

From Direction of Motion to Patterns of Motion: Hierarchies of Motion Analysis in the Visual Cortex

RICHARD A. ANDERSEN, STEFAN TREUE,
MICHAEL GRAZIANO, ROBERT J. SNOWDEN, AND
NING QIAN

Experiments on the visual cortex of primates have shown partially segregated pathways for the analysis of different attributes of visual perception such as motion, color, and form perception (DeYoe & Van Essen, 1988; Livingstone & Hubel, 1988; Schiller & Logothetis, 1990). Probably the best characterized of these pathways are those involved in the processing of motion information. The cortical areas involved in motion analysis appear to be connected anatomically in stages, which suggests that there is a hierarchy in the machinery for motion perception. This chapter presents experiments we have performed along the most prominent pathway for motion analysis which begins in the primary visual cortex (V1) and proceeds through the middle temporal area (MT) to the medial superior temporal area (MST). One goal of our research has been to determine the sequential processing steps at each anatomical stage in this presumed hierarchy. In the first section, we discuss how a special class of stimuli, transparent motions, have enabled us to show distinct and different mechanisms for the processing of directional information in areas V1 and MT. The second section covers recent work on the possible role of area MST in integrating motion-pattern information. Area MST neurons selective to translating, expanding, contracting, rotating, or spiraling velocity fields were examined. We found that these cells respond to particular patterns of stimuli irrespective of the location in the receptive field or, to a large extent, the size of the stimulus. These cells may represent a stage in higher level processing of, among other things, object motion.

MOTION TRANSPARENCY

Motion transparency occurs often in natural scenes. It exists whenever two different directions of motion occur at local points in an image. One example

of motion transparency is the view one receives looking through a windshield of a moving automobile with rain pouring down the glass. Other less obvious but still common instances of transparency are the motion of a shadow over a textured background, the specular reflections that remain stationary when an object rotates, and the motion discontinuities that occur at the boundaries between a moving object and the background against which it moves. Zucker, Iverson, and Hummel (1990) have pointed out that most computer algorithms developed to analyze moving video images have difficulty representing transparent motion and motion discontinuities because they generally allow only one velocity vector to be present at each pixel in the image. They argue from theoretical and psychophysical grounds that a different algorithm from the usual regularization approaches is required, which allows the simultaneous representation of different directions of motion at a single retinal locus. The fact that artificial systems display an inability to analyze motion transparency suggests that the primate visual system may provide some useful hints for machine vision.

Not only computer algorithms but also physiological models of direction selectivity cannot account for transparent motion perception. The currently prevalent model of motion-direction selectivity employs inhibitory interactions among groups of nerve cells (Barlow & Levick, 1965). These inhibitory interactions would suppress motion-selective cells in transparent conditions and render the visual system blind to motion, similar to the artificial systems mentioned above.

We recorded the activity of motion-selective cells to transparent motion stimuli in areas V1 and MT in behaving monkeys in order to examine how the primate visual system solves this problem (Erickson, Snowden, Andersen, & Treue, 1989; Snowden, Erickson, Treue, & Andersen, 1990; Snowden, Treue, Erickson, & Andersen, 1991). Each neuron's preferred direction of movement was assessed by random dot patterns drifting in different directions within a stationary aperture centered on the receptive field. The one eliciting the greatest discharge was termed the preferred direction and the opposite direction was designated the antipreferred direction. These two "single-surface" stimuli were then presented in an interleaved block of trials along with a stimulus which was comprised of these two stimuli superimposed on one another, producing the transparent condition (Figure 11-1). By and large, the direction-selective V1 cells gave a similar response to the preferred direction and to the two-surface stimulus (Figures 11-2 and 11-3). However, we found that the activity of MT cells stimulated with the two-surface stimulus was always suppressed below the level of that seen to the preferred stimulus alone (Figures 11-2, and 11-3). We concluded that area V1 direction-selective neurons are less suppressed under transparent conditions compared to MT neurons. Thus there appears to be a hierarchy for directional mechanisms in the motion pathway.

The function of inhibitory interactions in MT may be to average or smooth the velocity field. It may also be part of a mechanism for improving signal to noise, which inhibits dissimilar direction signals. We examined the nature of the inhibitory process in area MT in further experiments. By varying the dot density in each of the two directions we "titrated" the effect of excitatory and inhibitory influences. The inhibition appeared to be divisive, reducing the response not by




Figure 11-1 A cartoon representation of the stimuli used in this study. Dots were randomly placed within a window and shifted every frame by a set amount. Those falling outside the original window were wrapped to the opposite side of the display. For the single surface stimuli (A, B) all dots were displaced by the same amount and in the same direction. The two-surface stimuli (e.g., C) had two sets of dots which could undergo separate manipulations. In the stimulus portrayed in C the dots moved in the same speed in opposite directions, and thus this stimulus is equivalent to the superimposition of the stimuli portrayed in A and B. In other experiments the dots could be made to move at any arbitrary angle to one another, or to have different densities in each surface. (From Snowden et al., 1991.)


a set number of spikes per second, but rather by a fraction of the response (Figure 11-4). This divisive inhibition was found in all 15 cells tested in this way.

One consequence of a divisive inhibition is greater contrast insensitivity of the cells for two-surface patterns, assuming that the response to each individual surface is proportional to contrast. This is desirable because, perceptually, the degree of transparency of a pattern is not a sensitive function of its contrast.

In experiments that further explored this cross-direction inhibition, we examined the effect of varying the direction of the inhibiting surface. The one surface is always drifted in the preferred direction of the neuron, and the direction of the inhibiting surface is varied in 45° steps. Figure 11-5 shows the results collapsed over 11 neurons. The inhibition is present at 45° off the preferred direction and by 90° it is nearly maximal. It is interesting to note that the inhibiting surface, at 45° off the preferred direction, will normally excite the MT neurons when it is presented alone.

We also performed the reverse experiment of plotting tuning curves in the presence and absence of the antipreferred stimulus. In this case the inhibiting stimulus is held constant at 180° from the preferred direction, and the second surface is varied in direction to plot the tuning curve. Figure 11-6 shows two neurons examined in this fashion. The peak of the tuning curve remains in the same preferred direction in the presence of the antipreferred stimulus, but the magnitude of the response is greatly reduced. These data again demonstrate that the inhibition is divisive, since the peaks of the tuning curves are much more reduced than the flanks.

A: V1 cell



B: MT cell




Figure 11-2 Examples of the responses of a V1 and a middle temporal (MT) neuron to single and two surface stimuli. The circles containing arrows depict the type of stimulus shown on those trials and the dark bars below the response histograms indicate the time the stimulus was displayed. Each major division of the abscissa is 1,000 msec. (A) The V1 cell gives a directional response (favoring movement up and to the right) and gives a response to the two-surface stimulus which is similar in magnitude to the preferred stimulus alone. (B) The MT cell is also directional but gives a lower response to the two-surface stimulus than it does to the preferred stimulus alone. (From Snowden et al., 1991.)




Figure 11-3 For each of our cells on which the suppression test was performed we calculated a suppression index, I_s , and the frequency of occurrence of this index is plotted for our population of V1 cells (top graph) and middle temporal (MT) cells (bottom graph). An I_s of 0.0 (indicated by the vertical line) means that the cell responded in a similar fashion to the preferred stimulus alone and to the two-surface stimulus. Points to the left of this line indicate a greater response to the two-surface stimulus, while points to the right indicate a smaller response to the two-surface stimulus. The median index for V1 was 0.04 and for MT, 0.54. The distribution of the I_s in our population of V1 and MT cells is significantly different ($p < .0001$). (From Snowden et al., 1991.)

These experiments have interesting implications for mechanisms of direction selectivity for V1 neurons. The classic model comes from the study of direction-selective rabbit retinal ganglion cells. The model postulates that the presence of the antipreferred stimulus causes a wave of inhibition preceding the arrival of the excitatory signal, vetoing any response (Barlow & Levick, 1965). This model, transferred to primary visual cortex of the monkey, would predict a strong suppressive interaction for our transparent motion condition. The comparative lack of suppression in primary visual cortex suggests that many primate

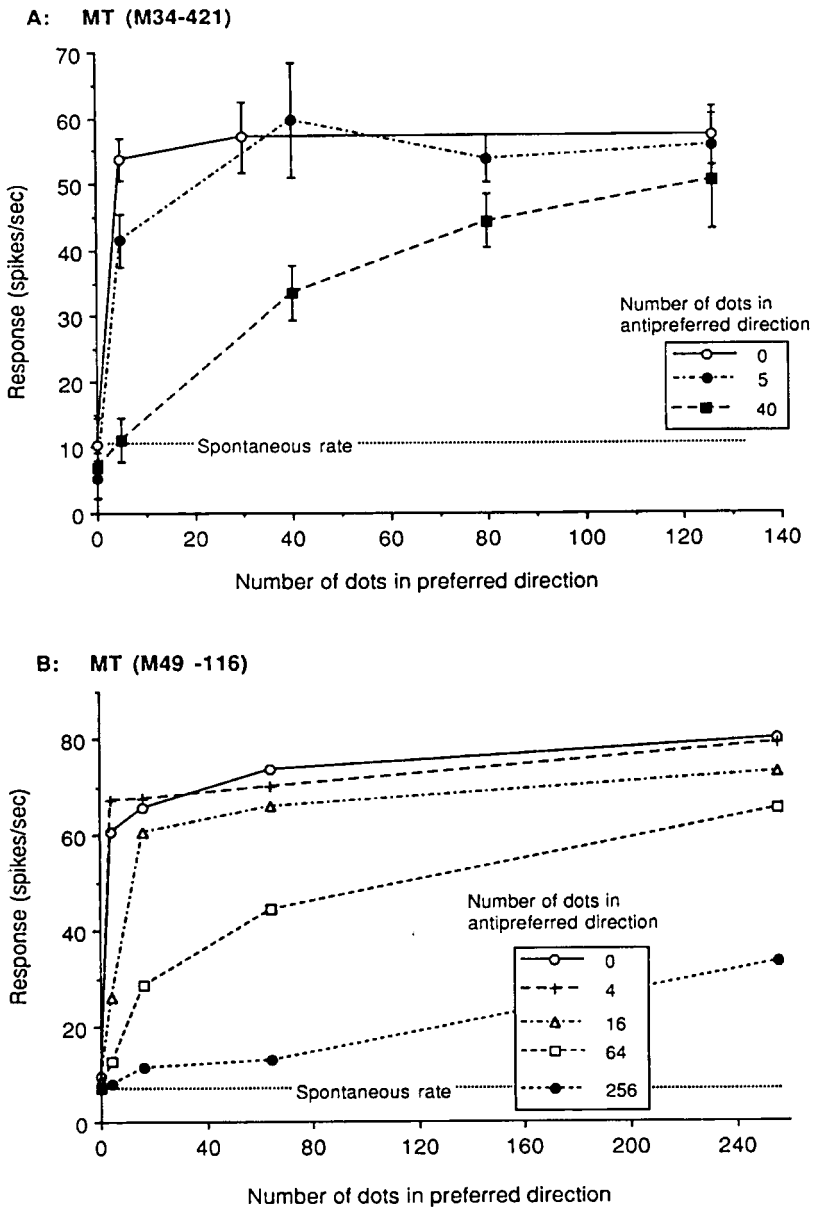


Figure 11-4 (A, B) The response of two middle temporal (MT) cells is plotted as a function of the number of dots in the preferred direction. Each curve corresponds to repeating this function with a certain number of dots always drifting in the antipreferred direction (the number is given on the figure). All trials for each cell were presented in a pseudorandom order. The standard deviation of each point is plotted in **A** but omitted in **B** in order to avoid cluttering. When no dots are drifted in the antipreferred direction the curves rise sharply and asymptote with increasing dot density of the preferred direction. The addition of dots in the antipreferred direction causes this rise to be more gradual but does not necessarily reduce the maximum response of the cell (though this is now reached at a greater dot density). Increasing the dot density in the antipreferred direction increases these effects. (From Snowden et al., 1991.)





Figure 11-5 The direction tuning of the suppressive effect across a population of 11 middle temporal (MT) neurons. The test for each neuron is described in the text. The response in each condition was then scaled to the response to the preferred direction, and the mean and standard error are portrayed in the figure. (From Snowden et al., 1991.)

V1 neurons might achieve directionality by other mechanisms, one possibility being facilitation in the preferred direction (Barlow & Levick, 1965). Facilitatory mechanisms have been proposed for directional selectivity in the cat striate cortex, although inhibitory mechanisms are generally believed to be more prominent (Movshon, Thompson, & Tolhurst, 1978; Douglas, Martin, & Whitteridge, 1988; Goodwin & Henry, 1975; Sillito, 1979).

One important concern that may influence the interpretation of these experiments has to do with the density of the dots and the relative sizes of V1 and MT receptive fields. It could be argued that the smaller degree of inhibition of V1 may be due to the fact that the smaller V1 receptive fields may be more likely to have fluctuations in the balance of preferred and antipreferred moving dots. These fluctuations may give rise to more activity when the number of preferred dots is greater than average, or the antipreferred dots less than average. We have performed two experiments to control for this potential artifact. In the first, we used counterphase gratings, which are equivalent to two sine-wave gratings moving in opposite directions (Qian, Andersen, & Adelson, 1991). Thus the counterphase grating is a physically transparent stimulus, although it does not lead to the perception of transparency. Since it is a continuous stimulus there are no local fluctuations that occurred for the random dot surfaces. Interestingly, V1 cells showed similar suppression indices for counterphase and transparent random dot patterns, arguing that the local fluctuations in the dot

A: MT (M49-136)



B: MT (M49-134)




Figure 11-6 The direction tuning of two middle temporal (MT) cells under conditions of just a single surface (solid line and black circles), and where an equal number of dots was always drifting in the antipreferred direction (180°, open dashed line and white circles). The spontaneous rate is also shown as a dotted line. The greatest suppression occurs at 0° when the cell would be most active. (From Snowden et al., 1991.)




Figure 11-7 The suppression index calculated when the surfaces in the transparent condition contained 64 points each is plotted against the same index when the surfaces contained 32 points each. The single-surface stimulus used to determine the response to the preferred direction alone contained 64 points in both cases. The top graph contains the results from V1 cells, and the bottom graph results from middle temporal (MT) cells. Both plots show a strong correlation between the two indexes and points fall with approximately equal preference above and below the dotted line at 45°, which indicates no systematic shift in the I_s with dot density. The V1 data have a greater scatter than that of the MT data and this scatter was found to increase with decreasing dot density. This scatter may reflect the statistical probabilities of the number of dots in the receptive fields of the V1 neurons. (From Snowden et al., 1991.)

patterns are not a concern. In a second control experiment we plotted the suppression index for 32 dot surfaces against the suppression index for 64 dot surfaces. If the degree of suppression is not influenced by the density of the dots, then these points should lie along a 45° line. Figure 11-7 shows that this is in fact the case for both V1 and MT neurons. There was a bit more scatter for the V1 cells, probably reflecting local fluctuations in balance within the receptive field, but these fluctuations were not biased toward greater suppression at the higher dot density. We further confirmed this result for dot densities ranging from .45 to 28.8 dots/deg² (a 12-fold difference) in 12 V1 and 8 MT neurons.

The suppression in MT may result from neurons receiving opposing inputs from different directional cells in the projection from V1 to MT, with perhaps an intervening interneuron to reverse the sign of one of the inputs. MT neurons may be arranged in a network in which cells inhibit one another. It is also possible that direction selectivity may be reproduced *de novo* in MT, using inhibitory mechanisms (Rodman & Albright, 1989) that are evident in the transparent condition. At present, our data cannot distinguish between these three alternatives. γ -Aminobutyric acid (GABA) is a likely inhibitory transmitter for producing this type of inhibition for direction selective neurons in light of experiments performed in the cat cortex and fly visual system (Sillito, 1979; Koch, Poggio, & Torre, 1986; Egelhaaf, Borst, & Pils, 1990).

The two stages demonstrated in these experiments have interesting parallels with prominent models of motion processing which also have two stages (Adelson & Bergen, 1985; Reichardt, 1961). The first stage of these models measures local motions, and the second stage performs some form of integration that averages or "smoothes" the velocity field. Similarly, the neurons in area V1 behave like directional filters, extracting their preferred direction of motion from the directions present in the image. As a result, the transparent motion stimulus activates two separate groups of neurons tuned to opposite directions of movement. This segregation could then form the basis for segmenting the stimulus into two separate moving surfaces (Zucker et al., 1990). Thus the V1 neurons may correspond to the first stage of the above models where local motion measurements are made.

The transparent motion stimulus would also activate two subpopulations of neurons in MT tuned for different directions of motion. Our results show, however, that this activation is considerably less than when each direction is presented alone. If area MT is a major site for the processing of the perception of motion (Siegel & Andersen, 1986; Newsome, Britten, & Movshon, 1989; Newsome & Paré, 1988), even under transparent conditions, then the recordings suggest that each direction of motion would be less detectable under transparent conditions, compared to when the direction of motion is presented in isolation. This result has been reported for transparent motions of orthogonal direction (Snowden, 1989).

RESPONSE OF MEDIAL SUPERIOR TEMPORAL AREA NEURONS TO MOTION PATTERNS

The Stimuli

Studies by Sakata and colleagues (Sakata, Shibutani, Kawano, & Harrington, 1985; Sakata, Shibutani, Ito, & Tsurugai, 1986) and Tanaka and Saito and

colleagues (Saito et al., 1986; Tanaka et al., 1986, Tanaka, Fukada, & Saito, 1989; Tanaka & Saito, 1989), and Duffy & Wurtz (1991a, 1991b) have reported neurons in area MST that are selective for rotating or expanding or contracting visual stimuli. There are several complicating parameters that often accompany these types of stimuli. In cases where a textured stimulus such as a random dot field is used, there is always the potential that the neurons are responding to shape cues in the stimulus and not to its motion. The expansion and contraction stimuli are sometimes generated, using dot patterns and zoom lenses, which produce changes in luminance, dot size, and dot density along with the movement. Even natural stimuli produce changes in texture size and density with motion toward or away from the animal. The individual elements of a rotating field have curved trajectories and, for expanding/contracting stimuli, they have accelerations. Moreover, it is generally difficult to equate motion speed distributions when testing neurons with rotation and expansion/contraction stimuli. These complications were appreciated by the investigators, and they initiated controls that indicated it was likely these neurons were in fact selective to the pattern of motion rather than any other, confounding variable (Tanaka et al., 1989; Tanaka & Saito, 1989).

In order to extend these studies, we found it necessary to develop new expansion/contraction and rotation stimuli, with the global pattern of the motion as the only difference. By using this approach we could begin to manipulate these displays without risking the possibility of introducing confounding variables. These new stimuli have dots of finite point lifetime which are randomly, asynchronously replotted so that the static shape cues are constantly, randomly changing. The use of limited point lifetimes also enable an even distribution of points at all times. Each short-lived velocity vector is linear and produces no curvature or acceleration cues. The speed distributions for expansion, contraction, and rotation are all equivalent. The only difference between the displays are the global patterns of individual motions. In the case of expansion the dots move outward from the center of the display, for rotation they move tangentially, and for contraction they move toward the center of the display.

Using these highly controlled stimuli, we confirmed that area MST neurons were often selective for expansion, contraction, clockwise or counterclockwise rotation, or pure translational movement (Graziano, Andersen, & Snowden, 1990). As reported earlier by Duffy and Wurtz (1991a) some cells responded to different combinations of these stimuli, whereas others only responded to one type of movement. Neurons that responded to all three classes of stimulus (rotation, expansion/contraction, and translation) are referred to as triple-component, two classes as double-component, and only one class as single-component. We also found a small number of rotation cells that responded to both clockwise and counterclockwise movement. These cells were recorded from the dorsal aspect of area MST (MSTd), consistent with previous findings (Tanaka et al., 1986; Saito et al., 1986; Duffy & Wurtz, 1991a).

Position Invariance

If area MSTd cells are truly selective for pattern, then they would show position invariance; they should be selective for the same stimulus regardless of its po-

sition in the large receptive field. Position invariance was tested by placing 10° diameter stimuli at five overlapping locations in the receptive field arranged in a cloverleaf pattern such that the local direction of motion in the areas of overlap was reversed in direction (Andersen, Graziano, & Snowden, 1990). In every case tested the cells retained the same stimulus selectivity, indicating that they were position invariant. These cells also showed a degree of size invariance, responding well to 10° and 20° diameter stimuli, although they often tended to respond better to the larger stimuli. We confirmed Duffy and Wurtz's finding that single-component cells tend to have a little higher degree of position invariance.

Spiral Tuning

The finding that some MSTd neurons were selective for only rotation, expansion/contraction, or translation suggested that MST may play a role in the analysis of "optical flow" generated by locomotion through the environment (Koenderink & van Doorn, 1981; Longuet-Higgins & Prazdny, 1980; Rieger & Lawton, 1985; Warren & Hannon, 1988). These signals form orthogonal components that can be used to decompose the optical flow and provide information about the direction of heading and organization of the environment. We decided to test this idea directly by examining the selectivity of MST neurons for rotation and expansion/contraction, and stimuli intermediate to these two classes, which comprise a family of spirals. The prediction is that if area MST cells are doing an orthogonal decomposition, then most cells should prefer expansion/contraction or rotation over spiraling stimuli.

Since our stimuli had balanced vector fields we could generate a complete continuum of stimuli on a polar plot in which 90° corresponds to expansion, 270° contraction, 0° clockwise rotation, and 180° counterclockwise rotation (see Figure 11-8). All intermediate angles represent spiral stimuli with, for example, angles between 0° and 90° representing various outward and clockwise spirals. By indicating the magnitude of the neural responses by the length of the radius in these graphs, we could plot "spiral" tuning curves. We generated tuning curves for 66 MST neurons. All but two cells showed single peak tuning curves, these two exceptions were cells that showed peaks for both clockwise and counterclockwise rotation. Figure 11-8 shows a cell that gave a best response for expansion. Figure 11-9 illustrates a cell that preferred inward, clockwise spiral. Of 33 cells tested with spiral stimuli, we also tested with up, down, left, and right translational motion. Only 8 of these cells responded better to the translational motion than the spiral, expansion/contraction, or rotation. Some cells had peaks for spiraling motion and gave virtually no response for translation.

Of the 66 neurons, 57 (86%) had tuning curves with good gaussian fits ($r > .9$). For these 57 neurons we calculated a best direction. These directions ranged widely over the whole spiral graph. If area MST was using orthogonal components, then the best directions should have clustered along the cardinal axes representing pure expansion/contraction and rotation, which they did not.

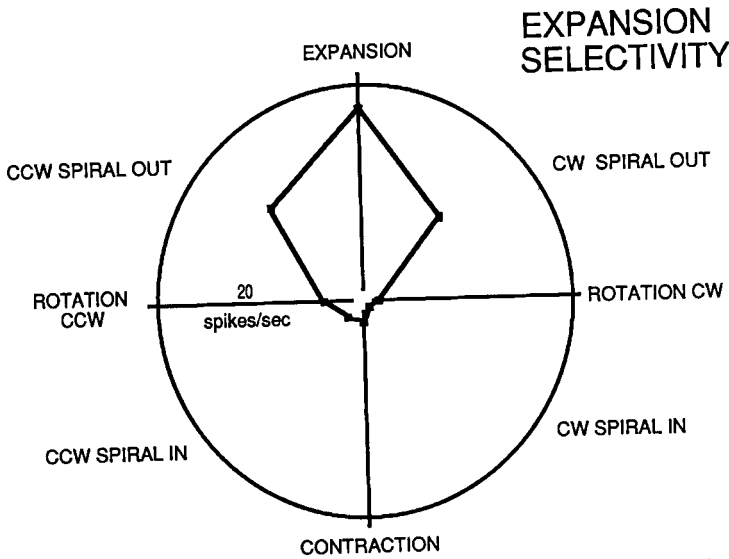


Figure 11-8 Tuning curve for a medial superior temporal (MST) cell which prefers expanding dot patterns. CCW, counterclockwise; CW, clockwise.

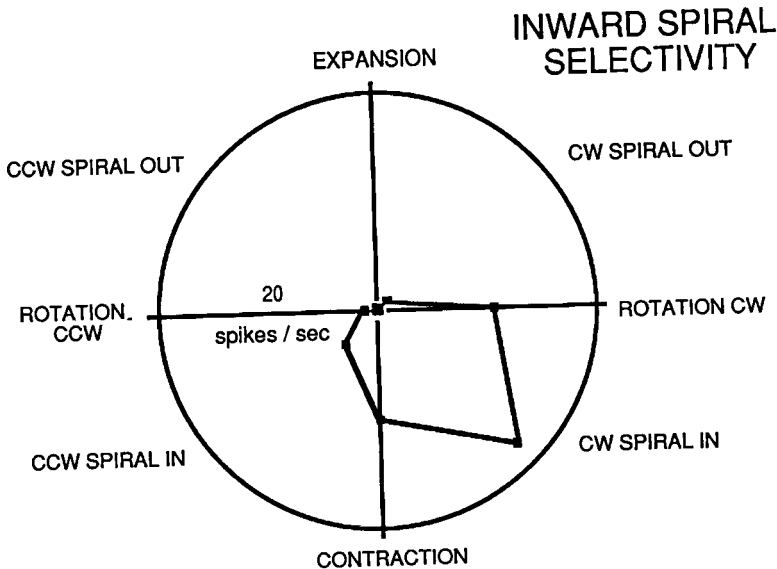


Figure 11-9 Tuning curve for a medial superior temporal (MST) cell which prefers inward-clockwise spirals. CCW, counterclockwise; CW, clockwise.

A Possible Explanation for Variations in Position Invariance Between Single-, Double-, and Triple-Component Cells

As mentioned above, single-component cells were found to give the highest degree of position invariance. The cells that are tuned best for spiraling stimuli are typically double-component, giving weaker responses for both contraction/

expansion and rotation. Some are also triple-component, responding to translational motion. We decided to test the hypothesis that multicomponent cells appear less position invariant because their preferred stimulus lies between the test stimuli. The prediction is that spiral tuned cells will be as position invariant as expansion or rotation tuned cells, if they are tested with their preferred, spiraling stimulus.

We tested position invariance by generating complete spiral tuning curves at two locations in the receptive field. The stimuli were overlapping and the two locations spanned 25° of the receptive field. For 24 cells that gave significant responses at the two locations, all showed a high degree of position invariance with similar best directions at the two locations. Twenty-two of the cells showed a good gaussian fit and best directions were computed; 8 of the cells were tuned to spiral motion and 14 to expansion/contraction or rotation. The average shift in preferred direction for the population of cells was 10.7 degrees. The average shift of the spiral tuned (11.2°) and expansion/contraction or rotation tuned cells (10.5°) was not significantly different for the two groups.

To test whether cells were less position invariant along the flanks of their tuning curves compared to the centers, we calculated a change in direction tuning metric in the best direction and 45° to either side of the best direction. The average change in direction across the population of cells was over twice as great on the flanks compared to the center of the tuning curve, and was statistically significant.

These results suggest that cells are equally position invariant, whether single- or double-component, when tested with their preferred or more nearly preferred stimulus. Double- and single-component cells were not different, at least when tested in the domain of expansion/contraction and rotation.

Size

Twenty cells were tested for the effect of stimulus size on selectivity. On average, cells maintained their selectivity over different sizes, but the average response for 10° stimuli was 76% of the response for 20° stimuli.

Implications of These Results for the Analysis of Optical Flow

These results suggest that area MST neurons do not perform an orthogonal decomposition of the optical flow into strictly rotation and expansion/contraction components. The experiments add another class of motion for which area MST cells are selective, spiral motions. Spiral motions are quite common in natural scenes, occurring under conditions of motion parallax during locomotion and for the motion of objects. The position invariance result complicates any mechanism for deriving direction of heading information for navigation. The task for navigation is to subtract out motions generated by eye movements to extract the expansion component of the flow, whose central focus corresponds to the heading direction. Since area MST expansion cells show position invariance anywhere in their large receptive fields, they do not provide an appropriate signal for localizing the focus of expansion. If area MST is important for navi-

gation, the encoding of the locus of expansion would have to be performed by a population of neurons using a course-coding style of computation.

Probably the best fit of the experimental data to function is to assign area MST a role in the analysis of pattern motion. Area MST could be thought of as the motion equivalent of area IT. Both have large receptive fields, which include the fovea. Both are tuned to a wide variety of patterns, although IT is specialized for static patterns and MST is specialized for patterns of movement. Finally, both show position invariance, giving a similar selectivity for their preferred stimulus over the large area of their receptive fields.

CONCLUDING REMARKS

The experiments outlined above provide evidence for hierarchical processing along the cortical visual motion pathway in macaque monkeys. Area V1 direction-selective cells appear to act as directional filters signalling motion in the preferred direction with the non-preferred direction having little or no effect. Transparent motion perception may be made possible by allowing separate and multiple independent measures of direction at the initial stage of motion analysis. On the other hand, for area MT neurons, the nonpreferred directions powerfully inhibit the response to preferred direction stimuli. The functional role for inhibitory interactions between different directions at this later stage is not clear, but it may perform a smoothing operation that could be important for interpolating surfaces from sparse data (Siegel & Andersen, 1988; Husain, Treue, & Andersen, 1989; Treue, Husain, & Andersen, 1990), reducing noise, or determining the direction of pattern motion from component motion measurements (Bülthoff, Little, & Poggio, 1989; Wang, Mathur, & Koch, 1989). Area MST appears to play a higher level role in the processing of motion information with cells selective for particular patterns of motion. The cells show position invariance giving similar response selectivity at different positions in their large receptive fields. MST cells do not respond only to rotation and expansion motions, but to a wide range of intermediate spiral patterns as well. This result suggests that MST does not perform a single, orthogonal decomposition of the flow field into rotation and expansion/contraction components.

ACKNOWLEDGMENTS

This work was supported by NIH grant EY07492. Our thanks go to Gail Robertson, Stephen Marchetti, and Richard Keough for technical assistance; and to Carol Andersen and Catherine Cooper for editorial assistance.

REFERENCES

- Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America A*, 2, 284-299.

- Andersen, R. A., Graziano, M., & Snowden, R. J. (1990). Translational invariance and attentional modulation of MST cells. *Society of Neuroscience Abstracts*, 16, 7.
- Barlow, H. B., & Levick, W. R. (1965). The mechanism of directionally selective units in the rabbit retina. *Journal of Physiology, London*, 178, 477-504.
- Bülthoff, H., Little, J., & Poggio, T. (1989). A parallel algorithm for real-time computation of optical flow. *Nature*, 337, 549-553.
- DeYoe, E. A., & Van Essen, D. C. (1988). Concurrent processing streams in monkey visual cortex. *Trends in Neurosciences*, 11, 219-226.
- Douglas, R. J., Martin, K. A. C., & Whitteridge, D. (1988). Selective responses of visual cortical cells do not depend on shunting inhibition. *Nature*, 332, 642-644.
- Duffy, C. J., & Wurtz, R. H. (1991a). Sensitivity of MST neurons to optic flow stimuli. I. A continuum of response selectivity to large-field stimuli. *Journal of Neurophysiology*, 65, 1329-1345.
- Duffy, C. J., & Wurtz, R. H. (1991b). Sensitivity of MST neurons to optic flow stimuli. II. Mechanisms of response selectivity revealed by small-field stimuli. *Journal of Neurophysiology*, 65, 1346-1359.
- Egelhaaf, M., Borst, A., & Pilsz, B. (1990). The role of GABA in detecting visual motion. *Brain Research*, 509, 156-160.
- Erickson, R. G., Snowden, R. J., Andersen, R. A., & Treue, S. (1989). Directional neurons in awake rhesus monkeys: Implications for motion transparency. *Society of Neuroscience Abstracts*, 15, 323.
- Goodwin, A. W., & Henry, G. H. (1975). Direction selectivity of complex cells in a comparison with simple cells. *Journal of Neurophysiology*, 38, 1524-1540.
- Graziano, M., Andersen, R. A., & Snowden, R. J. (1990). Stimulus selectivity of neurons in macaque MST. *Society of Neuroscience Abstracts*, 16, 6.
- Husain, M., Treue, S., & Andersen, R. A. (1989). Surface interpolation in three-dimensional structure-from-motion perception. *Neural Computation*, 1, 324-333.
- Koch, C., Poggio, T., & Torre, V. (1986). Computations in the vertebrate retina: gain enhancement, differentiation, and motion discrimination. *Trends in Neurosciences*, 9, 204-211.
- Koenderink, J. J., & van Doorn, A. J. (1981). Exterspecific component of the motion parallax field. *Journal of the Optical Society of America*, 71(8), 953-957.
- Livingstone, M. S., & Hubel, D. H. (1988). Segregation of form, color, movement and depth: Anatomy, physiology and perception. *Science*, 240, 740-749.
- Longuet-Higgins, H. C., & Prazdny, K. (1980). The interpretation of a moving retinal image. *Proceedings of the Royal Society, London B*, 208, 385-397.
- Movshon, J. A., Thompson, I. D., & Tolhurst, D. J. (1978). Receptive field organization of complex cells in the cat's striate cortex. *Journal of Physiology, London*, 283, 79-99.
- Newsome, W. T., Britten, K. H., & Movshon, J. A. (1989). Neuronal correlates of a perceptual decision. *Nature*, 341, 52-54.
- Newsome, W. T., & Paré, E. B. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *Journal of Neuroscience*, 8, 2201-2211.
- Qian, N., Andersen, R. A., & Adelson, E. (1991). V1 responses to two-surface transparent and non-transparent motion. *Society of Neuroscience Abstracts*, 17, 177.
- Reichardt, W. (1961). Autocorrelation: A principle of the evaluation of sensory information by the central nervous system. In ?? Rosenblith (Ed.), *Sensory communication*. New York: Wiley.
- Rieger, J. H., & Lawton, D. T. (1985). Processing differential image motion. *Journal of the Optical Society of America A*, 2, 354-360.

- Rodman, H. R., & Albright, T. D. (1989). Afferent basis of visual response properties in area MT of the macaque: 1. Effects of striate cortex removal. *Journal of Neuroscience*, *9*, 2033-2050.
- Saito, H., Yukio, M., Tanaka, K., Hikosaka, K., Fukada, Y., & Iwai, E. (1986). Integration of direction signals of image motion in the superior temporal sulcus of the macaque monkey. *Journal of Neuroscience*, *6*, 145-157.
- Sakata, H., Shibutani, H., Ito, Y., & Tsurugai, K. (1986). Parietal cortical neurons responding to rotary movement of visual stimuli in space. *Experimental Brain Research*, *61*, 658-663.
- Sakata, H., Shibutani, H., Kawano, K., & Harrington, T. (1985). Neural mechanisms of space vision in the parietal association cortex of the monkey. *Vision Research*, *25*, 453-464.
- Schiller, P., & Logothetis, N. K. (1990). The color-opponent and broad-band channels of the primate visual system. *Trends in Neurosciences*, *13*, 392-398.
- Siegel, R. M., & Andersen, R. A. (1986). Motion perceptual deficits following ibotenic acid lesions of the middle temporal area in the behaving rhesus monkey. *Society of Neuroscience Abstracts*, *12*, 1183.
- Siegel, R. M., & Andersen, R. A. (1988). Perception of three-dimensional structure from motion in monkey and man. *Nature*, *331*, 259-261.
- Sillito, A. M. (1979). Pharmacological approach to visual cortex. *Trends in Neurosciences*, *2*, 196-198.
- Snowden, R. J. (1989). Motions in orthogonal directions are mutually suppressive. *Journal of the Optical Society of America A*, *6*, 1096-1101.
- Snowden, R. J., Erickson, R. G., Treue, S., & Andersen, R. A. (1990). Transparent motion stimuli reveal divisive inhibition in area MT of macaque. *Investigative Ophthalmology and Visual Science* *31*(Supplement) 399.
- Snowden, R. J., Treue, S., Erickson, R. G., & Andersen, R. A. (1991). The response of area MT and V1 neurons to transparent motion. *Journal of Neuroscience*, *11*, 2766-2785.
- Tanaka, K., Fukuda, Y., & Saito, H. (1989). Underlying mechanisms of the response specificity of the expansion/contraction and rotation cells in the dorsal part of the medial superior temporal area of the macaque monkey. *Journal of Neurophysiology*, *62*, 642-656.
- Tanaka, K., Hikosaka, K., Saito, H., Yukie, M., Fukada, Y., & Iwai, E. (1986). Analysis of local and wide-field movements in the superior temporal visual areas of the macaque monkey. *Journal of Neuroscience*, *6*, 134-144.
- Tanaka, K., & Saito, H. (1989). Analysis of motion of the visual field by direction, expansion/contraction, and rotation cells clustered in the dorsal part of the medial superior temporal area of the macaque monkey. *Journal of Neurophysiology*, *62*, 626-641.
- Treue, S., & Andersen, R. A. (1990). 3-D structure from motion: Rigidity and surface interpolation. *Investigative Ophthalmology and Visual Science* *31*(Supplement), 172.
- Treue, S., Husain, M., & Andersen, R. A. (1991). Human perception of structure from motion. *Vision Research*, *31*, 59-75.
- Wang, H. T., Mathur, B., & Koch, C. (1989). Computing optical flow in the primate visual system. *Neural Computation*, *1*, 92-103.
- Warren, W. H., & Hannon, J. (1988). Direction of self-motion is perceived from optical flow. *Nature*, *336*, 162-163.
- Zucker, S. W., Iverson, L., & Hummel, Robert A. (1990). Coherent compound motion: Corners and nonrigid configurations. *Neural Computation*, *2*, 44-57.