B*: saccade latency as a function of trials of task. Each point represents average saccade latency (\pm SE) over 5 correct trials, in any stimulus configuration. The error bars are ± 1 SE. Circles correspond to the original training, and squares to reversal training.

a location other than that of the two individual images that made up the array.

Monkeys learned to saccade to the positive stimulus to a criterion of 80% correct in a median of 15 trials (mean of 25 trials), across the population.

After the monkeys were performing consistently at a level near 80% correct, the reward contingencies were reversed, and the monkeys were trained to saccade to the other picture in the array. Monkeys learned to reverse their behavior to a criterion of 80% correct saccades in a median of 28 trials (mean of 45 trials).

On average, performance rose above chance within 5 trials on the original training with each pair and within 15 trials with the reversal training (Fig. 3*A*). Across the entire period of data collection, the average percent correct was 85% during the original training, and 76% during the reversal training. Across the pairs of stimuli, average performance increased between the first 10 and last 10 trials for both the original and reversal training ($66 \pm 2.7\%$, mean \pm SE, to $92 \pm 1.3\%$ for original training, paired *t*-test, $P < 0.01$; $45 \pm 2.8\%$ to $85 \pm 1.8\%$ for the reversal training, paired *t*-test, $P < 0.01$). The mean saccade latency in the original training and reversal training was 214 ± 11.2 ms and 223 ± 11.4 ms, respectively. During reversal, saccade latency decreased with a time course similar to the increase in performance accuracy (Fig. 3*B*). The saccade

latency remained roughly constant over the course of the training session during the original learning.

Single-cell recordings in task 1

Each cell's response to the array when the good stimulus was the target was compared with its response to the same physical array when the poor stimulus was the target. For one-half of the cells the good stimulus was the target during the original learning and the distractor during the reversal, and for the other half it was the poor stimulus that was the target during original learning and the distractor during the reversal. For each cell, the target and distractor stimuli were switched during the reversal phase of the task. Because the neuronal results differed according to whether the two stimuli were located within the same hemifield, or across the vertical meridian from one another in opposite hemifields, we analyzed these two categories of stimulus configurations separately. The results when both stimuli were contained within the contralateral or ipsilateral hemifield were similar to one another, on the other hand, so these data were combined for the within-hemifield condition.

Responses to the array were larger when the good stimulus was the target than when the poor stimulus was the target, and this difference appeared before the saccadic eye movement was made to the target stimulus. An example of an experiment on a typical cell is illustrated in Fig. 4, *A–D*. The cell was first tested with an array of 24 stimuli while the animal performed

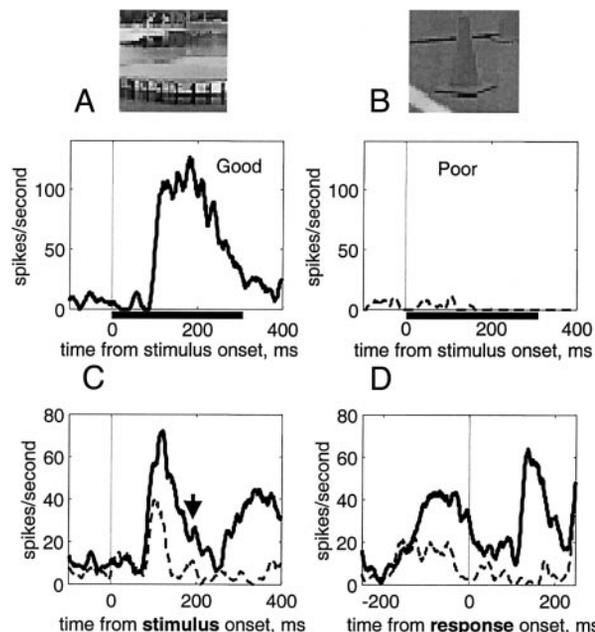


FIG. 4. *A* and *B*: response to good and poor stimulus presented at the fovea for a single cell. Each graph shows the average spike density histogram for 17 trials of stimulus presentation, time locked to the onset of the stimulus. The bar in the graph shows the stimulus duration of 300 ms. *C* and *D*: responses of a single cell to the stimulus pair shown in Fig. 4 in correct trials when the good stimulus was the target stimulus (solid line), or the poor stimulus was the target stimulus (dashed line). Stimuli were contained within a single hemifield (either ipsilateral or contralateral). Each graph shows the average spike density histogram for 40–50 trials of stimulus representation. The monkey was trained to saccade 1st to the poor stimulus, and then to the good stimulus. *C*: time locked to stimulus onset. Arrow at 200 ms represents the average saccade latency. *D*: time locked to saccade onset.

the passive fixation task, and a good (85 spikes/s) and poor (1.8 spikes/s) stimulus chosen from the set (Fig. 4, A and B). Based on these responses, the stimuli were paired together in the discrimination phase of the experiment. The monkey was then trained with the poor stimulus as the target. After sufficient physiological data were collected in this phase, the monkey was then trained with the good stimulus as the target, and comparison physiological data were collected in this condition (Fig. 4, C and D).

The cell started responding to the stimulus in both conditions (Fig. 4C), but the response was maintained at a high rate through the saccade period only when the good stimulus was the target. When the poor stimulus was the target, the cell's response to the good stimulus in the array was suppressed before the saccade, even though the sensory conditions remained the same. Thus by the time of onset of the saccade, the response of the cell appears to be dominated by the target stimulus.

In Fig. 4D, the same data are shown time locked to the onset of the saccade. The response to the array diverged according to which stimulus was the target, beginning about 100 ms before the onset of the saccade. The saccade brought the image of the target stimulus onto the fovea and moved the image of the nontarget stimulus into the periphery. After the saccade, the nontarget stimulus in the array was turned off, and therefore the activity in this time period was determined by the presence of the good or poor stimulus at fixation.

Population response in task 1

The same results are evident in the population of recorded cells ($n = 55$). Across the population (Fig. 5, A and B), the average response to the good stimulus presented alone at the fovea was 22.6 spikes/s, and the average response to the poor stimulus alone was 10.3 spikes/s (paired t -test, $P < 0.01$).

Figure 5C shows the responses of the population average to the paired stimuli on correctly performed discrimination trials, when the good versus poor stimulus was the target for a saccade. In the period just before the saccade, the response was greater on trials when the good stimulus was the target than when the good stimulus was the distractor. Across the population, the firing rates diverged 100 ms after stimulus onset (paired t -test on successive 20-ms bins, $P < 0.01$ for 3 consecutive bins). The response in the postsaccade time period was also much greater when the good stimulus was the target, but this simply reflects the monkey fixating the good versus the poor stimulus.

If the information provided by these neuronal responses had been used to guide the saccade, the response difference must have been present before the onset of the saccade. To examine this, we averaged the responses of the same population of cells, but time locked to the onset of the saccade (Fig. 5D). The firing rate with the good versus poor stimulus as the target diverged significantly at 80 ms before the saccade onset (paired t -test on successive 20-ms bins, $P < 0.01$ for 3 consecutive bins). Across the population of cells, the activity in the saccade interval (75 ms before saccade onset to 50 ms after) was significantly lower with the poor stimulus as target (10.7 spikes/s), than with the good stimulus as target (15.5 spikes/s; paired t -test, $P < 0.01$).

The modulation of the response according to the identity of

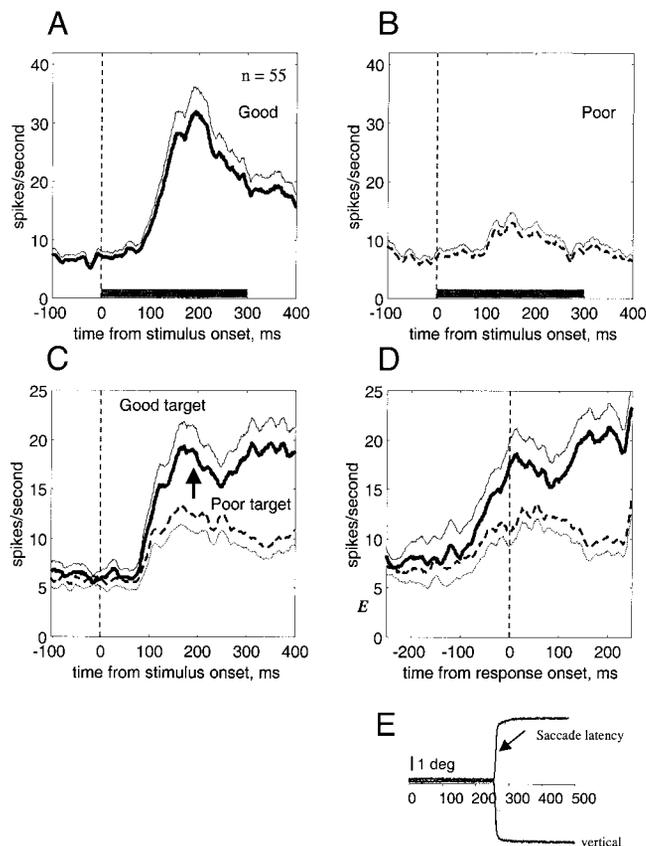


FIG. 5. A and B: responses of the population of cells ($n = 55$) to the good and poor stimuli presented at the fovea (thick lines). Thin lines show 1 SE above the mean response, calculated by smoothing individual cell responses with a gaussian of kernel (20 ms) and calculating the SE across the population. Stimuli subtend 1.5 – 2.5° of visual angle. Each graph shows the population spike density histogram. The bar in the graph shows the stimulus duration of 300 ms. The stimulus array presented was different for every cell. C and D: responses of population of cells ($n = 55$) to the visual array consisting of the good and poor stimulus (thick lines). Thin lines are 1 SE above or below the mean, calculated as described for A and B. Only correct trials are shown. The target was either the good stimulus (solid line), or the poor stimulus (dashed line). Stimuli contained within a single hemifield (either ipsilateral or contralateral). Each curve shows the population spike density histogram. The average saccade latency is shown as the arrow at 220 ms. C: same data locked to stimulus onset. D: locked to saccade onset. E: a single sample eye movement trace, showing the calculation of the saccade latency of 250 ms for the trial shown.

the target was significant in slightly over one-half of the cells. When we summed up the response that occurred in the saccade interval, the firing rate significantly varied according to which stimulus was the target for 29/55 cells in the sample (individual cell unpaired t -test, $P < 0.05$). For 27 of the cells that showed a significant effect, the response in the saccade interval was significantly higher when the good stimulus was chosen. For only two of the cells, the response in the saccade interval was significantly higher when the poor stimulus was chosen. Figure 6 shows the distribution of the index of response difference, defined as the normalized difference in response when the good versus poor stimulus was the target (the difference in response in the 2 conditions, divided by the sum). The indexes have been sorted by rank, and the filled circles represent cells for which the difference in activity was significant.

We randomized, on a cell by cell basis, whether the monkey learned to saccade first to the good or first to the poor stimulus

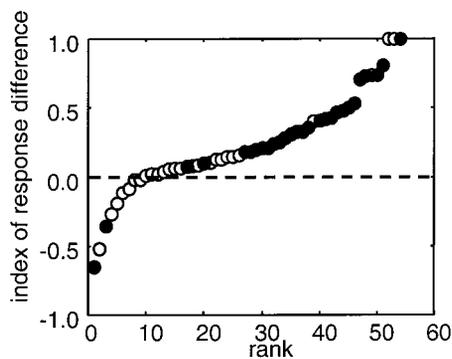


FIG. 6. Index of response difference between good and poor stimulus target arrays. The index is the difference in response to the good and poor target arrays divided by the sum of responses to the good and poor target arrays. The index was calculated for each cell using the mean response, across trials in the time interval 75 ms before saccade onset to 50 ms after saccade onset. The indexes have been sorted and plotted in rank order. The solid circles indicate cells for which the difference in response in the 2 conditions was significant for that cell (unpaired t -test, $P < 0.05$).

in the array. To verify that the order of the training had no effect, we examined the responses of the population of neurons in the original and reversal training blocks. The mean response during the saccade interval was 13.9 spikes/s during the original training and 13.2 spikes/s during reversal training (paired t -test, $P > 0.10$). Order of training therefore did not significantly affect the cell's responses to the array.

Correlation between neuronal responses and behavior

1) INCORRECT TRIALS. On correct trials, the effects of training on the neuronal response to the stimulus array are confounded with the effects of making an eye movement to the target. For that reason, it is not possible on correct trials to tell whether the target-related difference in neural activity is related to the choice the monkey will actually be making, or to the choice that it should be making, based on the training. However, incorrect trials can provide insights in differentiating the effects of training from the effects of choosing a stimulus. In some incorrect trials, the monkey incorrectly chose to saccade to the stimulus that was *not* paired with reward during the training. If the neuronal response differences we found on correct trials were caused by the saccade target choice made by the monkey rather than the effects of training per se, the neuronal response difference should reverse in direction on those error trials in which the monkey made a saccade to the incorrect stimulus. On the other hand, if the neuronal response difference on correct trials were caused by the training the monkey received, and the incorrect behavioral choices on error trials were made because the effects of learning failed to be expressed by aIT cells on those trials, then the response difference should be smaller, or absent on error trials.

The responses on incorrect trials, averaged across the population of 55 cells, is shown in Fig. 7, locked to the onset of the saccade. As shown in the figure, choosing the good versus poor stimulus as the target had no effect on the population activity in the period just before the onset of the saccade. That is, the average activity in the saccade time period was not significantly different when the good stimulus was chosen incorrectly compared with when the poor stimulus was chosen incorrectly, (12.9 vs. 14.5 spikes/s, respectively, paired t -test, $P > 0.10$). The firing rate was slightly higher in those trials when the poor

stimulus was chosen incorrectly, diverging shortly after the onset of the saccade. These results are in contrast to the results found on correct trials. Therefore the activity difference in the saccade interval seen on correct trials seems to reflect not simply the intention of the monkey to saccade to the good stimulus, but rather the training the monkey has received. An alternative explanation for the absence of target selection effects in error trials could be that the monkey merely did not select either stimulus in these trials (i.e., guessed). However, in that case, it must be proposed that saccading to one of the stimuli can occur without selection of a target. [See Seidemann and Newsome (1999) for description of error results in a different task, where the selection of a saccade target occurs separately from attentional selection of a spatial location.] Other types of errors could also affect the expected response differences in correct and error trials. If the errors were unrelated to activity in aIT (e.g., the error occurred because a later stage of saccade control failed), the neuronal responses in aIT should be the same on error and correct trials. However, this was not observed. Alternatively, if the errors result from a misperception of the nontarget stimulus as the target stimulus (an error of stimulus representation/selection), the neuronal responses in aIT would be expected to reverse direction on error trials. Again, this was not observed. A combination of different causes for errors would complicate the interpretation of neuronal responses differences on error trials. However, we can rule out the possibility that the neuronal responses reflect nothing other than the selection of the saccade target. The above results show that on correct trials, the response to the array is larger when the good stimulus is the target than when the poor stimulus is the target, whereas on incorrect trials the responses to the good and poor targets are nearly equivalent. A two-way repeated measures ANOVA across the population shows a significant interaction between the target identity (good vs. poor target, which also has a significant main effect, $P < 0.05$) and performance (correct vs. incorrect, $P < 0.05$). This interaction could result either from an enhancement of the response to the target stimulus on correct trials compared with incorrect trials, or from a suppression of the response to the nontarget stimulus on correct versus incorrect trials, or from a

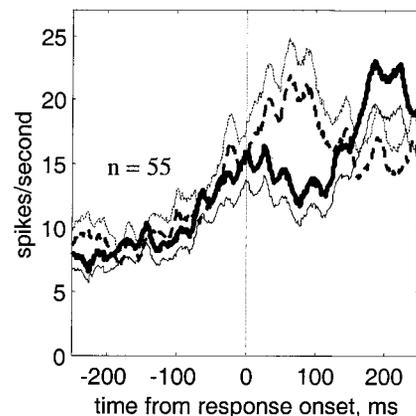


FIG. 7. Responses of population of cells ($n = 55$) on incorrect trials. The dashed line represents the trials when the good stimulus was the correct target stimulus, but the poor stimulus was incorrectly chosen. The solid line represents trials when the poor stimulus was the correct target stimulus, but the good stimulus was incorrectly chosen. Responses are locked to saccade onset. Each curve shows the average spike density histogram (thick lines). Thin lines are 1 SE above or below the mean, calculated as described for Fig. 5, A and B.

combination of both mechanisms. Comparing responses in the correct and incorrect trials separately revealed that the response to the array with the good stimulus as target did not differ significantly on correct versus incorrect trials (good stimulus chosen correctly, mean 15.5 spikes/s; poor stimulus chosen incorrectly, mean 14.5 spikes/s; Wilcoxon signed rank test, $P > 0.10$). By contrast, the response to the array with the poor stimulus as target (and the good stimulus as nontarget) was significantly smaller on correct trials compared with incorrect trials (poor stimulus chosen correctly, mean 10.7 spikes/s; good stimulus chosen incorrectly, mean 12.9 spikes/s; Wilcoxon signed rank test, $P < 0.01$). These results suggest that the effect of the training is to suppress the activity of cells that respond to the distractor or nontarget stimulus. This suppression fails to occur on incorrect trials.

2) RESPONSES EARLY AND LATE IN RECORDING SESSION. Unlike neuronal responses on correct trials, those on incorrect trials did not differ according to whether the good or poor stimulus was the target. This response difference between correct and incorrect trials suggests that modulation of aIT activity during target selection is correlated with the animal's behavioral performance. If so, then the effects of training on neuronal responses might increase over trials, as behavioral performance improves.

To test this, we compared the average firing rate in the saccade interval, during the first block versus the last block of 10 trials in the learning session. Figure 8A shows data collapsed across original and reversal learning trials, for 26 cells that showed a significant difference in response to the good and poor target arrays, had good/poor target indexes >0 , and at least 20 trials in each condition. Of the 29 cells that showed a significant difference in response to the good and poor target arrays, 2 were eliminated because they had good/poor target indexes <0 , and one cell was eliminated because it had <20 trials in each block. The response pattern was similar in both the original and reversal learning trials, as expected from the fact that there was no difference in the response in the original and reversal learning trials (see RESULTS above). The firing rate remains constant when the good stimulus is the target but is suppressed when the poor stimulus is the target. Because of the small size of the effects in some of these cells, the interaction fails to reach significance in this group of cells (2-way repeated measures ANOVA interaction term, $P = 0.10$). We then selected a subgroup of these cells that fit the following criterion: 1) at least 40 trials of response for both the good and poor target arrays and 2) good/poor target index (cf. Fig. 6) >0.10 . This yielded the 20 cells that are shown in Fig. 8, B and C. As shown in Fig. 8B, although the trends shown in the population of 26 cells shown in Fig. 8A remain, the response difference between trials with the good versus poor stimulus as target is larger in the last block of trials than it is in the first block for this subgroup of cells.

In this subgroup of cells, a two-way repeated measures ANOVA shows that there is no main effect of block (first vs. last, $P > 0.10$), but there is a main effect of target stimulus ($P < 0.01$), and that there is a significant interaction between target identity and block ($P < 0.05$). Although there was no difference in response to the good stimulus as target in the last block compared with the first block, the response when the poor stimulus was target was significantly more suppressed in

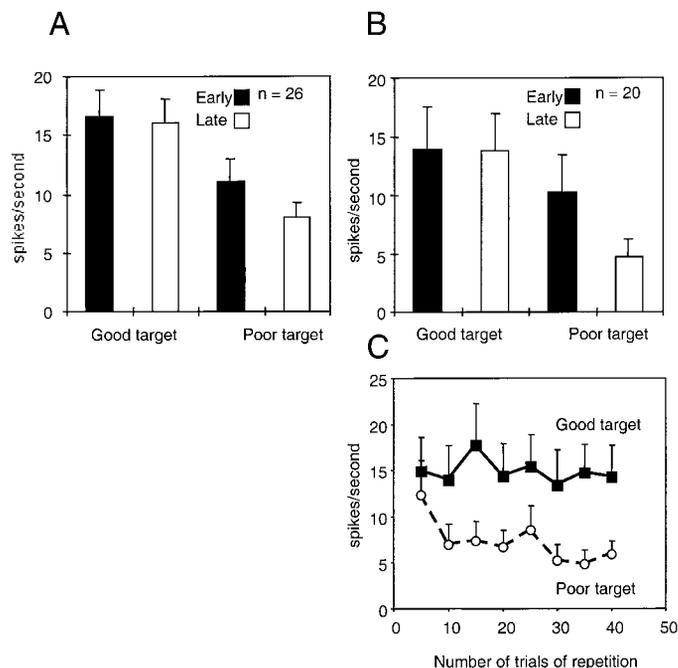


FIG. 8. A: average activity in saccade time interval (75 ms before saccade onset to 50 ms after saccade onset) for 1st 10 trials and last 10 trials for each of those stimulus sets, in which either the good stimulus is the target or the poor stimulus is the target. Only correct trials are shown. Only cells with a significantly larger response to good target arrays than to poor target arrays (1 cell with fewer than 20 trials for the poor target array is excluded). B: subgroup of 20 cells with a significant difference between the response to arrays with good and poor target, at least 40 trials of repetition for both good and poor target arrays, and good-poor target index >0.1 . The responses early and late in the recording session differ significantly when the poor stimulus was the target [(post hoc *t*-test, SPSS), $P < 0.05$], but not when the good stimulus was the target [(post hoc *t*-test, SPSS), $P > 0.10$]. Error bars are 1 SE. C: response to good and poor target arrays in blocks of 5 trials for the same subgroup of 20 cells. Error bars are 1 SE. Differences between 2 lines are not significant for the 1st 5 trials, but are significant for the following groups of trials ($P < 0.01$). For the poor target arrays (\circ), blocks of trials from 25–40 are significantly suppressed below the response in the 1st 5 trials. For the good target arrays (\bullet), the response does not change over the trials.

the last block compared with the first (Fig. 8B). For the smaller subgroup of cells in Fig. 8B, the mean index of response difference (cf. Fig. 6) increased from 0.25 to 0.45, in the same period when behavioral performance increased from 47 to 88% correct.

Data from the same set of cells are illustrated in Fig. 8C, showing the response to the good and poor stimulus in blocks of five trials during the course of the training. The response to the good stimulus remains constant as the stimulus pair is learned, but the response to the poor stimulus drops after the first five trials, and remains low during subsequent trials. The response was not significantly different during the first one block of five trials (paired *t*-test = 0.10), but response differences in the remaining blocks of five are all significant (paired *t*-test < 0.01). The responses are illustrated only for correct trials, during both initial and reversal training. Collapsed across initial and reversal learning, the median trials to a criterion of 80% correct for this subset of cells was 27. Although noisier, the trends are similar when data are separated into the groups of cells in which the good stimulus was the first target, and those in which the poor stimulus was the first target. There was no significant difference in response in the first block of five

trials ($P > 0.10$), but the next block, and five of six of the subsequent blocks were significantly different ($P < 0.05$).

In all of the cells in this subgroup (as well as the nonselected subgroup), the activity remained constant over the course of the session when the good stimulus was the target (early vs. late, post hoc *t*-test after significant ANOVA, $P > 0.10$; Fig. 8B). However, when the good stimulus was the distractor (and the poor stimulus was the target), the firing rate was suppressed later in the session compared with earlier in the session (early vs. late, post hoc *t*-test after significant ANOVA, $P < 0.05$). Thus the data from early and late in the session support the conclusion drawn from the data on correct and incorrect trials. In both analyses, the firing rate that changes with performance is the one in which the good stimulus is the distractor (and the poor stimulus is the target). Specifically, the responses of cells that respond to the good stimulus are suppressed when it is the distractor, and this suppression grows with learning.

Effects of stimulus location

There were statistically significant effects of learning on the neuronal response before the onset of the saccade when stimuli were positioned across the vertical meridian from one another, but these effects were less striking than when the stimuli were contained within the same hemifield. When the good stimulus was in the ipsilateral hemifield, the average response when the good stimulus was the target was 17.3 spikes/s, compared with 13.5 spikes/s when the poor stimulus was the target, which was a significant difference (paired *t*-test, $P < 0.01$). About one-half of the cells (23 cells) showed a significant difference in the response (either larger or smaller) depending on which stimulus was the target (unpaired *t*-test, $P < 0.05$, $n = 23$ cells), which is similar to the proportion of cells that showed significant effects when the array was confined to a single hemifield. However, only 18 of the 23 showed a *larger* response when the good stimulus was the target, a proportion smaller than the one that showed this effect when the stimuli were confined to the same hemifield.

When the good stimulus was in the contralateral hemifield and the poor stimulus was in the ipsilateral hemifield, the effects of learning on the firing rate were even smaller than in the preceding case, but were still statistically significant (good stimulus as target, 17.5 spikes/s; poor stimulus as target, 14.9 spikes/s; paired *t*-test, $P < 0.05$). A total of 22 cells showed a significant difference in the response (either larger or smaller) according to which stimulus was the target (unpaired *t*-test, $P < 0.05$, 22 cells), which again is similar to the proportion showing an effect when the array was contained within the same hemifield. However, only 12 of the 22 cells showed a *larger* response when the good stimulus was the target in this configuration, which was substantially fewer than when the array was contained within the same hemifield. Thus the effects of the training are less consistent when the stimuli are placed across the vertical meridian from one another.

Overall, only 18/55 cells had significantly higher activity when the target was the good stimulus, when the two stimuli were in opposite hemifields. On the other hand, when both stimuli were contained within the same hemifield, the response difference between the good and poor target arrays did not seem to depend significantly on whether that hemifield was ipsilateral or contralateral to the cell. The overall responses

were smaller when stimuli were contained within the ipsilateral hemifield (good stimulus as target, 15.1 spikes/s; poor stimulus as target, 10.2 spikes/s) than when the stimuli were contained within the contralateral hemifield (poor stimulus as target, 18.4 spikes/s; good stimulus as target, 12.8 spikes/s). However, the differences in response to the good and poor target arrays were significant in both hemifields (paired *t*-test, $P < 0.01$) and the index of response difference (cf. Fig. 6) for the good stimulus arrays was comparable in both hemifields.

Relationship between stimulus selectivity and effects of learning

We tested whether there was a correlation between the degree of selectivity of the cell for the good versus poor stimuli and the effects of choosing the good versus poor stimulus as the target. As a measure of selectivity, we calculated the difference in response to good and poor stimuli, presented at the fovea, divided by the sum. As a measure of learning, we calculated the difference in response to the two target arrays (good stimulus as target vs. poor stimulus as target) divided by the sum, as in Fig. 6 (Chelazzi et al. 1998). The two variables were correlated ($r = 0.56$, $z = 4.1$, $P < 0.01$); thus the greater the selectivity of the neuron for the individual stimuli that make up the pair, the greater the effect of learning and behavioral relevance (target identity) on the neuron's response to the pair.

Responses to individual stimuli during learning

A possible explanation for the effects of learning on the response to the search array is that the responses to the individual stimuli in the array change with learning. For example, the response to the good stimulus presented alone might increase or decrease over trials depending on whether it is the target, i.e., increasingly associated with reward, or the nontarget, i.e., increasingly associated with nonreward. If so, this could explain why the response to the pair was better when the good stimulus was the target than when the poor stimulus was the target. To test this, we recorded the responses of a subset of the cells to the individual stimuli presented alone on probe trials over the course of learning in *task 1*. On these probe trials, individual stimuli were presented at fixation, and the monkey was rewarded simply for maintaining fixation throughout the trial (see METHODS, passive fixation task).

Figure 9 shows the average response of eight cells to the good and poor stimuli presented alone on the probe trials, intermixed with trials in which they were paired together in the search task. Overall, there is no difference in the response to either the good or poor stimulus presented alone depending on whether they are target stimuli or nontarget stimuli (good stimulus, 42 spikes/s under both conditions, $P > 0.10$; poor stimulus, 12 spikes/s under both conditions, $P > 0.10$). Thus the change in response to the search array that occurs over the course of learning apparently depends on both stimuli being present together in the pair. Figure 9C shows the response during the discrimination task (*task 1*) for this subset of cells. Before the onset of the saccade, the response is higher in good target trials than in poor target trials. In these same cells, the mean index of selection for the stimulus arrays (as in Fig. 6) was 0.27, and the effects of target selection were significant in

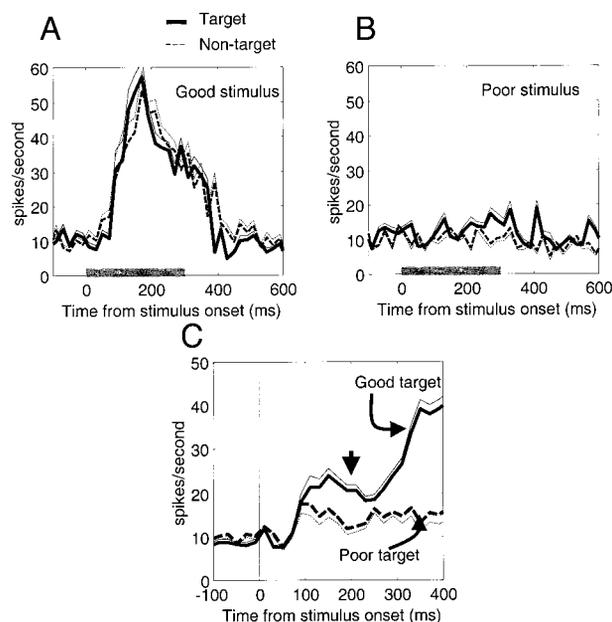


FIG. 9. Responses to single stimuli presented at fovea during the course of discrimination training. Data show responses collected during single stimulus presentation trials from *task 1*. These trials were intermixed with the presentation of the discrimination task in *task 1*. Solid lines show target stimuli, and dashed lines nontarget stimuli, presented intermixed with trials in which either the good or poor stimulus was the saccade target when it appeared within the array. Bar shows the stimulus duration of 300 ms. Data are average population ($n = 8$) histograms, with 20-ms bins. A: good stimuli. B: poor stimuli. C: response during discrimination task (*task 1*) for subset of cells shown in A and B. Arrow points to the average saccade latency of 200 ms. Thin lines show 1 SE above or below the mean for each time point.

seven of the eight cells, showing that selection did occur in this subset of neurons.

Firing rate in the baseline period

As in the present study, Chelazzi et al. (1993, 1998) also measured aIT responses when the monkey selected a target stimulus from an array. In those studies, however, the identity of the target stimulus varied randomly across trials, with the particular stimulus to be used as the target on a given trial being briefly presented to the animal as a cue stimulus at the start of the trial. It was found that the activity in the delay period between cue and choice array was higher when the cue was a good stimulus for the cell than when it was a poor stimulus. This difference in activity could even precede the cue, if the animal could predict that the good stimulus would be the cue on that trial. It was proposed that this differential activity both before and after the cue resulted from feedback to aIT cortex from structures involved in working memory, biasing activity in favor of cells coding the target stimulus.

In the present study, it was also possible that the animal used a working memory strategy, and that feedback from structures mediating working memory biased aIT activity in favor of the target stimulus on each trial, just as in the Chelazzi et al. task. We therefore tested whether, preceding the onset of the array, the population activity was higher in blocks of trials when the good stimulus was the target than in blocks when the poor stimulus was the target. We found no significant difference in activity (good stimulus as target, 6.6 spikes/s; poor stimulus as

target, 5.8 spikes/s; paired t -test, $P > 0.10$), which is inconsistent with the feedback hypothesis. This was tested more explicitly in the concurrent discrimination task described below.

Concurrent discrimination learning

To further rule out the possibility of mediation by working memory, one monkey was trained on a concurrent discrimination version of the search task, in which any of 10 different stimulus pairs could appear randomly on a given trial. We reasoned that 10 pairs would exceed the capacity of working memory and would prevent the animal from forming an expectation about the identity of the target stimulus on a given trial. Since it was not possible to train monkeys to reach a high level of performance on 10 different stimulus pairs in one recording session lasting 2–3 h, a consistent set of stimuli (the entire set of 10 stimulus pairs in Fig. 2) was used every day. The stimuli were initially arbitrarily chosen from a set of previously unused stimuli and assigned as targets and nontargets. The stimuli, their pairing, and the target-nontarget stimulus in each pair remained constant throughout the recording period for all cells. The stimuli appeared in the same configurations as in the single pair version of the task, and, as in that task, the location of the stimuli was unpredictable from trial to trial, requiring a discrimination of visual features, not spatial position.

The monkey learned to saccade reliably to the rewarded stimulus in each of the pairs over a period of 30–40 daily sessions, consisting of 200–300 trials each. At that point, the monkey was performing at an average of 91% correct and with an average saccade latency of 181 ms.

Unlike the case in *task 1*, the pairs were already learned, and therefore we could not customize the stimulus pairs so that they were composed of good and poor stimuli for each cell. We therefore used the responses of each cell to each of the 20 stimuli presented at the fovea, to determine whether or not an already-learned stimulus pair contained a “good” and “poor” stimulus. For each stimulus pair for each cell, an index was calculated, based on the difference in response between the target and nontarget stimulus (during a 200-ms period starting 100 ms after stimulus onset), divided by the sum of the responses. This index could range from -1 to $+1$. Pairs for which this index was greater than $+0.2$ were classified as having a good stimulus as the target. Pairs for which this index was less than -0.2 were classified as having the poor stimulus as the target. Pairs with indexes between -0.2 and $+0.2$ were

TABLE 1. Categorization of good and poor target arrays in concurrent discrimination task (*task 2*)

Category	Number of Pairs	Percent
1) No difference	226	58
2) Good target (selectivity index $> +0.2$)	79	20
3) Poor target (selectivity index < -0.2)	87	22

Number of pairs refers to the stimulus pairs across cells that fit into that category. A single cell can appear more than once, in either the same or different category. A single stimulus pair appears only once, and its categorization depends on the response to the cell’s response to the individual stimuli that make up the pair. The total number of cells was 49.

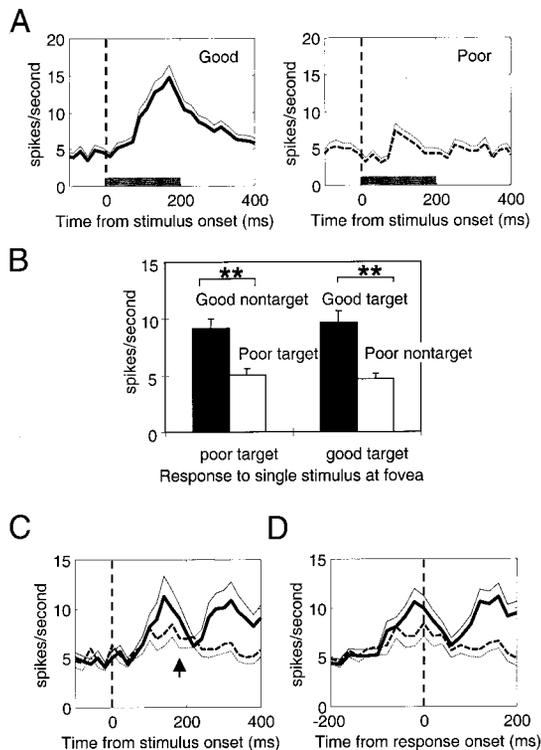


FIG. 10. Responses in the concurrent discrimination task (*task 2*). *A*: average response for good and poor stimuli (20-ms bins) on correct trials. Good and poor stimuli are as described in Table 1. Solid bar shows the 200 ms duration of the stimulus. Population average histograms, with bin size of 20 ms. *B*: mean response to good and poor stimulus at the fovea, for stimulus pairs where the poor or good stimulus was the target in the array. The response difference between good and poor is significant (*t*-test, $P < 0.01$) for both sets of pairs. The response difference between good and poor stimuli depending on whether they were target stimuli or nontarget stimuli is not significant (*t*-test, $P > 0.10$). *C*: response across the population of cells with differences in response to target and nontarget stimuli during concurrent discrimination performance (*task 2*). Average response (20-ms bins) when the good stimulus was the target (solid line) and when the poor stimulus was the target (dashed line). Cells and stimuli categorized as described in Table 1. Arrow points to the average saccade latency of 180 ms. *D*: same data as in *C*, time locked to the saccade onset in each trial. Thin lines show 1 SE above or below the mean for each time point.

assigned to the “no difference” category. Table 1 shows the distribution of stimulus pairs ($n = 392$) in the three categories. The table shows that there were roughly as many stimulus pairs in which the target stimulus caused a greater response than the nontarget as there were stimulus pairs in which the target stimulus caused a smaller response than the nontarget. The population average responses to the good and poor stimuli presented alone at the fovea are shown in Fig. 10, *A* and *B*. As expected, the response difference between good and poor stimuli presented alone was not as large as in *task 1* (Fig. 5, *A* and *B*). More importantly, there was no difference in response to the good and poor stimuli as a function of their status as a target or nontarget (Fig. 10*B*). Thus the stimulus selection appeared to be unbiased across the population.

With the stimuli categorized in this way, we computed the population average response to the array when the good or poor stimulus was the target (excluding responses to pairs in the no-difference category). The pattern of results (Fig. 10*C*) is the same as seen in Fig. 5*C*. In the period 100–160 ms after stimulus onset, the response is higher for pairs in which the good stimulus was the target than where the poor stimulus was

the target ($P < 0.05$). The response difference during performance was smaller than in *task 1*, probably because the response difference between the good and poor stimuli presented individually was also smaller than in *task 1*.

Two other similarities were noted between these results and those found during discrimination learning and reversal with single stimulus pairs. First, the baseline activity (i.e., the activity before the onset of the array) did not differ depending on which stimulus was the target (good stimulus chosen correctly, mean = 4.9 spikes/s; poor stimulus chosen correctly, mean = 4.6 spikes/s, $P > 0.20$). Second, there was no effect of behavioral relevance on the responses on incorrect trials (good stimulus chosen incorrectly, mean response 8.7 spikes/s; poor stimulus chosen incorrectly, mean response 9.3 spikes/s, $P > 0.20$). These similarities in responses in the two tasks suggest that the effects of learning on neuronal responses in these tasks is not due to the use of neuronal mechanisms for working memory. Because this task was learned before the recording session began, it was not possible to compare the neural responses early and late in the learning, as in the on-line learning task.

Similarity of responses to paired stimuli

Previous studies of inferior temporal cortex have reported that behavioral training with pairs of stimuli increases the probability that single cells will respond similarly to both members of the stimulus pair (Erickson and Desimone 1999; Sakai and Miyashita 1991). Because in our arrays, the target stimulus frequently appeared together with the nontarget stimulus in the array, we tested whether there was any correlation between the neuronal responses to individually presented target and nontarget stimuli. Each cell was tested by presenting, individually, at the fovea, each of the 20 target and nontarget stimuli taken from the concurrent discrimination. A correlation between the responses to the paired targets and nontargets was calculated for each cell. The distribution of correlation coeffi-

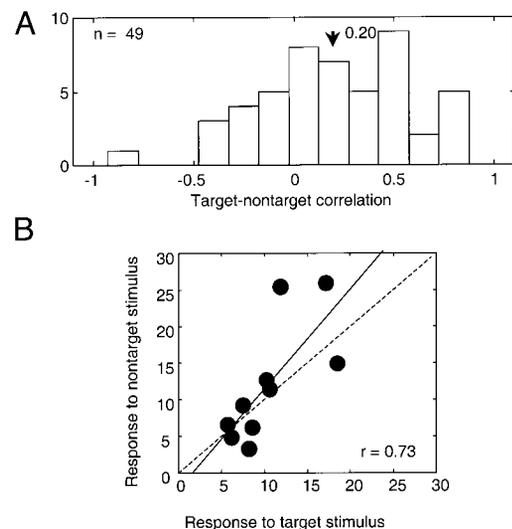


FIG. 11. Correlation between responses to target and nontarget stimuli. *A*: the distribution of correlation coefficients between responses to target and nontarget stimuli for cells with significant stimulus selectivity ($n = 49$). The distribution was significantly different from 0 (*t*-test on Fisher Z transformed correlation coefficients, $P < 0.01$). *B*: an example cell showing correlated responses to the paired stimuli. The response was measured in the time interval 75 ms after stimulus onset to 250 ms.

cients across the population of cells is shown in Fig. 11. The mean correlation coefficient was 0.196, which is significantly different from 0 (t -test on Fisher Z transformed r values, $P < 0.01$). An example of a single cell with a significant correlation between its responses to target and nontarget stimuli is shown in Fig. 11B. Thus stimulus pairing appears to alter the selectivity of the cells for the stimuli (although correlations may have also existed before training, by chance). The correlation is small, however, and obviously did not eliminate the cell's ability to discriminate between the target and nontarget stimuli.

DISCUSSION

The present results suggest that learning can increase the salience of stimuli in aIT cortex. When an animal is presented with an array of two stimuli, the responses of aIT cells appear to be largely dominated by whichever stimulus has become behaviorally relevant through stimulus-response-reward training. If the relevant stimulus is a good one for the cell, the cell responds well, but if the relevant stimulus is a poor one for the cell, the response to the identical array is poor.

This learning effect is consistent with the idea that learning directly modifies the neuronal representations of the stimuli in extrastriate cortex, giving the relevant stimulus a competitive advantage. By themselves, however, these data do not rule out the possibility that aIT responses were dynamically modulated by attentional selection of the target stimulus. When an animal attends to one of two stimuli in an array, neuronal responses in extrastriate cortex are determined primarily by the attended stimulus (Chelazzi et al. 1993, 1998). One could argue that the training in the present study simply taught the animal to attend to the relevant stimulus, and the site of learning was within structures concerned with attentional control, located outside of extrastriate cortex. According to this view, the effects of learning on aIT responses are due to feedback from these structures mediating stimulus selection.

Two lines of evidence weigh against this possibility. First, on incorrect trials, the response to the array did not vary according to which stimulus was the target, even though the animal selected a stimulus on these trials and used it for the target of an eye movement, just as it did on the correct trials. Second, the effects of learning on the response to the array were larger later in the recording session than earlier, even though a stimulus was selected on every trial, in both the early and late trials. Thus the results cannot easily be explained by feedback from sites mediating attentional-selection per se, although it is possible to argue that attentional selection parallels learning in such a way that the two phenomena are indistinguishable. One might argue, for example, that early in the session and during error trials, selection of the target for a saccade occurs without attention. Later in the session, when the behavioral response is consistent, targets are selected by an attentional mechanism that sends feedback to aIT. This hypothesis would be contrary to some ideas about the close relationship between attentional selection and saccades, however (Deubel and Schneider 1996; Kowler et al. 1995; Kustov and Robinson 1996; Sheliga et al. 1995).

Another related possibility suggested by previous experiments is that aIT responses may be dynamically "biased" on a trial-by-trial basis to respond to a particular stimulus, as a result of feedback from structures concerned with working

memory. Chelazzi et al. (1993, 1998) recorded from aIT using a visual search task, in which the target on each trial was specified by a cue stimulus briefly presented at the start of the trial. The animal had to use its memory of the cue to guide its selection of the target from the search array. They found that the baseline activity between the time of the cue and the time of the array was higher when a good stimulus for the recorded cell was used as a cue-target than when a poor stimulus was used as a cue-target. This higher baseline activity was taken to be direct evidence for the bias in favor of cells coding the target stimulus. A baseline difference was not seen in the present study, however. Although the positive, or target, stimulus remained consistent for a block of trials (during a discrimination learning or reversal; *task 1*), the baseline activity before the presentation of the array was similar, regardless of whether a good or poor stimulus was the expected target.

Nevertheless, to test further the possible role of working memory, we also recorded responses in a task in which any of 10 different pairs of stimuli could be presented on a given trial. This presumably prevented the animal from anticipating which target would be present and precluded a working memory strategy. The results in this multiple-pair stimulus version of the task were similar to those in the task with a single pair; these similarities included the difference in response to good and poor target arrays, the difference in response on error trials, and the absence of baseline activity differences. These similarities argue against the working memory hypothesis.

The most likely explanation for these results is that learning directly modifies the representations of stimuli in extrastriate cortex and thereby gives the behaviorally relevant stimuli a competitive advantage. Conclusively proving that hypothesis, however, would require a more complete analysis of the different structures that might play a role in this task, or alternatively, the demonstration of a permanent learning based change in the response of neurons in aIT (a structural change, or a biophysical change in their properties) as a result of the learning.

Regardless of the precise site of the plastic change, however, processing is largely limited to relevant stimuli in aIT, and cells representing irrelevant stimuli in a scene are suppressed. As a relevant object acquires salience, it may attract attention through its mere presence in the scene, without first requiring a conscious guiding of attention to a location or feature. Psychophysical data and models (Shiffrin and Czerwinski 1988; Shiffrin and Schneider 1977) suggest that an individual item may indeed attract attention automatically, when that stimulus has been repeatedly presented as a target stimulus. The interaction among top-down effects of attention, bottom-up effects of stimulus salience, and stimulus-specific learning is supported by psychophysical experiments.

The bottom-up effects that make an image salient are similar to the top-down influences wielded by active attention, and the benefits conferred by perceptual learning (Ito et al. 1998). A collinear flanking stimulus increases the perceived brightness (i.e., salience) of a central target stimulus (Ito et al. 1998; Kapadia et al. 1995). When attention is first directed to a specific location (a top-down process) or training results in perceptual learning, the flanking stimuli no longer cause increased salience of the target (Ito et al. 1998). Perceptual learning in the ability to detect targets in dense texture displays also suggest that low-level processes, like scene segmentations

can be modified by the top-down influences produced by learning (Ahissar and Hochstein 1995; Karni and Sagi 1991; Sireteanu and Rettenbach 1995).

Response differences to individual stimuli presented alone, and role of specific pairing

We considered the possibility that learning caused changes in the response to the individual stimuli that made up the choice array. For example, learning might have enhanced the response to the good stimulus alone, which would lead to a larger response to the array when the good stimulus was the target. However, we found no evidence that the response to individual stimuli changed with learning. Instead, the effects of learning on the response to the arrays appeared to depend on the simultaneous presentation of the target and nontarget stimuli. We suggest that learning asymmetrically strengthens the competitive interactions between the cells representing each of the two stimuli in the array, with the advantage favoring the relevant stimulus. Because it is only the competitive (suppressive) connections that are changed, the responses to each of the stimuli presented alone would be unaffected, as was found in this study (Fig. 9).

If so, one implication of this hypothesis is that the learning effects would be specific to the precise pairing of target and nontarget stimuli used in training. If the target stimulus were paired with a different nontarget following training, the neural effects should diminish and the animal should have difficulty choosing the target. Substituting a different negative stimulus for a well-learned distractor in the array does seem to impair the animal's performance in the task (Jagadeesh, Mishkin, and Desimone, unpublished observations). This hypothesis should be tested in future studies. The design of this task does make the strategy of learning to ignore the irrelevant stimulus in a specific pair efficient (because both stimuli always occurred together, it was not necessary to learn anything about the stimuli apart from their role as a pair). An altered version of the task where many different nontarget stimuli could appear with a single target stimulus, which would then become the target stimulus across many different contexts, might lead to different neuronal effects.

Changes in the responses to individual learned stimuli compared with other, un-learned stimuli have been demonstrated in previous studies. For example, Sakai and Miyashita (1994) showed that after long-term training in a discrimination task, cells were likely to respond better to stimuli that had been used in the task than to other, visually similar stimuli. Kobatake et al. (1998) compared the responses to the same sets of stimuli in trained and untrained animals and found cells were more likely to give an optimal response to stimuli in the trained stimulus set in the trained animals. Neither study, however, demonstrated short-term changes in the responses to individual stimuli as a function of training; in the present study, also, we were unable to demonstrate any short-term changes in the response of individual cells to single stimuli as a function of the learning during a single recording session.

In the concurrent discrimination task, which involved long-term learning, cells were more likely to respond similarly to stimuli that had been paired together in the discrimination task, than would be expected by chance (Fig. 11). This phenomenon, termed "pair coding," has been observed in several other studies

(Erickson and Desimone 1999; Sakai and Miyashita 1991, 1994). In the present study, however, unlike the others, where association of individual stimuli into pairs was either required (Sakai and Miyashita 1991, 1994) or beneficial (Erickson and Desimone 1999) to performance, similarity in the responses to the target and distractor stimuli developed even though this association was not reinforced and could tend to impair rather than improve performance in the task. An endpoint result of pair coding can be for a single cell to respond identically to two visually dissimilar stimuli (like the target-nontarget pair). Development of pair coding of this form could impair performance in the discrimination task because the degree of response suppression seen in the neurons was correlated with the selectivity that the neurons had for the target-nontarget stimulus (see RESULTS). Neurons that are not selective for the target-nontarget stimulus are presumed to not participate in the selection of the target (Chelazzi et al. 1993, 1998). It is possible to develop pair coding while still maintaining selectivity for the target nontarget stimuli, or alternatively, allowing a group of cells to drop out of population of cells participating in the discrimination, while serving some other purpose, such as distinguishing the target-nontarget stimulus pair from other stimuli. Since pair coding was neither required, or necessarily beneficial in this task, it is possible that "pair coding" can occur as a bottom-up process, simply as a result of spatiotemporal proximity in the presentation of individual object stimuli, as found in the Erickson and Desimone passive association task (Erickson and Desimone 1999). If so, this process could enable cells to respond to different transforms of the same image (for example, transformations of location

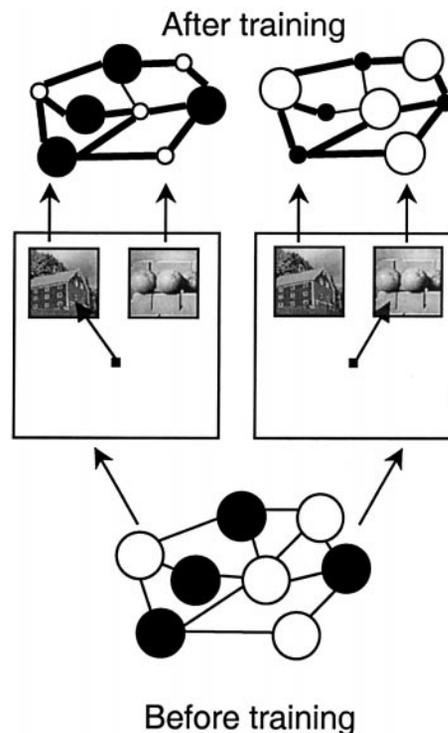


FIG. 12. Diagram of change in cortical circuit with training. Dark circles represent cells that respond better to the house than to the oranges; light circles represent cells that respond better to the oranges than the house. Before training, both groups of cells respond when the array is presented. After training, only those cells that respond to the relevant, target stimulus are active, and those cells that respond to the nontarget, irrelevant stimulus are suppressed. This is a simplified diagram that represents best the data shown for the subset of 15 cells in Fig. 8C.

or size) (Ito et al. 1995; Lueschow et al. 1994), as well as to different rotations of the same three-dimensional object (Logothetis and Pauls 1995; Logothetis et al. 1995).

Mechanism for implementing learned salience

Although the responses to individual stimuli did not, on the whole, change with the learning of behavioral relevance, the responses to arrays of stimuli did change (Figs. 7 and 8). A simplified summary of the effects of the short-term learning, extrapolating from the data shown in Figs. 5, 7, and 8, *B* and *C*, is diagrammed in Fig. 12. When monkeys are trained to saccade to a target stimulus in a scene, the competitive interactions between cells responding to the target and nontarget stimulus are biased in favor of cells responding to the target, resulting in suppression of the cells responding to the nontarget. Two pieces of evidence point to this suppressive mechanism. First, the neuronal response to the behaviorally irrelevant stimulus (distractor) was suppressed over the course of training, whereas the response to the relevant stimulus (target) remained unchanged (Fig. 8). Second, the activity was suppressed on correct trials compared with incorrect trials, when the good stimulus was also the behaviorally irrelevant stimulus, but was similar on both correct and incorrect trials when the good stimulus was the behaviorally relevant stimulus. In the population overall, then, the pattern of activity is biased in favor of the stimulus that is the relevant one in the task, i.e., the stimulus to which a saccade should be made.

This pattern, across the cortex, resembles that for stimuli that are more salient because of their image characteristics. For example, if we imagine the pattern of activity produced by dim and bright bars, the bright bars, too, will elicit greater activity in cortex than the dim bars, which are less salient. Cells responding to the bright bars will then have a competitive advantage.

An important point apparent from the diagram is that the activity of a single neuron does not determine the salience of the image in an array. The populations of cells that respond to the relevant image in the array do not change their responses as that stimulus becomes salient. Only the population of cells that respond to the irrelevant image in the scene is suppressed. Thus the information that codes the increased salience of the image is coded in the differential activity between the two populations of cells that code the relevant and irrelevant image in the scene. The choice of the behaviorally relevant stimulus must be made by comparing the activity of these two populations. Thus to accurately predict the use of this information by the monkey, we should compare the activities of the two populations of cells that respond to the relevant and irrelevant stimuli.

We thank Drs. Giuseppe Bertini and Peter De Weerd for discussions on the data.

This work was supported by the National Institute of Mental Health Intramural Research Program (NIMH-IRP).

REFERENCES

- AHISSAR M AND HOCHSTEIN S. Learning pop-out detection: specificities to stimulus characteristics. *Vision Res* 36: 3487–3500, 1995.
- CHELAZZI L, DUNCAN J, MILLER EK, AND DESIMONE R. Responses of neurons in inferior temporal cortex during memory-guided visual search. *J Neurophysiol* 80: 2918–2940, 1998.
- CHIELAZZI L, MILLER EK, DUNCAN J, AND DESIMONE R. A neural basis for visual search in inferior temporal cortex. *Nature* 363: 345–347, 1993.
- CONNOR CE, GALLANT JL, PREDDIE DC, AND VAN ESSEN DC. Responses in area V4 depend on the spatial relationship between stimulus and attention. *J Neurophysiol* 75: 1306–1308, 1996.
- CONNOR CE, PREDDIE DC, GALLANT JL, AND VAN ESSEN DC. Spatial attention effects in macaque area V4. *J Neurosci* 17: 3201–3214, 1997.
- DEUBEL H AND SCHNEIDER WX. Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Res* 36: 1827–1837, 1996.
- ELLISON A AND WALSH V. Perceptual learning in visual search: some evidence of specificities. *Vision Res* 38: 333–451, 1998.
- ERICKSON CA AND DESIMONE R. Responses of macaque perirhinal neurons during and after visual stimulus association learning. *J Neurosci* 19: 10404–10416, 1999.
- FUSTER JM. Inferotemporal units in selective visual attention and short-term memory. *J Neurophysiol* 64: 681–697, 1990.
- GAFFAN D. Dissociated effects of perirhinal cortex ablation, fornix transection and amygdalotomy: evidence for multiple memory systems in the primate temporal lobe. *Exp Brain Res* 99: 411–422, 1994.
- GROSS CG, BENDER DB, AND MISHKIN M. Contributions of the corpus callosum and the anterior commissure to visual activation of inferior temporal neurons. *Brain Res* 131: 227–239, 1977.
- ITO M, TAMURA H, FUJITA I, AND TANAKA K. Size and position invariance of neuronal responses in monkey inferotemporal cortex. *J Neurophysiol* 73: 218–226, 1995.
- ITO M, WESTHEIMER G, AND GILBERT CD. Attention and perceptual learning modulate contextual influences on visual perception. *Neuron* 20: 1191–1197, 1998.
- JUDGE SJ, RICHMOND BJ, AND CHU FC. Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res* 20: 535–538, 1980.
- KAPADIA MK, ITO M, GILBERT CD, AND WESTHEIMER G. Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys. *Neuron* 15: 843–856, 1995.
- KARNI A AND SAGI D. Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proc Natl Acad Sci USA* 88: 4966–4970, 1991.
- KAROL EA AND PANDYA DN. The distribution of the corpus callosum in the rhesus monkey. *Brain* 94: 471–486, 1971.
- KOBATAKE E, WANG G, AND TANAKA K. Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. *J Neurophysiol* 80: 324–330, 1998.
- KOWLER E, ANDERSON E, DOSHER B, AND BLASER E. The role of attention in the programming of saccades. *Vision Res* 35: 1897–1916, 1995.
- KUSTOV AA AND ROBINSON DL. Shared neural control of attentional shifts and eye movements. *Nature* 384: 74–77, 1996.
- LOGOTHETIS NK AND PAULS J. Psychophysical and physiological evidence for viewer-centered object representations in the primate. *Cereb Cortex* 5: 270–288, 1995.
- LOGOTHETIS NK, PAULS J, AND POGGIO T. Shape representation in the inferior temporal cortex of monkeys. *Curr Biol* 5: 552–563, 1995.
- LUCK SJ, CHELAZZI L, HILLYARD SA, AND DESIMONE R. Neural mechanisms of spatial selective attention in areas V1 V2 and V4 of the macaque visual cortex. *J Neurophysiol* 77: 24–42, 1997.
- LUESCHOW A, MILLER EK, AND DESIMONE R. Inferior temporal mechanisms for invariant object recognition. *Cereb Cortex* 4: 523–531, 1994.
- MEUNIER M, BACHEVALIER J, MISHKIN M, AND MURRAY EA. Effects on visual recognition of combined and separate ablations of the entorhinal and perirhinal cortex in rhesus monkeys. *J Neurosci* 13: 5418–5432, 1993.
- MILLER EK, LI L, AND DESIMONE R. Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* 13: 1460–1478, 1993.
- MORAN J AND DESIMONE R. Selective attention gates visual processing in the extrastriate cortex. *Science* 229: 782–784, 1985.
- MORRIS JS, FRISTON KJ, AND DOLAN RJ. Neural responses to salient visual stimuli. *Proc R Soc Lond B Biol Sci* 264: 769–775, 1997.
- MURRAY EA AND MISHKIN M. Visual recognition in monkeys following rhinal cortical ablations combined with either amygdalotomy or hippocampectomy. *J Neurosci* 6: 1991–2003, 1986.
- PANDYA DN, KAROL EA, AND HEILBRONN D. The topographical distribution of interhemispheric projections in the corpus callosum of the rhesus monkey. *Brain Res* 32: 31–43, 1971.

- RICHMOND BJ AND SATO T. Enhancement of inferior temporal neurons during visual discrimination. *J Neurophysiol* 58: 1292–1306, 1987.
- ROCHA-MIRANDA CE, BENDER DB, GROSS CG, AND MISHKIN M. Visual activation of neurons in inferotemporal cortex depends on striate cortex and forebrain commissures. *J Neurophysiol* 38: 475–491, 1975.
- SAKAI K AND MIYASHITA Y. Neural organization for the long-term memory of paired associates. *Nature* 354: 152–155, 1991.
- SAKAI K AND MIYASHITA Y. Neuronal tuning to learned complex forms in vision. *NeuroReport* 5: 829–832, 1994.
- SATO T. Effects of attention and stimulus interaction on visual responses of inferior temporal neurons in macaque. *J Neurophysiol* 60: 344–364, 1988.
- SATO T. Interactions of visual stimuli in the receptive fields of inferior temporal neurons in awake macaques. *Exp Brain Res* 77: 23–30, 1989.
- SEIDEMANN E AND NEWSOME WT. Effect of spatial attention on the responses of area MT neurons. *J Neurophysiol* 81: 1783–1794, 1999.
- SHELIGA BM, RIGGIO L, AND RIZZOLATTI G. Spatial attention and eye movements. *Exp Brain Res* 105: 261–275, 1995.
- SHIFFRIN RM AND CZERWINSKI MP. A model of automatic attention attraction when mapping is partially consistent. *J Exp Psychol Learn Mem Cogn* 14: 562–569, 1988.
- SHIFFRIN RM AND SCHNEIDER W. Controlled and automatic human information processing: 2. Perceptual learning, automatic attending, and a general theory. *Psychol Rev* 84: 127–190, 1977.
- SIRETEANU R AND RETTENBACH R. Perceptual learning in visual search: fast enduring but non-specific. *Vision Res* 35: 2037–2043, 1995.
- SPITZER H, DESIMONE R, AND MORAN J. Increased attention enhances both behavioral and neuronal performance. *Science* 240: 338–340, 1988.
- SUZUKI WA, ZOLA-MORGAN S, SQUIRE LR, AND AMARAL DG. Lesions of the perirhinal and parahippocampal cortices in the monkey produce long-lasting memory impairment in the visual and tactual modalities. *J Neurosci* 13: 2430–2451, 1993.
- ZEKI SM. Comparison of the cortical degeneration in the visual regions of the temporal lobe of the monkey following section of the anterior commissure and the splenium. *J Comp Neurol* 148: 167–175, 1973.
- ZOLA-MORGAN S, SQUIRE LR, AMARAL DG, AND SUZUKI WA. Lesions of perirhinal and parahippocampal cortex that spare the amygdala and hippocampal formation produce severe memory impairment. *J Neurosci* 9: 4355–4370, 1989.
- ZOLA-MORGAN S, SQUIRE LR, CLOWER RP, AND REMPEL NL. Damage to the perirhinal cortex exacerbates memory impairment following lesions to the hippocampal formation. *J Neurosci* 13: 251–265, 1993.