



## Contents

INTRODUCTION .....	158
MT WAS A KEY PART OF THE EARLY EXPLORATION OF EXTRASTRIATE CORTEX .....	158
CONNECTIONS .....	159
FUNCTIONAL ORGANIZATION ..	162
BASIC TUNING PROPERTIES .....	164
SURROUND MECHANISMS .....	165
THE COMPUTATION OF VELOCITY .....	168
NOISE REDUCTION .....	173
SEGMENTATION .....	173
THE COMPUTATION OF STRUCTURE .....	174
EXTRARETINAL EFFECTS .....	176
PERCEPTUAL CORRELATES AND POPULATION CODES ....	177
Single-Neuron Sensitivity .....	177
Vector Summation versus Winner-Take-All .....	178
Distributed Speed and Acceleration Codes .....	179
CONCLUSIONS .....	179

## INTRODUCTION

The middle temporal visual area (MT or V5) of the macaque monkey possesses a number of attributes that have made it particularly attractive to systems neuroscientists. This region is typical of extrastriate cortex but is still readily identifiable both anatomically and functionally. Though extrastriate, it is still quite close to the retina—its principle inputs as few as five synapses from the photoreceptors—a feature which means, among other things, that the mechanisms by which its receptive field properties arise can be profitably studied. And, although MT neurons are near enough to the inputs to be mechanistically tractable, they are also close enough to some outputs—in particular, those involved in eye movements—to provide an easily measurable, continuous readout of computations performed in this

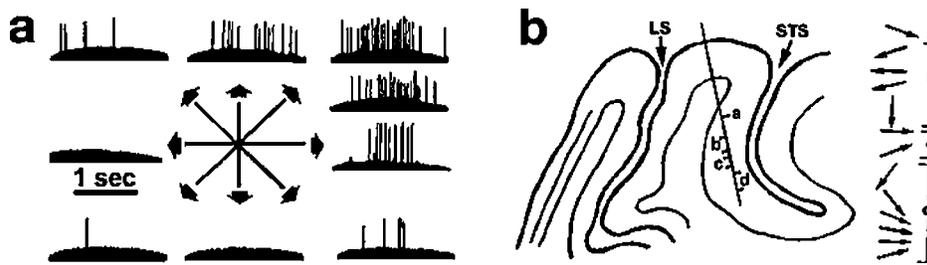
pathway. Finally, MT neurons are concerned with visual motion, which is of obvious ethological importance, which has been extensively characterized psychophysically, and for which there are well-defined mathematical descriptions. Much of the work on MT has focused on its role in visual motion processing, though, as we hope to make clear in what follows, MT plays a richer and more varied role in vision.

## MT WAS A KEY PART OF THE EARLY EXPLORATION OF EXTRASTRIATE CORTEX

Part of MT's significance is historical; it played an important role in the discovery of new extrastriate visual areas (Felleman & Van Essen 1991) and in the idea that they constitute specialized representations of the visual world (Zeki 1978, Barlow 1986).

At the beginning of the twentieth century, primate visual cortex was thought to consist of only three architectonically distinct fields (Brodmann 1909). Beginning in the late 1940s, however, it became clear that considerably more of the cortex was involved in vision. The first demonstration came from temporal lobe lesions that produced visual impairment (Mishkin 1954, Mishkin & Pribram 1954) unaccompanied by deficits in other sensory modalities (Weiskrantz & Mishkin 1958, Brown 1963). Mapping studies using surface electrodes also revealed visually responsive regions well anterior to those traditionally associated with vision (Talbot & Marshall 1941, Clare & Bishop 1954, Woolsey et al. 1955). In addition, new anatomical techniques (Nauta & Gyax 1954) permitted the labeling of connections after lesions of striate cortex (Kuypers et al. 1965, Cragg & Ainsworth 1969, Zeki 1969), which revealed a direct striate (V1) projection zone situated on the posterior bank of the superior temporal sulcus (STS).

MT was discovered at roughly the same time by two different groups. In England, Dubner & Zeki (1971) were able to record



**Figure 1**

First demonstration of direction selectivity in macaque MT/V5 by Dubner & Zeki (1971). (a) Neuronal responses to a bar of light swept across the receptive field in different directions (modified from figure 1 of Dubner & Zeki 1971). Each trace shows the spiking activity of the neuron as the bar was swept in the direction indicated by the arrow. The neuron's preferred direction was up and to the right. (b) Oblique penetration through MT (modified from **Figure 3** of Dubner & Zeki 1971) showing the shifts in preferred direction indicative of the direction columns subsequently demonstrated by Albright et al. (1984). See also **Figure 4**.

visual responses from the V1-projection zone in anesthetized macaques, in so doing establishing a number of physiological hallmarks, particularly their direction-selective responses (**Figure 1a**). Quite presciently, they also suggested a columnar organization for direction-selective neurons (**Figure 1b**) and a role for MT signals in guiding pursuit eye movements, both subsequently confirmed (Albright et al. 1984, Lisberger et al. 1987). Around the same time, Allman & Kaas (1971) were recording from owl monkeys and using a different approach. They made systematic rows of microelectrode penetrations across the entire cortex, mapping receptive fields as they went, thus discovering a large number of retinotopically organized maps. One of these, which they named MT for middle temporal, mapped onto a well-defined region of dense myelination in the lower layers and contained neurons that responded better to drifting bars than to flashed spots. The myelination was also later shown to be characteristic of the macaque motion area (Van Essen et al. 1981), which Zeki subsequently named V5. This histochemical feature has been an underappreciated factor in contributing to the detail with which MT has been studied because it has permitted reliable comparisons across different studies.

Following the first studies, a series of papers confirmed that MT contained a high concentration of direction-selective neurons in several species of both New and Old World monkeys (Zeki 1974, 1980; Baker et al. 1981; Van Essen et al. 1981; Maunsell & Van Essen 1983a,b; Felleman & Kaas 1984). These studies indicated that MT was both unique as a cortical area highly specialized for visual motion and, at the same time, common to a number of different primate species.

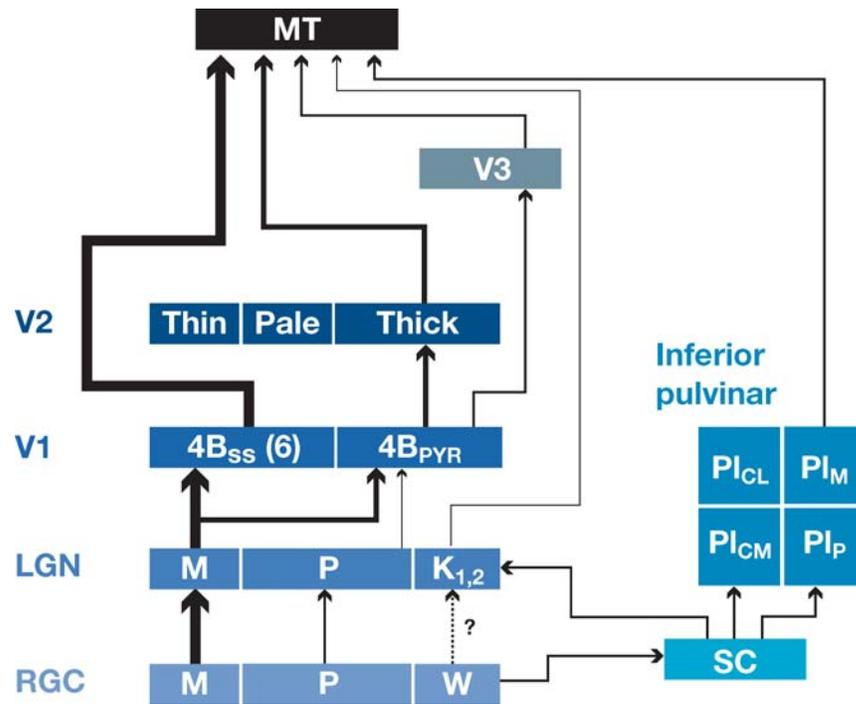
## CONNECTIONS

Like every other cortical area, MT has a rich set of interconnections with other regions of the cortex as well as with numerous subcortical structures. These connections have been discussed in previous publications (Felleman & Van Essen 1991, Orban 1997, Lewis & Van Essen 2000), so we do not recapitulate them here. From a broad perspective, MT's corticocortical connections identify it as one of the main inputs into the dorsal or posterior parietal processing stream (Ungerleider & Mishkin 1982, Maunsell & Newsome 1987), and its key outputs target structures that are implicated in the analysis of optic flow (e.g., MST, VIP) and the generation of eye movements (e.g., LIP, FEF, SC, dorsolateral

---

**preferred direction:**  
the direction of motion eliciting the greatest response from a given neuron

---



**Figure 2**

Gestalt map of major routes into MT in the manner of Felleman & Van Essen (1991). Line thickness is roughly proportional to the magnitude of the inputs, on the basis of a combination of projection neuron numbers and, where data are available, the characteristics of their axon terminals (see **Figure 3**). The thickest lines represent the direct cortical pathway emphasized in the text. Following are important caveats: The pathways shown are those discussed in the text and omit a number of known feedforward cortical inputs that appear lesser in magnitude (V3A, VP, PIP) as well as many subcortical inputs. The sources of the direct and indirect projections from V1 are probably not defined purely by cell morphology (i.e., spiny stellate versus pyramidal; see Elston & Rosa 1997), though they are largely distinct; the largest 4B cells contribute to the direct pathway (Sincich & Horton 2003). The precise nature of the retinal inputs to K1,2 is not known, though their response properties are W-like in the galago (Irvin et al. 1986). Also, the proposed input to MT from the SC via the pulvinar is rendered problematic by the finding that, in owl monkey pulvinar, the principle target of SC terminals ( $PI_{CM}$ ) is different from the main source of MT projections ( $PI_M$ ) (see Stepniewska et al. 1999). Abbreviations: 4B<sub>SS</sub>, spiny stellate neurons in layer 4B; 4B<sub>PYR</sub>, pyramidal neurons in layer 4B; LGN, lateral geniculate nucleus; M, magnocellular stream; P, parvocellular stream; K, koniocellular layers of LGN;  $PI_{CL}$ , central lateral nucleus of the inferior pulvinar;  $PI_{CM}$ , central medial nucleus of the inferior pulvinar;  $PI_M$ , medial nucleus of the inferior pulvinar;  $PI_P$ , posterior nucleus of the inferior pulvinar; RGC, retinal ganglion cells; SC, superior colliculus; VP, ventral posterior area.

pons). Because we desire to address how MT neurons acquire their unique visual response properties and discuss the role they play in motion computations, we focus the present discussion on MT's major inputs (**Figure 2**) and their functional implications. In particular, we argue that the most important

input to MT is from a magnocellular-dominated projection originating from layer 4B of V1.

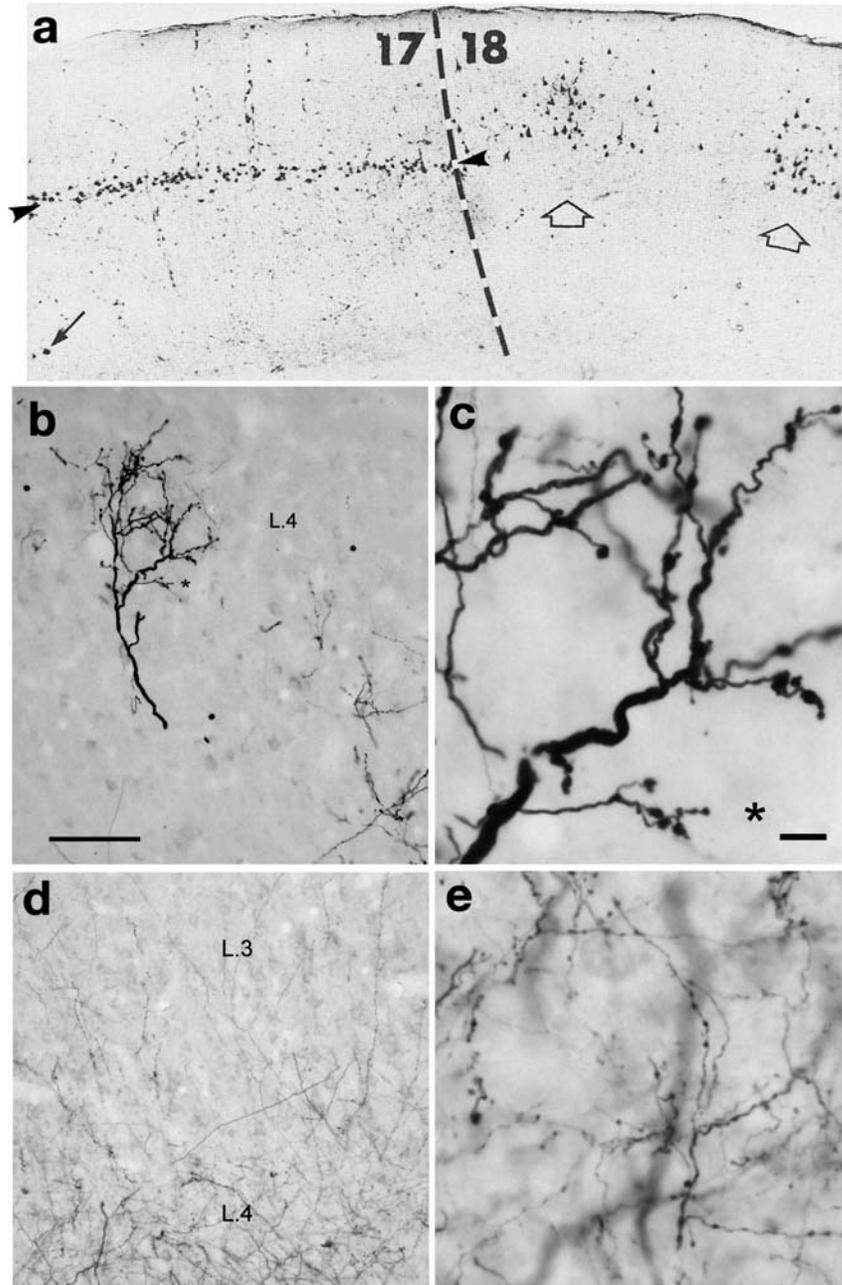
To a first approximation, MT is dominated by cortical rather than subcortical inputs. Nevertheless, unlike some other extrastriate areas such as V2 (Schiller & Malpeli



### FUNCTIONAL ORGANIZATION

MT is retinotopically organized, each hemisphere containing a more-or-less complete map of the contralateral visual hemi-field, with a marked emphasis on the fovea the

central [15° of the visual field occupies over half of MT's surface area (Van Essen et al. 1981)] and a bias toward the lower quadrant of the visual field (Maunsell & Van Essen 1987). Within this relatively crude retinotopic map,





**binocular disparity:** a difference in the relative position of a stimulus on each of the two retinas. This disparity is the basis of stereoscopic depth perception.

above. This clustering may occur predominantly in the upper (output) layers of macaque MT because a segregation of center-surround properties appears to exist in subdivisions of one of MT's main projection zones, MSTd versus MSTl (Tanaka et al. 1986, Komatsu & Wurtz 1988, Eifuku & Wurtz 1998). The finding that microstimulation in macaque MT has qualitatively different effects on smooth pursuit eye movements depending on the nature of center-surround interactions at the stimulation site (Born et al. 2000) is consistent with the idea of segregation and, to date, is the only direct experimental evidence to support a role for MT center-surround interactions in the type of figure-ground comparisons originally suggested by Allman and his colleagues (Allman et al. 1985b).

### BASIC TUNING PROPERTIES

The visual responses of MT neurons are determined principally by five properties of the stimulus: (*a*) retinal position (*b*) direction of motion, (*c*) speed of motion, (*d*) binocular disparity, and (*e*) stimulus size (due to surround suppression). In the following paragraphs we expand on selected aspects of the above description. The basic visual-response properties of MT neurons have been addressed extensively in excellent reviews by Orban (1997) and Britten (2003), so we focus on more recent studies and on those most relevant to the subsequent discussion. In particular, we attempt, where possible, to compare MT properties with those of its principle inputs, with an eye toward understanding what MT contributes to cortical motion processing.

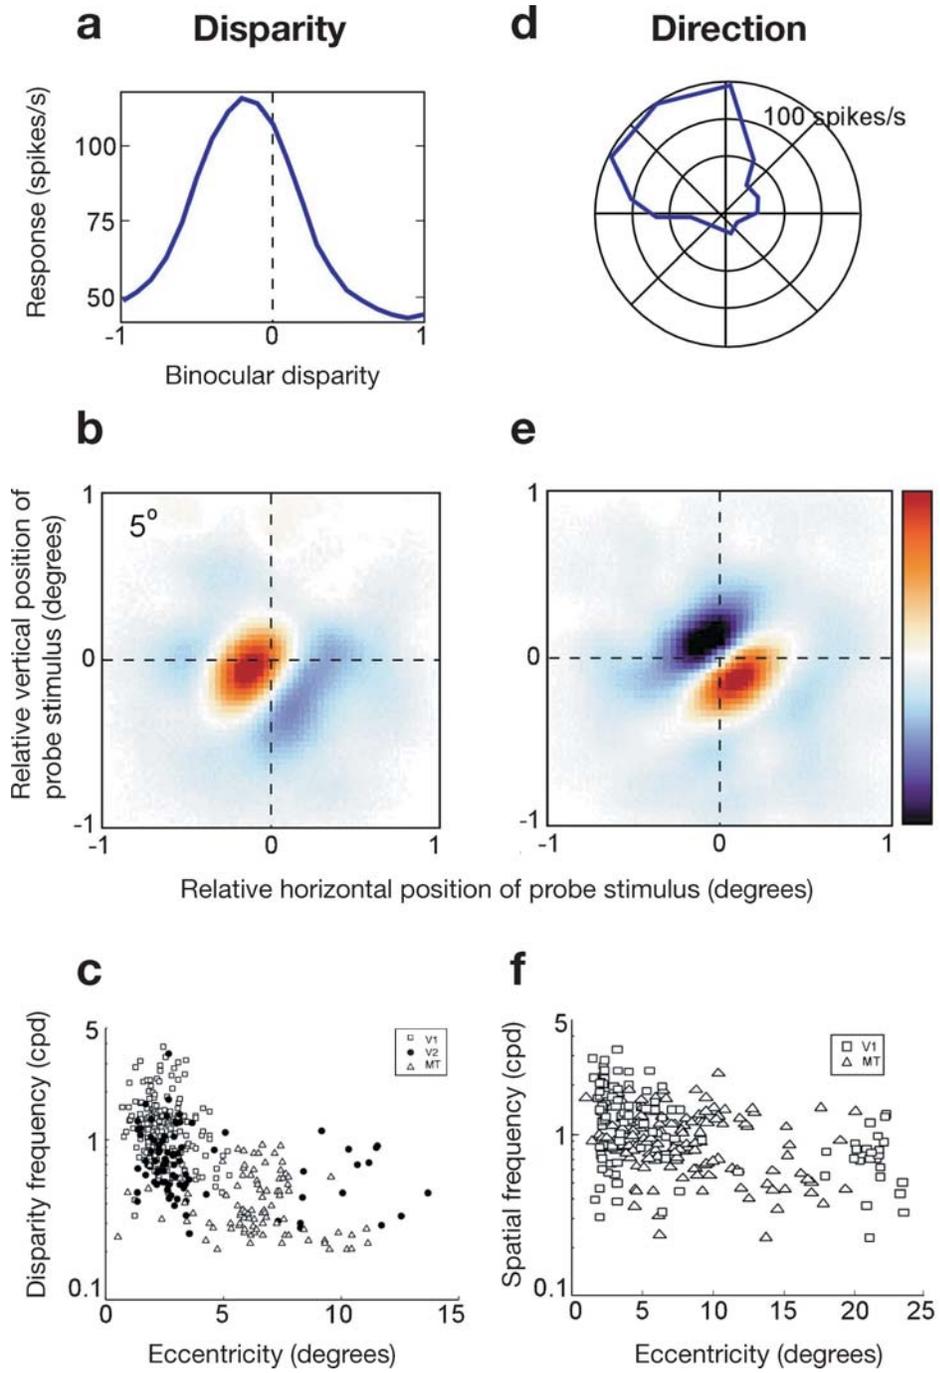
The principal V1 inputs to MT described above are themselves strongly direction selective (Movshon & Newsome 1996) and also are tuned for speed (Orban et al. 1986). Moreover, given that the most direction-selective V1 cells also tend to be highly selective for binocular disparity (Prince et al. 2000), many of MT's V1 inputs are also likely disparity tuned. Conceivably then, MT neurons in-

herit three of their most important properties from V1.

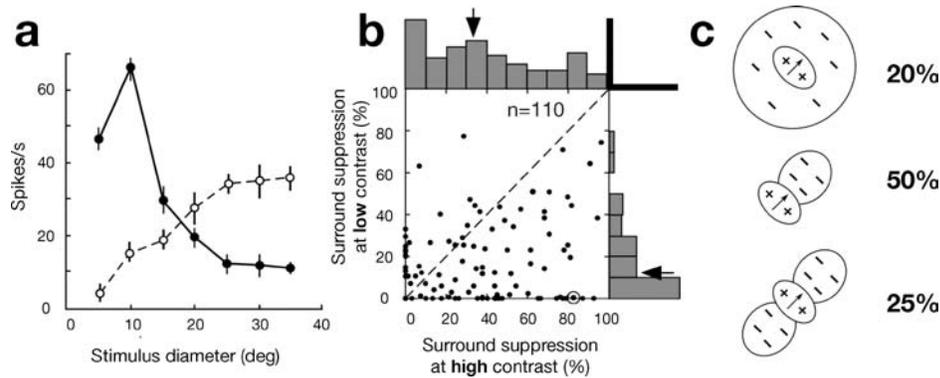
So what, then, does MT add? It is chastening to note up front that more than 30 years of physiology have not yielded a clear-cut answer to this question. In a number of cases, early evidence pointed to differences that were theoretically attractive but that, upon closer and more quantitative analysis, failed to materialize. Because these examples are informative, we first consider several of them, before moving on to other differences that have been borne out.

MT receptive fields are much larger than those in V1—a ballpark figure is tenfold greater in linear dimensions—so one might suppose that MT neurons can compute motion and disparity over larger spatial ranges than can V1 cells. This idea was particularly attractive for directional interactions because studies of human perception have revealed at least two different motion-sensing mechanisms that operate over different spatial scales and also differ with respect to other properties, such as contrast invariance (Nakayama 1985). As a result, investigators proposed that MT neurons might inherit short-range motion sensitivity from their V1 inputs but would, in addition, compute motion over longer spatial displacements, commensurate with the size of their receptive fields. Indeed, initial experiments by Mikami and colleagues (1986), using sequences of flashed bars, suggested that MT neurons did produce directional signals to larger spatial separations, on average, than did V1 cells. However, a recent reexamination of this issue, using more directly comparable stimuli revealed very similar upper limits for V1 and MT (Churchland et al. 2004). In fact, the directional interactions of neurons from both areas take place over extremely small spatial ranges—fractions of a degree in receptive fields that are, in MT, many degrees wide (**Figure 5d-f**)—and they reverse direction for contrast-inverting sequences (Livingstone et al. 2001, Livingstone & Conway 2003), both of which are characteristics of the short-range perceptual









**Figure 6**

Center-surround interactions in MT. (A) Effect of contrast on center-surround interactions for one MT neuron. When tested with high-contrast random dots (RMS contrast  $9.8 \text{ cd/m}^2$ ) the neuron responded optimally to a circular dot patch  $10^\circ$  in diameter and was strongly suppressed by larger patterns. The same test using a low-contrast dot pattern ( $0.7 \text{ cd/m}^2$ ) revealed strong area summation with increasing size. (B) Population of 110 MT neurons showing the strength of surround suppression measured at both high and low contrast. Surround suppression was quantified as the percent reduction in response between the largest dot patch ( $35^\circ$  diameter) and the stimulus eliciting the maximal response. Each dot represents data from one neuron; the dashed diagonal is the locus of points for which the surround suppression was unchanged by contrast. The circled dot is the cell from panel A. (C) Asymmetries in the spatial organization of the suppressive surround (after Xiao et al. 1997). Different kinds of surround geometry are potentially useful for calculating spatial changes in flow fields that may be involved in the computation of structure from motion. Neurons whose receptive fields have circularly symmetric surrounds (*top*) are postulated to underlie figure-ground segregation. The first- (*middle*) and second-order (*bottom*) directional derivatives can be used to determine surface tilt (or slant) and surface curvature, respectively (Buracas & Albright 1996). Panels A and B are from Pack et al. 2005.

order) or curvature (second order) (Droulez & Cornilleau-Peres 1990, Koenderink & van Doorn 1992, Buracas & Albright 1996). This potential role of the surround in structural computations is discussed further in the next section.

The source of MT surrounds remains unclear. One possibility is that the surrounds are already present in the inputs to MT. Though center-surround interactions for motion have been reported in V1 (Gulyas et al. 1987, Levitt & Lund 1997), the relative paucity of such interactions in the input layers of MT (Lagae et al. 1989, Raiguel et al. 1995, Born 2000) and the very large size of MT surrounds—at least several-fold larger than their centers (Allman et al. 1985a, Tanaka et al. 1986, Raiguel et al. 1995)—make this an unlikely explanation. It may be that surrounds reflect feedback from higher areas such as MST or are created by

horizontal connections within MT (Malach et al. 1997).

## THE COMPUTATION OF VELOCITY

By “velocity” we mean the vector representation of the direction and speed of retinal motion. As discussed above, MT adds little to the raw direction and speed tuning already found in V1, but researchers still think it plays a role in computing the motion of whole objects or patterns. The nature of that role is the subject of this section. We first discuss some theoretical considerations and outline the roles MT might play.

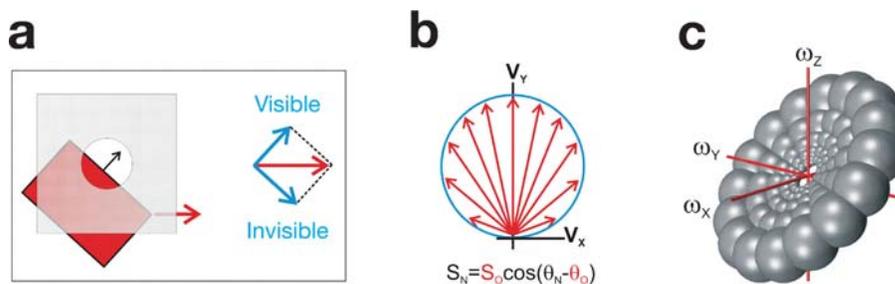
For a rigid object, it would seem trivial to compute pattern motion because one would expect every part of the object to have the same velocity. But the measurements obtained

**null direction:** the direction of motion opposite that eliciting the greatest response from a given neuron; e.g., if a neuron responds optimally to rightward motion, its preferred direction is right and its null direction is left.

depend critically on how the local motion is sampled. Because of the “aperture problem,” moving edges seen through small apertures appear to move orthogonally to their length; therefore, for an object sampled at high resolution (through small apertures), different parts of the object appear to move in different directions, depending on the orientation of the sampled edges (Fennema & Thompson 1979) (**Figure 7a**). In the primate, visual motion is first computed in V1, and these neurons have small receptive fields; therefore, investigators generally think that V1 neurons see normal (orthogonal) velocities (Hubel & Wiesel 1968, Movshon et al. 1985, Heeger 1987). The problem then is to compute two-dimensional (2D) pattern motion on the basis of these local, one-dimensional (1D) samples. In theory, two local samples are sufficient to compute the pattern motion; the geometric solution to the problem is called the intersection of constraints (IOC), which is illustrated in **Figure 7B**.

As we address below, some MT neurons report 2D motion, but the computation of 2D motion itself does not need to occur in MT; indeed, it could take place anywhere along a continuum between two extremes. At one extreme is the idea that the V1 stage is linear and the informative (nonlinear) computation occurs in MT (Adelson & Movshon 1982, Albright 1984, Movshon et al. 1985, Heeger 1987, Simoncelli & Heeger 1998). At the other extreme, 2D motion is extracted in V1 through nonlinearities, such as endstopping, and all that is required of MT is to pool the V1 inputs (Wilson et al. 1992, Barth & Watson 2000). Below, we discuss physiological evidence supporting each of these ideas. At present, there is no definitive evidence for either model, and indeed, we conclude by suggesting that both may be operative, depending on the stimulus conditions.

According to models of the first type, the aperture problem is built into the system by virtue of the linearity of the motion

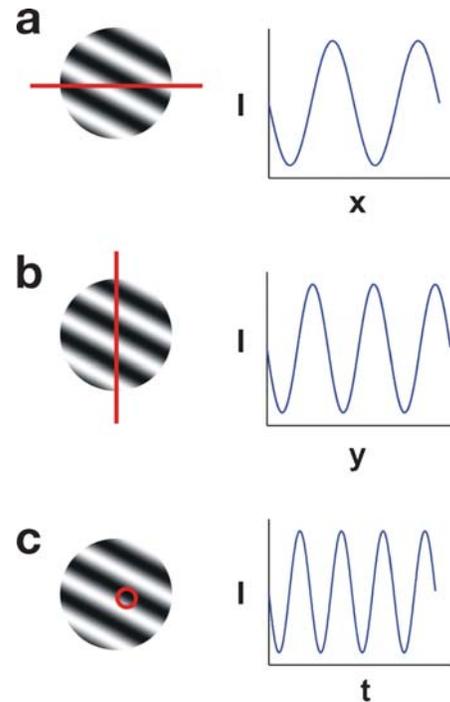


**Figure 7**

The problem of two-dimensional motion detection. (A) The aperture problem. A moving edge seen through an aperture appears to move perpendicularly to itself because the object's motion, in this case to the right, can be decomposed into vector components, one parallel to the edge and one perpendicular. The parallel component is invisible because there is no contrast parallel to an edge, so only the perpendicular component remains. (B) The relationship between the component (apertured) vector samples and the global motion of an object is cosinusoidal; that is, the speed of the samples is the object speed times the cosine of the difference between the object direction and the direction of the vector sample, as shown in the equation. In the equation,  $S_N$  and  $\theta_N$  are the normal (sampled) speed and direction, and  $\theta_O$  and  $S_O$  are the speed and direction of the object. Because there are two unknowns, the object direction and speed (*shown in red in equation*), two samples are needed to solve for the object velocity. This requirement is usually referred to as the intersection of constraints (IOC). It is often visualized differently (Movshon et al. 1985). (C) A possible neural algorithm for solving the IOC, as suggested by Simoncelli & Heeger (1998). For a rigid, translating object, all local spatiotemporal frequencies must lie on a plane in frequency space. To detect this, one could create a linear filter, represented by blobs in the diagram, for each location on the plane, then sum the energy passed through the filters.

detection stage of V1 simple cells. Consider the frequency representation of moving objects (Watson & Ahumada 1983). For any rigid moving object, the spatiotemporal frequency of all local samples must lie on a plane in frequency space (**Figure 7c**). V1 simple cells, to a first approximation, act like linear space-time filters (McLean & Palmer 1989, Reid et al. 1987), and complex cells are thought to differ mainly in their phase-insensitivity but not in their frequency selectivity (Movshon et al. 1978). So V1 neurons effectively measure the amount of motion energy in their passband (Adelson & Bergen 1985). In other words, such neurons do not really know what the stimulus direction is; they see only the motion component within their frequency band. This is the expression of the aperture problem in frequency space (see **Figure 8**).

If V1 neurons see component motion, they are ignorant, in a sense; MT neurons would need to be relatively intelligent in combining V1 inputs to recover the true, 2D direction of motion. As described above, the IOC construction is the basic rule needed to compute pattern velocity from component (local) velocities. But what physiological mechanism could do this? Heeger (1987) and Simoncelli & Heeger (1998), building on important theoretical (Adelson & Movshon 1982, Watson & Ahumada 1983, Movshon et al. 1985) and physiological (Hubel & Wiesel 1962, Movshon et al. 1985, Reid et al. 1987) groundwork, described what we refer to as the F-plane model. It assumes a front end made of linear V1 cells whose outputs are summed over a plane in frequency space by an MT pattern cell. This planar summation is an instantiation of the IOC rule (**Figure 7C**). The model is able to explain a number of physiological results, including MT responses to variable coherence stimuli (Newsome et al. 1989) and to plaid patterns (Adelson & Movshon 1982). There are other important models of MT computation, however, which we regrettably do not have space to discuss here (Wilson et al. 1992, Nowlan & Sejnowski 1995, Lisberger & Movshon 1999).



**Figure 8**

Cutting a horizontal slice through an oblique sine wave grating reveals a sinusoidal modulation of intensity vs.  $x$  position (*A*). Cutting a vertical slice through the same grating reveals an identical intensity modulation versus  $y$  position (*B*). Looking at a fixed location as the grating moves reveals a sinusoidal modulation of intensity versus time (*C*). Changing the angle (direction) of the grating changes the relative  $x$  and  $y$  sinusoidal frequencies, and changing the grating speed changes the temporal frequency. Thus, the velocity (direction and speed) of the grating is completely characterized by a single, three-dimensional frequency,  $(\omega_x, \omega_y, \omega_t)$ . Just as any sound is the sum of its harmonics, any moving object can be represented in terms of its component frequencies, each equivalent to a single, moving sinusoidal grating. If the object is rigid and not rotating, all of its component frequencies will lie on a plane in  $\omega_x, \omega_y, \omega_t$  space. The aperture “problem,” stated in these terms, arises insofar as V1 neurons are linear  $\omega_x, \omega_y, \omega_t$  filters, each detecting a certain component frequency. The relative blindness, or ignorance, implied by the aperture problem is in the linearity of these cells; they are presumably unaffected by frequencies outside their passband. Thus, the “component” cells of Movshon et al. (1985; see text) see only one grating—one frequency.



neurons respond well to line terminators moving in their preferred direction, regardless of the orientation of the contour (Pack et al. 2003). Although not tested directly, these cells likely show a similar preference for corners, dots, and other features compared with extended contours. Because features always move in the object direction, and because end-stopping suppresses neurons not responding to features, end-stopping tends to have the important effect of defeating the aperture problem. Whether end-stopping actually produces a pattern-motion computation in the naturally functioning visual system remains to be seen.

But some evidence suggests that end-stopping does play a critical role in solving the aperture problem. Pack & Born (2001) tested MT neurons with a pattern of line segments drifting coherently such that the motion components—the velocities normal to the segments—differed by  $45^\circ$  from the direction of the pattern as a whole. The neuronal responses were initially strongly biased by the component directions and then evolved to the pattern direction over the course of about 80 ms. In a separate study of V1 neurons, the end-stopping effect was shown to also develop over a similar time course (Pack et al. 2003). Thus, MT neurons may have computed the motion of the line pattern by integrating the output of a feature-based mechanism afforded by end-stopping in V1.

But the slow time course observed in MT by Pack & Born does not need to reflect end-stopping in V1. Smith et al. (2005) recently measured the time course over which MT neurons develop pattern selectivity and also found it to be gradual. Because they used sine wave plaids, the observed time course may have had nothing to do with end-stopping, reflecting instead mechanisms within MT or even top-down effects. Of course, one could argue that the overlap regions of the plaids constitute features that could be tracked by end-stopped (or other feature-selective) neurons; in that case, though, one would expect to find pattern cells in the V1 inputs to MT. Movshon & Newsome studied this in 12

MT-projecting V1 cells and did not find any (Movshon & Newsome 1996), but this sample may be too small to tell. Only 9 of these cells were tested with plaids, and only 6 of the 12 cells were in layer 4B. Because 90% of MT's V1 input comes from 4B, it is premature to conclude that the V1 cells feeding MT do not have pattern behavior. In fact, Tinsley et al. (2003) and Guo et al. (2004) both found a small number of pattern-selective cells in V1. Clearly, the issue of whether substantial pattern selectivity occurs in primate V1 remains unresolved. Still, under conditions where Movshon et al. did not find V1 pattern cells, they did find MT pattern cells (Movshon et al. 1985); therefore, the pattern mechanism under those conditions cannot be explained solely in terms of feature tracking.

In summary, there is substantial evidence for both feature tracking in V1 and a two-stage mechanism involving linear filtering in V1 followed by a nonlinear process, something like IOC or vector averaging, in MT. We note that IOC and feature tracking mechanisms are not necessarily incompatible; in fact, there is perceptual evidence that different sorts of computation are at play under different circumstances (Weiss et al. 2002). In the absence of end-stopping, V1 neurons are assumed to extract motion energy at the various spatiotemporal frequencies in the stimulus. For the F-plane model to be robust, it is best to have energy broadly distributed over the frequency plane defined by the object's velocity; otherwise, it is difficult for the MT population to determine the orientation of this plane. This may be why sinusoidal plaid stimuli, whose energy occupies only localized portions of the frequency plane, are not perceived as moving in the IOC direction unless the grating directions are roughly symmetric about the pattern direction (Yo & Wilson 1992). When end-stopping is operational, neurons with extended contours in their receptive fields are suppressed; so the remaining input consists mainly of features that have energy that is well dispersed over the frequency plane. Therefore, the overall effect

would be to provide MT with a well-balanced, broad-spectrum input. At low contrast, where end-stopping becomes weak or even nonexistent (Sceniak et al. 1999), the system would revert to a more linear mode of operation, using signals derived from moving contours as well as features. The idea of representing moving objects in terms of sparse features at high contrast and diffuse frequency components at low contrast is consistent with theoretical studies demonstrating the advantage of minimizing redundancy when signal-to-noise is large, and at the same time it emphasizes sensitivity and noise reduction using broad pooling mechanisms at low contrast (Field 1987, van Hateren 1992). Of course, other nonlinear mechanisms besides end-stopping could also be used to track features (Wilson et al. 1992, Bowns 1996).

## NOISE REDUCTION

Regardless of how pattern velocity is computed, the visual motion system is compelled at some point to filter noise from the processing stream. MT neurons appear to serve this function in a way that is akin to common mode rejection by a differential amplifier (Snowden et al. 1991, Qian & Andersen 1994, Qian et al. 1994, Bradley et al. 1995). Snowden et al. showed that MT neurons responding to dot patterns moving in their preferred direction were suppressed when dot patterns moving in nonpreferred directions were simultaneously introduced (Snowden et al. 1991). This suppression, termed motion opponency, is not a guaranteed result; V1 neurons, for instance, are not substantially affected by nonpreferred motion directions, consistent with their approximately linear behavior (Qian & Andersen 1995). Snowden et al. demonstrated that nonpreferred motions exhibited a suppression in MT that was roughly divisive; that is, the nonpreferred pattern tended to reduce the response gain for the preferred-direction pattern. Later studies by Qian & Andersen showed that suppression of MT by nonpreferred directions is particularly strong

when opposing motions occur within  $\sim 0.5^\circ$  of each other, roughly the scale of V1 receptive fields (Qian & Andersen 1994). The fact that V1 cells do not themselves show appreciable motion opponency suggests that there is strong mutual inhibition at the stage of V1 inputs onto the dendrites of MT neurons. Because a flash emits motion energy simultaneously in all directions, a mechanism that cancels opposite-direction signals on a local scale could be a particularly effective way of reducing responses to flicker.

There is no reason to assume that the only function of motion opponency in MT is to reduce noise. In fact, opponency is likely at least partly a manifestation of gain normalization, the process by which neural responses are scaled according to the total amount of neural activity in their immediate vicinity (Simoncelli & Heeger 1998, Heuer & Britten 2002). But there is little doubt that opponency would tend to reduce the responses to motion noise. Some behavioral evidence also exists: Perceptual studies in a patient with bilateral damage to a region corresponding to MT/V5 found that performance was normal for various motion tasks except when noise was added to the stimulus, in which case performance fell to chance (Zihl et al. 1983, Marcar et al. 1997). This, together with the suppressive effects of nonpreferred motion discussed above, suggests that MT has a basic role in noise reduction within the motion processing stream.

## SEGMENTATION

Although pooling is an important stage in motion processing, it introduces its own problems. It is not uncommon for more than one moving object to appear in the same part of visual space, and pooling these movements would obviously be inappropriate because there is nothing meaningful about the joint (say, average) velocity of two independent objects. Therefore, pooling mechanisms need to be accompanied by parsing mechanisms that distinguish groups of

signals to be integrated separately (Nowlan & Sejnowski 1995, Hildreth et al. 1995).

Several lines of evidence suggest that MT neurons are equipped with response properties that allow them to carry out, or at least contribute to, this parsing. For one thing, the suppression of MT responses normally exerted by nonpreferred motion is relaxed when the nonpreferred and preferred motions occur in separate depth planes (Bradley et al. 1995). Opposing movements in the same part of 2D visual space usually are separated in depth (otherwise they would collide), which would tend to prevent the pooling and cancellation of motion signals emanating from different objects. Also, most MT neurons have an antagonistic surround that is least suppressive under conditions where center and surround stimuli move in different directions and speeds and at different stereoscopic depths (See Surround Mechanisms). This could be a mechanism for emphasizing the relative motion of an object against its background. Other studies suggest that more complex mechanisms may be in place. For example, Albright and colleagues superimposed square wave gratings and manipulated the luminance at the overlaps, according to physical transmittance rules, to produce stimuli that looked either transparent (separate gratings) or coherent (single plaid) in their static state. They then demonstrated that single MT neurons could exhibit either pattern- or component-like direction tuning, depending on whether the stimulus was in its coherent or transparent configuration, respectively (Stoner & Albright 1992, Stoner & Albright 1996). These results could reflect external parsing mechanisms that influence MT, but they could also reflect low-level mechanisms where the overlap regions of the plaids are detected by nonlinear mechanisms in V1, leading to a change in the distribution of inputs to MT. Finally, MT neurons are better able to extract signal (motion) from noise when the two have different isoluminant colors (Croner & Albright 1997, 1999). All together the evidence rather strongly suggests

that MT neurons are critically involved in segmenting an image into separately moving parts.

## THE COMPUTATION OF STRUCTURE

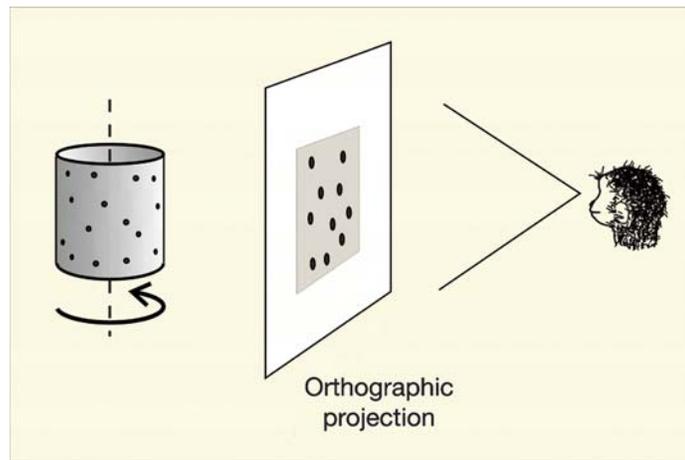
One of the most important discoveries since MT was first located was the binocular disparity tuning of its neurons. Maunsell & Van Essen (1983a) showed, using stereoscopic moving bars, that a majority of MT cells were disparity selective, although it is now known that moving dots, which have a broader frequency content, reveal selectivity in almost all MT cells. The tuning resembles sigmoids, Gaussians, and shapes in between. Curves sometimes peak near zero disparity, but usually they peak or flatten out well to the left or right. Therefore, with some exceptions, these neurons are tuned “near” or “far,” with about a twofold preponderance of the near-tuned (Bradley & Andersen 1998). Importantly, Maunsell & Van Essen showed that MT neurons are not tuned for motion through 3D space. They are simply tuned for a certain 2D direction in a plane a certain distance from the fixation point. This remained something of an anomaly until the mid-1990s when experimenters began searching for possible functions of disparity in MT. First, they showed that null-direction motion, which normally suppresses MT activity, becomes less suppressive if it occurs outside the preferred depth plane of the neuron. The suppression is itself thought to have a role in noise reduction, so this depth constraint provides a way of confining the filtering mechanism to a particular surface. Later it was shown that surround inhibition is also modulated by disparity, predictably in such a way as to minimize inhibition when the surround stimulus is outside the depth plane containing the center stimulus (Bradley & Andersen 1998). This could be a mechanism for segmenting an object from its background. So there is some evidence that disparity is wired into MT in a way that facilitates the processing of

visual motion, presumably for purposes of noise reduction, segmentation, and probably other functions.

However, subsequent studies suggested that MT is involved in the perception of depth itself (Bradley et al. 1998, DeAngelis et al. 1998, Dodd et al. 2001). Two groups of investigators trained monkeys to view revolving cylinders whose direction of revolution was unspecified in the stimulus, but which nevertheless appeared to rotate (**Figure 9**). The perceived rotation was bistable, and monkeys reported their percept on each trial. Both groups found a clear correlation between MT firing rates and the monkeys' judgments. In terms of motion, both percepts meant seeing two, opposite directions; the only difference was the order of these directions in depth. So these were the first neural correlates of the perception of depth. In fact, Dodd et al. (2001) computed choice probability, a measure of the correlation between perception and neural activity (Britten et al. 1996), and found the strongest neuro-perceptual correlation of any MT study, to date.

This role for MT in depth perception has only been strengthened by additional experiments. In fact, of the criteria linking neurons to perception proposed by Parker & Newsome (1998), the only important one that has yet to be met is that of MT lesions affecting disparity judgments. Thus far, this research has rigorously demonstrated (*a*) neuronal selectivity for disparity (Maunsell & Van Essen 1983b, DeAngelis & Uka 2003), (*b*) neuronal disparity sensitivity that is sufficient to account for the abilities of monkeys to perform coarse disparity discrimination tasks (Uka & DeAngelis 2003), (*c*) a predictive relationship between MT neuronal activity and monkeys' perceptual decisions concerning depth (Uka & DeAngelis 2004), and (*d*) the ability to predictably bias monkey's disparity judgments by microstimulation of disparity columns favoring a given depth (DeAngelis et al. 1998).

Another hint that MT neurons may be involved in the extraction of surface



**Figure 9**

Schematic illustrating cylinder experiments first used to link depth perception to MT. A monkey views the orthographic projection, or shadow, of dots revolving on a transparent, cylindrical surface. The planar image contains only dots moving in opposite directions, but the observer perceives a three-dimensional revolving cylinder. Monkeys are trained to report the perception of structure (Siegel & Andersen 1988) or the direction of revolution (Bradley et al. 1998, Dodd et al. 2001).

properties is the presence of inhomogeneities within the receptive field with respect to tuning for both speed (Treue & Andersen 1996) and binocular disparity (Nguyenkim & DeAngelis 2003). As a moving textured surface in the fronto-parallel plane is either tilted forward or backward or slanted to the right or left, the retinal projection of local motion and depth vectors will form a gradient in both speed and disparity. If MT neuronal receptive fields possessed systematic variation in these tuning properties as a function of position within their receptive fields, the resulting structure might serve as a template for a particular 3D surface orientation. Indeed, both of the above groups have demonstrated tuning for surface orientation on the basis of either cue alone. The group testing speed gradients confined their stimuli to the receptive field center, but it appears that, if asymmetries of the surround described above are included, the tuning to speed gradients may become even more marked (Xiao et al. 1997a). Furthermore, it now appears that at least some MT neurons are selective for surface orientation defined by both cues. In such cases, the selectivity is generally consistent and reinforcing, in that tuning is sharper to gradients defined by the combination than by gradients defined by either cue alone (Nguyenkim & DeAngelis 2004).

## EXTRARETINAL EFFECTS

For most behavior, visual information must be integrated with other, nonvisual information. An example of this is the use of copies of the signals used to generate eye movements, so-called efference copy or corollary discharge, to aid in disambiguating retinal motion caused by the eye movements themselves from that due to motion of an object. In ascending the hierarchy of visual areas, at least in the dorsal stream, such extraretinal influences tend to become increasingly powerful. Thus, whereas investigators have reported small-to-modest effects of eye position (Bremmer et al. 1997), saccadic eye movements (Thiele et al. 2002), and attention (Treue & Maunsell 1996) for some MT neurons, much larger and more prevalent effects of these signals have been found in areas above MT, particularly MST and LIP. So, at least to a first approximation, MT is a relatively faithful representation of events occurring on the retina.

The issue of extraretinal influences in the dorsal stream was first examined by Wurtz's group at the National Institutes of Health (NIH) in the context of smooth-pursuit eye movements (Newsome et al. 1988). MT neuronal signals are important for the initiation of pursuit (see Lisberger et al. 1987 for review), and many MT neurons with foveal receptive fields are also active during ongoing pursuit. However, when Newsome et al. (1988) eliminated retinal motion, either by briefly extinguishing the target or by using the recorded eye movements to stabilize the target's image on the retina, the MT responses disappeared. This result stood in marked contrast to many neurons in MST that showed continued directional responses under the same conditions.

Subsequent groups have found some tendency for the position of the eye in the orbit to affect the overall level of responsiveness of some MT neurons, without affecting their tuning for direction (Bremmer et al. 1997). But, in line with the idea of hierarchy, the effects in MT were both rarer and smaller in

magnitude than those found in MST, one of the next higher-tier motion-processing areas to which MT projects. And it is telling that this type of effect was initially discovered in an area still further up the hierarchy—7a of posterior parietal cortex (Andersen et al. 1985), where the effects of eye position are even more profound and where vestibular, auditory, and somatosensory information is also integrated (Andersen et al. 1997). None of these latter signals are known to influence neurons in MT, and anatomical data would suggest they are unlikely to do so.

Perhaps the most prevalent nonretinal influence on MT neurons is that of attention. This is manifest as an enhanced neuronal response to visual stimuli when an animal is attending to either the spatial location of the neuron's receptive field (Seidemann & Newsome 1999, Treue & Maunsell 1999) or to some preferred neuronal feature, such as a particular direction of motion (Treue & Martinez-Trujillo 1999). The gain increase appears to be a straightforward multiplication of the responses, without any changes in the shape of the direction tuning curve (Treue & Maunsell 1999) nor in the nature of the underlying motion computations (Cook & Maunsell 2004). However, the magnitude of the gain may change both as a function of stimulus contrast (Martinez-Trujillo & Treue 2002) and, for feature-based attention, as a function of the similarity between the attended feature and that preferred by the neuron, actually becoming negative as the attended feature approaches the antipreferred feature of the neuron (Martinez-Trujillo & Treue 2004). As for other extraretinal effects, however, the general magnitude of attentional modulation in MT appears to follow the cortical hierarchy: Attentional modulation strengths range from less than 10% in V1 (McAdams & Maunsell 1999), to a ballpark figure of 20%–30% in MT (Treue & Maunsell 1999), to values well over 50% in higher-tier areas, such as MST and 7a (Maunsell & Cook 2002).

## PERCEPTUAL CORRELATES AND POPULATION CODES

### Single-Neuron Sensitivity

In classic experiments at Stanford University, Newsome and colleagues trained macaques to watch moving dot patterns and then indicate which of two opposite directions they had seen (see Parker & Newsome 1998 for review). The strength of the motion signal was controlled by varying the fraction of dots moving coherently, and the remaining dots moved in random directions. The task was executed simultaneously with the recording of single MT neurons, in each case aligning the coherent motion axis with the preferred direction of the recorded neuron. Using methods from signal detection theory, Newsome and colleagues were able to compare neuronal sensitivity directly with that of the monkey itself. Remarkably, they discovered that most MT neurons were at least as sensitive as the monkey itself. Later, the same group defined the choice probability (CP), an extension of detection theory, which reflects the correlation between a subject's judgments and random fluctuations in a single neuron's firing rate. For their task, the mean CP in macaques was only 0.55; chance was 0.50 and perfect correlation was 1. In a computational analysis, Shadlen et al. (1996) concluded that roughly 70–100 neurons would have been involved in the decision pool for the task, far more than would seem necessary considering the exquisite sensitivity of single cells.

Shadlen et al. (1996) suggested several possible reasons for the discrepancy, including correlated noise in MT's inputs, which would limit the benefits obtained by pooling, noise in downstream decision processes, and the likelihood that, owing to limitations in the precision of cortical connectivity, signals from relatively insensitive neurons would also be included, thus degrading the calculation. But this coarse pooling may not occur under all conditions. For example, a recent study at

the University of Chicago showed that when monkeys performed fine direction discrimination, neuronal sensitivity was at best 2–3 times worse than the observer as a whole, and no amount of pooling could account for the psychophysical data unless the most sensitive neurons were selectively pooled (Purushothaman & Bradley 2005). The different results obtained by the two groups probably reflects important differences in the tasks they used. The Stanford task used a large direction difference embedded in noise, whereas the Chicago stimuli were noise-free but the monkeys had to discriminate very small direction differences. One would expect the former task to reflect the sensitivity and filtering capability of the neurons, whereas the latter task should depend more on the slopes of the direction tuning curves. Of course, both tasks probe critical aspects of MT processing. The point is that a neuron's sensitivity relative to the observer is likely to depend on the task, and in particular the computational role the neuron plays in the decision process.

Other examples reinforce this idea. Uka & DeAngelis (2004) trained monkeys in a near/far-depth-judgment task and compared their sensitivity with that of MT neurons. In each case the near and far depths were set at the worst and best values for the MT cell being studied, much like testing a neuron with its preferred and antipreferred directions. Though depth differences were large, the task was made difficult by adding noise to the stimulus. They found that single neurons were on average as sensitive as the observers. In contrast, Liu & Newsome (2003a, 2005) trained monkeys to discriminate small speed differences and found that MT neurons were much less sensitive than the observers. Overall, it would seem that single MT neurons are exquisitely sensitive when the task is to detect a large direction or disparity difference that is heavily corrupted by noise, but they are less sensitive (relative to the observer) when the task requires fine discrimination

of the relevant cue (e.g., direction or speed).

Yet another important factor in interpreting such studies is the role of temporal integration. The Stanford experiments used stimulus presentation intervals of two seconds, after which time monkeys indicated their decision. If the monkeys actually made their decisions much earlier, however, this could affect the ratio of neuron-to-observer sensitivity because the signal detection analysis used the entire two seconds to compute neuronal sensitivity (Mazurek & Shadlen 2002). Cook & Maunsell (2002) found some support for this idea in monkeys trained in a reaction time task with variable-coherence stimuli. When integration times were commensurate for neurons and observers, the neuronal sensitivity was found to be substantially less than observers. Some caution is warranted, however, because the task was not identical to that used in the Stanford experiments. In contrast, Britten et al. (1992) showed that psychophysical and neural thresholds declined similarly with shorter stimulus presentations, which would produce similar neuron/observer-sensitivity ratios regardless of the particular time window chosen. Uka & DeAngelis (2004) also found a similarly small effect of integration interval on neuronal sensitivity in their disparity task, probably because of serial correlation (Osborne et al. 2004). Overall, the bulk of the evidence suggests that observers probably do make decisions in less than two seconds, and that information continues to accumulate for two seconds in spike trains, assuming it is extracted with the type of signal detection analysis used in the detection and discrimination experiments discussed above. That said, the effects probably are not dramatic, and the basic claim of the early Stanford experiments, that monkeys could perform the task in question with a relatively small number of neurons—compared with the hundreds or thousands of neurons one might suspect—is probably correct.

### Vector Summation versus Winner-Take-All

For any simple moving stimulus, the response distribution in MT assumes a roughly Gaussian shape. By “simple” we mean that the stimulus is translating, rigid, and not rotating; by “distribution” we refer to the mean firing plotted on the ordinate versus the preferred direction of the neurons on the abscissa. Under the same visual circumstances that create such a distribution in MT, subjects accurately perceive the direction of the moving stimulus. Given the mountain of evidence linking MT to direction perception, we seek to determine how the response distribution is linked to the direction perception. More formally, we recognize that the response distribution is characterized by many numbers—the firing rates of many neurons—but the perceived direction is a single number. So there must be a code, a rule, for the conversion of many to one. But what is it?

Two simple possibilities come to mind. First, the direction percept may derive from the peak of the distribution; that is, the preferred direction of the most active neurons is taken by decision networks as the direction signal. Second, the overall distribution mean could constitute the direction signal. The Newsome lab has also attempted to probe these two potential mechanisms using microstimulation (Groh et al. 1997, Nichols & Newsome 2002). The basic experiments involved a moving visual target paired with the electrical stimulation of neurons tuned to a different direction, and animals were required either to pursue the visual target or to report simply the direction of motion they perceived. The authors’ rationale was that if direction percepts derive from activity peaks, then the animals’ answers should center on either the direction of the visual target or the direction encoded by the stimulated neurons. In comparison, if direction percepts depend on the activity distribution as a whole, then the perceived direction should lie somewhere in between. The general finding of these

experiments was that direction percepts are usually somewhere between the two extremes, suggesting that direction percepts derived from pooled activity rather than from activity peaks. But a potential confound in these experiments is that animals might have perceived two directions at one time. Forced to choose a single direction, they may have settled on a strategy of splitting the difference. There is, however, no evidence for this.

Besides the winner-take-all and vector-average hypotheses, Weiss et al. (2002) have proposed a Bayesian model of motion perception, which could be implemented with appropriate weighting of MT responses. Unlike vector-average and winner-take-all schemes, the model would address the MT response distribution probabilistically; as such, it could take into account the effects of noise on uncertainty and allow the introduction of priors (biases). The model can explain a remarkable array of psychophysical observations, so it will be interesting to test the model with MT recordings as well.

### Distributed Speed and Acceleration Codes

Maunsell & Van Essen (1983b) found that MT cells were generally broadly speed tuned but, by and large, tended to have distinct preferred speeds. Lagae et al. (1993) later distinguished MT neurons in terms of their speed tuning as being low-pass, high-pass, or band-pass. Because high-pass neurons generally increased firing as stimulus speed increased, they pointed out that the overall mean activity of these cells could be used to gauge the stimulus speed. Recently, both Churchland et al. (2001) and Priebe & Lisberger (2004) found that speed percepts in macaques were consistent with the firing rate-weighted average preferred speed of MT neurons. Lisberger & Movshon (1999) successfully used a similar population average to account for changes in stimulus speed in anesthetized monkeys. Because the population average can shift only if single neurons change their ac-

tivity, one would expect MT firing rates to correlate with speed judgments. Indeed, preliminary evidence suggests that they do (Liu & Newsome 2005).

Lisberger & Movshon (1999) studied MT responses to accelerating random dot patterns. Judging from the neurons' sustained firing rates, there was no evidence for single-neuron encoding of acceleration. However, using a measure of the neurons' transient responses and taking their weighted average, the authors could accurately predict the target's acceleration. This result was important not only because it was the first evidence for acceleration coding in MT, but also because it revealed a signal quite hidden in the population response. One wonders how many other kinds of information are encoded in such distributed forms.

## CONCLUSIONS

Cortical physiologists often claim—whether or not it is true—that in studying a particular area they hope to discover general principles of cortical function. One could argue that more information has been learned from MT than from any other visual area. We make this bold claim because MT has, in many ways, disappointed its explorers, who initially expected a number of obvious contributions to motion processing but did not find them. Although MT may slightly extend motion detection to longer ranges and higher speeds, these are clearly not its main functions; nor does it make the obvious step from speed tuning (as in V1) to acceleration tuning, at least in single neurons. So it is a testament to the persistence and ingenuity of researchers in the field that visual motion research has reexamined itself and reapproached the study of MT in a variety of novel and fruitful ways. These new approaches and their results have been the focus of this review.

Overall, MT does not appear to detect or measure visual motion; this computation occurs in V1. It also does not elaborate substantially on this basic signal; for example,



















