

This paper was presented at a colloquium entitled “Memory: Recording Experience in Cells and Circuits,” organized by Patricia S. Goldman-Rakic, held February 17–20, 1996, at the National Academy of Sciences in Irvine, CA.

Neural mechanisms for visual memory and their role in attention

(primate/inferior temporal cortex/neurophysiology/prefrontal cortex/visual search)

ROBERT DESIMONE

Laboratory of Neuropsychology, National Institute of Mental Health, 49 Convent Drive, Building 49, Room 1B80, Bethesda, MD 20892-4415

ABSTRACT Recent studies show that neuronal mechanisms for learning and memory both dynamically modulate and permanently alter the representations of visual stimuli in the adult monkey cortex. Three commonly observed neuronal effects in memory-demanding tasks are repetition suppression, enhancement, and delay activity. In repetition suppression, repeated experience with the same visual stimulus leads to both short- and long-term suppression of neuronal responses in subpopulations of visual neurons. Enhancement works in an opposite fashion, in that neuronal responses are enhanced for objects with learned behavioral relevance. Delay activity is found in tasks in which animals are required to actively hold specific information “on-line” for short periods. Repetition suppression appears to be an intrinsic property of visual cortical areas such as inferior temporal cortex and is thought to be important for perceptual learning and priming. By contrast, enhancement and delay activity may depend on feedback to temporal cortex from prefrontal cortex and are thought to be important for working memory. All of these mnemonic effects on neuronal responses bias the competitive interactions that take place between stimulus representations in the cortex when there is more than one stimulus in the visual field. As a result, memory will often determine the winner of these competitions and, thus, will determine which stimulus is attended.

Attention is often thought of as a gateway to learning and memory because we typically learn and remember much more about stimuli in the environment that we attend to than about stimuli we ignore. However, the converse is equally true. That is, mechanisms for learning and memory play a critical role in the selection process that determines which stimuli of the many stimuli in a complex environment are attended. In fact, some of the mechanisms for memory and attention are so intertwined that one might question whether they are even distinguishable. In the first part of this review I will summarize what is known about several cortical memory mechanisms, and in the second part I will take up a role for memory mechanisms in attention.

Fig. 1 summarizes in a diagrammatic fashion some of the cortical mechanisms for visual memory that appear to impact on attentional selection processes. It is important to note that all of the processes illustrated are related to what could be called “stimulus memory”—i.e., plastic changes induced in the cortex during the storage or expression of memories of visual stimuli. The types of memories normally associated with the hippocampus and medial temporal lobe structures, including memories of more complex behavioral events, maps, scenes, episodes, etc., are not considered.

Most neurophysiological studies of visual memory in primates have used some variation of the delayed matching-to-sample task (DMS). Typically in these studies, a sample stimulus is presented at the start of the trial, followed after by a delay by one or more test stimuli. The animal is rewarded for indicating which test stimulus matches the sample. To solve the DMS task the monkey must, in principle, solve three problems. First, it must discriminate among the different stimuli. Second, it must retain the memory of the sample for the length of the trial. Third, it must make a decision about whether the current test stimulus matches the sample held in memory. Neuronal mechanisms that may contribute to each of these operations are described in the next sections.

Most of the examples to be described come from recordings in the inferior temporal (IT) cortex of macaque monkeys. IT cortex is the last visual processing region in the “ventral” visual pathway that mediates object recognition and memory (2–4). It is comprised of at least two separate areas, area TE laterally and the perirhinal cortex anteromedially. IT cells have large receptive fields that typically include the center of gaze and have complex feature selectivity, and lesions of IT cortex cause impairments in both stimulus discrimination and recognition (2, 5–7).

Repetition Suppression

Repetition suppression (top of Fig. 1) was the first type of stimulus-induced change in the cortex that we found in our own studies in the perirhinal portion of IT cortex (8–11). In these studies, the monkey was performing a variation of the DMS task in which the sample stimulus was followed by a sequence of up to five different test stimuli, each separated by a short delay. The monkey was rewarded for indicating when a test stimulus in the sequence matched the sample. All of the stimuli and sequences were randomized, and a stimulus that appeared as a sample and matching test on one trial would appear as a nonmatching test item on another. For example, on one trial the monkey might see a stimulus sequence consisting of A . . . B . . . C . . . D . . . A and would respond to the final A, and on another trial it might see a sequence consisting of C . . . A . . . C and would respond to the final C. The stimuli were digitized pictures of complex objects, which are the sort of stimuli that typically activate IT neurons in a stimulus-selective fashion. For the initial study, the stimuli were always familiar to the animal, although the stimuli were often changed from cell to cell. Our goal was not to identify the “optimal” stimuli for a given cell; rather, we simply needed stimuli that would elicit a range of responses from the cells so that we could measure how the responses were affected by the mnemonic demands of the task.

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

Abbreviations: DMS, delayed matching-to-sample; IT, inferior temporal.

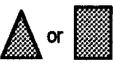
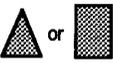
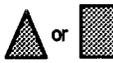
	A neuron that initially responds to	After experience with or instruction about	Now responds to
Repetition suppression (priming)			
Enhancement (selection/working memory)			
Delay/maintained activity (working memory)			
Tuning (perceptual learning)			
Association			

FIG. 1. Five ways in which neuronal activity is modified during the formation or expression of memory traces. (Adapted from ref. 1.)

Not surprisingly, most cells gave stimulus-selective responses, responding better to some stimuli than to others. IT stimulus selectivity presumably contributes to the first requirement for performing the DMS task described above—i.e., the ability to discriminate the stimuli.

For about half the cells, responses to the test items were apparently unaffected by any mnemonic aspect of the task. However, for the remaining cells, the responses to the test stimuli varied not only according to the features of those stimuli but also according to whether the test stimulus matched the sample held in memory. This influence of the sample memory on IT responses might contribute to the third requirement for performing the DMS task described above, namely a mechanism for making a decision about whether the current test stimulus matches the sample held in memory. Several other studies have found comparable match–nonmatch effects on responses of cells in IT cortex (12–18).

The most common type of interaction effect we found was suppressive: the more similar the test stimulus was to the sample, the more the response was suppressed. For example, if the cell was normally selective for stimulus A, the response to A when it was a matching test item was substantially smaller than when A appeared as a nonmatching item in the sequence. The degree of similarity between the sample and test stimulus that caused the suppression was at a more abstract level than simply a pixel by pixel comparison, because for many cells the suppressive effect was maintained even when the sample and test stimulus were presented at different retinal locations or differed in size (11).

We found a related suppressive effect when we used stimuli in the DMS task that were initially novel to the animal (9). For about a third of the cells in IT cortex, the responses to specific sample and test stimuli declined systematically over the course of the recording session, as they became familiar. These effects were highly stimulus specific and long lasting: the reduction in response with familiarity was maintained even when several minutes and more than a hundred other stimulus presentations intervened between repetitions of a given stimulus. However, the cells did not act as “novelty detectors,” in the sense of responding to any stimulus that is novel. Rather, both novelty and stimulus features determined the response. For example, a cell that responded well to stimulus A and poorly to stimulus B when they were both novel at the start of the session might respond poorly to both A and B when they had become

familiar. A study by Brown and colleagues (17) suggests that these effects may be permanent. They find that presentation of a set of stimuli on one day leads to a reduced incidence of cells activated by those stimuli on a subsequent day. Another way of describing these result is that stimulus repetition leads to a smaller populations of activated cells, or smaller stimulus representations in the cortex.

Selection/Enhancement

Because IT responses distinguish between matching and non-matching stimuli in the DMS task, we initially proposed that the repetition suppressive mechanism was part of the neural mechanism of “working memory” (8). In the field of human psychology, the concept of working memory includes aspects of short term memory, rehearsal, a visuo-spatial “sketchpad,” and the actions of a central executive (19). Some of the important features of working memory are that it holds information only briefly before it is discarded, that it is voluntary, that it requires effort, and that it is limited in capacity. A good example is silently rehearsing a new phone number while waiting to dial it. In primates, working memory is usually studied in tasks such as DMS or delayed response, in which the monkey actively holds a specific item in memory for a short delay and then makes a behavioral judgment based on the item in memory. We reasoned that the act of holding the sample stimulus in memory decreased the sensitivity of a subpopulation of IT cells to a reoccurrence of that stimulus. When the matching stimulus occurred in the sequence, the suppressed responses of this population (perhaps in comparison to the responses of other IT cells that were unaffected by the mnemonic demands of the task) might then trigger the animal’s decision to respond to the match stimulus.

However, the picture became more complicated when we studied IT cells in a variant of our standard DMS design (10). In the standard design, a given nonmatch test stimulus was presented only once within a given trial. For example, a particular stimulus sequence in a trial might be A . . . B . . . C . . . A, in which the B and the C appeared only once in the sequence. In the new variation, one of the nonmatching test stimuli was repeated within the sequence. For example, a stimulus sequence might be A . . . B . . . B . . . A. We termed this version of the task the ABBA task. The question we asked was whether only responses to the matching test (e.g., A) would be suppressed, or would responses also be suppressed for the behaviorally irrelevant repeated nonmatch (e.g., B). In other words, was active maintenance of the sample memory necessary for suppression to occur, or would suppression occur automatically with any stimulus repetition.

The first surprise in the experiment was that when monkeys trained on the standard task were tested in the ABBA task, the monkeys incorrectly responded to the repeated nonmatch (e.g., B), not waiting for the final match. This indicated that the monkeys had always thought that the task requirement was to respond to any repeated stimulus in the sequence. The monkeys were apparently detecting simple repetition and were not actively maintaining the memory of the sample. In the original task design, only the sample stimulus was repeated in the trial, so that this erroneous strategy actually met the task requirements and was rewarded. We therefore had to train the monkeys to respond only to the matching test (e.g., A) in the ABBA task.

Once the monkeys were performing the ABBA task correctly and responded only to the matching test stimulus, we again recorded from cells in the perirhinal portion of IT cortex. We found that, as before, the responses of many cells were suppressed by the matching test stimulus (e.g., A) in the sequence. However, there was equal suppression for the repeated nonmatch (e.g., B), which was behaviorally irrelevant. This indicated that the repetition suppression effect was

caused by any type of stimulus repetition and was not specifically linked to active maintenance of the sample memory trace. Consistent with this, repetition suppression is also found when stimuli are repeated during passive fixation and in monkeys under anesthesia (20, 21). Repetition suppression is also found in monkeys given systemic scopolamine, which greatly impairs their behavioral performance in the DMS task (22).

In contrast to the cells showing match-suppression, we found that the responses of another class of IT cells were enhanced for the matching test stimulus in the ABBA sequence. For example, a cell that normally responded well to stimulus A gave a stronger response when A was a match to the sample than when it was an intervening nonmatch stimulus. Most importantly, the responses of these cells were enhanced only by the stimulus that matched the sample and not by the repetition of the nonmatch stimuli in the sequence. Thus, these cells uniquely signaled the stimulus for which the monkey was actively searching.

Parallel Memory Mechanisms

Based on these results, the perirhinal portion of IT cortex contains at least two parallel mechanisms that might mediate performance in DMS tasks, an enhancement mechanism for active working memory (Fig. 1) and a suppressive mechanism that is engaged automatically by stimulus repetition (10). The latter might be considered a type of "recency" or "novelty" memory. In monkeys performing the DMS task, both enhancement and suppression occur (for different cells) at the time of presentation of the test stimuli. Since this is the precise time at which the animal must make a decision about whether the test stimulus is matching or nonmatching, the animal might utilize either the suppression mechanism, the enhancement mechanism, or both mechanisms, depending on the specific requirements of the task. In principle, the ABBA task cannot be solved by an animal relying only on the suppressive mechanism, whereas the standard version of the DMS task might be solved by a monkey relying on either enhancement or suppression. The fact that our monkeys initially trained on the standard version initially failed to distinguish between the stimulus matching the sample and irrelevant repetitions of the nonmatch stimuli in the ABBA task suggests they may have utilized the suppression mechanism when performing the standard task. More generally, the results imply that the kind of short-term memory typically studied in behavioral and lesion studies in animals and humans may appear superficially to be a single phenomenon but in fact might be mediated by different mechanisms depending on the specific requirements of the task, training history, and perhaps individual variables.

What role does the suppressive mechanism play in memory formation? Brain imaging studies in humans using positron-emission tomography or functional MRI suggest that it plays a role in repetition priming, a type of implicit memory (23). In repetition priming, experience with an item leads to faster and better performance when subjects are required to name or identify the item at a later time (24). For example, if a subject is asked to identify briefly presented drawings of objects, they will be faster if they have seen the drawings before. Moreover, these effects occur regardless of whether the subject actually remembers seeing the drawings before. Imaging studies have found that under conditions that lead to priming, repetition of either visually presented objects or words leads to reduced activation of cortical areas, compared with stimuli that were not seen before (25, 26). It seems likely that the reduction in cortical activation is due to the repetition suppression effect and the shrinkage of the pool of activated cells. The fact that a smaller population of activated cells is associated with better task performance suggests that a smaller representation of a stimulus is a better representation, particularly if it is due to sharpened stimulus selectivity of the remaining cells. In this

view, repetition suppression is a by-product of sharpening stimulus representations in the cortex. In addition, repetition suppression may contribute to behavioral habituation and, conversely, automatic orientation to novel stimuli (see below).

Delay Activity

There remains one important requirement for solving the DMS task that is not accounted for by the properties of IT neurons described so far in this review—namely, a mechanism for maintaining the memory of the sample. One possibility is that some IT cells continue to respond to a sample stimulus throughout the delay when the stimulus is no longer present. If such cells existed, they might provide a chronic input to the other IT cells selective for the same stimuli, leading to the enhancement effect when the matching test stimulus appears. Consistent with this possibility, several studies have reported IT cells that respond selectively to particular sample stimuli and show higher maintained activity during the delay following the preferred samples (27–33).

We found counter-evidence for this idea, however, in our studies using the DMS task with multiple nonmatch test stimuli intervening between the sample and the matching test (31). In both the standard and ABBA version of the task, many cells had stimulus-selective maintained activity in the delay following the sample but this activity was "reset" to a different level following the first intervening stimulus in the sequence. Consider, for example, a cell that responds selectively to stimulus A. In an A . . . B . . . C . . . A sequence, such a cell might have high maintained activity following A, which might then be reset to a low firing level following the intervening B and a different firing level following intervening stimulus C. After the first intervening stimulus, there was no clear link between the magnitude of delay activity and the identity of the sample stimulus that the animal was actively maintaining in memory.

This result explains an otherwise contradictory report from Nakamura and Kubota (33), who found that cells in the anterior and temporal polar regions of IT cortex had stimulus-specific maintained activity in the delays following the sample in a delayed nonmatch-to-sample task. In their task, the sample stimulus was presented repeatedly in a sequence, and the monkey was rewarded for indicating when the sequence was interrupted by a stimulus that differed from the sample. For example, a trial sequence might consist of A . . . A . . . B. Thus, in their task, delay activity could be reset after each intervening stimulus as we found but would nonetheless appear to be maintained throughout the trial.

Our failure to find delay activity that bridged intervening stimuli in IT cortex suggests that delay activity, in itself, is not the mechanism that maintains the memory of the sample, at least not in IT cortex. Yet, other work indicates that IT delay activity is at least under the control of a short-term memory mechanism. Miyashita and colleagues (27, 34) recorded from IT neurons in a paired-associate task, in which the monkey was presented with a sample stimulus at the start of the trial and was rewarded for matching it to an arbitrary associate stimulus, following a delay. They found that IT delay activity anticipated the paired-associate stimulus—e.g., if the animal was taught an A–B associate pair, and the neuron was selective for stimulus B, then the neuron showed elevated activity in the delay between A and B. Likewise, we found in a DMS task that when the same sample stimulus was used for several consecutive trials in a block, the maintained activity at the start of the trial varied according to the sample that the animal learned to expect (unpublished data). Activity was highest on trials in which the animal anticipated seeing a stimulus that was the preferred stimulus for the cell. If the same cell was tested in trials with randomized samples, such that the animal could not predict which sample would be used on a given trial, the differential activity at the start of the trial was abolished.

Putting these results together indicates that maintained or delay activity in IT cortex may provide a representation of whatever object is expected or behaviorally important but only at a given moment of time.

Prefrontal Cortex

Having failed to find delay activity that bridged intervening items in IT cortex, we next turned to prefrontal cortex. There were several reasons for suspecting that cells with the appropriate properties might exist in prefrontal cortex. First, prefrontal cells have stimulus-specific delay activity in delayed response tasks (32, 35–42). Cells in the dorsolateral portion of prefrontal cortex tend to have delay activity specific for remembered locations in space whereas cells on the ventral convexity of prefrontal cortex tend to have activity specific for complex objects (42). Second, lesions or deactivation of prefrontal cortex impair performance on working memory tasks in monkeys (43–49). Third, anatomical studies reveal that prefrontal cortex has reciprocal connections with most or all of extrastriate visual cortex, and thus the anatomy is at least consistent with the idea that prefrontal cortex is a source of biasing activity in visual cortex (50–53). Finally, if the idea is correct that visual working memory requires feedback to visual cortex, biasing responses in favor of stimuli that are relevant or expected, then it makes sense for this feedback to come from structures that are not purely visual. Behavioral relevance and expectations are often defined by the behavioral context rather than by inputs into a single sensory system.

We recorded from neurons in the ventral convexity of prefrontal cortex using both our standard DMS and ABBA task (54). Consistent with previous studies (32, 35–42), it was found that many prefrontal cells responded to the visual stimuli in the task, and many of these showed stimulus specificity. Many cells also had high maintained activity during the delay intervals and, for many cells, this activity was stimulus-specific. Most importantly, stimulus-specific delay activity in prefrontal cortex did bridge the intervening stimuli in the DMS task, unlike delay activity in IT cortex. For example, a cell that had higher maintained activity in the delay immediately following sample A than it did following sample B also showed higher maintained activity in all of the subsequent delay intervals in a trial sequence consisting of A . . . B . . . C . . . D . . . A (see Fig. 2). For some cells, the overall amount of delay activity either increased or decreased with each subsequent delay interval, but the selectivity of the delay activity for a given stimulus was maintained. Thus, prefrontal cells might, in principle, provide feedback to IT cortex, biasing activity in favor of the stimulus matching the sample held in memory on a given trial (Fig. 1).

Interestingly, some prefrontal cells also showed repetition suppression and match-enhancement, like we found in IT cortex. The suppression and enhancement effects exactly paralleled our results in IT cortex, in that prefrontal cells with repetition suppression showed suppression for not only the stimulus matching the sample but also for repetitions of the nonmatching test stimuli in the ABBA task whereas prefrontal cells showing match-enhancement were only enhanced for the stimuli that matched the sample. Overall, prefrontal cells were less likely to show stimulus selective visual responses but were more likely to have stimulus-selective delay activity and had a much stronger match-enhancement effect than cells in IT cortex. Together, the results point to a more central role of prefrontal cortex in working memory (ref. 54; see ref. 32 for a review).

The results support the idea that prefrontal cortex is a source of feedback to extrastriate cortex, biasing or priming cells to respond preferentially to stimuli held in short-term memory that are expected or behaviorally relevant. In this view, feedback from prefrontal cortex contributes to the

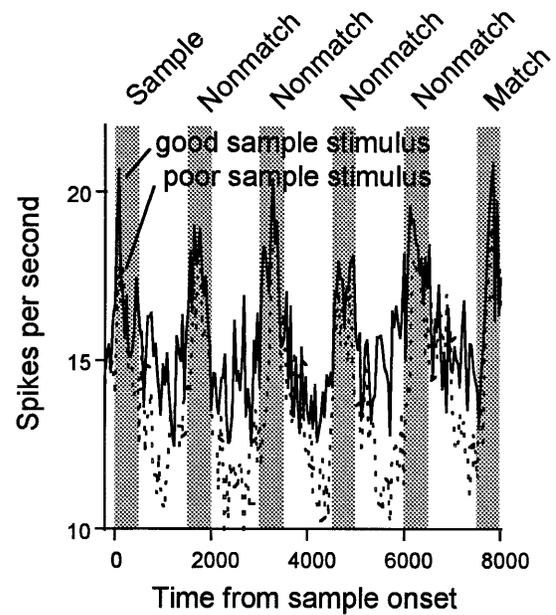


FIG. 2. Response histograms averaged from a population of 40 prefrontal neurons that had significant sample-selective delay activity. Responses are shown separately for trials in which a preferred or “good” stimulus was used as the sample and trials in which a nonpreferred or “poor” stimulus was used as the sample (bin width, 40 ms). (Adapted from ref. 54.)

match-enhancement effect in IT cortex and likely plays a role in generating delay activity in IT cortex as well. Although prefrontal cortex may not be the only source of such inputs, a direct test of its role would be to record from IT neurons during reversible deactivation of the prefrontal cortex and ask whether match-enhancement is affected. This experiment has not yet been done, but Fuster *et al.* (47) have recorded delay activity in IT cortex during cooling of prefrontal cortex. IT delay activity was not eliminated during prefrontal cooling but it became less selective. Since it appears that prefrontal delay activity is modulated by dopaminergic inputs to prefrontal cortex, it may be possible in the future to test prefrontal-temporal interactions with pharmacological manipulations (45, 55–57).

There is a striking parallel between these results on delay activity in IT and prefrontal cortex and results in the posterior parietal and prefrontal cortex. Neurons in both posterior parietal cortex and in the more dorsal portion of prefrontal cortex (area 46) have delay activity that is selective for spatial location (38, 41, 47, 58–61). Constantinidis and Steinmetz (62) have recorded from posterior parietal cells in a spatial analog of our standard DMS task, in which the animal is presented with a sample at one spatial location, followed by a sequence of stimuli at other spatial location. The animal was rewarded for indicating when a stimulus in the sequence appeared at the same location as the sample. They found that many cells had delay activity following the presentation of the sample but this activity did not survive the first intervening stimulus in the sequence. By contrast, delay activity specific for spatial location is maintained following intervening stimuli in prefrontal cortex (60, 61). Thus, prefrontal but not parietal cortex appears to maintain an explicit representation of the remembered spatial location for the length of the trial. This distinction may not apply to presaccadic activity, however, as posterior parietal activity immediately preceding a saccade to a remembered target is not disrupted by the presentation of a second target in a double saccade task (63, 64).

Biased Competition

Only a small amount of the information available on the retina can be fully processed at any given moment in time. Thus, at

some point between stimulus and response, the objects in a typical crowded scene must compete for representation, analysis, and control over behavior. This competition is biased in favor of information that is currently relevant for behavior, and the winner of the competition is said to be “attended.”

In our work, we have extended the biased competition account of attentional selection to the visual cortex itself (65). According to this hypothesis, the multiple objects present in a complex scene will activate, in parallel, their respective neural representations in the cortex. Any overlap between the neuronal populations participating in different representations presumably results in degraded information available about each individual object; however, the representations should never be completely overlapping even if some objects are identical to others, as each object will occupy a different location on the retina. The cells participating in the different representations then enter into some type of competitive interaction, in which each representation tries to suppress other representations. The competition is biased in favor of some representations, or neural populations, at the expense of others depending on the influence of many different processes, some of which are intrinsic to the visual cortex itself (bottom-up), and some of which depend on extrinsic feedback to visual cortex from structures outside of the traditional visual system (top-down). All of the memory mechanisms described in the first part of this review influence the firing rate or sensitivity of cells in visual cortex and would therefore be expected to influence any competitive interactions or attentional selection of objects.

Visual Search

A good example of the role that memory mechanisms play in attentional selection is provided by visual search. The visual search task is a commonly used attentional paradigm in psychology; in one form it is akin to finding a “face in the crowd.” The problem is to find an object in a scene when the object’s location is not known in advance. In our monkey version of visual search, the monkey was given a cue at the start of the trial, followed by a blank screen for a short delay period (66). The monkey was required to fixate the cue at the center of the display and to maintain fixation on a fixation spot at the same location during the delay. At the end of the delay, an array of two to five stimuli were presented extrafoveally, and the monkey was rewarded for making a saccadic eye movement to a stimulus (target) matching the prior cue. The location of the target in the array was varied from trial to trial, so that the monkey had to find it based on its features. The stimuli were digitized pictures of complex objects, similar to those used to study memory in IT cortex.

We recorded from cells in the same perirhinal portion of IT cortex in which we had studied the memory mechanisms described above. According to the biased competition hypothesis, the onset of the array should activate, in parallel, IT cells selective for any of the objects in the array. Eventually, however, cells that participate in the representations of nontarget objects should be suppressed.

Because IT receptive fields are typically very large, all of the objects in the array might potentially activate a given cell. Therefore, it was necessary to “label” IT responses to the different objects, so that we would know which specific object was responsible for a given cell’s response. We did this by first testing the cell’s response to each stimulus individually and then by choosing the stimuli in the array so that only one of them activated the cell strongly when presented by itself. In a two-item array, for example, we would choose one highly preferred stimulus (good stimulus) for the cell we were recording and one highly nonpreferred stimulus (poor stimulus). We could then examine the response to the array when the good stimulus was the target compared with trials when the

good stimulus was the nontarget. The physical arrays would be identical on the two types of trials; the only difference was which object was the target.

When the array was presented, cells responded initially to their preferred stimulus in the array, regardless of whether it was the target. However, about 100 ms after the onset of the response, responses to the good stimulus in the array were suppressed if it was the nontarget on that trial (Fig. 3). By contrast, when the good stimulus was the target, responses remained at a high rate up until the time of the eye movement, which occurred about 100 ms or so after responses to the target and nontarget diverged. That is, by the time the eye movement was made, only cells participating in the representation of the target remained active. This is just what the biased competition idea would predict: an initial parallel activation of cells participating in the representations of all object in the scene, followed by suppression of cells that represent nontarget (irrelevant) objects.

In addition to evidence for the competition, there was also evidence for the “bias.” In the delay between the offset of the cue and the onset of the array, many IT cells showed higher delay activity if their preferred stimulus was the cue-target. These cells acted as though they had been “primed” or

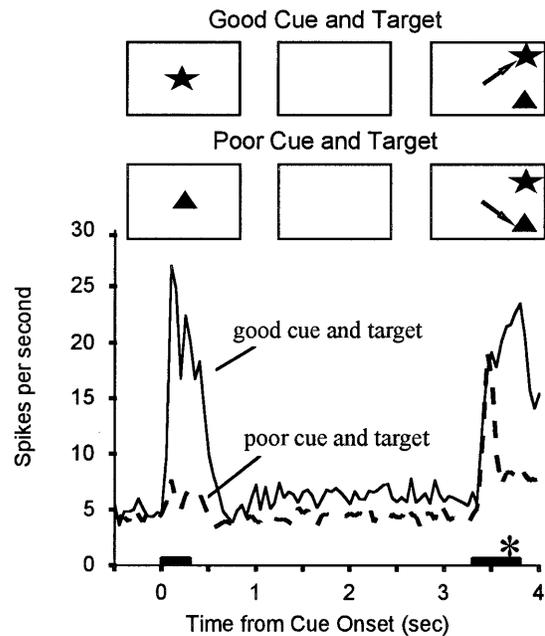


FIG. 3. Response histograms averaged from a population of 22 IT neurons recorded while the monkey performed a visual search task. A cue stimulus was briefly presented at the start of the trial, followed by a blank delay period during which the animal maintained fixation at the center of the display. At the end of the delay, a choice array containing two stimuli at random locations was presented and the animal was rewarded from making a saccadic eye movement to the stimulus that matched the cue (target). The presentation periods for the cue and array are indicated by horizontal bars. The cues were chosen for each cell such that one would activate a cell when presented alone (“good” cue) and one would only poorly activate the cell when presented alone (“poor” cue). Responses to identical choice arrays are shown separately for trials with the good versus poor stimulus for the recorded cells used as the cue. When the good cue was used, activity was higher during the delay period, and responses to the choice array remained high. When the poor cue was used, activity was lower during the delay period, and responses to the identical choice array were suppressed approximately 200 ms after the onset of the array. Thus, only cells selective for the target stimulus remained active, all other cells being suppressed. The saccadic eye movement to the target (asterisk) began about 300 ms after the onset of the array, well after responses to the nontarget stimuli were suppressed. (Adapted from refs. 65 and 66.)

sensitized to respond to their preferred stimulus when it would appear as the target in the array. Consistent with this, some cells showed an enhanced response to the onset of the array when their preferred stimulus was the target (unpublished data). This is the same “enhancement effect” that we found in our studies of working memory in IT cortex, that we believed is caused at least in part by feedback from prefrontal cortex.

Interestingly, some cells showed a reduced response to the onset of the array when their preferred stimulus was the target (unpublished data). This is presumably the same “repetition suppression” that we found in our short-term memory studies, as the target stimulus was presented twice on each trial—once as the cue and once in the array. Because the suppression mechanism biased responses in favor of the nontarget stimulus (the only unique stimulus in the trial), it worked against both the enhancement mechanism and, presumably, the animal’s goal in the task.

Conclusions

Mechanisms for learning and memory are continually modulating and altering the representations of stimuli in the cortex. We have considered only three short-term memory mechanisms in this brief review, namely repetition suppression, enhancement, and delay activity. These are hardly exhaustive, and, in particular, we have not touched on known mechanisms for long-term plasticity and perceptual learning. Besides the contributions these mechanisms make to performance of traditional memory demanding tasks, such as DMS, there are perhaps equally important contributions to tasks that have been traditionally considered to be attentional. Indeed, we have argued that “attention” derives, at least in part, from the impact of memory mechanisms on cortical sensory mechanisms, which are intrinsically competitive in nature.

1. Desimone, R. (1992) *Science* **258**, 245–246.
2. Desimone, R. & Ungerleider, L. G. (1989) in *Neural Mechanisms of Visual Processing in Monkeys*, eds. Boller, F. & Grafman, J. (Elsevier, New York), Vol. 14, pp. 267–299.
3. Ungerleider, L. G. & Mishkin, M. (1982) in *Two Cortical Visual Systems*, eds. Ingle, D., Goodale, M. A. & Mansfield, R. J. W. (MIT Press, Cambridge, MA), pp. 549–586.
4. Ungerleider, L. G. & Haxby, J. V. (1994) *Curr. Opin. Neurobiol.* **4**, 157–165.
5. Desimone, R., Albright, T. D., Gross, C. G. & Bruce, C. (1984) *J. Neurosci.* **4**, 2051–2062.
6. Desimone, R. & Gross, C. G. (1979) *Brain Res.* **178**, 363–380.
7. Tanaka, K. (1996) *Annu. Rev. Neurosci.* **19**, 109–139.
8. Miller, E. K., Li, L. & Desimone, R. (1991) *Science* **254**, 1377–1379.
9. Li, L., Miller, E. K. & Desimone, R. (1993) *J. Neurophysiol.* **69**, 1918–1929.
10. Miller, E. K. & Desimone, R. (1994) *Science* **263**, 520–522.
11. Lueschow, A., Miller, E. K. & Desimone, R. (1994) *Cereb. Cortex* **5**, 523–531.
12. Riches, I. P., Wilson, F. A. & Brown, M. W. (1991) *J. Neurosci.* **11**, 1763–1779.
13. Eskandar, E. N., Richmond, B. J. & Optican, L. M. (1992) *J. Neurophysiol.* **68**, 1277–1295.
14. Eskandar, E. N., Optican, L. M. & Richmond, B. J. (1992) *J. Neurophysiol.* **68**, 1296–1306.
15. Vogels, R. & Orban, G. A. (1993) *Prog. Brain Res.* **95**, 427–444.
16. Baylis, G. C. & Rolls, E. T. (1987) *Exp. Brain Res.* **65**, 614–622.
17. Fahy, F. L., Riches, I. P. & Brown, M. W. (1993) *Exp. Brain Res.* **96**, 457–472.
18. Vogels, R. & Orban, G. A. (1994) *J. Neurophysiol.* **71**, 1428–1451.
19. Baddeley, A. (1986) *Working Memory* (Clarendon, Oxford).
20. Miller, E. K., Gochin, P. M. & Gross, C. G. (1991) *Vis. Neurosci.* **7**, 357–362.
21. Vogels, R., Sary, G. & Orban, G. A. (1995) *Vis. Neurosci.* **12**, 207–214.
22. Miller, E. K. & Desimone, R. (1993) *NeuroReport* **4**, 81–84.
23. Ungerleider, L. G. (1995) *Science* **270**, 769–775.
24. Tulving, E. & Schacter, D. L. (1990) *Science* **247**, 301–306.
25. Squire, L. R., Ojemann, J. G., Miezin, F. M., Petersen, S. E., Videen, T. O. & Raichle, M. E. (1992) *Proc. Natl. Acad. Sci. USA* **89**, 1837–1841.
26. Buckner, R. L., Petersen, S. E., Ojemann, J. G., Miezin, F. M., Squire, L. R. & Raichle, M. E. (1995) *J. Neurosci.* **15**, 12–29.
27. Sakai, K. & Miyashita, Y. (1991) *Nature (London)* **354**, 152–155.
28. Miyashita, Y. (1988) *Nature (London)* **335**, 817–820.
29. Miyashita, Y. & Chang, H. S. (1988) *Nature (London)* **331**, 68–70.
30. Fuster, J. M. & Jervey, J. P. (1981) *Science* **212**, 952–955.
31. Miller, E. K., Li, L. & Desimone, R. (1993) *J. Neurosci.* **13**, 1460–1478.
32. Fuster, J. M. (1995) *Memory in the Cerebral Cortex* (MIT Press, Cambridge, MA).
33. Nakamura, K. & Kubota, K. (1995) *J. Neurophysiol.* **74**, 162–178.
34. Sakai, K. & Miyashita, Y. (1993) *Curr. Opin. Neurobiol.* **3**, 166–170.
35. Sawaguchi, T., Matsumura, M. & Kubota, K. (1988) *Behav. Brain Res.* **31**, 193–198.
36. Sawaguchi, T., Matsumura, M. & Kubota, K. (1988) *Neurosci. Res.* **5**, 465–473.
37. Fuster, J. M. (1973) *J. Neurophysiol.* **36**, 61–78.
38. Funahashi, S., Bruce, C. J. & Goldman-Rakic, P. S. (1989) *J. Neurophysiol.* **61**, 331–349.
39. Fuster, J. M., Bauer, R. H. & Jervey, J. P. (1982) *Exp. Neurol.* **77**, 679–694.
40. Funahashi, S., Bruce, C. J. & Goldman-Rakic, P. S. (1993) *Biomed. Res.* **14**, 85–88.
41. Funahashi, S., Chafee, M. V. & Goldman-Rakic, P. S. (1993) *Nature (London)* **365**, 753–756.
42. Wilson, F. A. W., O Scalaidhe, S. P. & Goldman-Rakic, P. S. (1993) *Science* **260**, 1955–1958.
43. Mishkin, M. & Manning, F. J. (1978) *Brain Res.* **143**, 313–323.
44. Mishkin, M. (1957) *J. Neurophysiol.* **20**, 615–622.
45. Sawaguchi, T. & Goldman-Rakic, P. S. (1991) *Science* **251**, 947–950.
46. Funahashi, S., Bruce, C. J. & Goldman-Rakic, P. S. (1993) *J. Neurosci.* **13**, 1479–1497.
47. Fuster, J. M., Bauer, R. H. & Jervey, J. P. (1985) *Brain Res.* **330**, 299–307.
48. Shindy, W. W., Posley, K. A. & Fuster, J. M. (1994) *Cereb. Cortex* **4**, 443–450.
49. Bauer, R. H. & Fuster, J. M. (1976) *J. Comp. Physiol. Psychol.* **90**, 293–302.
50. Barbas, H. & Pandya, D. N. (1989) *J. Comp. Neurol.* **286**, 353–375.
51. Barbas, H. (1988) *J. Comp. Neurol.* **276**, 313–342.
52. Ungerleider, L. G., Gaffan, D. & Pelak, V. S. (1989) *Exp. Brain Res.* **76**, 473–484.
53. Webster, M. J., Bachevalier, J. & Ungerleider, L. G. (1994) *Cereb. Cortex* **4**, 470–483.
54. Miller, E. K., Erickson, C. A. & Desimone, R. (1996) *J. Neurosci.* **16**, 5154–5167.
55. Sawaguchi, T., Matsumura, M. & Kubota, K. (1990) *J. Neurophysiol.* **63**, 1401–1412.
56. Sawaguchi, T., Matsumura, M. & Kubota, K. (1990) *J. Neurophysiol.* **63**, 1385–1400.
57. Williams, G. V. & Goldman-Rakic, P. S. (1996) *Nature (London)* **376**, 549–550.
58. Kojima, S. & Goldman-Rakic, P. S. (1982) *Brain Res.* **248**, 43–49.
59. Gnadt, J. W. & Andersen, R. A. (1988) *Exp. Brain Res.* **70**, 216–220.
60. di Pellegrino, G. & Wise, S. P. (1993) *J. Neurosci.* **13**, 1227–1243.
61. di Pellegrino, G. & Wise, S. P. (1993) *Somatosens. Motor Res.* **10**, 245–262.
62. Constantinidis, C. & Steinmetz, M. A. (1996) *J. Neurophysiol.* **76**, 1352–1355.
63. Barash, S., Bracewell, R. M., Fogassi, L., Gnadt, J. W. & Andersen, R. A. (1991) *J. Neurophysiol.* **66**, 1109–1124.
64. Barash, S., Bracewell, R. M., Fogassi, L., Gnadt, J. W. & Andersen, R. A. (1991) *J. Neurophysiol.* **66**, 1095–1108.
65. Desimone, R. & Duncan, J. (1995) *Annu. Rev. Neurosci.* **18**, 193–222.
66. Chelazzi, L., Miller, E. K., Duncan, J. & Desimone, R. (1993) *Nature (London)* **363**, 345–347.