

Activity of Neurons in Anterior Inferior Temporal Cortex during a Short-Term Memory Task

Earl K. Miller, Lin Li, and Robert Desimone

Laboratory of Neuropsychology, National Institute of Mental Health, Bethesda, Maryland 20892

Inferior temporal (IT) cortex of primates is known to play an important role in visual memory. Previous studies of IT neurons during performance of working memory tasks have found modulation of responses when a current stimulus matched an item in memory; however, this effect was lost if other stimuli intervened in the retention interval. To examine how IT cortex retains memories while new stimuli are activating the cells, we recorded from IT neurons while monkeys performed a delayed matching-to-sample task, with multiple intervening items between the sample and matching test stimulus. About half of the cells responded differently to a test stimulus if it matched the sample, and this difference was maintained following intervening stimuli. For most of the affected cells, the responses to matching stimuli were suppressed; however, for a few cells the opposite effect was seen. Temporal contiguity alone could not explain the results, as there was no modulation of responses when a stimulus on one trial was repeated on the next trial. Thus, an active reset mechanism appears to restrict the memory comparison to just the stimuli presented within a trial. The suppressive effects appear to be generated within or before IT cortex since the suppression of response to matching stimuli began almost immediately with the onset of the visual response. The memory of the sample stimulus affected not only the responses to matching stimuli but also those to nonmatching stimuli. There was suggestive evidence that the more similar a nonmatching stimulus to the sample, the more the response was suppressed. About a quarter of the cells showed stimulus-selective activity in the delay interval following the sample. However, this activity appeared to be eliminated by intervening stimuli. Thus, it is unlikely that delay-interval activity in IT contributed to the performance of this particular version of delayed matching to sample.

To determine how much information about the match–nonmatch status of the stimulus was conveyed by individual neurons, we analyzed the responses with discriminant analysis. The responses of an individual IT neuron could be used to classify a stimulus as matching or nonmatching on about

60% of the trials. To achieve the same performance as the animal would require averaging the responses of a minimum of 25 IT neurons. There was no evidence that mnemonic information was carried by temporal variations in the spike trains. By contrast, there was a modest amount of temporal variation in sensory responses to different visual stimuli. This variation appeared to be due to different stimuli having different effects on the early and late (transient and sustained) portions of the response. We propose that two populations of IT cells contribute to memory. One functions as adaptive mnemonic filters and the other provides a sensory referent. The difference in response between the two populations is a measure of the difference between the current stimulus and stored memory traces. A temporal “figure-ground” mechanism such as this could contribute to performance of a variety of mnemonic tasks, including working memory tasks.

[Key words: macaque monkey, memory, inferior temporal cortex]

Normal perception and memory of visual form depends on the neural circuitry in inferior temporal (IT) cortex (Gross, 1973; Mishkin, 1982). IT cortex receives information from prestriate visual areas, particularly area V4, and is directly or indirectly connected with other medial temporal and prefrontal structures important for memory (Van Hoesen and Pandya, 1975a,b; Aggleton et al., 1980; Desimone et al., 1980; Ungerleider and Mishkin, 1982; Amaral and Price, 1984; Insausti et al., 1987a,b; Baizer et al., 1991; Webster et al., 1991; Martin-Elkins and Horel, 1992). Congruous with its role in high-order processing and storage of visual information, IT cortex has neurons with large receptive fields and complex stimulus-selective properties (Desimone and Gross, 1979; Desimone et al., 1984).

Anterior IT cortex has been implicated in both long-term and short-term storage of visual information (Mishkin, 1982; Horel et al., 1987; Murray et al., 1989; Zola-Morgan et al., 1989, 1993; Meunier et al., 1990; Suzuki et al., 1991, 1993; Gaffan and Murray, 1992). Following deactivation of IT cortex, animals are impaired on short-term memory tasks, such as delayed match-to-sample (DMS) tasks with small stimulus sets (Fuster et al., 1981; Horel and Pytko, 1982; Horel et al., 1987). In DMS tasks of this type, memory of the sample on one trial may interfere with performance on subsequent trials. Thus, the task involves memory spanning just the length of the behavioral trial.

IT neurons have been studied during performance of DMS tasks, and their properties seem consistent with a role in storage of information. Some neurons show elevated activity during the retention interval on DMS tasks, as if they are actively holding

Received July 8, 1992; revised October 1, 1992; accepted October 7, 1992.

We are indebted to Mortimer Mishkin for his advice and support throughout the study and to Thomas White for writing many of the computer programs. We also thank Michelle Adams and John Sewell for help with histology, Verity Brown, Leonardo Chelazzi, Timothy Gawne, Andreas Lueschow, William Newsome, Barry Richmond, and Leslie Ungerleider for valuable comments on an earlier version of the manuscript, and Yadin Dudai for stimulating discussions of the results.

Correspondence should be addressed to Earl Miller, Laboratory of Neuropsychology, National Institute of Mental Health, Building 9, Room 1E104, Bethesda, MD 20892.

Copyright © 1993 Society for Neuroscience 0270-6474/93/131460-19\$05.00/0

memory of the stimulus (Fuster and Jervey, 1981; Miyashita and Chang, 1988; Fuster, 1990) and some respond differently to the same stimulus depending on the behavioral context (sample, match, or nonmatch) in which the stimulus appears (Gross et al., 1979; Mikami and Kubota, 1980; Baylis and Rolls, 1987; Vogels and Orban, 1990a; Miller et al., 1991b; Riches et al., 1991; Eskandar et al., 1992).

Behavioral tasks employed in the laboratory, though, can often fall short of the demands of the real world. Most investigations of the role of IT neurons in memory have employed tasks that require the monkey to retain memories across a "blank" retention interval, during which no visual stimuli are presented. Outside the laboratory, however, memories frequently must be retained while new information is being processed. It is not understood how IT neurons can participate in memory storage during times when they are concurrently responding to new visual stimuli. One study of neurons in lateral IT cortex of monkeys performing a serial recognition task found that responses were affected by stimulus recency, but this effect was lost if one or two other stimuli appeared in the retention interval (Baylis and Rolls, 1987). However, Riches et al. (1991) tested four cells with intervening stimuli in the anterior ventral portion of IT cortex and found that three of them responded differently to a stimulus if it had recently been seen before, even when more than two other stimuli intervened.

In the present study, we tested responses of neurons in the anterior ventral portion of IT cortex in monkeys that were holding a stimulus in memory while concurrently processing other stimuli. We describe several new experiments indicating that neuronal responses in this region carry mnemonic information that could be used in performance of a recency-memory task, even when several stimuli appear in the retention interval.

Some of these findings have been reported in a preliminary form (Miller et al., 1991b).

Materials and Methods

Subjects and surgical procedures

Two rhesus monkeys (*Macaca mulatta*) weighing 8–9 kg were used. Prior to the implantation of the recording chamber, the animals were placed in a plastic stereotaxic machine and scanned with magnetic resonance imaging (MRI). A head restraint post, recording chamber, and scleral eye coil for monitoring eye position (Robinson, 1963) were implanted under aseptic conditions while the animal was anesthetized with sodium pentobarbital. Using the coordinates derived from the MRI images, the recording chamber was centered on the dorsal surface of the skull above anterior-ventral IT and oriented in the Horsley–Clark stereotaxic planes. The animals received antibiotics and analgesics post-operatively.

Recording techniques

Activity was recorded using tungsten microelectrodes (Micro Probe, Rockville, MD). Following penetration of the dura with a guide tube, the electrode was advanced in the vertical plane down to the cortex on the ventral surface. Single-neuron activity was isolated from a multi-neuron signal using an online spike-sorting system (Signal Processing Systems, Prospect, Australia) that employs a template-matching algorithm. The spike-sorting system produced a pulse output that was stored as a spike event by our computer with a 1 msec time resolution. Two separate neurons were frequently studied simultaneously.

While the monkey performed the task, the electrode was advanced through IT cortex until the activity of one or two single neurons could be isolated. Responsiveness of the neuron(s) was assessed through use of an audio monitor. If a neuron exhibited a visual response (a change in activity during presentation of one or more of the stimuli) while the animal performed the behavioral task, data collection was initiated. If no visual response was detected, the electrode was advanced farther.

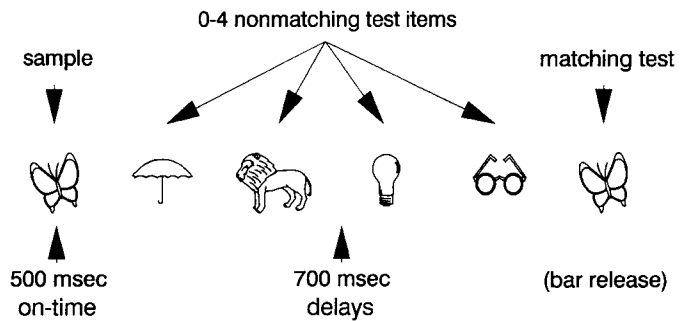


Figure 1. Outline of the DMS task. The matching test stimulus was terminated when the animal made its response, typically 350 msec after stimulus onset.

Behavioral task

The task (Fig. 1) was a modified version of delayed matching-to-sample (DMS), similar to one used by Maunsell et al. (1991) in their study of neurons in area V4. At the start of each trial the monkey grasped a metal bar after which a small (0.2°) fixation target appeared, which remained on throughout the trial. The monkey was then required to fixate the target for 300 msec before presentation of the first stimulus and to maintain fixation throughout the trial.

The first stimulus presented on each trial was the sample, which was followed by one to five test stimuli. When one of the test stimuli matched the sample stimulus, the animal was required to release the bar within 900 msec of stimulus onset to receive a drop of orange juice. Each stimulus was presented for 500 msec, and there was a 700 msec delay between the disappearance of one stimulus and the onset of the next. The match stimulus was the last stimulus presented on each trial, and it was extinguished as soon as the animal made its response, which was generally around 350 msec after stimulus onset. During the intertrial interval, there were no behavioral constraints on the animals. Each trial usually began with a saccade to the fixation target.

Stimuli

The stimulus set consisted of over 500 complex, two-dimensional, multicolored pictures presented on a computer graphics display. The pictures were digitized from magazine photographs and objects around the laboratory. Some were identifiable objects, such as faces or fruit, and others were colored textures and patterns. This stimulus set was chosen because the pictures seemed easily "memorable" to the experimenters, they spanned a range of two-dimensional and three-dimensional features known to be effective for producing stimulus-selective responses from IT neurons, and they represented the sort of objects that a person might actually hold in memory outside the laboratory. The stimuli ranged in size between 1° and 3° on a side and were presented at the center of gaze. For each recording session, six stimuli were chosen at random, and the animal was given 0.5–1 hr of experience with the stimuli in the behavioral task before data collection began, thus ensuring that the stimuli were familiar. Once a stimulus was used for a session it was not used in another session until the entire stimulus set had been exhausted, typically after 2 months. Each of the six stimuli appeared as a sample and match on some trials and as a nonmatch on other trials. Within a trial, the only stimulus that was repeated was the stimulus used as a sample and as a match; all the nonmatching stimuli differed from each other and from the sample.

Data analysis

For most analyses, responses were calculated over a 200 msec time interval beginning 75 msec following stimulus onset. The beginning of this time interval was chosen to coincide with the typical response latencies of IT neurons, and the end was chosen to occur before the animals' behavioral responses to the match stimulus (mean reaction time was 350 msec). For discriminant analysis (see below), the time interval used was slightly longer, 225 msec, so that it could be evenly divided into three 75 msec subintervals for the multivariate analysis described below.

Each trial with a correct behavioral response ended with a matching stimulus. Thus, there was a maximum number of three intervening stimuli preceding a given nonmatching stimulus but a maximum of four preceding a matching stimulus. Therefore, for all statistical and summary analyses of nonmatching and matching responses, the responses to the match stimulus on trials with four intervening stimuli were dropped so that the average numbers of intervening stimuli preceding matching and nonmatching stimuli were equal.

Statistical analyses. Responses to stimuli presented as samples and as matching and nonmatching test stimuli were appraised with both analysis of variance (ANOVA) and *t* tests, evaluated at the $p < 0.05$ level of significance. However, the fact that a response difference is statistically significant does not, in itself, indicate how potentially useful the difference is in discriminating among the stimuli. We were particularly interested in how well the responses of IT neurons distinguished whether a test stimulus matched the sample stimulus held in memory. That is, how well could one classify a stimulus as either matching or nonmatching on each trial of the task, if the only information available was the response of an individual IT neuron? To answer this question for a given cell, one needs to know the distributions for the two classes of responses, that is, matching and nonmatching. Given the distributions, one could design an optimal (Bayesian) classifier. Classification success depends on the amount of overlap of the two distributions of responses; the greater the overlap, the poorer the classification performance.

Because the underlying distributions cannot be measured directly, they must be estimated from the data. There are several different methods available for estimating distributions and constructing a classifier, each of which makes different assumptions about the distributions. In a previous study (Miller et al., 1991b), we used logistic regression. However, discriminant analysis often does better at discriminating among more than two classes and thus was employed in the present study. We also made use of a neural network classification technique. We applied the methods in a number of different ways and to different sorts of data; the details are given in the appropriate sections of the Results. Below we give a brief overview of the methods.

Discriminant analysis is typically used to classify individuals into two or more discrete classes based on one or more continuous or discrete variables. The analysis fits normal distributions to each of the different response classes (e.g., match and nonmatch) and estimates the probability that the two sets of responses came from the same distribution. Responses are classified into groups based on a decision rule. When there are two groups, the decision rule is typically whether the response falls above or below the midpoint between the means of the two distributions. Multivariate discriminant analysis is an extension of the univariate case, except that multivariate normal distributions are fit to the multiple response measures. In the multivariate case, the different response measures are weighted and combined, resulting in a "canonical variable" that can be used to discriminate between the groups. The number of canonical variables is equal to either the number of groups or the number of response measures, whichever is smaller. The first canonical variable accounts for most of the difference among the groups, followed by the second, third, and so on. The different canonical variables are uncorrelated and the weightings of the underlying response measures are optimal for distinguishing among the groups. The discriminant analysis was performed using the SYSTAT statistical package.

The function relating neuronal response to match versus nonmatch status of the stimulus was also estimated using a three-layer backpropagation neural network, which is a nonlinear optimization method. After "training," the network was presented with responses from the data set, and the network's prediction of class membership (matching or nonmatching) was compared with the actual class, as a measure of classification performance. An advantage of the neural network over conventional statistical methods is that it does not assume normally distributed data and can be trained to classify data from arbitrarily complex distributions, including ones in which class membership is a nonmonotonic function of the variable(s). Disadvantages are that network parameters, such as number of hidden units and rate of learning, are essentially empirically determined, with no assurance that any particular set of parameters is optimal.

When any statistical classification procedure is applied to the same data on which the parameters of the classification algorithm are fit, the classification success rate will be optimistically biased. The success of the classification procedure is inversely proportional to the overlap in the probability distributions of the response classes. The smaller the

number of data used in estimation, the less the overlap and the more the classification procedure can take advantage of chance differences between the response classes in achieving an artificially high success rate. The problem is most severe when classifiers are applied to multivariate data, in which the space over which the distribution is estimated is far larger (each variable adds a dimension to the space) and, therefore, the data points are sparser.

To correct for this bias we used the cross-validation method. The discriminant functions and neural network were first estimated, or trained, on half of the data (chosen randomly). The functions, or network weights, were then fixed and the classification procedure applied to the remaining half of the data. This gives an unbiased estimate of how successfully a new set of data could be classified.

Univariate and multivariate response measures. The univariate response measure was simply the firing rate in an interval beginning 75 msec after stimulus onset and ending 300 msec after onset. We used two different multivariate measures. For both, we first "smoothed" the spike train on each trial by convolving it with a Gaussian with a standard deviation of 10 msec (Richmond and Optican, 1987). For one of the multivariate measures, we calculated the first three principal components (PCs) of all of the smoothed spike trains, using the SYSTAT statistical program. For matching versus nonmatching responses, there were only two sources of variance due to stimulus effects. Therefore, to maximize the variance due to the stimuli, we calculated PCs on the spike trains for all six stimuli, matching and nonmatching, taken together. In fact, whether the components were calculated over all stimuli or just one seemed to make little difference in either the shape of the PCs or the ultimate classification success.

As a comparison to the principal component method, we divided the smooth spike train into three successive time intervals of 75 msec duration beginning 75 msec after stimulus onset. The firing rates during these "early," "middle," and "late" intervals were then used as three different response measures in the multivariate analyses.

Histology

At the conclusion of the experiments, fluorescent dyes were injected at the boundaries of the recording area, through a cannula inserted in place of the electrode. A few days later, the animals were killed with sodium pentobarbital and perfused transcardially with formalin. Sections were cut, stained with thionin, and examined for electrode tracks and the dyes. Because of the length of the study, most tracks could not be visualized, but their locations could be inferred from those of the dyes and other identified tracks.

Results

Anatomical location of penetrations

All of the recording sites were located on the inferior temporal gyrus, just lateral to the rhinal sulcus and medial to the anterior medial temporal sulcus. Figure 2 shows the location of recording sites in the monkeys along with representative coronal sections.

General properties

A total of 146 visually responsive neurons were recorded in the DMS task with up to three intervening stimuli. All of the neurons sampled gave excitatory responses. The lack of inhibitory responses may reflect a sampling bias due to selecting neurons for study by assessing responsiveness with an audio monitor. Because the performance of the animals was better than 90% correct in the task, there were too few error trials to analyze. Consequently, all analyses were performed using data from correctly performed trials.

To determine whether a response to a sample stimulus was significant, we compared the firing rate to the sample with the spontaneous (prestimulus) firing rate using a paired *t* test. Of the 846 stimuli used as a sample (146 neurons \times 6 stimuli each), 75% (635) elicited a visual response. To assess the stimulus selectivity of each neuron, we compared the responses to the

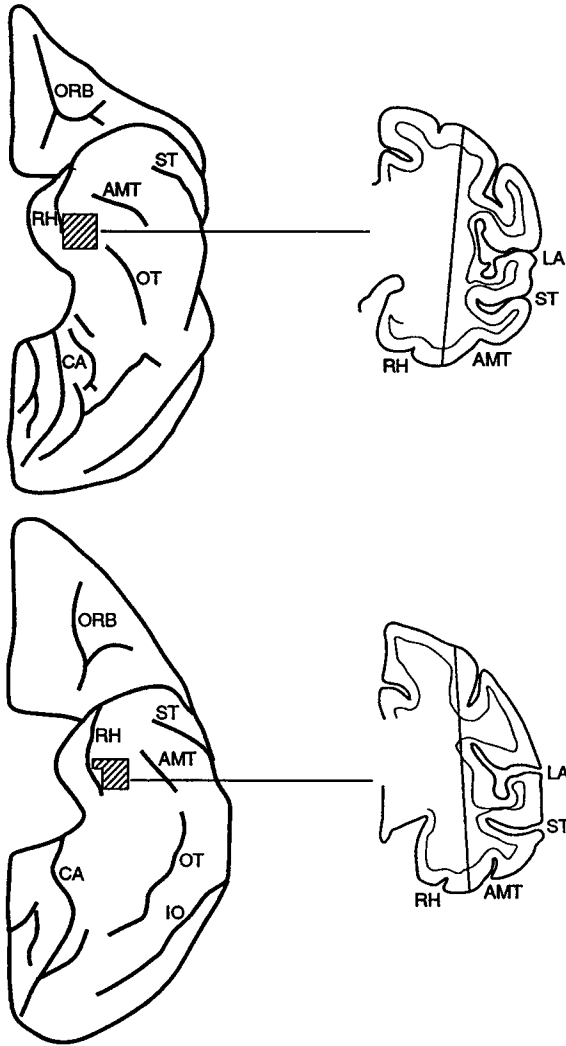


Figure 2. Location of recording sites. The shaded areas on the ventral views indicate the recording areas in the two animals, and the vertical lines through the coronal sections illustrate representative penetrations. AMT, anterior middle temporal sulcus; CA, calcarine sulcus; IO, inferior occipital sulcus; LA, lateral sulcus; ORB, orbital sulcus; OT, occipito-temporal sulcus; RH, rhinal sulcus; ST, superior temporal sulcus.

six sample stimuli separately for each neuron using an ANOVA. The ANOVA was significant ($p < 0.05$) for 85% of the neurons (124 of 146). Thus, most of the neurons were stimulus selective for the sample stimuli.

Effect of context on responses to stimuli

Many IT neurons responded differently to the same stimulus depending on whether it appeared as a sample, match, or non-match. Figure 3 shows the average response (across all intervening stimuli) to a single stimulus appearing as a sample, non-match, and match for one neuron. For this cell, the response was strongest when the stimulus appeared as a sample, intermediate when the stimulus appeared as a nonmatch, and weakest when the stimulus appeared as a match.

To study the neural information available for matching-nonmatching classification of the test stimuli, we conducted two general classes of analyses. The first focused on how many cells distinguished among matches and nonmatches, whereas the sec-




Figure 3. Average responses to stimuli appearing as samples, non-matches, and matches for one IT neuron. Histograms were computed by summing responses across all six stimuli, in 10 msec bins. The bar beneath each histogram shows the time of the stimulus presentation. The y-axis represents the firing rate in spikes per second.

ond focused on how much information was available in the response to individual stimuli.

To determine whether a cell distinguished between matching and nonmatching stimuli, we performed a two-way ANOVA on the responses to all six test stimuli, with stimulus and matching-nonmatching status as factors. Responses were averaged over all the intervening stimulus presentations. Based on the ANOVA performed on each cell, 48% (68 of 141) of the cells distinguished between matching and nonmatching stimuli. Nearly all of these (92%, or 62 of 68) responded better to non-matching than matching stimuli. In the Discussion, we consider

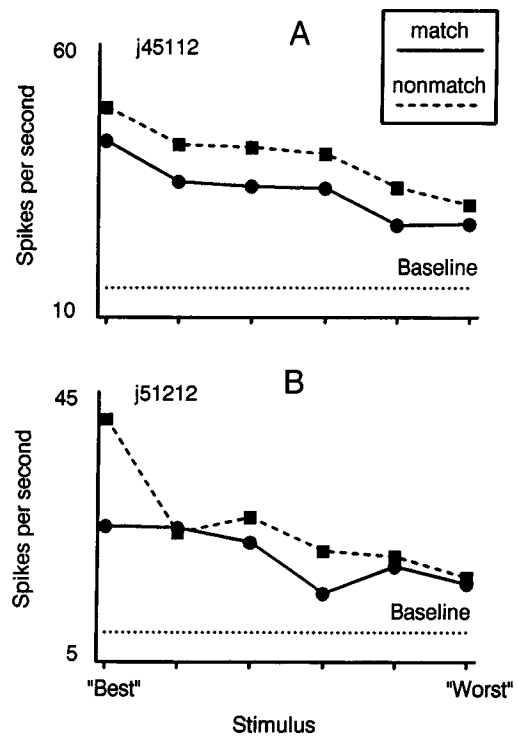


Figure 4. Examples of two IT neurons that showed different match-nonmatch effects. Six stimuli are represented along the x-axis, ranked from best to worst. Dotted line shows the neuron's baseline or spontaneous firing rate. A, Cell for which the match-nonmatch effect was proportional to the magnitude of the response. B, Cell for which the effect was disproportionately large for the best stimulus.

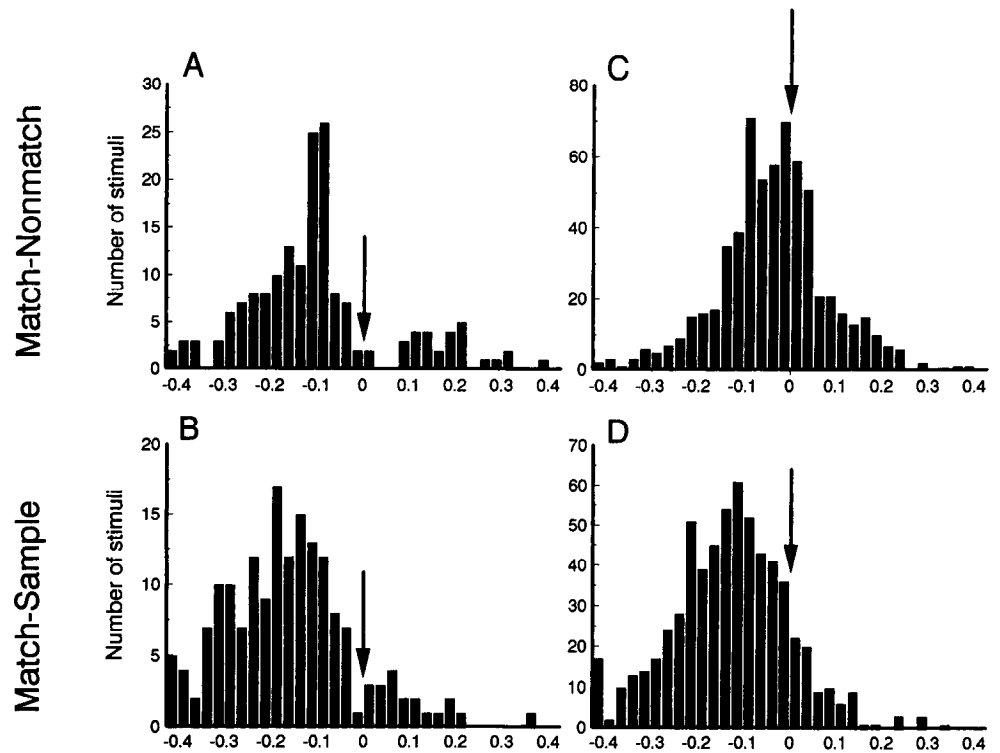


Figure 5. Distribution of indices showing the strength of the response to matching stimuli relative to the response to the same stimulus as a non-match (*A* and *C*) and as a sample (*B* and *D*). Values below zero (arrow) indicate weaker responses when the stimulus appeared as a match. Distributions *A* and *B* show the indices for the stimuli for which there was a significant match–nonmatch effect, while distributions *C* and *D* show the indices for the entire population of stimuli that elicited a visual response.

the role that the half of the cells that did not distinguish matching and nonmatching might play in a memory mechanism.

Of 25 pairs of cells recorded simultaneously with the same electrode in IT, both cells were responsive in only 12 pairs. In five of the 12 pairs, both cells showed significant match–nonmatch effects, in four pairs neither cell showed significant effects, and in the three remaining pairs just one cell showed a significant effect.

Based on the two-way ANOVA, almost all of the cells (94%, or 64 of 68) that exhibited a match–nonmatch difference were also stimulus selective (significant effect of stimulus, $p < 0.05$). For most of these cells (67%, or 43 of 64) there was no significant interaction between stimulus and matching–nonmatching factors, that is, the magnitude of the match–nonmatch effect was about equal for all stimuli. Figure 4*A* shows the responses of one such neuron. The stimuli were rank ordered from “best” to “worst” based on the magnitude of response to the same stimulus appearing as a sample. The lines corresponding to the average match and nonmatch responses are roughly parallel, indicating that the magnitude of the effect is similar for all six stimuli. Thus, for this type of cell, the suppression of the match response was proportional to response magnitude.

For the other cells that showed a match–nonmatch difference (33%, or 21 of 64), the magnitude of the difference was disproportionately larger for some stimuli (i.e., there was a significant interaction between matching–nonmatching and stimulus factors). Figure 4*B* shows an example of such a cell. For this cell, there is a larger difference between match and nonmatch responses for the best stimulus of the set compared to the other five stimuli. Indeed, for most of the cells that exhibited interactions between match–nonmatch and stimulus factors (71%, or 15 of 21), the match–nonmatch effect was disproportionately largest for the best stimulus of the set.

Having determined that nearly half of the cells gave a significantly different response to all matching and nonmatching stimuli taken together, we then asked for how many stimuli was the effect significant. Based on a *t* test applied to each stimulus (evaluated at $p < 0.05$), about one-fourth (27%, 171 of 635) of the effective stimuli showed significant match–nonmatch response differences. The stimuli with significant effects derived from 99 cells, of which 68 showed significant match–nonmatch effects across all stimuli based on the omnibus ANOVA described above. (Although the *per-comparison* type I error rate is 0.05 for both the *t* test and the ANOVA, the cumulative *per-cell* error rate is higher with the *t* test; thus, 68 is the more conservative estimate of the number of significant cells.)

To quantify the magnitude of the match–nonmatch effect, we computed an index for each stimulus by subtracting the mean match response from the mean nonmatch response and dividing the result by the sum of the two means. Figure 5*A* shows the distribution of indices for the stimuli with significant match–nonmatch effects. Values below zero indicate match responses weaker than nonmatch, which was true for 83% (142 of 171) of the stimuli. Only 17% (29 of 171) of stimuli showed the opposite effect (match greater than nonmatch), and these were largely from a separate population of neurons: only four cells showed mixed effects, that is, significant match suppression for one or more stimuli and match enhancement for one or more other stimuli.

Figure 5*B* shows response indices for the same stimuli as in Figure 4*A*, but in this case the comparison was made between the match and sample responses. Values below zero indicate responses that were weaker to matching test stimuli than to samples. In Figure 5, *C* and *D* show the distribution of the match–nonmatch and match–sample indices for the entire population of effective stimuli. The means of both distributions are

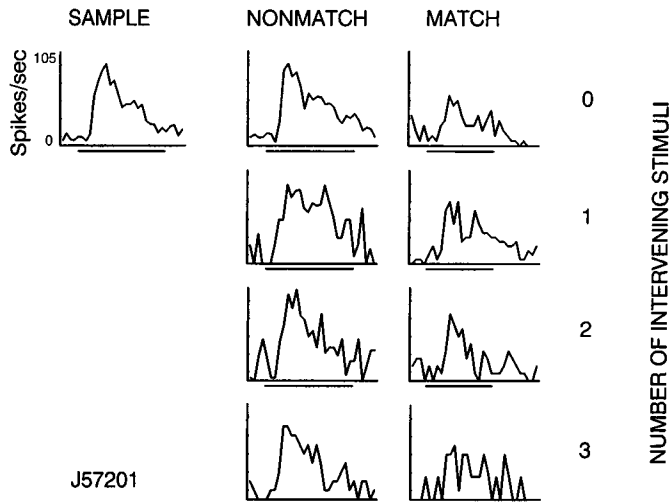


Figure 6. Responses of an IT neuron to a single stimulus appearing either as a sample or as a match or nonmatch item following different numbers of intervening stimuli. Bar beneath each histogram indicates when the stimulus was on.

shifted below zero, indicating weaker responses to match stimuli than to either sample or nonmatching stimuli. However, the match-sample distribution (Fig. 5D) shows a larger negative shift than the match-nonmatch distribution (Fig. 5C). Thus, neuronal responses to samples were typically strongest, followed by responses to nonmatching and then to match stimuli.

Retention of sample information across intervening stimuli. Not only could IT neurons detect whether a test stimulus was a match or a nonmatch, but they could do so even when several stimuli intervened in the retention interval. The responses of such a neuron are shown in Figure 6. Average responses of the population of IT neurons to samples and to test stimuli following intervening stimuli are shown in Figure 7. Figure 7A shows responses to all of the stimuli that had a significant match-nonmatch effect when the responses were averaged over all intervening stimuli. All responses were included in the average, regardless of whether the match or nonmatch responses were larger. The figure shows that the responses to the sample were largest, followed by the nonmatching and matching stimuli, respectively. Based on a paired *t* test on the population, there was a significant difference between match and nonmatch responses even with three intervening stimuli ($t = 3.32, p = 0.0013$). For comparison, Figure 7B shows the average responses to sample and test stimuli for the entire population of effective stimuli, that is, regardless of whether there was a statistically significant match-nonmatch response difference. Even when stimuli without significant effects are included in the average, match responses are smaller than nonmatch, and this difference is maintained after two intervening stimuli ($t = 8.349, p < 0.0001$).

To test further the “memory span” of the cells, an additional group of 18 cells in one monkey was tested with up to five intervening stimuli before the final matching stimulus. The animal required additional training to perform this version of the task. Of the cells tested, 44% showed significant overall match-nonmatch effects according to an ANOVA ($p < 0.05$). When each stimulus was tested for significant match-nonmatch effects separately with a *t* test, 28% showed significant effects. The average response to these stimuli is shown in Figure 7C. Based

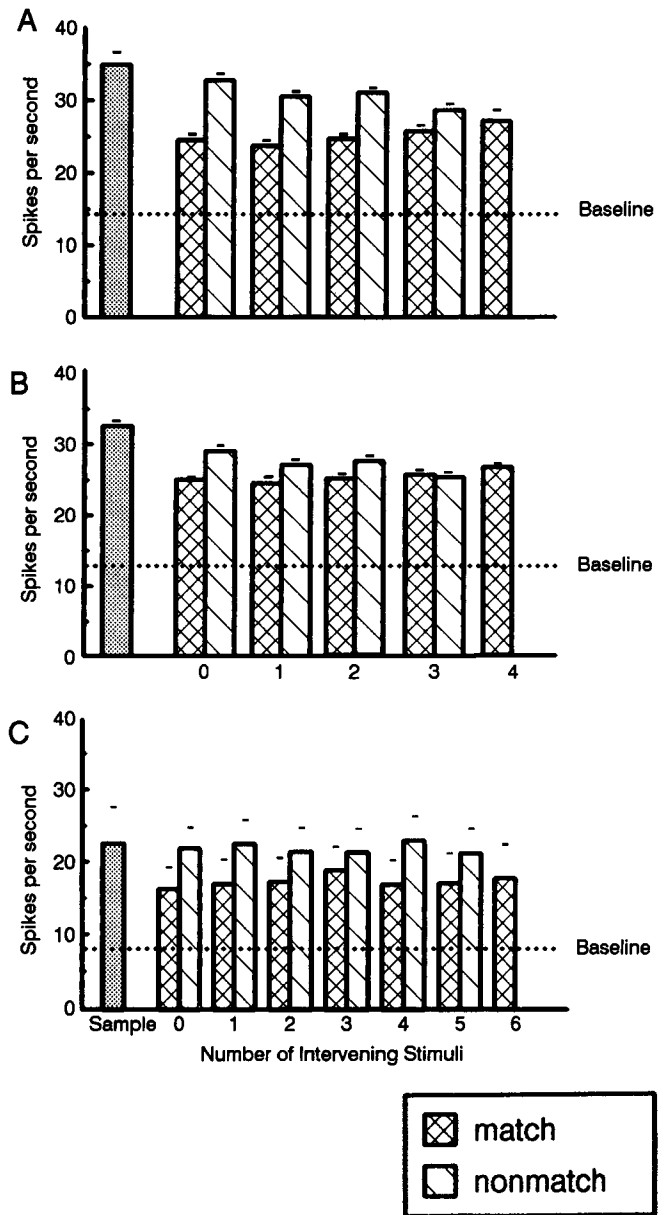


Figure 7. Average responses across cells to the same set of stimuli appearing as samples and as matches and nonmatches following different numbers of intervening stimuli. The *Small horizontal line* above each bar indicates the SEM. *A*, Average responses for the population of stimuli that elicited a significant match-nonmatch effect. *B*, Average responses for the population of stimuli that elicited a visual response. *C*, Average responses for the population of stimuli that elicited a significant match-nonmatch effect and were tested with up to five intervening stimuli before the final match.

on a paired *t* test of the population response, there was a significant difference between the match and nonmatch responses even with five intervening stimuli ($t = 2.33, p = 0.02$). Because there were only a few cells tested in this group, the remainder of the Results will be devoted to the larger group of cells tested with up to three intervening stimuli before the final match.

As indicated earlier in the Results, some (17%) of the stimuli with significant match-nonmatch effects actually elicited larger matching than nonmatching responses. Figure 8 shows the av-

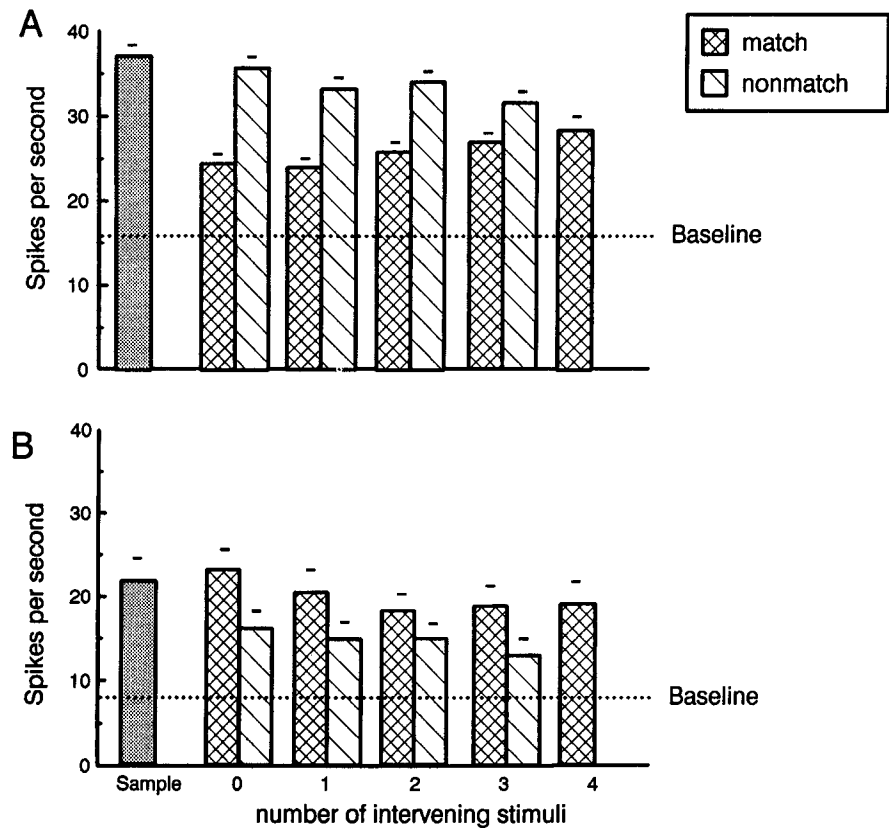


Figure 8. *A*, Average responses for stimuli for which match responses were weaker than nonmatch responses. *B*, Average responses for which match responses were stronger than nonmatch responses. The line above each bar indicates the SEM.

verage sample and test stimulus responses for the same stimuli included in Figure 7*A*, except that the averages are given separately for stimuli showing nonmatch responses greater than match (Fig. 8*A*) and stimuli showing the opposite effect (Fig. 8*B*). For the stimuli that elicited stronger match responses (Fig. 8*B*), the sample and match responses were similar and both were larger than the nonmatch responses; that is, nonmatch responses appeared to be suppressed (compared to the sample) rather than the match responses being enhanced. As was the case for stimuli with larger nonmatch responses, the match–nonmatch effect was maintained across all intervening stimuli (paired *t* test, $p < 0.01$).

The difference in response to matching and nonmatching stimuli tended to decrease with increasing numbers of intervening stimuli. Furthermore, the animals' ability to perform the task showed a corresponding decrease. Figure 9 shows the difference in response to matching and nonmatching stimuli plotted along with the animals' percentage of correct performance. Both the response difference and the animals' ability to perform the task decreased with increasing numbers of intervening stimuli.

To determine whether the memory of stimuli from one behavioral trial could carry across the intertrial interval to modulate the responses on subsequent trials, we examined the responses to sample stimuli when the same sample stimulus had been used on successive trials. We included in the task several instances in which one trial with a given stimulus as a sample and a match was followed, after a 1–2 sec intertrial interval, by another trial with the same stimulus as a sample. If close temporal contiguity was all that was important for producing the effect, then we would expect the responses to sample stimuli to

be reduced when they followed a trial with the same stimulus as a sample. We compared the responses to sample stimuli when the same sample had appeared on the previous trial to the responses when a different sample had appeared on the previous trial. This comparison was made for the neurons that exhibited a matching–nonmatching effect based on the two-way ANOVA. The mean response across cells to sample stimuli that were the same as the sample on the previous trial (32.08 spikes/sec) was not significantly different from the mean response to sample stimuli that were different than the sample on the previous trial (32.36 spikes/sec), according to a paired *t* test ($t = 0.51$, $p = 0.61$). Thus, the match–nonmatch effect did not appear to bridge the 1–2 sec intertrial interval even though the within-trial effect was observed when up to 5 sec elapsed between the sample and the matching stimulus. This rules out temporal contiguity alone as the explanation of the sample–match effect and, by extension, the match–nonmatch effect (see Discussion).

Time course of the match–nonmatch effect. To examine the time course of the match–nonmatch effect across the population of cells, we computed average histograms of the match and nonmatch responses across all stimuli that elicited a significant suppression of matching responses. Figure 10 shows the average response to matching and nonmatching stimuli as well as a plot of the difference between the histograms. The match and nonmatch histograms have the same shape but show a shift in absolute magnitude, suggesting that the mechanism underlying match suppression works in a sustained fashion throughout the response. This conclusion is further supported by an analysis of temporal variation in response, presented later in the Results.

To determine how soon after the onset of the visual response the suppression began, each 10 msec bin of the nonmatch re-

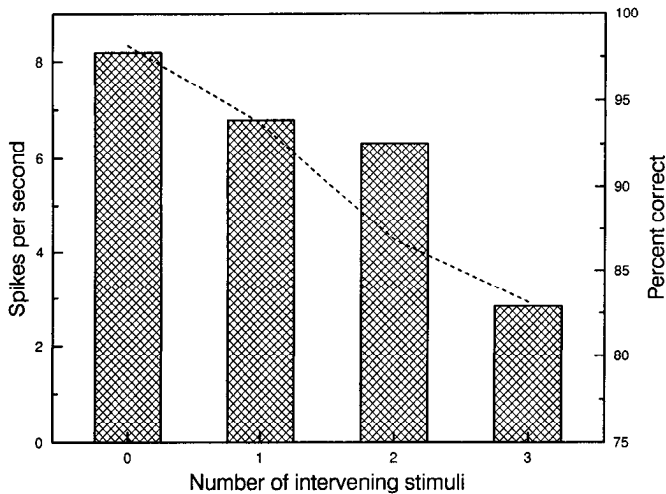


Figure 9. Average difference in response to matching and nonmatching stimuli (bars) following different numbers of intervening stimuli compared with percentage of correct performance by the animals on corresponding trials (line).

sponse histogram was first tested against the average prestimulus activity using a *t* test. The first significant response occurred 70–80 msec after stimulus onset. The match and nonmatch histograms were then compared bin by bin using a paired *t* test, to determine when the two histograms became significantly different. The difference was near significance in the first bin showing a significant visual response ($p = 0.08$), and remained highly significant ($p < 0.01$) for the remainder of the visual response. Thus, sustained suppression of the match stimulus responses occurred within 10 msec of the onset of the visual response, that is, with virtually no additional latency, suggesting that the suppression originates either within or before IT cortex.

Control for behavioral response and expectation of reward. An additional 14 cells were tested on a modified version of the task in which the animal was only required to maintain fixation. The response bar was removed from the animal's chair, and the animal was rewarded simply for fixating throughout the trial. The temporal order of the matching and nonmatching stimuli within the trial was also varied so that the match did not always appear at the end of the trial and thus did not predict the reward. Under these conditions, differential responses to match and nonmatch stimuli were still observed: 32% (21 of 66) of the effective stimuli exhibited a significant matching–nonmatching effect. Thus, the attenuation of the match responses was not related to the behavioral response (bar release) or to the fact that the match predicted the reward.

Neural classification of test stimuli as matching or nonmatching. The above analyses demonstrated that, for many stimuli, the responses of IT neurons differed significantly depending on whether the current stimulus matched the sample stimulus held in memory. To determine how useful this response difference would be in actually classifying a stimulus as match or nonmatch and to relate the neuronal performance to that of the monkey, we applied discriminant analysis to the match–nonmatch data (see Materials and Methods).

The discriminant analysis was calculated for each cell and each stimulus separately. The analysis gave us a decision rule to classify each stimulus as a match or nonmatch, based on the strength of response to the stimulus on each trial. This predicted




Figure 10. Population average histograms for matching and nonmatching stimuli. The difference line plots the difference between the two histograms. Bin width, 10 msec.

class could then be compared with the actual class (match or nonmatch) on each trial. To control for bias in the apparent predictive success, we tested the prediction by cross-validation; that is, we estimated the parameters of the discriminant function on half of the data (the training set) and applied the function to the other half (the test set).

On the average, 41 match responses and 85 nonmatch responses were available for each stimulus. Figure 11A shows the distribution of the classification scores for the full data set (i.e., without cross-validation) for the 173 stimuli for which the discriminant analysis was significant ($p < 0.05$). (The significant stimuli were virtually identical to those identified with the *t* test described earlier, as both methods evaluate significance in the same way.) The mean successful classification rate was 60.3%, with a range of 45–75.5% (compared to chance, 50%). The mean square root of the Mahalanobis distance between the two distributions (analogous to d' in signal detection theory) was 0.9. When the same data were split into two sets (i.e., cross-validated), the mean classification rate on the test set was similar (60.1%), which indicates that the discriminant model provided a good fit of the data and was not strongly biased. For comparison, Figure 11B shows the distribution of classification scores on the full data set for the entire population of effective stimuli (i.e., including stimuli that did not elicit a significant match–nonmatch response difference). The mean was 54.8% with a range of 37–76% (compared to chance, 50%).

As a comparison to the discriminant analysis method, we also analyzed the data using the receiver operating characteristic (ROC) approach, a psychophysical method that has also been applied to neural data (Bradley et al., 1987; Vogels and Orban, 1990b; Britten et al., 1992). For each stimulus that showed a significant match–nonmatch effect based on the *t* test, we plotted the hit and false alarm rates for several different firing rate criteria. We fit an ROC curve to the hit and false alarm rates and measured the area under the curve, which is a measure of the percentage of correct performance of an unbiased observer. The area under the curve corresponded to a mean successful classification rate of 64% correct, which was somewhat better performance than that predicted from the discriminant analysis. However, the ROC performance was not corrected for opti-

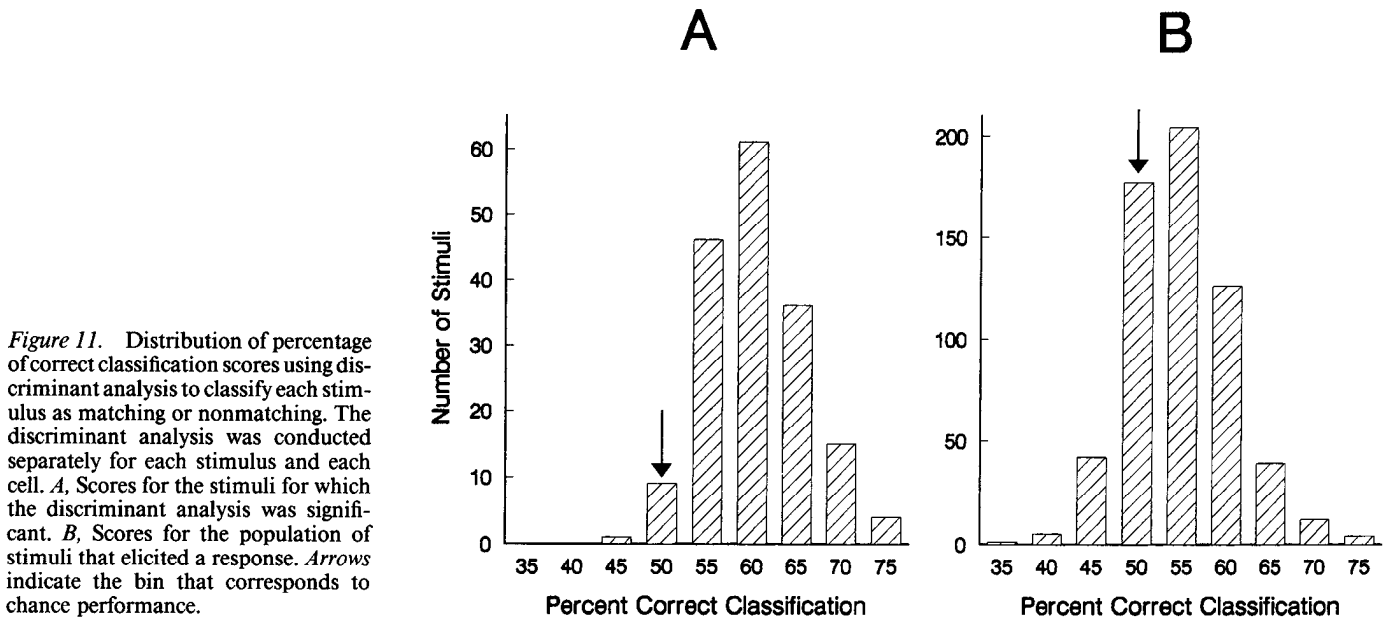


Figure 11. Distribution of percentage of correct classification scores using discriminant analysis to classify each stimulus as matching or nonmatching. The discriminant analysis was conducted separately for each stimulus and each cell. *A*, Scores for the stimuli for which the discriminant analysis was significant. *B*, Scores for the population of stimuli that elicited a response. Arrows indicate the bin that corresponds to chance performance.

mistic bias. Although chance performance should equal 50% with very large sample sizes, it will tend to be better than 50% as the sample size is decreased. As a rough measure of bias in the estimate, we recomputed the ROC curve after randomizing the data, that is, randomly reassigning the responses to the match or nonmatch classes. The mean area under the curve corresponded to a classification rate of 54% correct, indicating that up to 4% of the success rate was due to chance variations in the means of the match and nonmatch distributions. Subtracting 4% from the success rate derived from the unrandomized data gave a 60% successful classification rate, which was the same

as the performance of the discriminant analysis. Again, the similarity of the ROC results to those from the discriminant analysis indicates that the discriminant model provided a good fit to the data.

Effects of sample memory on nonmatch responses

The finding that the responses of IT neurons to a current stimulus were suppressed if the stimulus matched the stimulus held in memory raised the question of whether *all* responses were influenced by the memory of the sample. That is, perhaps even the responses to the nonmatching stimuli were differentially affected by the different sample stimuli used on a given trial. To determine whether the responses to a nonmatch stimulus varied significantly depending on which of the five sample stimuli it followed, we computed for each cell an ANOVA on the responses to the nonmatching stimuli using sample as the factor. Most of the cells (45 of 68, or 67%) that showed a significant match–nonmatch effect (based on the two-way ANOVA described previously) also showed significant modulation of the nonmatch responses by the sample stimulus ($p < 0.05$). Modulation of nonmatch responses by sample stimuli has also been reported by Eskandar et al. (1992).

We used discriminant analysis to determine how much information about the sample was carried in the neuronal responses to the nonmatching test stimuli. In this case, the discriminant analysis classified which of the five samples (i.e., excluding the one sample that matched the stimulus) appeared on each trial, on the basis of the average firing rate for each nonmatching stimulus. There were enough data to compute the discriminant analysis for 408 stimuli. On the average, there were 16 responses for each test stimulus available for the analysis. The discriminant analysis was significant for 24% (98 of 408) of the stimuli, and the classification success rates for these stimuli are shown in Figure 12. The mean classification rate was 30.2% (chance = 20%). There were not enough data to divide them into training and test sets for cross-validation.

Most often, responses to stimuli that matched the sample in memory appeared to be suppressed. Moreover, even the responses to nonmatching stimuli were influenced by the partic-




Figure 12. Distribution of percentage of correct classification scores using discriminant analysis to classify the sample stimulus based on the response to a nonmatching test stimulus. Discriminant analysis was conducted separately for each cell and each nonmatching test stimulus. Arrow indicates the bin that corresponds to chance performance.

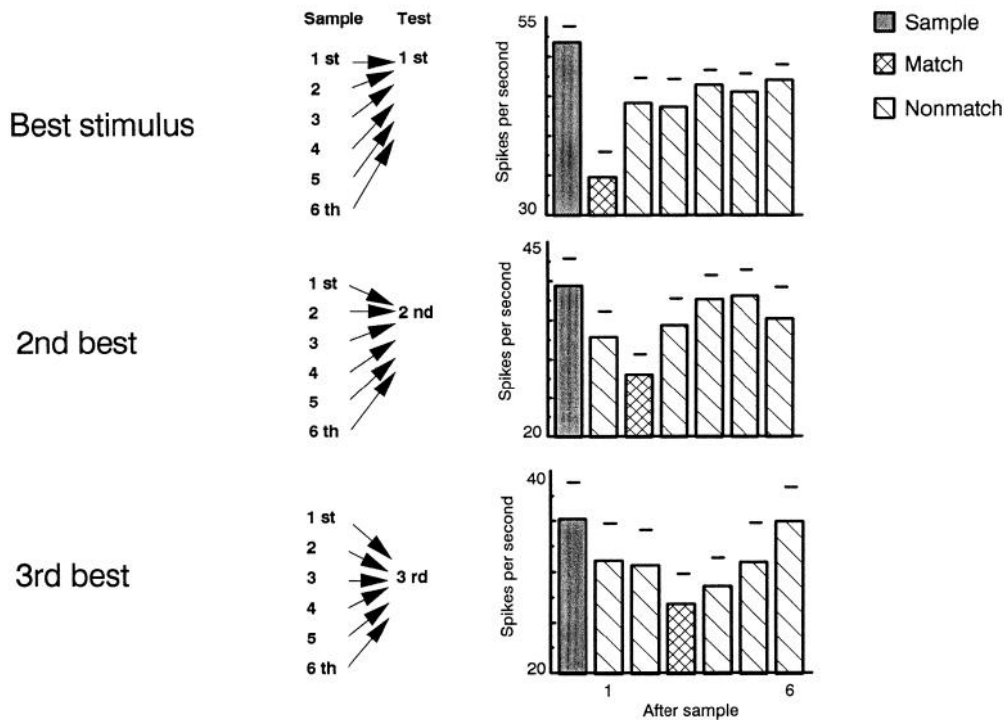


Figure 13. Average responses across cells to the first, second, and third ranked stimuli appearing as samples and as test stimuli following each of the (ranked) samples. The leftmost bar on each graph shows the average response to the stimulus when it had appeared as a sample, and the other bars show the average responses to the stimuli when they had appeared as test stimuli. Position along the x-axis indicates the ranking of the sample stimulus preceding the test stimulus. The cross-hatched bar indicates the response to a test stimulus preceded by itself as a sample (match trial). Test responses are averaged across all intervening stimuli.

ular stimulus held in memory. These two facts suggested that the responses to particular nonmatch stimuli were suppressed according to how *similar* they were to the sample. If so, this could explain the modulation of the nonmatch responses by the memory of the sample. Unfortunately, the stimuli used were complex and did not lend themselves to any objective measure of similarity to one another. Therefore, to test this hypothesis, we used the cells' relative preferences for the stimuli as a crude measure of similarity.

For each cell, we rank ordered the six stimuli from "best" to "worst" according to the magnitude of response to the stimuli appearing as samples. We assumed that similar stimuli would have similar rankings (i.e., produce similar responses), but this assumption would be invalid if similarity according to a given cell's response differed from similarity according to the visual system as a whole. For example, red and green stimuli might have similar rankings based on the responses of a cell nonselective for color. With this qualification in mind, we then examined the average response of the best, second-best, and third-best test stimulus for the population of cells that exhibited match-nonmatch effects based on the two-way ANOVA described earlier in the Results. We examined the responses to these test stimuli following each of the sample stimuli. Figure 13 shows the average responses for all comparisons.

If responses to the nonmatch stimuli were not determined by similarity to the sample (or if our stimulus rankings were an inadequate measure of similarity), the average responses to all nonmatching stimuli should have been the same. However, we found a weak tendency for responses to be smaller when a nonmatch stimulus followed samples that were more similar to it (i.e., closer in rank order) than when the stimuli followed samples that were less similar to it (i.e., more distant in rank order), although this tendency did not reach significance based on an ANOVA of the population data. For example, the average response to the third-ranked test stimulus (Fig. 13, bottom graph) was weaker when it followed either the second- or the fourth-

ranked sample than when it followed the sixth-ranked sample. In all comparisons, the responses to the match stimulus were weakest, confirming the match-suppression effect. Thus, there is suggestive evidence that stimulus similarity explains the effect of the sample on responses to the nonmatching test stimuli. The weakness of the effect may be due to the inadequacy of our ranking measure and, thus, the hypothesis should be tested using ordered stimulus sets, such as stimuli varying in color, orientation, or size.

Memory versus sensory classification ability

The suggestive evidence that the degree of modulation of responses of IT neurons was related to the relative similarity between the sample and test stimuli raised the possibility that there was also a relationship between a neuron's mnemonic and sensory abilities (e.g., between its ability to distinguish whether test stimulus A was preceded by sample A or sample B and its ability to distinguish between A and B themselves). We tested this relationship by computing two discriminant analyses for each cell: (1) a cell's ability to distinguish the best from worst sample stimulus (sensory classification), and (2) the cell's ability to distinguish between the best test stimulus preceded by itself as the sample from the same test stimulus preceded by the worst stimulus as a sample (memory classification). The discriminant analyses were performed on the 65 cells that exhibited both stimulus selectivity and a match-nonmatch effect based on the two-way ANOVA.

Figure 14 shows a plot of each neuron's sensory classification score against the memory classification score. There was no relationship ($r^2 = 0.038$) between the mnemonic and sensory abilities. This could be explained if either the mnemonic ability of cells is an independent random variable, or the detection of matching test stimuli is performed by the network in which a given IT neuron is embedded and not any individual cell on its own.

The diagonal line in the graph of Figure 14 represents equal




Figure 14. Memory classification ability of IT neurons versus sensory classification ability. Data are from 65 neurons that were stimulus selective and had a significant match–nonmatch effect. *Diagonal* indicates equal performance.

performance on the memory and sensory classifications. The population is clearly biased toward better performance on the sensory classification. The sensory classification scores (mean = 74.8% correct) were significantly better than the memory classification scores (mean = 62.3% correct) for the same stimuli, according to a paired *t* test ($t = 6.37$, $p < 0.0001$). Thus, the responses of IT neurons carry more sensory than mnemonic information. This conclusion was also reached by Eskandar et al. (1992).

Activity in the delay intervals

In the delay interval immediately following the sample, some IT neurons exhibited differential activity depending on which sample had appeared. We computed an ANOVA of the firing rate in the first delay interval following the six samples. To ensure that this activity was not influenced by responses to the offset of the sample, we used the firing rate in the last 200 msec of the 700 msec delay. Based on this ANOVA, 25% (37 of 146) of the cells exhibited differential firing rates in the delay interval depending on which stimulus had been used as a sample. Sample-selective delay activity by IT neurons has been hypothesized to reflect retention of information about the sample stimulus (Fuster and Jervey, 1981; Miyashita and Chang, 1988; Fuster, 1990).

If neuronal activity in the delay intervals is a retention mechanism for the samples, then this activity should not be unduly influenced by stimuli that intervene between the sample and the final match stimulus on each trial. To test this, we examined the activity in the delay interval following each nonmatching test stimulus.

Figure 15 shows the average delay activity following the best and worst samples for the cells that showed significantly different delay activity according to the ANOVA. In the delay interval immediately following the sample stimulus, the average level of activity across cells following the best sample is significantly greater than the activity following the worst sample ($t = 5.17$,

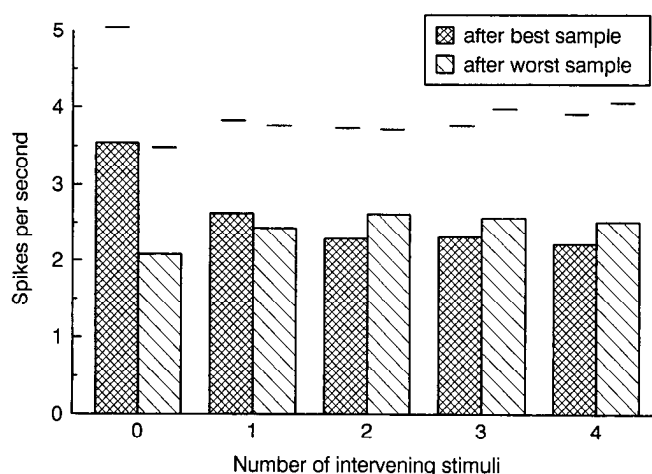


Figure 15. Average activity in the delay intervals when the “best” stimulus had been used as the sample and when the “worst” stimulus had been used as the sample. *Number of intervening stimuli* refers to the number of intervening nonmatching stimuli that had appeared before that delay. The averages are based on the 37 cells that showed significant differences in the first delay (0 intervening stimuli), depending on which sample had been used (ANOVA, $p < 0.05$). *Line* above each bar indicates the SEM.

$p < 0.0001$), confirming the results of prior studies (Fuster and Jervey, 1981; Miyashita and Chang, 1988; Fuster, 1990). However, after the first intervening (nonmatch) stimulus, this difference is no longer significant ($t = 1.48$, $p = 0.1482$). Following the remaining intervening stimuli, the relationship is actually reversed: the average delay activity following the worst sample is greater than the average activity following the best sample, although these differences did not reach significance. Thus, introduction of the intervening nonmatching test stimuli appears to have eliminated any differential delay interval activity due to the particular sample. The reversal of the effect following two to four intervening stimuli can be explained if the activity in the delay is determined by the immediately preceding stimulus rather than the memory of the sample. According to our task design, the best stimulus could not appear as an intervening stimulus if it had appeared as the sample on a given trial. Consequently, the intervening nonmatch stimuli following the worst sample sometimes included the best stimulus. This best stimulus may have accounted for higher levels of activity in the delays following the worst sample.

To determine the relative amount of information carried in the delay intervals, we computed two discriminant analyses. For both, the independent variable was the level of activity in the delay following the first intervening stimulus. We used this activity to predict (1) which sample stimulus had appeared at the start of the trial, and (2) which test stimulus had appeared immediately preceding that delay. We examined the percentage of correct classifications on the (cross-validated) test set for each analysis for the cells for which the ANOVA on the sample delay was significant. The mean classification success (21.56%) was above chance (16.7%) in predicting the immediately preceding nonmatch stimulus but at chance (16.89%) in predicting the sample stimulus. Thus, in our behavioral paradigm, the delay activity appears to carry information about the immediately preceding stimulus rather than about the sample throughout the trial.

Multivariate analysis of responses

Classification of matching and nonmatching stimuli. It was recently reported that IT neurons recorded in a DMS paradigm (without intervening stimuli between the sample and test stimuli) distinguished between matching and nonmatching stimuli at least in part on the basis of a temporal code, as measured by the first three PCs of the spike trains (Eskandar et al., 1992). To test this in our task, we performed a multivariate discriminant analysis on the values of the first three PCs of the spike trains for matching and nonmatching stimuli. For comparison, we also performed a multivariate discriminant analysis on the average firing rate during three successive 75 msec time intervals beginning 75 msec following onset of the test stimulus. As in the univariate case, the discriminant analysis was calculated separately for each stimulus. To ensure that the comparisons were not biased toward better performance on any type of model (temporal or rate coding), the comparisons were made for any response that showed a significant effect for *any* discriminant analysis: a significant match–nonmatch effect on the univariate discriminant analysis or either of the two multivariate discriminant analyses. The discriminant analysis was significant for the responses to 172, 155, and 159 stimuli for the average firing rate, PCs, and three intervals, respectively, resulting in a combined pool of 201 stimuli. Thus, there was a large degree of overlap in which stimuli showed a significant effect based on the three response measures.

The PCs of the responses to matching and nonmatching stimuli resembled those described by Richmond and Optican (1987). The first component appeared largely to reflect the DC level, or average firing rate over the interval. The second component typically had one zero crossing, usually with an initial positive period followed by a negative period. The third component typically had two zero crossings.

The top panels of Figure 16 show the classification scores based on the PCs plotted against the scores based on average firing rate over the entire response interval. The diagonal lines indicate equal performance. The top left panel shows the scores obtained from the entire data set. The values tended to cluster around the diagonal line, but more scores fall below the diagonal, indicating better classification (i.e., more information) from the PCs. Indeed, a paired *t* test showed that the classification scores based on the PCs (mean = 62.03%) were significantly better than the scores based on the average firing rate (mean = 59.28%, $p < 0.0001$). However, these results do not take into account the optimistic bias of the classification scores resulting from small data sets, a problem known to be more severe for multivariate analyses than for univariate. The magnitude of the bias is suggested by a comparison of univariate and multivariate discriminant analysis applied to randomized data. For each stimulus, we randomly reassigned both the firing rates and PC values for matching and nonmatching stimuli. Following the randomization, we recomputed the discriminant analysis and applied it to the training data. The average classification success for the randomized data across the population was 56.6% for the PCs and 53.2% for the average firing rate, which was a significant difference based on a paired *t* test ($t = 20.36$, $p < 0.0001$). Thus, the multivariate analysis apparently did better at capitalizing on random variations in the responses. Performing the cross validation test on the randomized data completely eliminated the difference in performance. When the discriminant analysis was calculated on half of the randomized data and

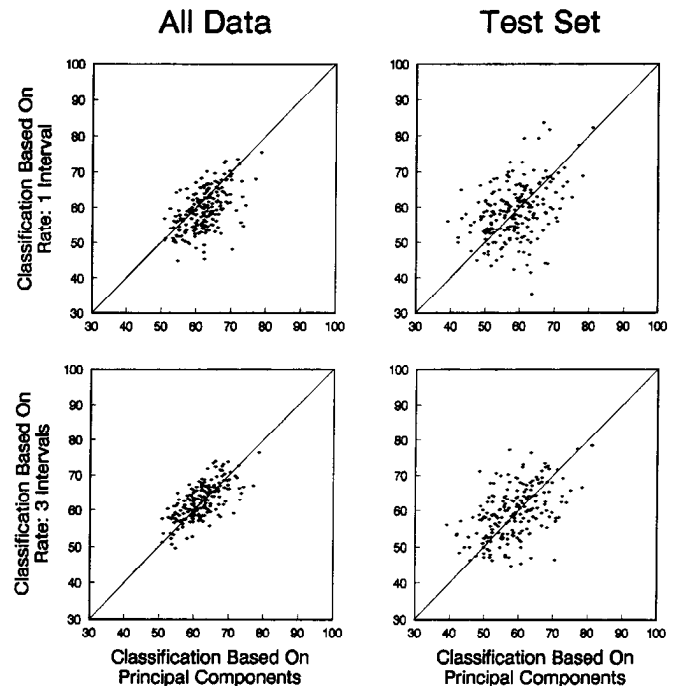


Figure 16. Performance of discriminant analysis applied to classification of stimuli as match or nonmatch, based on either average firing rate over the entire response interval, three PCs, or average firing rate during three subintervals of the response. Each point on each plot represents the performance on one stimulus for one neuron. Plots on the left were derived from applying the discriminant analysis to the entire data set for each stimulus. Plots on the right are derived from computing the discriminant analysis on half the data and applying it to the other half (test set). Diagonal line in each graph represents equal performance.

used to classify the other half, performance for the multivariate PCs and the univariate firing rate was equal, at 50%.

The top right panel of Figure 16 shows the results of the cross-validation test of the discriminant analysis applied to actual (nonrandomized) PC and average firing rate values. The discriminant analysis was calculated on half of the data (randomly selected) and used to classify the other half. The mean classification success based on PCs (58.5%) was no different from that based on the average firing rate (58.8%), according to a paired *t* test ($p = 0.6314$).

For comparison, we also calculated the multivariate discriminant analysis on the three firing rate intervals. The bottom panels of Figure 16 show the comparison between the multivariate spike rate classifications and the classifications based on PCs. The classifications based on multivariate spike rates on the entire data set (mean = 62.73%) were slightly, but significantly, better than the PC classification scores (mean = 62.03%) according to a paired *t* test ($p = 0.0013$). On the test set, the multivariate spike rate scores (mean = 59.64%) were slightly, but significantly, better than the scores based on the PCs (mean = 58.5%) based on a paired *t* test ($p = 0.0286$).

Although the discriminant analysis did not find additional information conveyed by the PCs compared to the average firing rate, it was possible that the discriminant analysis was missing important information in the data. The discriminant model assumes normal distributions of responses and, furthermore, that the probability of membership in a given class increases monotonically with the magnitude of the response variable. However, it is possible that the relation is actually not mono-

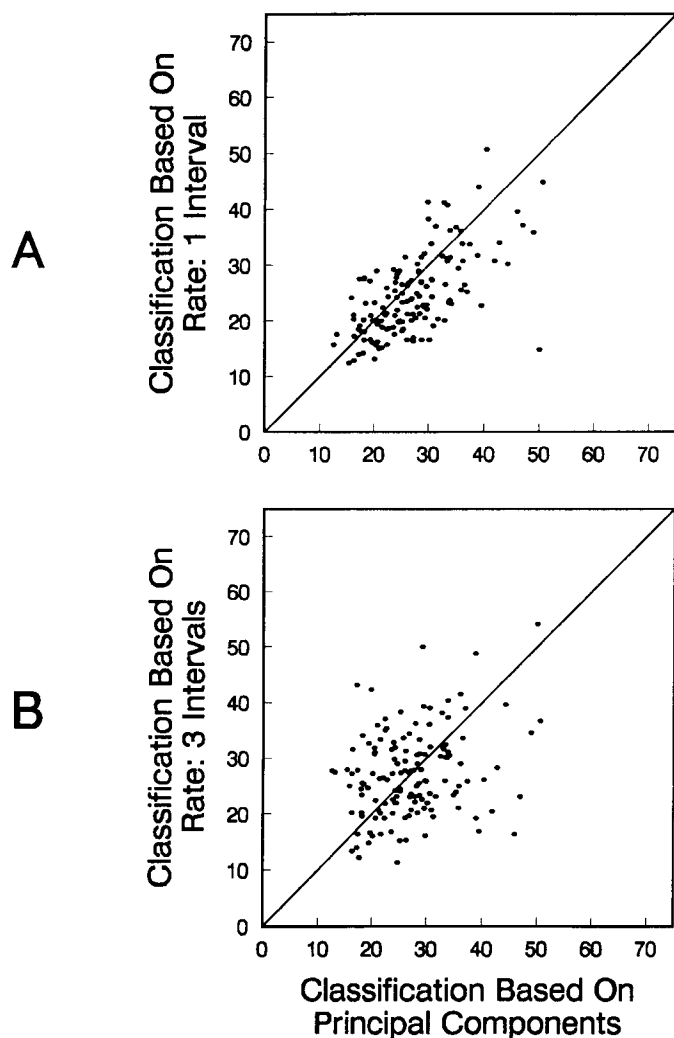


Figure 17. Performance of discriminant analysis applied to classification of the six sample stimuli. The data points show the classification scores for the cross-validated test set. For conventions, see Figure 16.

tonic and might even be discontinuous. For example, a matching stimulus might sometimes cause a large response and sometimes a small response but never an intermediate response, whereas the opposite might be true of the nonmatching stimulus. That is, the two distributions might be nonoverlapping but have identical means, which would result in poor classification by the discriminant analysis. The same could be true of discriminant analysis applied to PC values. To control for these possibilities, we used a three-layer neural network, trained by backpropagation, to classify the responses as matching or nonmatching based on the response rate in the three time intervals or PCs. Such a network should, in principle, find any nonlinear mapping of the stimulus to response.

For classifying stimuli based on firing rate in the three intervals, the network had three input units, three hidden units, and one output unit. For classifying stimuli based on three PCs, the network had six input units, three for positive PC values and three for negative PC values, five hidden units, and one output unit.

The classification scores from both networks were similar to each other and also to the scores obtained from the discriminant analysis. Because the neural network provides no test of statis-

tical significance, we examined the mean classification scores (with cross-validation) for the pool of 201 stimuli that showed significance effects on *any* discriminant analysis. The mean success rates of the neural network for the PCs and multivariate spike rates were 58.0% and 57.9%, respectively, which was about the same as those from the discriminant analysis. The absence of any significant temporal variation in response due to match–nonmatch status is consistent with our earlier analysis of the population histogram, which showed a sustained suppression of the match response.

Classification of sample stimuli. The fact that we found no evidence for temporal coding in the mnemonic domain in our paradigm raised the question of whether there was any evidence for temporal coding in the sensory domain. To investigate this, we applied discriminant analysis to the classification of the six different sample stimuli used for each cell. In this case, the discriminant analysis was used to classify each response into one of six stimulus classes rather than two.

The average firing rate and PCs were calculated in the same manner as in the match–nonmatch analysis. Based on the three PCs, the multivariate discriminant analysis was significant for 128 of the cells, which was almost the same result we found based on the ANOVA of the firing rates described earlier in the Results (124 cells). The percentage of cells showing significant stimulus selectivity is therefore about the same whether one considers temporal variation in the response or just the average firing rate.

On the full data set before cross-validation, the PCs did significantly better (mean of 31.68% correct) than the average firing rate (mean of 26.95% correct) in classifying the sample stimuli, which was expected (paired t test, $t = 26.95$, $p < 0.0001$). However, unlike in the case of the match–nonmatch response data, an advantage was also maintained in the test set data. The mean classification success based on PCs was 27.17% correct, compared to 24.63% correct with average firing rate, which was a significant difference according to a paired t test ($t = 4.60$, $p < 0.0001$). The relative performance of the two methods is plotted in Figure 17A. Considering that chance performance was 16.6%, another way of looking at this difference is that classifying the stimuli based on firing rate resulted in an 8% improvement over chance, whereas classifying them based on the three PCs resulted in a 10.5% improvement over chance, which was about 30% better performance compared to average firing rate. We checked these results with a neural network trained to classify the stimuli using the PCs and found no improvement in classification success of the neural network over the discriminant analysis. Thus, in using the neuronal response to distinguish among the samples, taking temporal variation into account results in a modest improvement in discriminability.

Because there were six sample stimuli, there were five different misclassification errors possible for each sample. If the errors are not random, they might themselves carry information, which is taken into account by the information measure used by Optican and Richmond (1987). To test the possibility that the error distribution based on PCs contained more information than that based on firing rate, we calculated the information contained in the entire confusion matrix (six samples by six response classes) produced by the discriminant analysis for each cell, using the same transmitted information measure as Optican and Richmond (1987). We found the mean information in the matrix to be the same for both PCs and average firing rate, namely, 0.35 bits per stimulus.

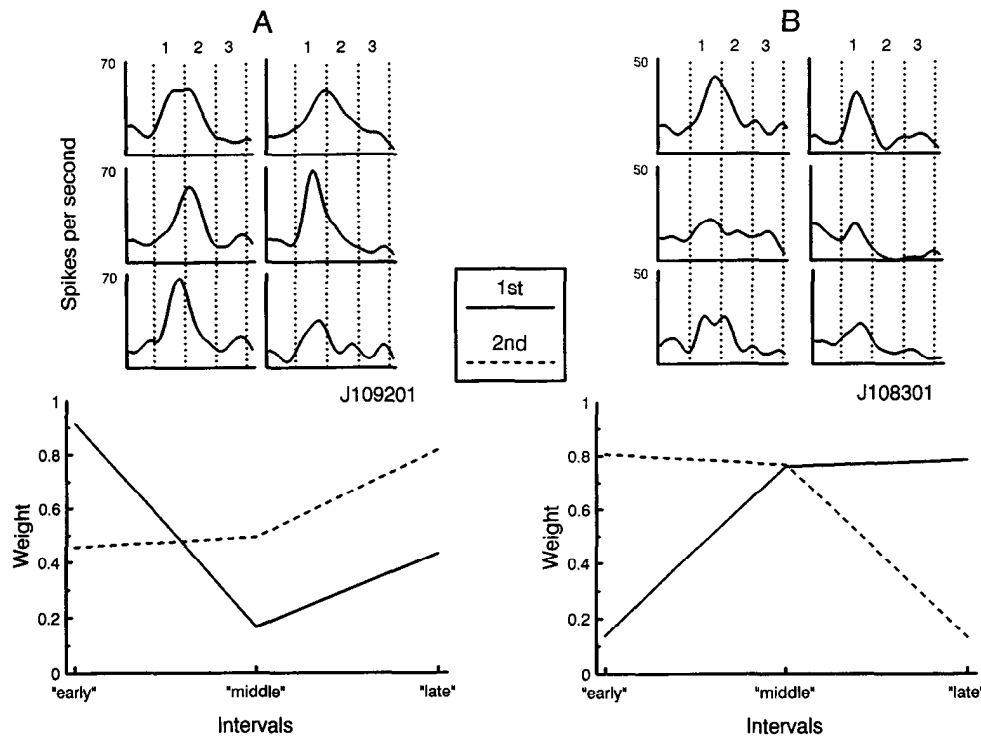


Figure 18. Weights assigned "early," "middle," and "late" response intervals by a discriminant analysis of sample responses for two IT neurons and histograms showing the responses to the six sample stimuli. Histograms were smoothed by convolving the spike trains with a Gaussian with an SD of 10 msec. Dotted lines through each histogram show the "early," "middle," and "late" subintervals, marked 1, 2, and 3, respectively. Weights are shown for the two significant canonical variables.

We next explored the nature of the different temporal variations in responses that gave the PCs an advantage in classifying the stimuli. Did the PCs extract a complex temporal code from the spike trains, or would a simpler measure of temporal variation do as well in distinguishing among the stimuli? To help answer this, we compared the performance of the discriminant analysis based on PCs with that of an analysis based on the same three firing rate intervals (early, middle, and late response) used to study match–nonmatch differences. Figure 17B shows the classification success scores based on three firing rate intervals compared with those based on the PCs. The classification success rate using the three rate intervals on the entire data set (mean = 31.68%) were slightly, but significantly, better than the success rates using the PCs (mean = 31.28%; $t = 2.56$, $p = 0.011$). On the test set, the scores were not significantly different: the means for the three rate intervals and the PCs were 27.25% and 27.17%, respectively ($t = 0.02$, $p = 0.9352$). Thus, simply breaking up the response interval into early, middle, and late components seems to capture all of the temporal information available from the three PCs.

We next returned to the analysis of the three firing rate intervals, to see if the canonical variables arrived at by the discriminant analysis gave any insight into the temporal variation in response caused by the stimuli. The discriminant analysis for each cell produced three uncorrelated canonical variables, each of which had a different set of loadings on the three response intervals. The relative loadings are a measure of how much each interval contributed to a discrimination. For 40% (50 of 125) of the stimulus-selective cells, only the first canonical variable was significant; that is, only a single set of weights applied to the three intervals was needed. This group included, for example, cells that simply had transient responses to all stimuli. For such cells, weighting the early part of the response more than the later part improved the discriminability of the stimuli over that based on average firing rate over the entire interval.

For the remainder of the cells (75 of 125), the second canonical variable was also significant; that is, for these cells it was necessary to apply two independent sets of weights to the three intervals to achieve the best discriminability. The third canonical variable was significant for only a few cells (5 of 125, or 4%).

To determine whether there were characteristic patterns to the relative weightings of the three response intervals, we performed a cluster analysis on weights for the first canonical variable across all cells. All cells fell into one of two clusters. In the first, which consisted of over half of the cells (75 of 125, or 60%), the early, or transient, parts of the response were weighted the most. The average weights for the three response intervals were 0.619, 0.583, and 0.264, respectively. We also examined the average weights for the second canonical variable for this group and found the opposite trend. Thus, for this group, the early and late portions of the response appeared to carry independent information about the sample stimuli, with the early portion the most important. An example of a cell from this cluster is shown in Figure 18A. For this cell, differences in latency to the peak response, a property that was noted for other cells as well, accounted for some of the early variation in response to the different stimuli.

The second cluster of cells (50 of 125, or 40%) showed the opposite trend; that is, the late portions of the response were weighted the most in the discrimination. The average weights on the response intervals for the first canonical variable were 0.350, 0.464, and 0.738, respectively, and the second canonical variable tended to show the opposite trend. An example of such a cell is shown in Figure 18B. Thus, as for the first group of cells, the early and late portions of the response tended to carry independent information about the samples but, unlike the first group, the late portion was more important. These early and late components roughly correspond to a difference between transient and sustained responses.

Discussion

The responses of many IT neurons to a current stimulus are influenced by the traces of stimuli held in memory. When the current stimulus matches a stimulus held in memory, the response appears to be suppressed. Thus, a substantial portion of IT neurons may function as adaptive mnemonic “filters” that give their strongest response to stimuli with the appropriate visual features, but which are unexpected or not recently seen. We reached this conclusion in an earlier report (Miller et al., 1991b), but the analyses of the present study allow us to expand upon the previous one. We now have evidence for the exact time course of the suppressive effects, the role of stimulus similarity in the suppressive phenomenon, the existence of cells that show the opposite behavior, the relationship between mnemonic and sensory discriminability, the relationship between the magnitude of suppression and the animal’s performance in the task, the role of delay activity, the presence of an active reset mechanism between trials, and the role of temporal response variation in both mnemonic and sensory discriminability.

Behavioral studies in both normal and brain-damaged humans and animals have identified many different types of memory, both long and short term (Mishkin, 1978; Squire, 1982; Malamut et al., 1984; Zola-Morgan and Squire, 1984, 1985; Overman et al., 1990). Some studies distinguish declarative (explicit) memory from procedural (habit) memory as well as from the types of memory that underlie priming and simple forms of learning such as behavioral habituation. Evidence from lesion and cooling studies suggests that IT cortex participates in most or all of these different memory systems. Likewise, the ability of IT neurons to compare the current stimulus to stored memory traces might contribute to any or all of the memory systems defined in behavioral studies, depending on how the memory system is implemented neurally. In the present study, the memory traces were actively changed from trial to trial, suggesting a role of IT neurons in working memory. Because the neuronal response to a repeated stimulus was smaller, there is a suggestive link to behavioral habituation. In other contexts, the memory traces might be long-term ones. Indeed, when we tested IT neurons with novel stimuli that were repeatedly presented until they became familiar, we found that many cells developed stimulus-specific suppression lasting long periods (Miller et al., 1991b). Such a mechanism could play a role in long-term recognition memory or, perhaps, priming.

Relationship to prior studies. A number of recent studies indicate that the anterior-ventral portion of IT cortex is particularly important for memory. Lesions or reversible cooling of the cortex in the vicinity of the rhinal sulcus cause a devastating impairment in a variety of memory tasks, which may be even more severe than those caused by lesions of the remainder of IT cortex (Horel et al., 1987; Murray et al., 1989; Zola-Morgan et al., 1989; Meunier et al., 1990; Suzuki et al., 1991, 1993; Gaffan and Murray, 1992; Murray, 1992; Zola-Morgan et al., 1993). Behavioral studies that have failed to find long-lasting impairments on DMS tasks following IT lesions have spared the anterior-ventral portion (Dean, 1974; Gaffan and Weiskrantz, 1980). Likewise, anatomical studies have found differences in connections of the anterior-ventral region compared to the remainder of IT cortex (Insausti et al., 1987a,b; Martin-Elkins and Horel, 1992). On the basis of cytoarchitecture and anatomical connections, some studies have labeled a portion of

the cortex between the anterior middle temporal and rhinal sulci as area 36 (Amaral et al., 1987; Suzuki et al., 1993), a portion of which is included in our recording area. In the future, it will be important to compare the mnemonic information communicated by individual cells in this region to that communicated by cells in the remainder of IT cortex as well as that by cells in adjacent areas such as area 35 in the rhinal sulcus, areas TF/TH in the parahippocampal gyrus, and the temporal pole.

Modulation of responses in IT cortex during memory-related tasks has been reported in several prior studies (Gross et al., 1979; Mikami and Kubota, 1980; Fuster and Jervey, 1981; Baylis and Rolls, 1987; Miyashita and Chang, 1988; Fuster, 1990; Vogels and Orban, 1990a; Miller et al., 1991b; Riches et al., 1991; Eskandar et al., 1992). Consistent with the present study, the two studies reporting match–nonmatch differences found, on the average, smaller responses to matching stimuli for those cells showing a significant match–nonmatch effect (Baylis and Rolls, 1987; Eskandar et al., 1992). Further, both Miller et al. (1991b) and Riches et al. (1991) found that IT neurons typically give smaller responses to recently seen stimuli. Thus, there is additional support for the adaptive memory filter idea. Suppression of responses to stimuli that match memory traces is also consistent with some neural network architectures for memory storage (Carpenter and Grossberg, 1987; Kohonen, 1988) as well as with the results of a recent positron emission tomography study of cortical activation in humans. In the latter study, subjects performing a visual word-stem completion task showed less activation of temporal cortex when they had recently seen the same words than when the words had not been seen (Squire et al., 1992).

The only other study of the effect of several intervening stimuli on IT responses in a memory task was by Baylis and Rolls (1987). In a running recognition task, they found that for most cells the match–nonmatch effect was eliminated by even a single intervening stimulus, and for all cells it was gone entirely following two intervening stimuli. On this basis, they suggested that IT cortex was involved in memory lasting only a second or so, that is, primary memory. By contrast, we found effects in IT cortex that bridged at least five intervening stimuli. Further, we have previously found modulation of responses by stimulus familiarity in IT cortex that bridges more than a hundred stimuli (Miller et al., 1991b). The results suggest that the “memory span” of IT neurons may be equivalent to those of the monkey.

There are at least three important differences between the Baylis and Rolls study and ours. First, the stimuli in their study were all initially novel and were presented only twice in a session, whereas the stimuli we used were familiar to the monkeys and were used repeatedly. Second, the monkeys in their study were performing a serial recognition task, in which they were required to remember all of the stimuli in a trial, whereas ours needed to remember only the sample. These differences in either the stimuli or the task might account for the different results. The most likely explanation, however, is that their recordings in IT cortex did not include the anterior-ventral region included in the present study, supporting the notion that this part of IT cortex is especially important for longer-lasting memories.

Several previous studies have found stimulus-selective activity in IT cortex during the delay interval of DMS tasks, suggesting that this maintained activity mediates the memory of the sample stimulus (Fuster and Jervey, 1981; Miyashita and Chang, 1988; Fuster, 1990). About a quarter of the cells we

recorded showed stimulus-selective activity in the delay interval following the sample, supporting the findings of these studies. However, because this activity was eliminated by intervening stimuli, it seems unlikely that delay-interval activity in IT cortex contributed to the performance of this particular version of delayed matching to sample.

Responses of neurons in area V4, which provides a major input to IT, are also modulated by the memory of a sample stimulus (Haenny et al., 1988; Maunsell et al., 1991). Using a task similar to ours but with a fixed set of four oriented gratings as stimuli, Maunsell et al. (1991) found that a given sample stimulus affected the responses to all subsequent test stimuli in a trial, that is, that the responses to test stimuli were "tuned" to the sample. There was no tendency for the responses to matching stimuli to be either enhanced or suppressed. Further, the modulation of responses did not appear to be related to any interaction between particular test stimulus representations and the sample memory trace; rather, a given sample caused a non-specific suppression or enhancement of all test stimuli responses on that trial. Although these V4 effects might be caused by feedback from IT, they are so different from those in IT that they may be independent, for example, reflecting mechanisms that come into play only after extensive training on a small set of stimuli.

Active or passive processes in memory. The suppression to matching test stimuli might have been caused, in principle, by either a strictly passive (or automatic) mechanism sensitive to stimulus recency or repetition, or an active (or voluntary) mechanism that restricted the comparison of the test stimulus to just the actively held memory of the sample, or by some combination of active and passive processes. Stimulus repetition alone cannot explain our results in the DMS task, because sample responses were not suppressed when a sample on one trial was identical to the final (matching) stimulus on the previous trial. However, in a different control condition, we did observe suppressive effects that appeared to be caused just by stimulus repetition. In this condition, a test stimulus matching the sample was simply presented within a trial while the monkey was maintaining fixation, and the response to this match stimulus was suppressed.

Taking the results of these two conditions together, it seems likely that both a passive process related to stimulus recency and an active process related to the memory trace of the sample play a role in the mnemonic modulation of IT responses. In the absence of the active process, stimulus repetition alone probably leads to response suppression. This view is consistent with recent results of Miller et al. (1991a), who found suppressed responses of IT neurons to repeated stimuli in both anesthetized animals and passively fixating animals, and of Riches et al. (1991), who also found such effects in a passive viewing condition. The fact that IT cells differentiate between matching and nonmatching stimuli within 10 msec of response onset, that is, practically on the first action potential, suggests that the passive recency mechanism could not be mediated by feedback to IT cortex during the time of the match stimulus presentation. Rather, the mechanism that causes the suppression must be located within or before IT cortex.

In addition to any passive recency mechanism, the failure of the suppressive effects to carry across trials in the DMS task indicates that an active process either "clears out" or "resets" the memory traces of stimuli from one trial to the next (to avoid cross-trial interference) or "primes" IT cells with the memory

of the sample so that they are prepared to detect its reoccurrence. Such a reset signal might arise from prefrontal cortex, since behavioral evidence suggests that prefrontal cortex is critical for working memory tasks with short delays (Mishkin and Manning, 1978; Bachevalier and Mishkin, 1986) and Fuster et al. (1985) has described interactions between prefrontal and IT cortex during short-term memory tasks.

Figure-ground in space and time. Many cells at all levels of the visual system respond best to contrast of some sort. In the retina, contrast is created by luminance differences, whereas in the neocortex some cells can detect differences in orientation, color, spatial frequency, and direction of motion between stimuli within their classically defined receptive field and stimuli in broad surrounding regions. For these cells, the greater the similarity between the stimulus in the surround and the stimulus in the receptive field, the more the response to the receptive field stimulus is suppressed. As one advances through the visual system, the stimulus features that are contrasted may become more sophisticated and the spatial areas over which the interactions occur may become larger. Based on these properties, it has been conjectured that one of the functions of the visual system is to separate figures from background (Allman et al., 1985; Desimone et al., 1985).

The properties of IT neurons suggest that figure-ground separation occurs in the temporal as well as in the spatial domain. Many IT neurons respond best to stimuli that are dissimilar to those that have been seen in the recent past. Thus, the past may function as the surround, which is compared with the current stimulus in the receptive field.

If the analogy with spatial receptive fields is valid, IT cortex should not be unique in its temporal properties. Figure-ground separation in the spatial domain appears to build incrementally as one moves through the visual system, and there may be a comparable buildup of temporal processing in the visual system as well. Nelson (1991), for example, found that cells in striate cortex of the cat show orientation-specific suppression lasting a few hundred milliseconds. As described above, orientation-specific temporal interactions have also been found in area V4 (Haenny et al., 1988; Maunsell et al., 1991), although these interactions did not appear to be suppressive. Presumably, the temporal interval over which stimuli are compared is much smaller (and less able to span intervening stimuli) in earlier visual areas, and the features that are compared are less complex than in IT cortex.

Pooling of IT responses. Because no individual IT neuron distinguished between matching and nonmatching stimuli as well as the animal as a whole, the responses of multiple IT neurons would have to be pooled to support the behavioral level of performance. Our discriminant analysis of neuronal performance assumes that the neural circuitry that interprets the output of IT cortex "knows" the magnitude and variances of the match and nonmatch responses for a given IT neuron. Performance of the decision network in this case is limited by the distance between the means of the response distributions. This distance can be used to estimate, in principle, how many IT cells the decision network would need to sample to achieve a given level of behavioral performance. The square root of the Mahalanobis distance (D^2) is the distance between the means of the distributions measured in units of the SD of the distributions (equivalent to d' in signal detection theory). For an unbiased criterion, the average error in classification is given by the area to the left of $-D/2$ under the standard normal curve.

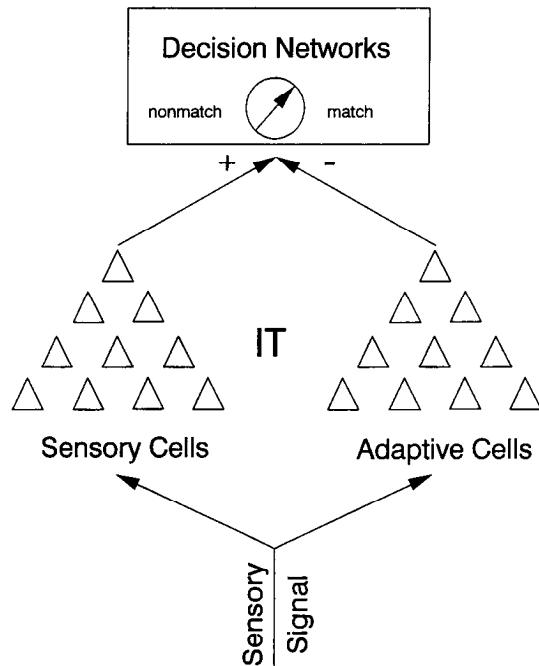


Figure 19. Model of how sensory cells and adaptive mnemonic filter cells in IT cortex contribute to the match–nonmatch decision. *Plus* and *minus* indicate that the outputs of the adaptive and sensory cells work in opposition, and are not meant to imply excitation or inhibition.

For stimuli that elicited a significant match–nonmatch difference, we found the estimated mean D to be 0.94. This estimate is probably overly optimistic, however, as a distance of 0.94 should have resulted in a neural classification success of 68% correct, whereas the measured success rate was only 60% correct on average. Small discrepancies between the predicted and actual classification success rates can result from response distributions that are not strictly normal. Based on the measured classification success, a more conservative estimate of D would be 0.5. By contrast, a behavioral performance of 90% correct corresponds to a D of 2.56, which is five times larger than that achieved by the average IT neuron. To achieve a five-fold improvement in D would require a comparable reduction in the SD of the response classes, which could be achieved by averaging the responses of 25 IT neurons that exhibit a match–nonmatch effect whose responses are independent, or uncorrelated. If responses were correlated, this would increase the minimum size of the pool (Britten et al., 1992).

However, it seems unlikely that the neural circuitry that interprets the output of IT cortex would have prior knowledge of the distributions of match and nonmatch responses, especially since absolute response magnitude might fluctuate depending on extraneous factors such as arousal or the relative effectiveness of the stimuli. A more robust mechanism would be one that compared inputs from two opposing classes of IT cells, an “adaptive mnemonic filter” class and a “sensory” class (Fig. 19). The response of the sensory class to the current stimulus would provide a referent, not affected by past experience, against which to compare the responses of the adaptive cells. The average difference in response between the two classes of cells would be a measure of similarity between the current stimulus and a stimulus held in memory. According to this model, the

sensory class of IT cells would play as important a role as the adaptive class in the DMS tasks. In fact, from an efficiency standpoint, it would be best if the two classes were equal in number, which fits our findings in IT. We do not yet know if the two populations are anatomically distinguishable, but our data from a few simultaneously recorded pairs of IT neurons suggest that there may be clustering of the two classes.

Such a differencing model would be analogous to one used by Britten et al. (1992) to compare the behavioral performance of monkeys performing direction of motion discriminations to the performance of MT neurons in the same task. In their model, a decision network compares the outputs of MT cells tuned to opposite directions of motion. In IT, where there is no real opposite to the adaptive cells, the comparable comparison would be to the sensory cells that are unaffected by memory. Casting the few adaptive cells that showed relative match enhancement as “sensory” cells would enhance the difference between the memory and sensory classes.

The exact size of the IT neuronal pool necessary to achieve a given level of behavioral performance probably has more statistical than biological meaning, yet the surprisingly small size of the estimated numbers suggests that information equivalent to that held by the animal as a whole is distributed down to the level of small neuronal populations. The same conclusion was reached by Britten et al. (1992) based on a study of direction of motion in area MT. An important question is how this information is extracted from just the relevant cells in order to guide behavior.

Temporal coding versus rate coding. Based on both discriminant and neural network analyses of match and nonmatch responses, we found no evidence for significant temporal coding. Rather, cells differentiated between matching and nonmatching stimuli practically by the first action potential, and the suppression of the match response appeared to be sustained throughout the response interval. These results stand in apparent contrast to the conclusions of Eskandar et al. (1992) that mnemonic information is communicated, at least in part, by temporal coding in IT cortex.

There are at least four possible reasons for the different findings of the two studies, namely, differences in the stimulus sets, the recording areas in IT cortex, the methods used to compare temporal and rate coding, and the use of intervening stimuli in the DMS task. Of these, we believe the most important difference is the fact that we used multiple intervening stimuli between the sample and matching test stimulus in our task, whereas Eskandar et al. (1992) did not. The delay interval in their study was only 500 msec, and we have sometimes observed either postresponse suppression or rebound activity lasting for that amount of time. Thus, the response to one stimulus may interact, temporally, with the response to a subsequent one, but these effects may not persist for many seconds and across several intervening stimuli. The only long-lasting effect we found in our data was a sustained suppression of the match stimulus response.

Unlike the case with the matching–nonmatching data, we did find some evidence for temporal variation in response elicited by the sample stimuli. Using the first three PCs of the spike train, we found a modest improvement (27.17% vs. 24.63% classification success) in the ability to classify the samples over that achieved by using just the average firing rate, confirming the findings of Richmond et al. (1987) that stimuli elicit temporal variation in responses of IT neurons, as well as changes

in the mean firing rate. From the ratio of the logs of the error rates using the two methods, one can calculate the relative number of neurons that would be needed to classify the stimulus, assuming a maximum-likelihood detector (Optican and Richmond, 1987). The probabilities of classification error using PCs and average firing rate were 0.728 and 0.754, respectively, and $\log(0.728)/\log(0.754) = 1.12$. Thus, it would take about 12% more neurons to classify a stimulus using firing rate than using three PCs. The magnitude of the difference was smaller than that found by Optican and Richmond (1987) and Eskandar et al. (1992), which might be due to differences in either the stimulus sets, recording sites, or methods of analysis. Other recent studies have investigated the role of temporal coding in the response of IT neurons to faces and have concluded that most of the information is transmitted in the early phase of the response (Oram and Perrett, 1992; Rolls et al., 1992), with only modest opportunity for improvement by temporal coding (Rolls et al., 1992).

The PCs are optimal linear features for representing variance and, in principle, might have captured complex patterns in the spike trains. In fact, we have found the shapes of the PCs to be arbitrary, in the sense that they depend strongly on the interval over which they were calculated (E. K. Miller, L. Li, and R. Desimone, unpublished observations). Further, we achieved almost precisely the same classification success from using three relatively large, successive time intervals of the response, that is, early, middle, and late response. Thus, the underlying temporal variation in response appears to be relatively coarse.

Our analysis of this underlying variation showed that, for some cells, one portion of the response, either the transient or the sustained portion, appeared to carry most of the information while the other portion appeared to add noise. For other cells, the transient and sustained portions were differentially affected by different stimuli. For example, two stimuli might elicit a smaller mean response but one might cause greater activity in the early phase of the response and the other might cause greater activity in the late phase. Different stimuli appeared to shift the balance somewhat between the transient and sustained portions of the response. The question is whether such temporal variations are simply a reflection of activity in dynamic neural networks, changing state over time, or whether they are evidence that neurons send and decode temporally coded messages about specific visual stimuli. In either case, a full accounting of neural properties requires that these temporal variations be explained. To determine whether they represent a "code" might require that they be tested for invariance over temporal variations in the stimulus (e.g., flicker or motion) or that they be purposely manipulated through, for example, electrical stimulation.

References

- Aggleton JP, Burton MJ, Passingham RE (1980) Cortical and subcortical afferents to the amygdala of the rhesus monkey (*Macaca mulatta*). *Brain Res* 36:243–248.
- Allman J, Miezin F, McGuinness E (1985) Direction- and velocity-specific responses from beyond the classical receptive field in the middle temporal area (MT). *Perception* 14:105–126.
- Amaral DG, Price JL (1984) Amygdalo-cortical projections in the monkey (*Macaca fascicularis*). *J Comp Neurol* 230:465–496.
- Amaral DG, Insausti R, Cowan WM (1987) The entorhinal cortex of the monkey: I. Cytoarchitectonic organization. *J Comp Neurol* 264:326–355.
- Bachevalier J, Mishkin M (1986) Visual recognition impairment follows ventromedial but not dorsolateral prefrontal lesions in monkeys. *Behav Brain Res* 20:249–261.
- Baizer JS, Ungerleider LG, Desimone R (1991) Organization of visual inputs to the inferior temporal and posterior parietal cortex in macaques. *J Neurosci* 11:168–190.
- Baylis GC, Rolls ET (1987) Responses of neurons in the inferior temporal cortex in short term and serial recognition memory tasks. *Exp Brain Res* 65:614–622.
- Bradley A, Skottun BC, Ohzawa I, Sclar G, Freeman RD (1987) Visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *J Neurophysiol* 57:755–772.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.
- Carpenter GA, Grossberg S (1987) A massively parallel architecture for a self-organizing neural pattern recognition machine. *Comput Vis Graphics Image Processing* 37:54–115.
- Dean P (1974) The effect of inferotemporal lesions on memory for visual stimuli in the rhesus monkeys. *Brain Res* 77:451–469.
- Desimone R, Gross CG (1979) Visual areas in the temporal cortex of the macaque. *Brain Res* 178:363–380.
- Desimone R, Fleming J, Gross CG (1980) Prestriate afferents to inferior temporal cortex: an HRP study. *Brain Res* 184:41–55.
- Desimone R, Albright TD, Gross CG, Bruce C (1984) Stimulus-selective properties of inferior temporal neurons in the macaque. *J Neurosci* 4:2051–2062.
- Desimone R, Schein SJ, Moran J, Ungerleider LG (1985) Contour, color and shape analysis beyond the striate cortex. *Vision Res* 25:441–452.
- Eskandar EN, Richmond BJ, Optican LM (1992) Role of inferior temporal neurons in visual memory: I. Temporal encoding of information about visual images, recalled images, and behavioral context. *J Neurophysiol* 68:1277–1295.
- Fuster JM (1990) Inferotemporal units in selective visual attention and short-term memory. *J Neurophysiol* 64:681–697.
- Fuster JM, Jervey JP (1981) Inferotemporal neurons distinguish and retain behaviorally relevant features of visual stimuli. *Science* 212:952–955.
- Fuster JM, Bauer RH, Jervey JP (1981) Effects of cooling inferotemporal cortex on performance of visual memory tasks. *Exp Neurol* 71:398–409.
- Fuster JM, Bauer RH, Jervey JP (1985) Functional interactions between inferotemporal and prefrontal cortex in a cognitive task. *Brain Res* 330:299–307.
- Gaffan D, Murray E (1992) Monkeys (*Macaca fascicularis*) with rhinal cortex ablations succeed in object discrimination learning despite 24-hr intertrial intervals and fail at matching to sample despite double sample presentations. *Behav Neurosci* 106:30–38.
- Gaffan D, Weiskrantz L (1980) Recency effects and lesion effects in delayed non-matching to randomly baited samples by monkeys. *Brain Res* 196:373–386.
- Gross CG (1973) Visual functions of inferotemporal cortex. In: *Handbook of sensory physiology* (Jung R, ed). Berlin: Springer.
- Gross CG, Bender DB, Gerstein GL (1979) Activity of inferior temporal neurons in behaving monkeys. *Neuropsychologia* 17:215–229.
- Haenny PE, Maunsell JHR, Schiller PH (1988) State dependent activity in monkey visual cortex. II. Retinal and extraretinal factors in V4. *Exp Brain Res* 69:245–259.
- Horel JA, Pytko DE (1982) Behavioral effect of local cooling in temporal lobe of monkeys. *J Neurophysiol* 47:11–22.
- Horel JA, Pytko-Joiner DE, Voytko ML, Salsbury K (1987) The performance of visual tasks while segments of the inferotemporal cortex are suppressed by cold. *Behav Brain Res* 23:29–42.
- Insausti R, Amaral DG, Cowan WM (1987a) The entorhinal cortex of the monkey: II. Cortical afferents. *J Comp Neurol* 264:356–395.
- Insausti R, Amaral DG, Cowan WM (1987b) The entorhinal cortex of the monkey: III. Subcortical afferents. *J Comp Neurol* 264:396–408.
- Kohonen T (1988) *Self-organization and associative memory*. New York: Springer-Verlag.
- Malamut BL, Saunders RC, Mishkin M (1984) Monkeys with combined amygdalo-hippocampal lesions succeed in object discrimination learning despite 24-hour intertrial intervals. *Behav Neurosci* 98:759–769.
- Martin-Elkins CL, Horel JA (1992) Cortical afferents to behaviorally defined regions of the inferior temporal and parahippocampal gyri as demonstrated by WGA-HRP. *J Comp Neurol* 321:177–192.

- Maunsell JHR, Sclar G, Nealey TA, DePriest DD (1991) Extraretinal representations in area V4 in the macaque monkey. *Vis Neurosci* 7:561–573.
- Meunier M, Murray EA, Bachevalier J, Mishkin M (1990) Effects of perirhinal cortical lesions on visual recognition in monkeys. *Soc Neurosci Abstr* 16:616.
- Mikami A, Kubota K (1980) Inferotemporal neuron activities and color discrimination with delay. *Brain Res* 182:65–78.
- Miller EK, Gochin PM, Gross CG (1991a) Habituation-like decrease in the responses of neurons in inferior temporal cortex of the macaque. *Vis Neurosci* 7:357–362.
- Miller EK, Li L, Desimone R (1991b) A neural mechanism for working and recognition memory in inferior temporal cortex. *Science* 254:1377–1379.
- Mishkin M (1978) Memory in monkeys severely impaired by combined but not by separate removal of amygdala and hippocampus. *Nature* 273:297–298.
- Mishkin M (1982) A memory system in the monkey. *Philos Trans R Soc Lond [Biol]* 298:83–95.
- Mishkin M, Manning FJ (1978) Non-spatial memory after selective prefrontal lesions in monkeys. *Brain Res* 143:313–323.
- Miyashita Y, Chang HS (1988) Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature* 331:68–70.
- Murray EA (1992) Medial temporal lobe structures contributing to recognition memory: the amygdaloid complex versus the rhinal cortex. In: *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction*, pp 453–470. New York: Wiley-Liss.
- Murray EA, Bachevalier J, Mishkin M (1989) Effects of rhinal cortical lesions on visual recognition memory in rhesus monkeys. *Soc Neurosci Abstr* 15:342.
- Nelson SB (1991) Temporal interactions in the cat visual system. I. Orientation-selective suppression in the visual cortex. *J Neurosci* 11:344–356.
- Optican LM, Richmond BJ (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J Neurophysiol* 57:162–178.
- Oram MW, Perrett DI (1992) Time course of neural responses discriminating different views of the face and head. *J Neurophysiol* 68:70–84.
- Overman WH, Ormsby G, Mishkin M (1990) Picture recognition vs. picture discrimination learning in monkeys with medial temporal removals. *Exp Brain Res* 79:18–24.
- Riches IP, Wilson FA, Brown MW (1991) The effects of visual stimulation and memory on neurons of the hippocampal formation and the neighboring parahippocampal gyrus and inferior temporal cortex of the primate. *J Neurosci* 11:1763–1779.
- Richmond BJ, Optican LM (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex: II. Quantification of response waveform. *J Neurophysiol* 57:147–161.
- Richmond BJ, Optican LM, Spitzer H (1987) Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex: I. Response characteristics. *J Neurophysiol* 57:132–146.
- Robinson DA (1963) A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE Trans Biomed Eng* 10:137–145.
- Rolls ET, Tovee MJ, Treves A (1992) A significant proportion of the total information is encoded in a short period at the beginning of a spike train of neurons recorded in the primate temporal visual cortex. *Soc Neurosci Abstr* 18:1275.
- Squire LR (1982) *The neuropsychology of the human mind*. Annu Rev Neurosci 5:241–273.
- Squire LR, Ojemann JG, Miezin FM, Petersen SE, Videen TO, Raichle ME (1992) Activation of the hippocampus in normal humans: a functional anatomical study of memory. *Proc Natl Acad Sci USA* 89:1837–1841.
- Suzuki WA, Zola-Morgan S, Squire LR, Amaral DG (1991) Lesions of the perirhinal and parahippocampal cortices in monkeys produce a modality general and long lasting memory impairment. *Soc Neurosci Abstr* 17:339.
- Suzuki WA, Zola-Morgan S, Squire L, Amaral DG (1993) Lesions of the perirhinal and parahippocampal cortices in the monkey produce long-lasting memory impairment in the visual and tactile modalities. *J Neurosci*, in press.
- Ungerleider LG, Mishkin M (1982) Two cortical visual systems. In: *Analysis of visual behavior* (Ingle J, Goodale MA, Mansfield RJW, eds), pp 549–586. Cambridge, MA: MIT Press.
- Van Hoesen GW, Pandya DN (1975a) Some connections of the entorhinal (area 28) and perirhinal (area 35) cortices of the rhesus monkey. I. Temporal lobe afferents. *Brain Res* 95:1–24.
- Van Hoesen GW, Pandya DN (1975b) Some connections of the entorhinal (area 28) and perirhinal (area 35) cortices of the rhesus monkey. III. Efferent connections. *Brain Res* 95:48–67.
- Vogels R, Orban GA (1990a) Effects of task related stimulus attributes on infero-temporal neurons studied in the discriminating monkey. *Soc Neurosci Abstr* 16:621.
- Vogels R, Orban GA (1990b) How well do response changes of striate neurons signal differences in orientation: a study in the discriminating monkey. *J Neurosci* 10:3543–3558.
- Webster MJ, Ungerleider LG, Bachevalier J (1991) Connections of inferior temporal areas TE and TEO with medial temporal-lobe structures in infant and adult monkeys. *J Neurosci* 11:1095–1116.
- Zola-Morgan S, Squire LR (1984) Preserved learning in monkeys with medial temporal lesions: sparing of motor and cognitive skills. *J Neurosci* 4:1072–1085.
- Zola-Morgan S, Squire LR (1985) Medial temporal lesions in monkeys impair memory on a variety of tasks sensitive to human amnesia. *Behav Neurosci* 99:22–34.
- Zola-Morgan S, Squire LR, Amaral DG, Suzuki WA (1989) Lesions of perirhinal and parahippocampal cortex that spare the amygdala and hippocampal formation produce severe memory impairment. *J Neurosci* 9:4355–4370.
- Zola-Morgan S, Squire LR, Clower RP, Rempel NL (1993) Damage to the perirhinal cortex exacerbates memory impairment following lesions to the hippocampal formation. *J Neurosci* 13:251–265.