Hierarchical Processing of Motion in the Visual Cortex of Monkey

R.A. ANDERSEN, R.J. SNOWDEN, S. TREUE, AND M. GRAZIANO
Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

Recent work on the visual system of primates has delineated several cortical fields involved in the processing of visual motion. These cortical areas appear to be connected anatomically in stages, which suggests that there is a hierarchy in the machinery for motion perception. In this paper, we outline experiments that we have performed along the most prominent pathway for motion analysis, which begins in area V1 and proceeds through the middle temporal area (MT) to the medial superior temporal area (MST). Our goal has been to demonstrate what sequential processing steps take place at each anatomical stage in this presumably hierarchical system. In the first section, we show how a special class of stimuli, transparent motions, have enabled us to separate distinct and different mechanisms for the processing of directional information in areas V1 and MT. In the second section, we discuss psychophysical experiments that explore the ability to perceive shape from motion. These results indicate that the brain interpolates motion data to form neural representations of moving surfaces, and they have interesting implications for our recording experiments. The last section covers recent work on the possible role of area MST in integrating motion-pattern information. Area MST neurons selective to expanding, contracting, or rotating 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 motion, including the perception of shape from motion.

Motion Transparency

Motion transparency is quite common 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. Computer algorithms developed to analyze moving video images have difficulty representing transparent motion because they generally allow only one velocity vector to be present at each pixel in the image. The fact that artificial systems display an inability to analyze motion transparency suggests that the primate visual system has developed a specialized method for its perception.

Motion transparency presents a problem not only to computer algorithms, but also to models of how the nervous system processes motion. The currently prevalent model of motion-direction selectivity employs inhibitory interactions among groups of nerve cells (Barlow and Levick 1965). These inhibitory interactions would suppress motion-selective cells under transparent conditions and render the visual system blind to motion, similar to the artificial systems mentioned above.

To examine how the primate visual system solves this problem, we recorded the activity of motion-selective cells to transparent motion stimuli in areas V1 and MT in behaving monkeys (Erickson et al. 1989; Snowden et al. 1990). Each neuron's preferred direction of movement was assessed by drifting random-dot patterns in different directions through the receptive field. The one eliciting the greatest discharge was termed the preferred direction, and the opposite direction was designated the antipreferred direction (Fig. 1). These two "single-surface" stimuli were then presented in an interleaved block of trials along with a stimulus that was composed of these two stimuli superimposed on one another, producing the transparent condition (Fig. 1). By and large, the direction-selective V1 cells gave a similar response to the preferred direction and to the two-surface stimulus. 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. We concluded that area V1 direction-selective neurons are only minimally affected under transparent conditions, whereas the activity of MT neurons is suppressed. Thus, there appears to be a hierarchy for directional mechanisms in the motion pathway.

In further experiments, we examined the nature of the inhibitory process in area MT. We "titrated" the effect of excitatory and inhibitory influences by varying the dot density in each of the two directions. It was found that the inhibition appeared to be divisive, reducing the response not by a set number of spikes per second, but rather by a number that increases with the increased firing of the neuron. These interactions may

Cold Spring Harbor Symposia on Quantitative Biology, Volume LV. © 1990 Cold Spring Harbor Laboratory Press 0-87969-059-3/90 $1.00

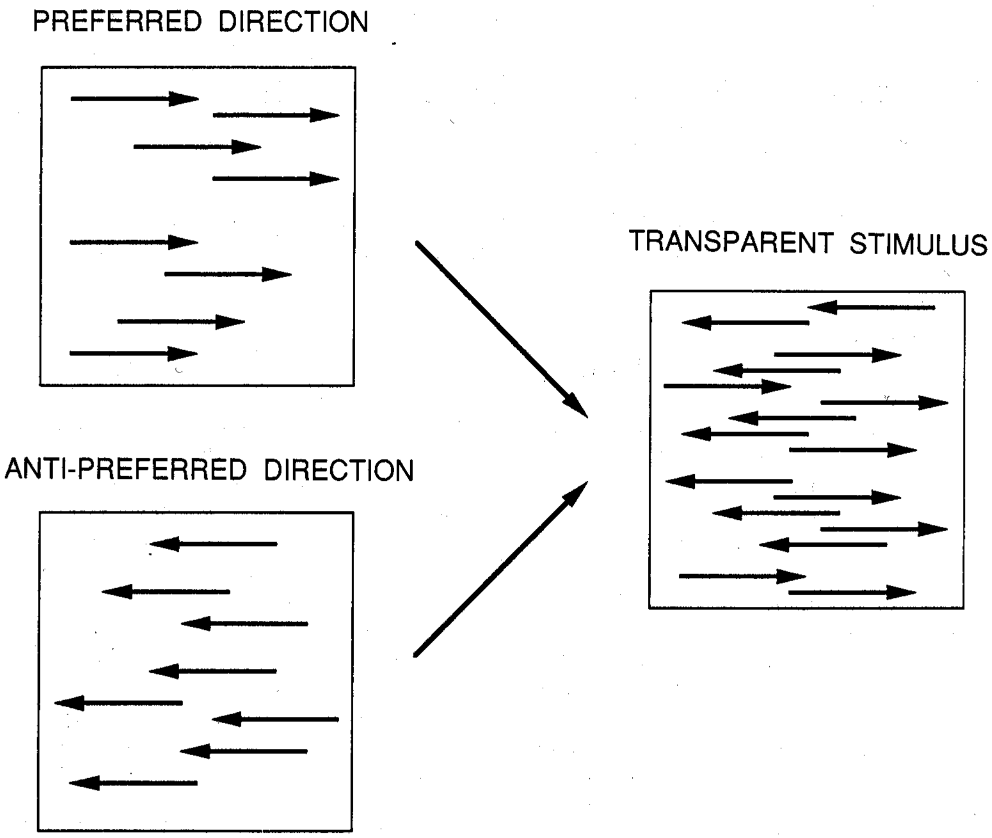

Figure 1. Example of the transparent motion stimulus used for recording experiments in areas V1 and MT. This stimulus is composed of two sets of dots: one set drifting in the preferred direction of the neuron being recorded from and the other set drifting in the opposite (antipreferred) direction.

be considered a kind of averaging or smoothing of the velocity field.

These results bear upon the issue of how area V1 neurons achieve a directional response. A classic model of how rabbit retinal ganglion cells achieve directional responses postulates that the presence of the antipreferred stimulus causes a wave of inhibition preceding the arrival of the excitatory signal, vetoing any response (Barlow and Levick 1965). Since this model predicts strong suppressive interactions for our transparent motion condition, the comparative lack of such suppression in primary visual cortex suggests that many primate V1 neurons might achieve directionality by other mechanisms such as facilitation in the preferred direction (Barlow and Levick 1965). Facilitatory mechanisms have occasionally been proposed for directional selectivity in the cat striate cortex, although inhibitory mechanisms are generally believed to predominate there (Goodwin and Henry 1975; Movshon et al. 1978; Sillito 1979; Douglas et al. 1988).

The suppression of the directional response in area MT by nonpreferred directions of motion may result from single MT neurons receiving opposing inputs from different directional cells in the projection from V1 to MT. It is also possible that direction selectivity may be reproduced de novo in MT using inhibitory mechanisms (Rodman and Albright 1989). At present, our data cannot distinguish between these alternatives. The divisive nature of the inhibition may result from an increase in chloride conductance producing a shunting inhibition (Koch et al. 1986). GABA has been proposed as a likely inhibitory transmitter for producing this type of inhibition for direction-selective neurons in the cat cortex and fly visual system (Sillito 1979; Koch et al. 1986; Egelhaaf et al. 1990).

The two-stage analysis demonstrated in these data has interesting parallels for some models of motion processing in humans and flies (Reichardt 1961; Adelson and Bergen 1985). The first stage of the models for both species measures local motions and the second stage then performs some form of integration, the most common being to average or "smooth" the velocity field. The neurons of area V1 behave like directional filters, extracting their preferred direction of motion from the directions present in the image. Therefore, the transparent motion stimulus appears to activate 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. As such, the V1 neurons may correspond to stage 1 of the above models, where local motion measurements are made.

In area MT, the transparent motion stimulus would also activate two subpopulations of neurons tuned for different directions of motion. However, our results show 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 and Andersen 1986; Newsome and Paré 1988; Newsome et al. 1989), 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. Such a result has been recently reported for transparent motions of orthogonal direction (Snowden 1989). Alternatively, one of several possible parallel motion pathways, such as V1 to V2, V3, or the parieto-occipital area, may be responsible for maintaining the segregation established in V1.

The results point to a two-stage hierarchical process in primate motion perception. The first stage, area V1, takes local velocity measurements and segments the image, and the second, area MT, spatially integrates these measurements. Transparent motion perception may be made possible by allowing separate and multiple independent measures of direction at the initial stage of motion analysis. The reason for the apparent combination of these independent measures with divisive interactions at a later stage in area MT is not clear, but it may perform a smoothing operation that could be important for interpolating surfaces from sparse data (Siegel and Andersen 1988; Husain et al. 1989; Treue et al. 1990), reducing noise, or determining the direction of pattern motion from component motion measurements (Bülthoff et al. 1989; Wang et al. 1989).

Surface Reconstruction for Structure-from-motion Perception

Humans can easily perceive the three-dimensional shape of objects from motion information. For example, the shadow of a wire frame projected onto a flat surface gives little information about the shape of the frame; however, when it is rotated, the three-dimensional shape becomes immediately apparent (Wallach and O'Connell 1953). This ability to recover structure from motion is not a trivial one since, generally, an infinite number of three-dimensional shapes can be given to a two-dimensional image. This structure-from-motion problem can be solved with computer algorithms by introducing certain reasonable constraints, such as the assumption that objects tend to be rigid or continuous. How the brain solves the problem, of course, is another matter. Recent psychophysical work from our laboratory suggests that one aspect of the neural mechanism for computing structure-from-motion is to reconstruct surfaces by interpolating motion cues. We have also extended this idea of surface interpolation to computational algorithms for structure-from-motion analysis.

Psychophysical studies. A task we developed for examining structure-from-motion perception is illustrated in Figure 2. The subjects are to indicate when they see a transition from an unstructured motion field to a motion field that gives the perception of a revolving, hollow cylinder. The time at which the transition occurs is random and unpredicted, and the subjects must release a key as quickly as they can once they perceive the change. To allow for a smooth transition from no structure to structure, the points are on for only a short period of time before being replotted at a new, random location on the screen. When the change in structure occurs, only new points assume the new structure, thus preventing any artifactual cues that would occur if the dots changed speed in midtrajectory. These finite point lifetimes also control for many other potential artifacts such as density and shape cues (see below and Siegel and Andersen 1988, 1990; Andersen and Siegel 1990).

As is illustrated in Figure 3, we found the reaction times in this task to be quite long, usually extending over several point lifetimes (Siegel and Andersen 1988, 1990; Husain et al. 1989; Andersen and Siegel 1990; Treue et al. 1990). This result suggested that the percept of structure from motion builds up over time and that new points were integrated into a representation partially computed from the old points. It is not possible, however, to determine from these data how much of the reaction time is needed as visual input and how much is computation time in the brain or motor reaction time. The surface interpolation hypothesis would predict that more visual input time would improve performance.

To test this idea directly, we presented equal numbers of structured and unstructured stimuli of 40–1700 msec duration in random order and asked subjects to indicate in a two-alternative forced-choice paradigm whether they saw a rotating cylinder or an unstructured noise pattern (Husain et al. 1989; Treue et al. 1990). Figure 4 (closed symbols) shows that performance peaked only after viewing stimuli for more than eight point lifetimes, a result consistent with the surface interpolation idea.

To control for the possibility that the buildup in performance is not due to presentation of new points but to some other process, we performed another experiment (Husain et al. 1989; Treue et al. 1990). After the first point lifetime, the dots were replotted at the beginning of the old trajectory instead of some new, random location. This procedure was repeated in all subsequent point lifetimes, resulting in the repeated presentation of the same motion data that was provided in the first point lifetime. Figure 4 (open symbols) shows that under these conditions, where no new information is supplied, the performance does not build up and remains at chance levels.

Another example supporting the use of surface interpolation comes from experiments in which individual dots in the cylinder were oscillated randomly back and forth along their horizontal trajectory (Treue and Andersen 1990). If the motion system were to follow the individual features of the stimulus, there would be no coherent interpretation of these motions. However,

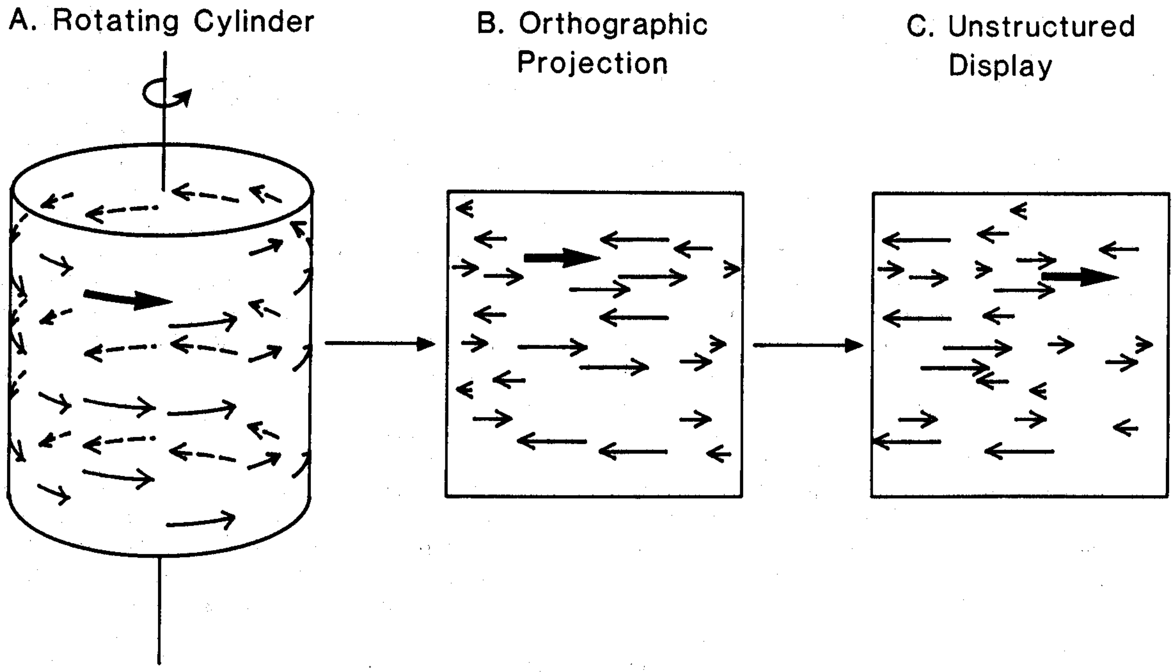

Figure 2. Demonstration of structure-from-motion reaction time task. The velocity of each point is indicated by the length of the arrows. (*A*) The stimuli are generated by first computing the location of points on the surface of the rotating cylinder for each instant in time. (*B*) The points are then parallel (orthographically) projected onto a plane perpendicular to the observer's line of sight. This is the "structured display." (*C*) The unstructured display is computed by taking each point's motion trajectories from the test display and displacing them randomly in a window equal to the width of the display. The fate of an individual motion trajectory for a point crossing the front of the display is shown by the bold arrow. The point density on the surface of the display is kept constant; each point is displayed for a finite amount of time. At the end of this time, the point vanishes and reappears randomly on the screen to follow a new motion trajectory. (Reprinted, with permission, from Siegel and Andersen 1988.)

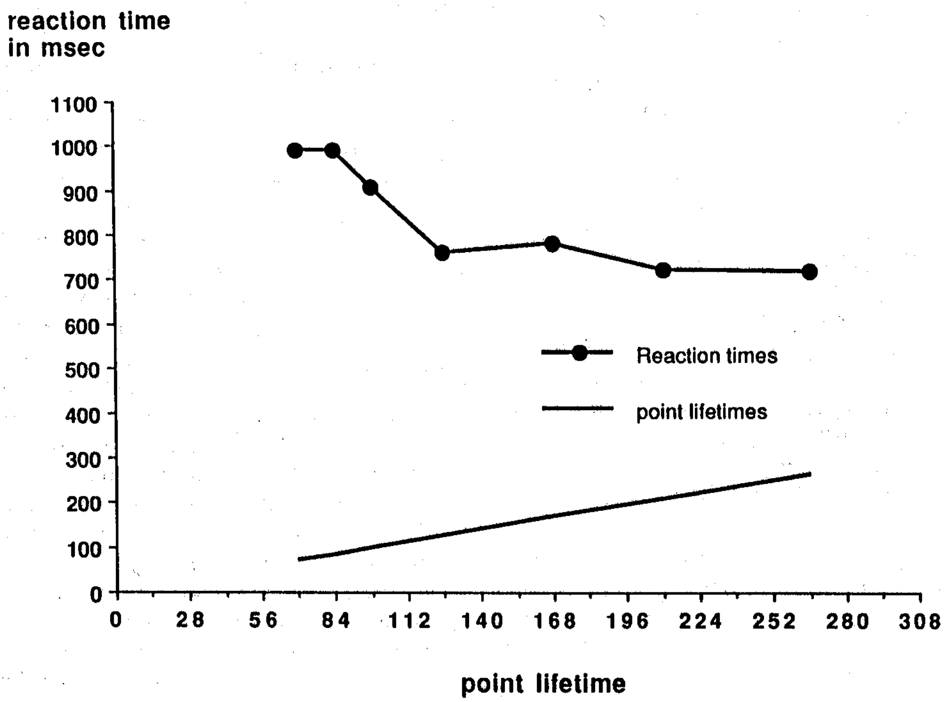

Figure 3. Reaction time as a function of point lifetime. Point lifetime is also plotted to allow easy comparison with reaction time. Note that the reaction time is always many times longer than the point lifetime. (Reprinted, with permission, from Treue et al. 1990.)

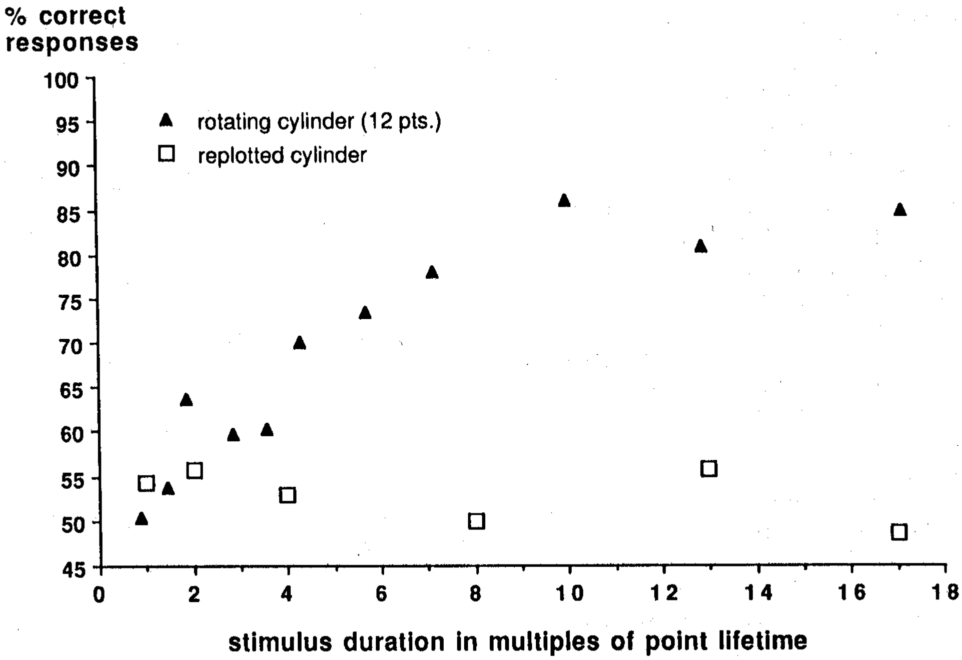

Figure 4. Percentage of correct responses in two alternative forced-choice tasks plotted as a function of stimulus duration. Note the long build up of performance. (□) Percentage of correct responses using 12 points; (▲) percentage of correct responses when points were repeatedly traveling along the same path. (Reprinted, with permission, from Treue et al. 1990.)

subjects simply see a revolving hollow cylinder. This percept can only be easily explained if the structure-from-motion process represents the observed object as two surfaces moving in opposite directions, rather than a collection of individual elements. A given point contributes information to each surface, depending on its direction of motion.

The neural basis for this segmentation of surfaces is suggested by the transparency experiments outlined above. The direction-selective cells in V1 appear to be direction-selective filters that to a large degree ignore motion in the nonpreferred directions. Our stimulus would activate two distinct populations of V1 cells, with each population corresponding to one of the surfaces. As the individual dots oscillated back and forth, they would alternately activate cells in one group and then the other.

Computational studies. Previous computational algorithms developed for analyzing structure-from-motion do not show buildup using short-lived data points. This observation has relevance not only for the biological plausibility of these models, but also for their general performance as image analyzers. In natural situations of object or observer motion, there is often a great deal of occlusion and disocclusion, which is more or less equivalent to the situation of finite point lifetimes. In addition, there is the computationally expensive problem of tracking many individual features over many frames.

We have recently developed a model that combines an extension of Ullman's (1984) incremental rigidity scheme for computing structure-from-motion with a surface interpolation process that reconstructs full three-dimensional surfaces using sparse motion data with finite point lifetimes (Ando et al. 1990). This new algorithm can operate on either displacements or velocities of moving points in an image. Using sequential inputs over time, the algorithm builds up a representation of three-dimensional structure. Each new estimate of the structure is derived by minimizing the overall deviation in structure that is consistent with the present and previous data. As the three-dimensional structure is derived at the locations of the moving points, a smooth surface is filled in between these locations. This constructed surface is saved between views, and individual image points can therefore appear and disappear without affecting the recovery of structure.

This new model has several interesting features from a biological perspective. It initially segregates surfaces based on their two-dimensional image motion before constructing the three-dimensional image. This step is necessary because without it the interpolated surface would average together the back and front surfaces of the cylinder and as a result would average the depth to zero. This "front end" segregation is reminiscent of the segregation in V1 by direction that was found in the motion transparency experiments. This model can also account for the psychophysical data, outlined above, using limited point lifetimes. A similar degradation in performance with fewer points is found in both our model and our psychophysical experiments (Treue et al. 1990), and similar interactions are seen between multiple transparent surfaces in motion and interactions with boundary shape as reported in human experiments by Ramachandran et al. (1988).

Response of Area MST Neurons to Motion Patterns

Several previous studies have reported neurons in area MST that were selective for rotating or expanding or contracting visual stimuli (Sakata et al. 1985, 1986; Saito et al. 1986; Tanaka et al. 1986, 1989; Tanaka and Saito 1989). Examples of these types of stimuli are shown in Figure 5. The stimuli used in these investigations had several potential artifacts. In cases where textured stimuli were used, there was always the potential that the neurons were responding to shape cues in the stimuli and not to their motion. The expansion and contraction stimuli were often 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 was not possible to equate the motion speed distributions when testing neurons with these two classes of stimuli. These complications were appreciated by several investigators, and they initiated controls that indicated that it was likely that these neurons were in fact selective to the pattern of motion, rather than the confounding potential artifacts listed above (Tanaka et al. 1989; Tanaka and Saito 1989). However, their presence in the test stimuli have made it difficult to extend these studies due to the linkage of these potential variables with the experimental variables one would like to investigate.

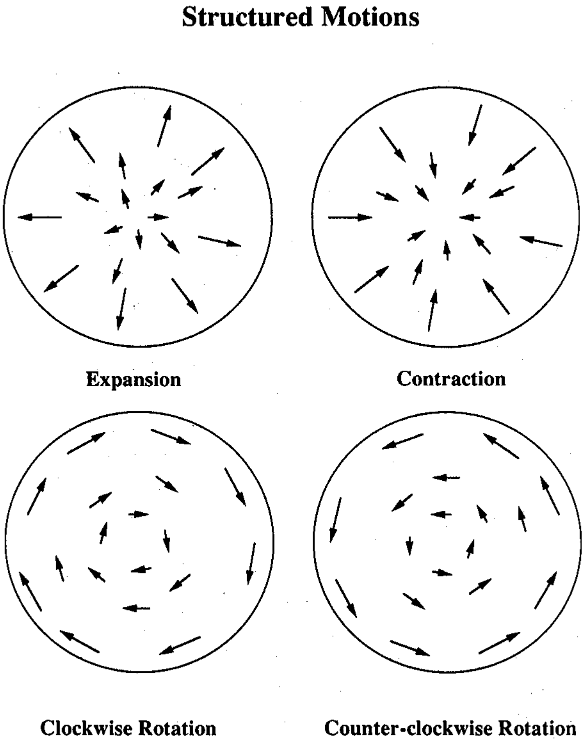
Structured Motions

Expansion

Contraction

Clockwise Rotation

Counter-clockwise Rotation

Figure 5. Illustration of expansion, contraction, and rotation stimuli used in recording experiments in area MST by Graziano et al. (1990) and Andersen et al. (1990).

To circumvent this problem, we have recently designed special rotation and expansion stimuli, generated on a high-speed graphics system, which contain only motion-pattern cues (Andersen et al. 1990; Graziano et al. 1990). These 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 thus there are no curvature or acceleration cues. The speed distributions for expansion, contraction, and rotation are equivalent. In short, the only difference between the displays are the local directions of the movement. For 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 et al. 1990). Some cells responded to different combinations of these stimuli, whereas others only responded to one type of movement. Some rotation cells were found to respond to both clockwise and counterclockwise movements.

Using these stimuli, we have been able to examine this motion-pattern selectivity in more detail. If these cells are truly selective for pattern, then they should show position invariance; i.e., the cell should show the same stimulus selectivity regardless of the position the stimulus occupies in the cell's receptive field. Area MST neurons have very large receptive fields, often exceeding 40° in diameter. We tested for position invariance by positioning 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 et al. 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.

Another parameter we examined was the effect of attention on the response of area MST neurons (Andersen et al. 1990). Attentional modulation was examined using a reaction time task in which the animal was required to release a key in response either to a change in the motion of the stimulus or to a small dimming of the fixation point. The two trial types were randomly interleaved, and the color of the fixation point at the beginning of each trial indicated to the animal which detection was required in a particular trial. Many area MST cells altered their response when the animal attended to the motion cue with a facilitated response to the preferred direction of motion, and often a reduced response to the nonpreferred direction of motion. This result suggests that in many cases, the tuning of area MST neurons is sharpened by attention.

CONCLUSIONS

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, signaling motion in the preferred direction, with the nonpreferred direction having little or no effect. On the other hand, for area MT neurons, the nonpreferred directions powerfully inhibit the response to preferred direction stimuli. We are currently studying the role this inhibitory interaction may play in the spatial integration of motion in area MT.

Our psychophysical experiments suggest that an important step in the analysis of structure-from-motion perception involves surface reconstruction. We are now investigating the role areas MT and MST might play in this surface reconstruction process.

Finally, area MST may play a higher-level role in the processing of motion information with cells selective for particular patterns of motion. The responses of these neurons show pattern position invariance and are also strongly modulated by attention. We are currently asking whether rotation and expansion/contraction are the primary channels used to analyze the structure of the velocity field. An alternative possibility is that there may be a wider range of patterns for which area MST neurons are selective. Thus, in the one case, structure in the velocity field would be decomposed into two basic channels; in the other case, each pattern of motion might activate a large array of different pattern cells by varying amounts.

ACKNOWLEDGMENTS

This work was supported by National Institutes of Health grant EY-07492. 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. and J.R. Bergen. 1985. Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am.* **2:** 284.

Andersen, R.A. and R.M. Siegel. 1990. Motion processing in primate cortex. In *Signal and sense: Local and global order in perceptual maps* (ed. G. Edelman et al.). Wiley, New York. (In press.)

Andersen, R.A., M. Graziano, and R.J. Snowden. 1990. Translational invariance and attentional modulation of MST cells. *Soc. Neurosci. Abstr.* (in press).

Ando, H., E.C. Hildreth, S. Treue, and R.A. Andersen. 1990. Recovering 3-D structure from motion with surface reconstruction. *Soc. Neurosci. Abstr.* (in press).

Barlow, H.B. and W.R. Levick. 1965. The mechanism of directionally selective units in the rabbit retina. *J. Physiol.* **178:** 477.

Bülthoff, H., J. Little, and T. Poggio. 1989. A parallel algorithm for real-time computation of optical flow. *Nature* **337:** 549.

Douglas, R.J., K.A.C. Martin, and D. Whitteridge. 1988. Selective responses of visual cortical cells do not depend on shunting inhibition. *Nature* **332:** 642.

Egelhaaf, M., A. Borst, and B. Pilz. 1990. The role of GABA in detecting visual motion. *Brain Res.* **509:** 156.

Erickson, R.G., R.J. Snowden, R.A. Andersen, and S. Treue. 1989. Directional neurons in awake rhesus monkeys: Implications for motion transparency. *Soc. Neurosci. Abstr.* **15:** 323.

Goodwin, A.W. and G.H. Henry. 1975. Direction selectivity of complex cells in a comparison with simple cells. *J. Neurophysiol.* **38:** 1524.

Graziano, M., R.A. Andersen, and R.J. Snowden. 1990. Stimulus selectivity of neurons in macaque MST. *Soc. Neurosci. Abstr.* (in press).

Husain, M., S. Treue, and R.A. Andersen. 1989. Surface interpolation in three-dimensional structure-from-motion perception. *Neural Computat.* **1:** 324.

Koch, C., T. Poggio, and V. Torre. 1986. Computations in the vertebrate retina: Gain enhancement, differentiation and motion discrimination. *Trends Neurosci.* **9:** 204.

Movshon, J.A., I.D. Thompson, and D.J. Tolhurst. 1978. Receptive field organization of complex cells in the cat's striate cortex. *J. Physiol.* **283:** 79.

Newsome, W.T. and E.B. Paré. 1988. A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J. Neurosci.* **8:** 2201.

Newsome, W.T., K.H. Britten, and I.A. Movshon. 1989. Neuronal correlates of a perceptual decision. *Nature* **341:** 52.

Ramachandran, V.S., S. Cobb, and D. Rogers-Ramachandran. 1988. Perception of 3-D structure from motion: The role of velocity gradients and segmentation boundaries. *Percept. Psychophys.* **44:** 390.

Reichardt, W. 1961. Autocorrelation: A principle of the evaluation of sensory information by the central nervous system. In *Sensory communication* (ed. W.A. Rosenblith), p. 303. Wiley, New York.

Rodman, H.R. and T.D. Albright. 1989. Afferent basis of visual response properties in area MT of the macaque. I. Effects of striate cortex removal. *J. Neurosci.* **9:** 2033.

Saito, H., M. Yukio, K. Tanaka, K. Hikosaka, Y. Fukada, and E. Iwai. 1986. Integration of direction signals of image motion in the superior temporal sulcus of the macaque monkey. *J. Neurosci.* **6:** 145.

Sakata, H., H. Shibutani, Y. Ito, and K. Tsurugai. 1986. Parietal cortical neurons responding to rotary movement of visual stimuli in space. *Exp. Brain Res.* **61:** 658.

Sakata, H., H. Shibutani, K. Kawano, and T. Harrington. 1985. Neural mechanisms of space vision in the parietal association cortex of the monkey. *Vision Res.* **25:** 453.

Siegel, R.M. and R.A. Andersen. 1986. Motion perceptual deficits following ibotenic acid lesions of the middle temporal area in the behaving rhesus monkey. *Soc. Neurosci. Abstr.* **12:** 1183.

———. 1988. Perception of three-dimensional structure from motion in monkey and man. *Nature* **331:** 259.

———. 1990. The perception of structure from visual motion in monkey and man. *J. Cognit. Neurosci.* (in press).

Sillito, A.M. 1979. Pharmological approach to visual cortex. *Trends Neurosci.* **2:** 196.

Snowden, R.J. 1989. Motions in orthogonal directions are mutually suppressive. *J. Opt. Soc. Am.* **6:** 1096.

Snowden, R.J., R.G. Erickson, S. Treue, and R.A. Andersen. 1990. Transparent motion stimuli reveal divisive inhibition in area MT of macaque. *Invest. Opthalmol. Vis. Sci.* (suppl.) **31:** 399.

Tanaka, K. and H. Saito. 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. *J. Neurophysiol.* **62:** 626.

Tanaka, K., Y. Fukada, and H. Saito. 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. *J. Neurophysiol.* **62:** 642.

Tanaka, K., K. Hikosaka, H. Saito, M. Yukie, Y. Fukada, and E. Iwai. 1986. Analysis of local and wide-field movements in the superior temporal visual areas of the macaque monkey. *J. Neurosci.* **6:** 134.

Treue, S. and R.A. Andersen. 1990. 3-D structure from motion: Rigidity and surface interpolation. *Invest. Opthalmol. Vis. Sci.* (suppl.) **31:** 172.

Treue, S., M. Husain, and R.A. Andersen. 1990. Human perception of structure from motion. *Vision Res.* (in press).

Ullman, S. 1984. Maximizing rigidity: The incremental recovery of 3-D structure from rigid and nonrigid motion. *Perception* **13:** 255.

Wallach, H. and D.N. O'Connell. 1953. The kinetic depth effect. *J. Exp. Psychol.* **45:** 205.

Wang, H.T., B. Mather, and C. Koch. 1989. Computing optical flow in the primate visual system. *Neural Computat.* **1:** 92.