Neuron
Article

Remembering Visual Motion: Neural Correlates of Associative Plasticity and Motion Recall in Cortical Area MT

Anja Schlack[1],* and Thomas D. Albright[1],*
[1] Systems Neurobiology Laboratories, The Salk Institute for Biological Studies, La Jolla, CA 92037, USA
*Correspondence: anja@salk.edu (A.S.), tom@salk.edu (T.D.A.)
DOI 10.1016/j.neuron.2007.02.028

SUMMARY

The pictorial content of visual memories recalled by association is embodied by neuronal activity at the highest processing stages of primate visual cortex. This activity is elicited by top-down signals from the frontal lobe and recapitulates the bottom-up pattern normally obtained by the recalled stimulus. To explore the generality and mechanisms of this phenomenon, we recorded motion-sensitive neurons at an early stage of cortical processing. After monkeys learned to associate directions of motion with static shapes, these neurons exhibited unprecedented selectivity for the shapes. This emergent shape selectivity reflects activation of neurons representing the motion stimuli recalled by association, and it suggests that recall-related activity may be a general feature of neurons in visual cortex.

INTRODUCTION

Objects that are frequently seen together become associated in memory, such that the sight of one object readily brings to mind the image of the other. The acquisition of such memories is believed to result from the establishment or strengthening of connections between neurons that represent the associated stimuli (Hebb, 1949; James, 1910; Miyashita, 1993). By this means, a neuron that represents a specific visual object, via ''bottom-up'' signaling from the sensory periphery, may also be activated by a second object that elicits recall of the first, via ''top-down'' signaling through newly established connections. Tests of this hypothesis have focused on the primate inferior temporal (IT) cortex because of its position at the pinnacle of the cortical visual processing stream (Felleman and Van Essen, 1991), the selectivity of its neurons for complex objects (Desimone et al., 1984), and its close connections with medial temporal lobe structures critical for learning and memory (Lavenex and Amaral, 2000; Lavenex et al., 2002). The selectivity of many IT neurons does change as predicted during associative learning (Erickson and Desimone, 1999; Messinger et al., 2001; Sakai and Miyashita, 1991).

We investigated whether associative recall-related activity might be a general property of cortical visual neurons, including those at early processing stages, which represent simple attributes of visual objects, such as direction of motion, rather than the objects themselves. To this aim, we trained rhesus monkeys to associate specific directions of moving dot patterns with specific directions of a static arrow shape (Figure 1). After animals learned these pairings, we recorded responses to both stimulus types from neurons in cortical visual area MT, which are known to be highly selective for the direction of stimulus motion but largely unresponsive to static stimuli (Albright, 1984). We predicted that the associative learning would be paralleled by changes in the sensitivity of directionally selective MT neurons, such that these neurons would become activated by the static arrows associated with preferred directions of motion.

RESULTS

Pretraining Neuronal Assessment

Our hypothesis presumes that MT neurons are not selective for the direction of static arrow stimuli prior to associative training. Although nothing in the extensive literature on MT response properties would suggest otherwise (for review see Albright, 1993; Born and Bradley, 2005), we nonetheless began our investigation by assessing whether MT neurons selective for motion direction were also selective for arrow direction. This assessment was made on a pretraining (''baseline'') sample of 103 MT neurons (see Experimental Procedures for details). Each neuron was tested with four directions of motion (up, down, left, right) and the corresponding four arrow directions. As expected, the vast majority (86%) of MT neurons were tuned for direction of motion. Less than 4% of the same population, a level consistent with chance probability, exhibited significant selectivity for the direction of static arrows.

Posttraining Neuronal Assessment

Animals were trained to associate two directions of motion (up and down) with two static arrow directions (Figure 1A).

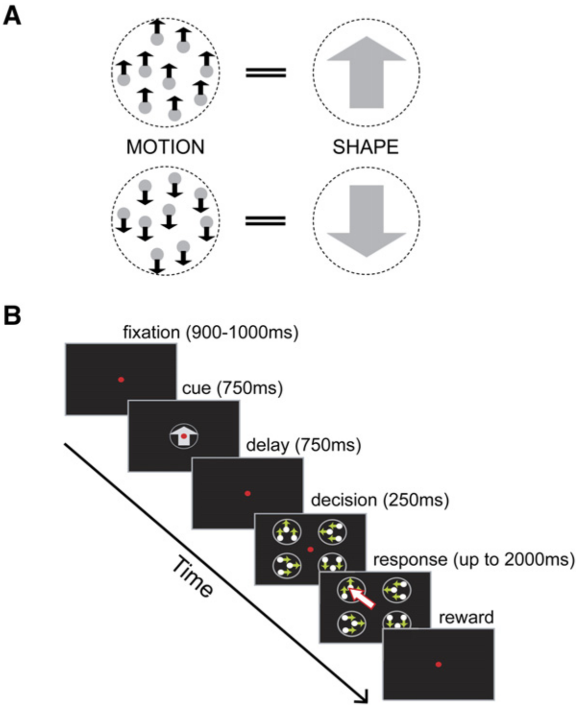

Figure 1. Visual Stimuli and Behavioral Paradigm for Association Learning

(A) Monkey A learned to associate up and down motions with up and down arrows, as illustrated. Monkey B learned the opposite pairings, i.e., arrows inverted relative to motion directions.

(B) Paired association task used to train motion-arrow pairings. Trial sequence is portrayed as a series of temporal frames. Each frame represents the video display and operant response (eye movement to chosen stimulus). Either motion or arrow stimuli could serve as the cue. Choice stimuli were always of the other type.

Training proceeded using a standard visual paired association task (Figure 1B) (Gaffan and Bolton, 1983; Messinger et al., 2001, 2005). Once animals learned the specified associations to criterion performance (85% correct for five consecutive days), we reassessed neuronal responses to motion and arrow direction using the same neurophysiological procedures that were applied prior to training. The prevalence of directional motion selectivity was unaffected by the training regime. Directional selectivity for the static arrow stimuli, however, became far more common among MT neurons (39/210, 19%). Moreover, among neurons that exhibited significant selectivity for arrow direction, the strength of that selectivity following training (quantified as average response vector magnitude; see Experimental Procedures) greatly exceeded (0.44, n = 39) that observed prior to training (0.17, n = 4).

Example Neuron Data

Patterns of selectivity for one MT neuron recorded after the monkey learned the motion-arrow association are illustrated in Figure 2. The top row in panel (A) contains trial-by-trial spike rasters and cumulative spike-density functions corresponding to the four directions of stimulus motion. The mean firing rate for each direction is plotted in polar form in Figure 2B (red curve). As expected, the cell was strongly tuned for direction of motion, with a preferred direction lying between leftward and downward (preferred direction = 216°).

The bottom row in Figure 2A contains data for the corresponding four static arrows. The animal had learned to explicitly associate up- and downward pointing arrows with up- and downward motions, respectively. Remarkably, the neuronal responses to these two static stimuli (bottom row, first and third columns) closely mirrored the responses to the associated motions: the arrow direction paired with downward motion elicited a strong response, whereas the arrow direction paired with upward motion elicited a weak response. The firing rates for each of the four arrow stimuli are plotted in Figure 2B (blue curve), with polar angle now representing the direction of the corresponding motion stimulus. The arrow direction tuning curve and preferred direction (223°) were nearly identical to those for direction of motion.

Population Data: Comparison of Strengths of Response and Directional Selectivity

To examine the relative response magnitudes and strengths of directional selectivity for motion and arrow stimuli, we performed two analyses. First, relative magnitudes of responses to motions and arrows were assessed by computing averaged response functions for preferred and null directions along the vertical axis. This analysis was restricted to neurons selective for both stimulus types (see Experimental Procedures), and the response functions (Figure 3) naturally reveal significant directional selectivity for both motions and arrows. However, the preferred direction responses elicited by arrows tended to be smaller, on average, than those elicited by motion. Although this method of examination is advantageous because it directly conveys response magnitudes to the relevant stimuli, it is limited because those stimuli (upward and downward directions) do not necessarily correspond to the preferred direction of a given neuron. It follows that a comparison of the response rates to up- versus downward directions will likely yield an underestimate of the strength of directional selectivity for each stimulus type.

We thus adopted a second procedure to assess strength of selectivity along the estimated preferred axis of motion. This procedure treats the response rate in each of the four stimulus directions as a vector magnitude and then computes the average of these four response vectors for each neuron. The magnitudes of these mean response vectors for each neuron are plotted in Figure 4, for neurons selective for both stimulus types. This comparison reveals that the population average vector magnitude for arrows (0.44) was only slightly less pronounced than that for motion (0.67). The aforementioned comparison of mean response rates to motion and arrow stimuli

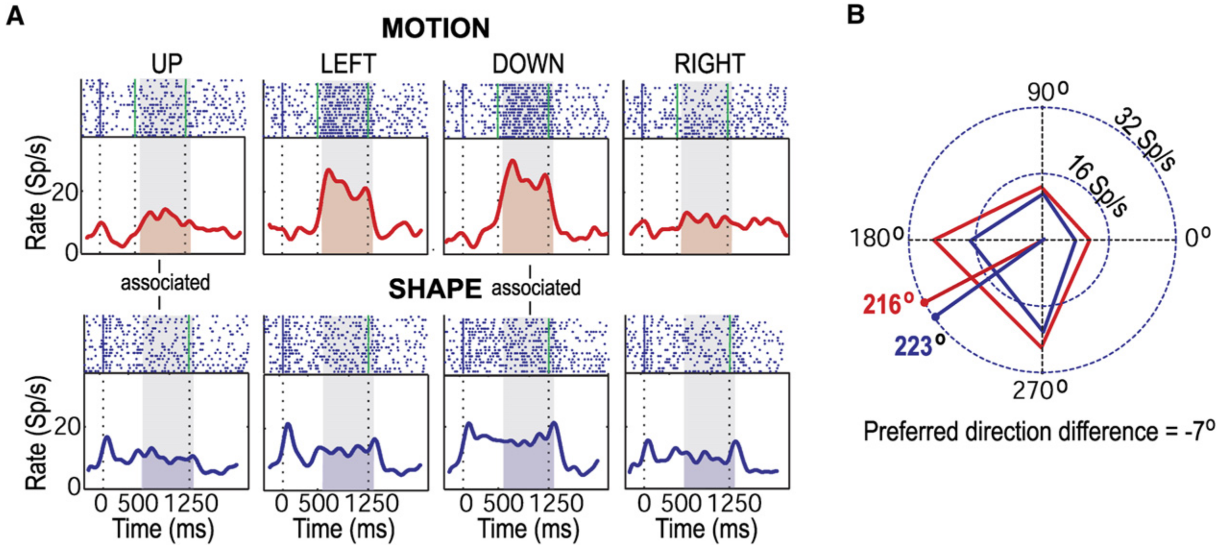

Figure 2. Selectivity of MT Neurons following Paired Association Learning

All illustrated neuronal data were collected following associative training on the task shown in Figure 1B, but during additional behavioral trials in which monkeys were simply required to fixate a central target (see Experimental Procedures).

(A) Data from representative MT neuron. (Top row) Responses to four motion directions. Spike raster displays of individual trial responses are plotted above cumulative spike-density functions. Vertical dashed lines correspond from left to right to stimulus onset, motion onset, and stimulus offset. Gray rectangle indicates analysis window (see Experimental Procedures). The cell was highly directionally selective (Rayleigh test, $p < 0.001$). (Bottom row) Responses to four static arrows. The animal previously learned to associate up and down pointing arrows with up and down motions, respectively (first and third columns of spike-density functions). Plotting conventions are same as in upper row. The cell was highly selective for arrow direction (Rayleigh test, $p < 0.001$).

(B) Mean responses of example neuron from panel (A) to motion directions (red curve) and corresponding static arrow directions (blue curve), indicated in polar format. Preferred directions for the two stimulus types (red and blue vectors) are nearly identical (7° difference).

suggests that the weaker directional selectivity for arrows is largely a consequence of the smaller responses elicited by these stimuli.

Population Data: Comparison of Preferred Directions

The marked parallels between motion and shape tuning for the neuron highlighted in Figure 2 suggest that the emergent tuning for arrow direction indeed reflects the learned associations. To test this hypothesis more generally, we first calculated the preferred directions for motion and arrow stimuli for each neuron that exhibited significant tuning for both. We then computed the angular difference between these measures. Our hypothesis predicts similar tuning for arrows and motions following learning, which should be manifested as a distribution of motion-arrow angular differences centered at 0°. Figure 5A illustrates the observed distribution of difference measures (plotted as unit vectors in polar format) for the monkey that learned the association illustrated in Figure 1A (upward and downward motions paired with upward and downward pointing arrows, respectively). Our prediction is borne out in the distribution, which is clustered around 0° (mean = 9°, v-test, $p < 0.001$, Batschelet, 1981). Importantly, values rarely exceeded ±90°, which means that the emergent tuning for static arrows nearly always respected the sign of the trained motion-arrow pairs.

Tuning Depends upon Learned Associations, Not Intrinsic Stimulus Properties

In this study, we paired motion directions with directional arrows, in part because that pairing conveniently exploits a graphical association familiar to the human experimenters. Our hypothesis, however, predicts that the observed plasticity in area MT should be independent of the specific shape stimulus used. We tested this hypothesis by training one of our monkeys on the reverse association, in which the direction of the static arrow was inverted relative to the direction of motion. If neuronal tuning reflects learning of this reversed association, we should expect preferred directions for motion and arrow stimuli to differ by 180°. Indeed, the distribution of motion-shape angular differences for this experiment (Figure 5B) clustered around the predicted value of 180° (mean = 143°, v-test, $p < 0.05$, Batschelet, 1981). This result demonstrates that tuning for static arrows reflects the learned association with motion and is not dependent upon the intrinsic properties of those shapes.

"Rule" Generalization to Untrained Arrows

The arrow and motion stimuli used in these experiments were drawn from two parallel periodic dimensions that are linked by an arbitrary pairing rule. Such parallel sensory dimensions are ubiquitous in normal human experience, and observers naturally map continuous values

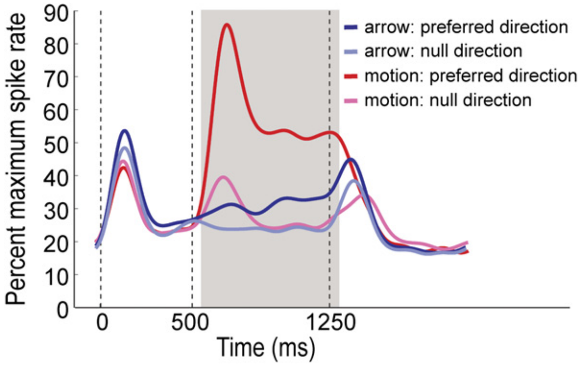

Figure 3. Population Data from MT (n = 39)
Cumulative spike-density functions reflect average responses to preferred and null directions for neurons selective for both motion and arrow direction (curves normalized to maximum response). Plotting conventions similar to individual panels in Figure 2A. Initial transient responses (at t = ~100 ms) reflect stimulus onset and are similar across stimulus type and direction. Stimulus motion began at 500 ms, causing a second response transient. Arrow stimuli remained static throughout presentation interval. Note that this analysis of arrow tuning favors neurons with significant tuning within the analysis window (580–1330 ms after stimulus onset), as opposed to other time windows following stimulus onset.

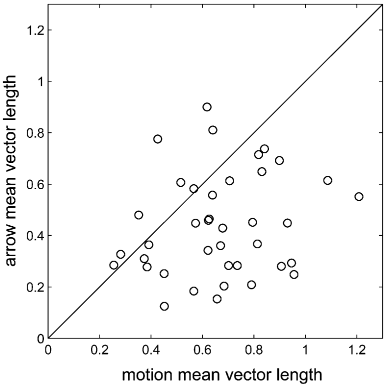

Figure 4. Population Comparison of Strength of Directional Selectivity for Motion versus Arrow Stimuli
Each data point represents the average direction tuning vector magnitude for motion versus arrow stimuli for an individual MT neuron. These direction tuning vectors were computed for each neuron by averaging four vectors with direction equal to stimulus direction and magnitude equal to mean response rate (see Experimental Procedures). Each vector magnitude was normalized relative to the maximum response rate for the cell and stimulus type. Data are shown only for neurons that were significantly selective for both motion and arrow stimuli (n = 39). This measure of selectivity reflects the estimated response bias along the preferred axis of motion and is, on average, slightly greater for motion (0.67) versus arrow (0.44) stimuli. This difference appears largely due to the difference in response magnitudes to these stimuli (see Figure 3).

from one dimension onto another following an imposed rule (Goldstone, 1998)—think, for example, of the relationship between the volume of sound and the rotational angle of a knob, or the depth of the gas pedal and the forward velocity of the car. In view of these considerations, we asked whether the pattern of neuronal responses to arrow stimuli generalized to arrow directions that were not explicitly paired with motion directions in our training regime, but were nonetheless "linked" by the pairing rule that defined the explicitly trained motion-arrow pairs.

To illustrate, consider the example cell shown in Figure 2. Under our training regime, the animal was explicitly trained to associate up- and downward motions with the up- and downward directed arrows (Figure 1A). As we have noted, the arrow directional selectivity observed for this example neuron reflects the explicit arrow-motion pairings (Figure 2A, bottom row, compare first and third columns). By contrast, no pairings of horizontally directed arrows and motions were explicitly trained. The responses of our example neuron were nonetheless tuned for the horizontally directed arrows (Figure 2A, bottom row, compare second and fourth columns). Moreover, the horizontal arrow direction preference corresponds to what one would expect if the neuron generalized to 90° rotated versions of the explicitly trained motion-arrow pairs (Figure 1), i.e., if arrow tuning reflected the rule that characterized the explicitly trained pairings.

We wondered whether generalization to nontrained stimulus pairs was common among our sampled MT

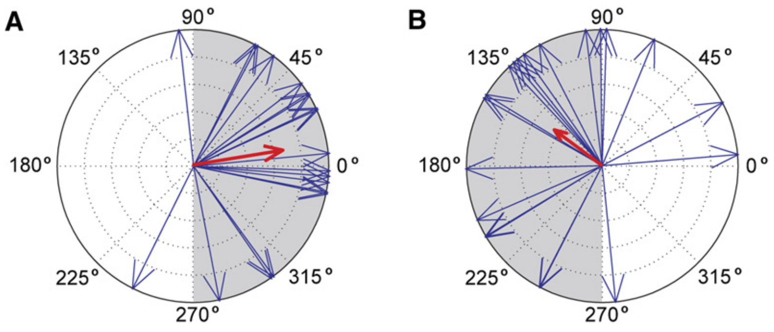

Figure 5. Comparison of Motion and Arrow Direction Preferences
Each vector represents the angular difference between preferred directions for motion and arrow stimuli, for each neuron selective for both.
(A) Monkey A learned the motion-arrow association illustrated in Figure 1A. Average of distribution (n = 19) is 9° (red vector). Shaded hemicircle indicates values that respect the sign (i.e., preferred motion direction same as preferred arrow direction) of the trained pairings.
(B) Monkey B learned the motion-arrow association opposite from monkey A. Average of distribution (n = 20) is 143°. Plotting conventions same as panel (A).

neurons. Specifically, we asked whether preference for left- versus rightward pointing arrows was congruent with preference for left- versus rightward motions, in accordance with generalization over a 90° rotation from the trained (i.e., up- and downward) motion-arrow pairs. As a reference point for this analysis, we first examined data from the subset of neurons selective for both motion and arrow direction in the vertical (i.e., trained) dimension (see Experimental Procedures for details). The high incidence (82%) of congruent direction preferences in this population (Figure 6A, left bar) honors the effect shown in Figure 5 and reinforces the conclusion that explicitly learned motion-arrow pairings lead to predictable shape tuning in MT.

The bar graph in Figure 6B illustrates data from the subset of neurons selective for both motion and arrow direction in the horizontal dimension. Surprisingly, despite the fact that horizontal motion-arrow pairings were not explicitly trained, 69% of the selective neurons exhibited an arrow direction preference congruent with the motion preference. These observations demonstrate that neuronal response changes resulting from learned stimulus pairings can generalize to other stimulus pairings of the same type. In other words, these neurons appeared to manifest the rule behind the learned associations.

Emergent Arrow Selectivity Was Robust over Subjects, Time, and Receptive Field Location

The incidence of congruent selectivity for motion and static arrow direction was highly consistent across animal subjects (both of whom learned the behavioral motion-arrow associations to criterion). Of the 33 neurons identified as selective for motion direction and static arrow direction along the vertical axis (Figure 6A), 19 were recorded in monkey A and 14 in monkey B. Of those 27/33 (82%) neurons that exhibited congruent vertical preferences for motion and arrow (i.e., both up or both down for monkey A, opposite for monkey B), 16/19 (84%) were from monkey A and 11/14 (79%) were from monkey B. Similarly, of the 32 neurons identified as selective for motion direction and static arrow direction along the horizontal axis (Figure 6B), 22 were from monkey A and 10 from monkey B. Of those 22/32 (69%) neurons that exhibited congruent horizontal preferences for motion and arrow (i.e., both left or both right for monkey A, opposite for monkey B), 15/22 (68%) were from monkey A and 7/10 (70%) were from monkey B.

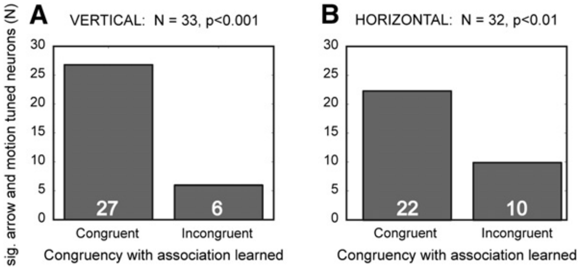

Figure 6. Generalization of Neuronal Associative Plasticity to Untrained Motion-Arrow Pairs

Data from both monkeys were similar and pooled (see Results for details).

(A) Data from trained motion-arrow pairs (vertical stimuli). Neurons were selected for significantly different responses to up- versus downward motion and significantly different responses to up- versus downward arrows (Wilcoxon test, $p < 0.05$). Bar graph indicates fraction for which up versus down direction preferences for motion and arrow stimuli were congruent (Sign test, $p < 0.001$), consistent with explicitly trained pairs.

(B) Data from untrained motion-arrow pairs (horizontal stimuli). Neurons were selected for significantly different responses to left- versus rightward motion and significantly different responses to left- versus rightward arrows (Wilcoxon test, $p < 0.05$). Bar graph indicates fraction for which left versus right direction preferences for motion and arrow stimuli were congruent (Sign test, $p < 0.001$), consistent with a rotated form of the motion-arrow pairing "rule" that applied to the trained (i.e., vertical) pairs.

Our assessment of postlearning neuronal selectivity for motion direction and static arrow direction took place during a 4–5 week period in which the animals no longer performed the motion-arrow association task. We examined the possibility that the incidence of arrow tuning declined over this period in which the learned associations were no longer behaviorally maintained. We partitioned data from each animal into two groups corresponding to the first and second halves of the neurophysiological assessment period. The numbers of significantly arrow tuned neurons were nearly identical during the two halves, and this was true for both animals (monkey A: 11 first, 9 second half; monkey B: 9 first, 10 second half).

We assessed the incidence of associative learning effects across the range of visual field eccentricities of the recorded neurons. Our sampled receptive field centers ranged from 1° to 14° from the center of gaze. The fraction of recorded MT neurons exhibiting postlearning static arrow tuning was uniform across this range (χ^2, $p < 0.05$).

Consideration of Potential Eye-Movement Artifacts

The arrow stimuli used in our experiments were presented in the CRF during nominal ocular fixation of a central target, but it is impossible to eliminate small-amplitude saccadic and smooth eye movements by these behavioral means. MT neurons exhibit directionally selective responses to retinal slip, and they do not distinguish between image displacement resulting from motion of an object in the visual scene or movement of the eyes across a static scene (Newsome et al., 1988). In principle, therefore, the observed responses to static arrow stimuli could have resulted from small eye movements that yielded retinal motion in the receptive field. Because the arrow responses were selective for arrow direction, such confounding eye movements would need to have been present for some arrow directions and not others. Although we consider this an unlikely possibility, we nonetheless attempted to rule it out empirically by analyzing the magnitude and directional properties of small eye movements that occurred during presentation of arrow stimuli.

For each neuron exhibiting postlearning arrow directional selectivity, we measured four eye-movement parameters during each stimulus presentation (see Experimental Procedures): (1) mean eye position, (2) standard deviation of eye position, (3) direction of fixational saccades, and (4) direction of smooth drift eye movements. (Note that trials were aborted and data not analyzed in cases where eye position deviated more than 1.0° from the central fixation target.) Neither the average eye position nor the position variance varied significantly as a function of arrow direction (rank sum test, $p < 0.05$), suggesting that retinal slip tied to these variables cannot account for the observed neuronal tuning. We also found that there was no significant effect of arrow direction on direction of small fixational saccades or smooth ocular drift (ANOVA, $p < 0.05$). In other words, when they occurred, the direction of small eye movements was unrelated to the direction of the arrow stimulus in the CRF, demonstrating that the associated retinal slip cannot account for the observed neuronal tuning.

DISCUSSION

After monkeys learned to associate directions of motion with directions of a static pointing arrow, we observed unprecedented selectivity for arrow direction among a subset of neurons in cortical area MT of the trained animals. This pattern of selectivity corresponded to that predicted from the associated motion stimuli. The apparent emergence of arrow tuning constitutes evidence for neuronal plasticity, which we conclude to be a product of associative learning. This associative neuronal plasticity is particularly notable for its presence in an early cortical visual area and for the extraordinary rule-based generalization it reveals. We consider the implications of these discoveries in greater detail below, along with a discussion of the functional significance of these properties and mechanisms that may give rise to them.

Does Neuronal Activity Reflect Recalled Motion?

Our discovery of emergent selectivity for static patterns in visual area MT naturally raises questions about what the activity represents. Two possibilities warrant consideration. First, it may be that through our behavioral conditioning paradigm MT neurons become co-opted for the processing of novel stimulus features. According to this argument, motion processing is the default operation in MT, but the inherent plasticity of cortex allows these neurons to take on other roles as dictated by the statistics of the observer's environment. At its core, this view supposes that the observed activity elicited in MT by a static arrow constitutes a neuronal representation of the arrow. Although this possibility deserves further study, we note that it defies the not unreasonable belief that properties of early visual neurons must remain stable in order to yield a stable interpretation of the world (Van Wezel and Britten, 2002).

The second possibility is that the emergent responses of MT neurons to static shapes play no direct role in the processing of those shapes, but rather represent the motions recalled by association. This is essentially the interpretation that has been offered to account for the effects of associative learning on visual response properties of neurons in inferior temporal cortex (Messinger et al., 2001; Sakai and Miyashita, 1991; Miyashita, 1993). According to this view, the static arrow stimuli are at all times represented by regions of cortex other than area MT. Through a learning-induced chain of connectivity, however, a given arrow stimulus ultimately elicits selective "top-down" activation of those MT neurons that represent the associated motion. We propose that this activation is the neuronal embodiment of pictorially recalled motion—motion imagery—which is represented in the same cortical region and by the same neuronal code as the original motion stimulus. Importantly, this view maintains that area MT remains stably committed to motion processing, with recognition that the same motion representations may become activated by either bottom-up or top-down signals.

Support for the view that postlearning responses to static arrows in area MT reflect motion recall comes from studies of the human perceptual phenomenon known as "implied motion," in which a static image drawn from a moving sequence—such as an animal in a predatory pounce or a swinging hammer—elicits pictorial recall of the complete motion sequence. These top-down effects can be objectified by their ability to interact perceptually with bottom-up motion signals (Freyd, 1987), and they presumably develop through experience, in which the static elements of a motion sequence are naturally associated with the movement itself. Consistent with the results reported herein and with the motion imagery interpretation, Kourtzi and Kanwisher (2001) found that static images that imply motion elicited activity in human cortical area MT+.

Finally, we acknowledge that it is impossible to rule strongly on these different interpretations, given our results thus far. We present them as logical alternatives with hope that they may stimulate further investigation and, ultimately, resolution.

Function of Recall-Related Activity in Visual Cortex: Drawing on Learned Environmental and Behavioral Regularities to Interpret Ambiguous Sensory Signals

Following a line of argument traceable to Helmholtz' (Helmholtz and Southall, 1924) concept of "unconscious inference," raised again by James (1890), Hebb (1949, 1968), Barlow (1990), and others (e.g., Davies et al., 1982) and pressed recently by Backus (Haijiang et al., 2006), we suggest that perceptual recapitulation of a stimulus recalled by association confers substantial functional benefit by generating an unambiguous (if not always correct) perceptual interpretation of noisy, incomplete, or ambiguous sensory input. Indeed, as rightly noted long ago by William James, "Perception is of probable things."

We often see what we expect to see, regardless of whether the experience reflects physical reality.

This familiar axiom is backed by a long-standing literature addressing the influence of associative experience on perception (Ball and Sekuler, 1980; Bartleson, 1960; Bruner et al., 1951; Farah, 1985; Hurlbert and Ling, 2005; Ishai and Sagi, 1995, 1997a, 1997b; Mast et al., 2001; Siple and Springer, 1983), which dates at least to Ewald Hering's (Hering, 1878) concept of memory colors—e.g., perceived color should be biased toward yellow if the color originates from a banana. The argument can also be understood in Bayesian terms (e.g., Kersten et al., 2004; Knill and Richards, 1996): learned associations constitute information about the statistics of the observer's environment, which come into play lawfully as the visual system attempts to identify the environmental causes of retinal stimulation (see also Brunswik, 1956).

These arguments regarding the influence of associative learning on perception were put to the test recently in an elegant experiment by Backus and colleagues (Haijiang et al., 2006), which supports the interpretation we have applied to our results in MT. Briefly, Haijiang et al. used classical conditioning to train associations between direction of motion and two values of a second cue (e.g., stimulus position). Following learning, human subjects were presented with ambiguous (bistable) motion stimuli along with one or the other cue value. Subjects exhibited a marked bias in the direction of perceived motion, which was dictated by the associated cue. Our discovery of recall-related activity in area MT leads us to hypothesize that such effects of association-based recall on perception are mediated by the integration of bottom-up and top-down signals at the level of individual neurons in visual cortex.

Mechanisms of Visual Associative Memory

Prior neurophysiological studies of associative memory implicated the inferior temporal (IT) cortex (Erickson and Desimone, 1999; Messinger et al., 2001; Sakai and Miyashita, 1991). In particular, a previously ineffective stimulus for a given IT neuron can become effective if it is associatively paired with an effective stimulus (Messinger et al., 2001). These findings have been interpreted as evidence for recall-related activity: The ineffective stimulus elicits recall of the effective stimulus and, hence, a neuronal response (Miyashita, 1993). Mechanistic insight for this interpretation is provided by an experiment by Miyashita and colleagues (Tomita et al., 1999), in which recall of one visual pattern by its paired associate leads to reactivation under the control of feedback from the prefrontal cortex.

Our discovery of emergent selectivity for static patterns in area MT adds to an understanding of the mechanisms of visual associative memory in two important ways. First, motion and static arrow directions are simple periodic variables that naturally parallel one another. Use of these stimuli thus afforded a unique opportunity to quantify the tuning of recall-related activity using arrow direction as a proxy for the associated motion direction. By this means we revealed a remarkable generalization of neuronal selectivity, consistent with the graphical "rule" that characterizes the relationship between our motion and shape variables. This finding is surprising in view of the fact that monkeys in a laboratory setting often fail to generalize behavior across changes in stimuli or task requirements. Although it remains to be seen whether our animals exhibit behavioral rule generalization that parallels the observed neuronal generalization, we speculate that the observed effects could underlie the perceptual learning phenomenon in which one sensory dimension is mapped onto another (Goldstone, 1998). This phenomenon is ubiquitous in normal human experience and of considerable behavioral significance—if you learn, for example, that the specific position of a dial or scale is associated with sound of a specific volume, you can readily generalize to predict the volume at other dial positions.

Second, our results provide unprecedented evidence that associative plasticity holds for neurons at an early processing stage that represents a highly specific visual attribute, i.e., motion. Additional support for this view comes from functional brain imaging studies, in which associative recall of visual memories and imagery is correlated with selective patterns of cerebral blood flow in early visual cortex (Buckner and Wheeler, 2001; Kosslyn, 1994; Roland and Gulyas, 1994), including area MT (Kourtzi and Kanwisher, 2000; Shulman et al., 1999). The associative plasticity that we observed in MT may thus be a general feature of sensory neurons at many hierarchical levels, which makes possible the manifold and highly flexible associations that underlie perceptual experience.

EXPERIMENTAL PROCEDURES

Subjects and Surgical Preparation

Two male rhesus monkeys (*Macaca mulatta*), weighing 8.5–9.5 kg, were used in accordance with procedures approved by the Salk Institute Animal Care and Use Committee and in compliance with USDA regulations and NIH guidelines for the care and use of laboratory animals. The approximate location of area MT was identified from sulcal landmarks in magnetic resonance images of each monkey's brain. A recording chamber was then implanted on each monkey's skull dorsal to area MT using standard surgical techniques (e.g., Thiele et al., 1999).

Visual Stimuli

Visual stimuli appeared on a CRT video display (21 inch, 1024 × 768, 75 Hz) and were viewed from a distance of 57 cm. Two stimulus types were used (see Figure 1): motion and static shape. Motion stimuli consisted of random dot patterns (dot diameter: 0.09°; dot density: 6.4 dots/deg^2) viewed within a circular aperture. Dot contrast (Michelson) with background was 95%. All dots moved continuously and in the same direction on each trial. Directions were left, right, up, and down. Dot speed was 16°/s. Shape stimuli (see Figure 1) were static arrows created from randomly positioned dots in an arrow-shaped aperture (dot density: 19.2 dots/deg^2). Arrows could be pointed in the same four directions as dot motions.

When presented as receptive field (RF) stimuli during fixation task (see below), both stimulus types were scaled to 70% of RF diameter.

During pair-association training (see below), both stimulus types subtended approximately 2.5°.

Behavior

Animals were seated in a standard primate chair. Eye position was monitored continuously using an infrared video-based device. Two types of behavioral tasks were used: fixation and pair association. The fixation task was used for all electrophysiological data acquisition, i.e., during the pre- and postlearning assessments of neuronal selectivity. The pair-association task was used to train the designated associations between static and moving stimuli. The current task regimen was identified at the outset of each trial by the use of differently colored fixation targets. Correct performance on both tasks was rewarded with a small drop of juice.

Fixation Task

In the fixation task, the monkey was simply required to maintain fixation on a centrally placed spot (within 2.0°) before and during presentation of a single visual stimulus at the previously determined RF location. Motion stimuli appeared in static form for 500 ms and then moved for an additional 750 ms. Static arrow stimuli appeared for the same total duration (1250 ms) without moving.

Pair-Association Task

In the pair-association task (Figure 1B), animals were additionally required to select the target stimulus that was paired with a preceding cue stimulus. Each pair-association trial began with the appearance of the fixation spot. Once fixation was achieved, the cue stimulus appeared at the center of gaze for 750 ms. Following a brief delay (750 ms), the choice display appeared, which contained four stimuli, one in each visual quadrant and all centered equidistant (4.25°) from the fixation point. One of these four stimuli was the matching stimulus and the other three were distractors. The position of the correct choice target varied randomly between trials.

Only vertical stimuli were used as cues/matches in pair-association training: (1) upward motion, (2) downward motion, (3) upward arrow, and (4) downward arrow. Each of these stimuli appeared as the cue, on a pseudorandom basis. The choice display contained stimuli that were of the type different from the cue, in which case the match was the stimulus that we chose to associate with the cue (see Figure 1). For example, if the cue was the upward pointing arrow, the match stimulus was upward motion, and vice versa. On all trials, the distractors consisted of three stimuli: two horizontal motion stimuli (left- and rightward motion) and the remaining stimulus of the same type as the match stimulus (i.e., the other direction of motion, or arrow from the list above). Animals selected one of the stimuli in the choice display by a saccadic eye movement directed to the stimulus.

Electrophysiology

During each recording session, a sharpened tungsten electrode (~3.0 MΩ impedance) was lowered through a 23 ga. stainless-steel guide tube into the dorsal telencephalon, reaching area MT in the lower bank of the superior temporal sulcus. Electrode placement was guided by magnetic resonance images of each monkey's brain and the distribution in depth of spontaneous neuronal activity. CORTEX software (http://www.cortex.salk.edu/) controlled stimulus presentation and data collection. Visually responsive neurons were isolated using the Plexon system.

All neuronal recordings reported herein were made while the animal performed the fixation task. These recordings were made both before (103 neurons in monkey A) and beginning immediately after (97 neurons in monkey A, 113 neurons in monkey B) animals were trained to criterion performance (85% correct on at least 5 consecutive days) on the pair-association task. Postlearning electrophysiological recording continued for 4 weeks in monkey A and 5 weeks in monkey B. RFs of isolated MT neurons were located in each quadrant of the contralateral hemifield, at eccentricities ranging from 1° to 14° (mean = 7.9°).

Analysis

Neuronal Responsivity and Selectivity

We assessed the responses of each neuron to four directions of motion and four directions of the static arrow stimuli. Each stimulus (direction and type) was presented in the RF for a minimum of 20 trials. For motion stimuli, responses were computed during a window extending 80–830 ms after motion onset (580–1330 ms after stimulus appearance). Responses to static arrows were computed during the same window (580–1330 ms) following stimulus appearance (times chosen to keep the analysis windows of the same duration for both stimulus types and to exclude luminance-onset responses from the analysis). The baseline neuronal firing rate was calculated over the interval -300 to 0 ms, relative to stimulus appearance.

To evaluate directional selectivity for each stimulus type (motion and arrows), we used responses (computed during the aforementioned analysis windows) to the four stimulus directions on each trial number to compute the average response vector for that trial. The directions of these trial-based response vectors ($n = \sim 20$ trials) were then used to compute the circular mean. Neurons were considered selective for direction if the Rayleigh test (Batschelet, 1981) revealed that the distribution of trial-based directional measures was concentrated around the circular mean ($p < 0.05$). When found significant by this test, the circular mean served as our estimate of the preferred direction for the relevant stimulus type and neuron. All statistical tests were evaluated at the $p = 0.05$ level and were two tailed.

To evaluate the strength of directional selectivity for motions and arrows along the estimated preferred axis, we first computed four response vectors for each neuron. The directions of these vectors corresponded to the four stimulus directions (up, down, left, right), and the vector magnitudes corresponded to the observed mean response rates for the corresponding directions. We then computed the average of these four response vectors for each neuron. To facilitate comparisons across neurons, the length of the averaged vector for each neuron was normalized by the maximum response of the neuron to the four stimulus directions. The magnitude of the resultant normalized averaged response vector thus reflects the strength of directional selectivity, and we used this measure to compare directional selectivity strengths for motion and arrow stimuli.

Polar Analysis of Preferred Directions for Motion and Arrow Stimuli

Preferred directions for motion and arrow stimuli were estimated as indicated above (Neuronal Responsivity and Selectivity) and were compared by subtracting the respective direction angles. These angular difference measures were then analyzed using the v-test (Batschelet, 1981) in order to determine whether the observed angles cluster around the predicted difference angles (0° in monkey A, 180° in monkey B). The value of the v statistic is inversely proportional to variance. By definition, the values of v that reach our criterion level of significance ($p < 0.05$) are those for which the variance is small enough to reject the null hypothesis (uniformity), i.e., to accept clustering around the predicted angle.

Analysis of Eye Movements for Recordings from Arrow-Tuned Neurons

To address the possible influence of eye movements on the observed postlearning neuronal selectivity for static arrow direction, we measured four variables during each stimulus presentation: (1) mean eye position, (2) standard deviation of eye position, (3) direction of fixational saccades (isolated during epochs in which they occurred), and (4) direction of smooth drift eye movements (isolated during epochs in which they occurred).

For variables (1) and (2), we determined whether the observed values across recordings differed as a function of the direction of the arrow in the CRF. To analyze saccade and smooth drift periods (variables 3 and 4), we first had to detect fixational saccade and smooth drift periods that occurred during stimulus presentation (see below). For each trial and condition, we then determined the mean direction of the occurring fixational saccades or smooth drift. We used ANOVA to assess whether the frequency of directional (up, down, left, right) eye movements was a function of arrow direction across the different recordings. This analysis enabled us to determine, for example, whether upward eye movements were more frequent on trials in which arrows were pointed upward, etc.

Saccade and Smooth Drift Detection Procedure

For the detection of saccades, we used a procedure similar to that described by Krauzlis and Miles (Krauzlis and Miles, 1996). Briefly, we first applied a 15 point finite response filter to the eye position traces to obtain eye velocity. Horizontal and vertical velocities were then combined to obtain radial eye velocity. By applying the same filter again to this velocity signal, we obtained eye acceleration. Saccades were detected using a three-step automated procedure. (1) Candidate saccades were first identified as points in time passing a velocity threshold (15°/s). (2) For each candidate saccade, the acceleration signal prior to and after the saccade was scanned. If an adjacent data point exceeded the acceleration threshold ($550^\circ/s^2$), the eye position at that point was flagged and added to the candidate saccade. (3) The duration of the candidate saccade was compared to our minimum duration criterion. Candidate saccades that both exceeded this criterion and passed visual inspection were confirmed for further analysis. The minimum amplitude of fixational saccades we were able to detect with this procedure was 0.42°. We defined smooth drift periods as those portions of the eye traces that did not exceed our velocity criterion for saccades.

ACKNOWLEDGMENTS

We thank T. Pasternak for many helpful discussions; R. Krauzlis and Z. Hafed for help with the eye movement analysis; and G. Horwitz, B. Krekelberg, A. Messinger, J. Reynolds, G. Stoner, and P. Rayudu for comments on the manuscript. This work was supported by a grant from the NEI. A.S. was supported by the DFG.

Received: December 14, 2006
Revised: January 25, 2007
Accepted: February 26, 2007
Published: March 14, 2007

REFERENCES

Albright, T.D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. J. Neurophysiol. *52*, 1106–1130.

Albright, T.D. (1993). Cortical processing of visual motion. Rev. Oculomot. Res. *5*, 177–201.

Ball, K., and Sekuler, R. (1980). Models of stimulus uncertainty in motion perception. Psychol. Rev. *87*, 435–469.

Barlow, H. (1990). Conditions for versatile learning, Helmholtz's unconscious inference, and the task of perception. Vision Res. *30*, 1561–1571.

Bartleson, C.J. (1960). Memory colors of familiar objects. J. Opt. Soc. Am. *50*, 73–77.

Batschelet, E. (1981). Circular Statistics in Biology (London, New York: Academic Press).

Born, R.T., and Bradley, D.C. (2005). Structure and function of visual area MT. Annu. Rev. Neurosci. *28*, 157–189.

Bruner, J.S., Postman, L., and Rodrigues, J. (1951). Expectation and the perception of color. Am. J. Psychol. *64*, 216–227.

Brunswik, E. (1956). Perception and the Representative Design of Psychological Experiments (Berkeley, CA: University of California Press).

Buckner, R.L., and Wheeler, M.E. (2001). The cognitive neuroscience of remembering. Nat. Rev. Neurosci. *2*, 624–634.

Davies, P., Davies, G.L., and Bennett, S. (1982). An effective paradigm for conditioning visual perception in human subjects. Perception *11*, 663–669.

Desimone, R., Albright, T.D., Gross, C.G., and Bruce, C.J. (1984). Stimulus selective properties of inferior temporal neurons in the macaque. J. Neurosci. *8*, 2051–2062.

Erickson, C.A., and Desimone, R. (1999). Responses of macaque perirhinal neurons during and after visual stimulus association learning. J. Neurosci. *19*, 10404–10416.

Farah, M.J. (1985). Psychophysical evidence for a shared representational medium for mental images and percepts. J. Exp. Psychol. *114*, 91–103.

Felleman, D.J., and Van Essen, D.C. (1991). Distributed hierarchical processing in the primate cerebral cortex. Cereb. Cortex *1*, 1–47.

Freyd, J.J. (1987). Dynamic mental representations. Psychol. Rev. *94*, 427–438.

Gaffan, D., and Bolton, J. (1983). Learning of object-object associations by monkeys. Q. J. Exp. Psychol. B *35*, 149–155.

Goldstone, R.L. (1998). Perceptual learning. Annu. Rev. Psychol. *49*, 585–612.

Haijiang, Q., Saunders, J.A., Stone, R.W., and Backus, B.T. (2006). Demonstration of cue recruitment: change in visual appearance by means of Pavlovian conditioning. Proc. Natl. Acad. Sci. USA *103*, 483–488.

Hebb, D.O. (1949). The Organization of Behavior; A Neuropsychological Theory (New York: Wiley).

Hebb, D.O. (1968). Concerning imagery. Psychol. Rev. *75*, 466–477.

Helmholtz, H.v., and Southall, J.P.C. (1924). Helmholtz's Treatise on Physiological Optics (Rochester, NY: The Optical Society of America).

Hering, E. (1878). Zur Lehre vom Lichtsinne (Principles of a New Theory of the color Sense), tanslated by K. Butler and partially reprinted in Treevan, R.C., and Birney, R.C. (1961) (Color Vision, Selected Readings., Van Nostrand Reinhold).

Hurlbert, A.C., and Ling, Y. (2005). If it's a banana, it must be yellow: The role of memory colors in color constancy. J. Vis. *5*, 787a.

Ishai, A., and Sagi, D. (1995). Common mechanisms of visual imagery and perception. Science *268*, 1772–1774.

Ishai, A., and Sagi, D. (1997a). Visual imagery: Effects of short- and long-term memory. J. Cogn. Neurosci. *9*, 734–742.

Ishai, A., and Sagi, D. (1997b). Visual imagery facilitates visual perception: Psychophysical evidence. J. Cogn. Neurosci. *9*, 476–489.

James, W. (1890). Principles of Psychology (New York: Henry Holt).

James, W. (1910). Principles of Psychology (New York: Holt).

Kersten, D., Mamassian, P., and Yuille, A. (2004). Object perception as Bayesian inference. Annu. Rev. Psychol. *55*, 271–304.

Knill, D.C., and Richards, W. (1996). Perception as Bayesian inference (Cambridge, England; New York: Cambridge University Press).

Kosslyn, S.M. (1994). Image and Brain (Cambridge, MA: The MIT Press).

Kourtzi, Z., and Kanwisher, N. (2000). Activation in human MT/MST by static images with implied motion. J. Cogn. Neurosci. *12*, 48–55.

Kourtzi, Z., and Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. Science *293*, 1506–1509.

Krauzlis, R.J., and Miles, F.A. (1996). Release of fixation for pursuit and saccades in humans: evidence for shared inputs acting on different neural substrates. J. Neurophysiol. *76*, 2822–2833.

Lavenex, P., and Amaral, D.G. (2000). Hippocampal-neocortical interaction: a hierarchy of associativity. Hippocampus *10*, 420–430.

Lavenex, P., Suzuki, W.A., and Amaral, D.G. (2002). Perirhinal and parahippocampal cortices of the macaque monkey: projections to the neocortex. J. Comp. Neurol. *447*, 394–420.

Mast, F.W., Berthoz, A., and Kosslyn, S.M. (2001). Mental imagery of visual motion modifies the perception of roll-vection stimulation. Perception *30*, 945–957.

Messinger, A., Squire, L.R., Zola, S.M., and Albright, T.D. (2001). Neuronal representations of stimulus associations develop in the temporal lobe during learning. Proc. Natl. Acad. Sci. USA *98*, 12239–12244.

Messinger, A., Squire, L.R., Zola, S.M., and Albright, T.D. (2005). Neural correlates of knowledge: stable representation of stimulus associations across variations in behavioral performance. Neuron *48*, 359–371.

Miyashita, Y. (1993). Inferior temporal cortex: where visual perception meets memory. Annu. Rev. Neurosci. *16*, 245–263.

Newsome, W.T., Wurtz, R.H., and Komatsu, H. (1988). Relation of cortical areas MT and MST to pursuit eye movements. II. Differentiation of retinal from extraretinal inputs. J. Neurophysiol. *60*, 604–620.

Roland, P.E., and Gulyas, B. (1994). Visual imagery and visual representation. Trends Neurosci. *17*, 281–287.

Sakai, K., and Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. Nature *354*, 152–155.

Shulman, G.L., Ollinger, J.M., Akbudak, E., Conturo, T.E., Snyder, A.Z., Petersen, S.E., and Corbetta, M. (1999). Areas involved in encoding and applying directional expectations to moving objects. J. Neurosci. *19*, 9480–9496.

Siple, P., and Springer, R.M. (1983). Memory and preference for the colors of objects. Percept. Psychophys. *34*, 363–370.

Thiele, A., Dobkins, K.R., and Albright, T.D. (1999). The contribution of color to motion processing in Macaque middle temporal area. J. Neurosci. *19*, 6571–6587.

Tomita, H., Ohbayashi, M., Nakahara, K., Hasegawa, I., and Miyashita, Y. (1999). Top-down signal from prefrontal cortex in executive control of memory retrieval. Nature *401*, 699–703.

Van Wezel, R.J., and Britten, K.H. (2002). Multiple uses of visual motion. The case for stability in sensory cortex. Neuroscience *111*, 739–759.