Two-Dimensional Substructure of MT Receptive Fields

Margaret S. Livingstone,[1] Christopher C. Pack,
and Richard T. Born
Department of Neurobiology
Harvard Medical School
220 Longwood Avenue
Boston, Massachusetts 02115

Summary

Neurons at progressively higher levels of the visual system have progressively larger, more complicated receptive fields, presumably constructed from simpler antecedent receptive fields. To study this hierarchical organization, we used sparse white noise to map receptive-field substructure (second order Wiener-like kernels) in an extrastriate motion processing area (MT) of alert monkeys. The maps revealed a clear substructure, on a spatial scale comparable to the receptive fields of the V1 inputs. There were both facilitatory and suppressive interactions that differed in spatial organization and time course. Directional interactions were remarkably precise over a very small spatial range, and reversed when successive stimuli reversed contrast—a neural correlate of "reverse phi" motion perception. The maps of some cells had an unexpected, curved shape, which challenges existing models for direction selectivity.

Introduction

Classical studies of the physiology of the visual system gave rise to the concept of "hierarchical elaboration of receptive fields"—that is, that neurons with more complicated types of receptive fields arise through the combination of inputs from neurons with simpler receptive fields. If this idea is correct, these simpler inputs should be discernible as substructure of the more complex receptive field, as proposed, for example, in Hubel and Wiesel's model of the formation of simple cells from LGN inputs and of complex cells from simple-cell inputs (Hubel and Wiesel, 1962). While this is an attractively simple model, direct evidence for such substructure has been relatively rare, largely for technical reasons. Two-bar interaction experiments have revealed subunit organization in direction-selective retinal ganglion cells of the rabbit (Barlow and Levick, 1965) and in complex cells of cat striate cortex (Movshon et al., 1978); more direct evidence for the latter was shown using "white noise" and reverse correlation (Emerson et al., 1987; Szulborski and Palmer, 1990). Because these techniques require considerable time and very precise localization of visual stimuli on the retina, they have heretofore been used only in anesthetized preparations and at very early stages of the visual system. Here, we use a modification of sparse white noise mapping that reveals, for the first time, substructure in the receptive

fields of neurons in an extrastriate visual area of alert monkeys.

The middle temporal visual area (MT or V5) is an extrastriate visual area in which most cells are selective for the direction of stimulus motion (Dubner and Zeki, 1971). MT receives input from V1, and, at any given eccentricity, receptive fields of MT neurons are about ten times larger than those in V1. MT must therefore summate inputs from many V1 cells with receptive fields distributed across visual space. There are at least two ways in which the direction selectivity in MT neurons could arise. It could be generated de novo from nondirectional inputs; it could be inherited from direction-selective inputs; or both. Direct inheritance from V1 is very likely because at least some MT-projecting V1 cells are themselves direction selective (Movshon and Newsome, 1996), and because directional interactions in MT act over a shorter range than the receptive-field size (Mikami et al., 1986). It remains possible, however, that MT cells perform additional operation(s) on direction-selective V1 inputs to generate longer-range motion interactions.

We used one- and two-dimensional sparse white noise stimuli to map directional interactions within MT receptive fields. We will first describe the results of mapping directional interactions in one dimension (1D) using pairs of oriented bars presented at various positions along each neuron's axis of preferred motion. The two-dimensional (2D) experiments use small, unoriented, spot stimuli to map the average spatial structure of a cell's paired-stimulus interactions without making any assumptions about their orientation or shape.

Results

One-Dimensional Interactions

We recorded from 48 single units in area MT of two macaque monkeys. For each unit, one-dimensional (1D) directional interactions were mapped using a modification of the 1D reverse-correlation mapping technique (deBoer, 1968; Eggermont et al., 1983; Jones and Palmer, 1987) (Figure 1, top). Our modification of this technique is to correct for eye position at each stimulus presentation, allowing us to correlate responses with stimulus position on the retina, rather than stimulus position on the screen (Livingstone, 1998). We will refer to interaction plots generated using this kind of one-dimensional stimulus range as "1D" maps to distinguish them from maps generated using a two-dimensional stimulus range. We used a one-dimensional sparse white noise stimulus consisting of pairs of optimally oriented bars, one black, one white, flashed at 75 Hz along a stimulus range parallel to the cell's axis of preferred motion. Because each stimulus presentation contains a black and a white bar, each sequence of two stimulus frames generates four possible pairwise combinations: white to white, black to black, black to white, and white to black. Maps for all four sequences from the same spike trains are shown for two cells in the bottom half of Figure 1. The position of the second (reference) bar is mapped along the x axis

[1] Correspondence: mlivingstone@hms.harvard.edu
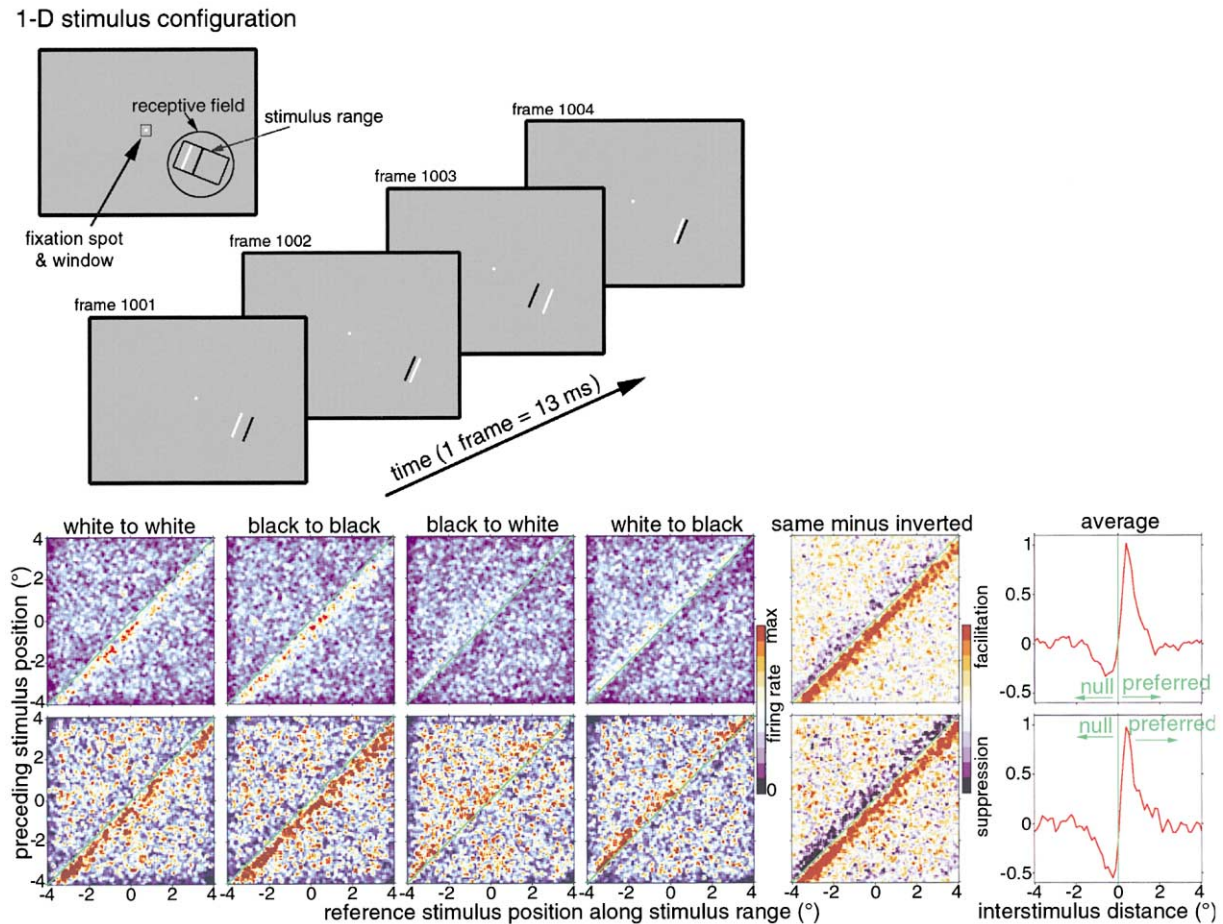
1-D stimulus configuration

Figure 1. One-Dimensional Directional Interactions for an MT Cell

(Top) Diagram of the stimulus configuration for 1D mapping. (Bottom) 1D interaction maps for two MT cells, receptive-field eccentricity 30° (top cell) and 19° (bottom cell). Pairs of 0.125° wide bars, one white, one black against an intermediate gray background were presented at 75 Hz at random positions along an 8°-long stimulus range. For each sequential-contrast condition, spikes were reverse correlated (at a delay of 45 to 55 ms) with each eye-position corrected bar position (plotted along the x axis) and the immediately preceding eye-position corrected bar position (plotted along the y axis). For both cells the positive direction on each axis corresponds to the preferred direction of stimulus motion. The green diagonal indicates occasions when the two bars appeared in precisely the same retinal location. Locations in stimulus space to the right of the green diagonal indicate preferred-direction sequences, and locations to the left of the diagonal indicate null-direction sequences. The first four panels show responses as a function of each bar-contrast condition. The fifth panels show the same-contrast conditions (white-to-white and black-to-black) minus the inverted-contrast conditions (black-to-white and white-to-black). The last panels show the average of the fifth panels collapsed along the green diagonal. Color scale for first four columns represents firing rate; color scale for difference maps ranges from maximal facilitatory interaction ($+1$) to maximal suppressive interaction (-1). The magnitude of the facilitation in the first (white to white) map for the top cell corresponds to an average peak firing rate that is 62% larger than the response to the two stimuli presented independently; the magnitude of the suppression is 12%. For the bottom cell, the facilitation in the first panel is 63% larger than the response to the two stimuli presented independently, and the suppression is 37%.

and the position of the immediately preceding stimulus is mapped along the y axis. On both axes the preferred direction of motion is from negative toward positive (i.e., right on the x axis and up on the y axis), with 0 corresponding to the center of the stimulus range. Thus, neural activity, as indicated by the color code, is mapped as a function of paired-stimulus position, not visual space. The $+45°$ diagonal green line indicates occasions when the reference stimulus and the immediately preceding stimulus appeared at exactly the same retinal location. Positions below/right of the green diagonal represent occasions when the stimulus sequence was in the preferred direction, and positions above/left indicate sequences in the opposite, or "null," direction.

The two cells illustrated in Figure 1 had receptive fields that were centered at 30° (upper cell) and 19° (lower cell) from the fovea, and were thus considerably larger in diameter than the 8° stimulus range. For the two *same-contrast* sequences (white to white and black to black) the two-bar interactions show narrow regions, lying below/right of the same-position diagonal, that indicate higher responsiveness than over the rest of the two-bar interaction space. This region corresponds to occasions when the two sequential bars were presented a small distance apart and in the preferred direction. This means that pairs of bars less than a degree apart presented in the preferred direction give larger responses than two bars presented in the opposite direction or farther apart.

We will refer to this type of interaction as "facilitation," without any implications as to the site of interaction or underlying synaptic mechanisms.

The fact that the region of enhanced responsiveness is elongated diagonally indicates that the directional interactions are uniform across the entire stimulus range; the narrowness of this band indicates that the interactions are quite local. Less obviously, there is also diminished activity to the left of the green diagonal, indicating that responses to local stimulus sequences in the null direction are smaller than if the two stimuli were presented farther apart or in the preferred direction. We will refer to this type of interaction as "suppression," again, without any implications as to the site of interaction or underlying synaptic mechanisms. Conversely, the *inverting-contrast* stimulus sequences (black to white and white to black) show facilitation in the null direction and suppression in the preferred direction. This reversal of the directional response with contrast-inverting stimuli has a perceptual correlate in humans: reversed motion (reverse phi) is seen for sequential contrast-inverting stimuli (Anstis and Rogers, 1975).

The first four maps for each cell in Figure 1 show responses to each possible two-stimulus combination. We can generate a single map that reflects only the two-stimulus *interactions* by subtracting the inverting-contrast maps from the same-contrast maps to get a *difference* map ("same minus inverted" column in Figure 1). Subtracting the inverting-contrast maps from the same-contrast maps gives us a map of only the directional *interactions,* i.e., only that part of the response that depends on the relative position of sequential bars, because all the independent (non-interacting) responses to each stimulus cancel out (see Experimental Procedures). This calculation is equivalent to the method of Emerson et al. (1987) for generating second order Wiener-like kernels.

In the *difference* maps (column 5 of Figure 1), same-contrast facilitation and inverting-contrast suppression are both depicted as positive interactions, and same-contrast suppression and inverting-contrast facilitation are mapped as negative. Interactions below the same-position diagonal all correspond to preferred-direction sequences, and interactions above correspond to null-direction sequences. Therefore preferred-direction interactions are positive, consisting of both same-contrast facilitation and inverting-contrast suppression, and null-direction interactions are negative, consisting of both same-contrast suppression and inverting-contrast facilitation. In these *difference* maps, both the positive preferred-direction interactions and the negative null-direction interactions are fairly homogeneous across the entire 8° stimulus range.

The graphs on the far right of Figure 1 show the difference interactions for each cell as a function of the distance between sequential stimuli, averaged across all positions along the stimulus range. These graphs reveal that the directional interactions are exceedingly precise. Even at these large eccentricities of 30° and 19°, the interactions show directionality at the smallest interstimulus distance we generated, one pixel (0.0625°), in either direction.

1D *difference* maps for interactions between successive stimuli presented every 13 ms are shown in the top row of Figure 2 for five other representative MT cells whose receptive fields were located at various eccentricities. All of the cells mapped showed similar characteristics: positive preferred-direction interactions and negative null-direction interactions that were relatively homogeneous across the entire stimulus range, were quite local, and were very precise spatially. These interactions were so precise that, on average, stimulus sequences only one pixel (0.065°) in the preferred direction gave responses twice the size of responses to sequences one pixel in the null direction.

To quantify the precision of the directional interactions, we took averages across the stimulus range for all 48 cells in our sample, as shown in the rightmost graphs of Figure 1. The scatter plot in Figure 3 shows the interstimulus distance giving the maximum positive preferred-direction interaction (x's) and the interstimulus distance giving the most negative null-direction interaction (filled circles) as a function of receptive-field eccentricity. For comparison the average V1 receptive-field diameter is shown as a solid line (Van Essen et al., 1984). By inspection, the optimal distance for direction-selective interactions is smaller than the average V1 receptive-field size, consistent with the idea that direction selectivity is generated within subunits of complex-cell receptive fields.

Both the interstimulus distance giving the most positive preferred-direction interaction and the interstimulus distance giving most negative null-direction interaction are small compared to the average V1 receptive-field diameter, and the latter interstimulus distances are, on average, smaller than the former. The interstimulus distance giving the most positive preferred-direction interaction is on average larger than the distance giving most negative null-direction interaction (one-tailed, paired t test, $p < 0.00004$; mean difference $\pm$ SEM = 0.12 $\pm$ 0.01°).

Short versus Long-Range Motion

Psychophysicists have described two (sometimes three) kinds of motion perception that are thought to represent hierarchical stages in motion processing (Braddick, 1974; Lu and Sperling, 1995). It is generally agreed that motion perception acting over short spatial and temporal ranges corresponds to the early stages of cortical processing. Short-range motion perception is characterized by reversal of perceived direction with inverting-contrast stimuli ("reverse phi") (Anstis and Rogers, 1975), failure to transfer between eyes (Braddick, 1974), and insensitivity to color and stimulus shape (Ramachandran and Gregory, 1978; Cavanagh et al., 1984; Lu and Sperling, 1995). The spatial, temporal, and contrast-dependent characteristics of short-range motion are consistent with the interactions revealed in Figure 1 and in the top row of Figure 2.

Conversely motion perception acting over longer spatial and temporal ranges is selective for stimulus color and shape, can transfer between eyes, and does not reverse with inverting-contrast stimuli (Braddick, 1974). It seems logical to suppose that short-range motion perception represents V1 direction selectivity and that long-range motion might be generated in extrastriate areas, like MT, with their larger, integrating receptive fields. We looked for longer-range, longer-lasting direc-
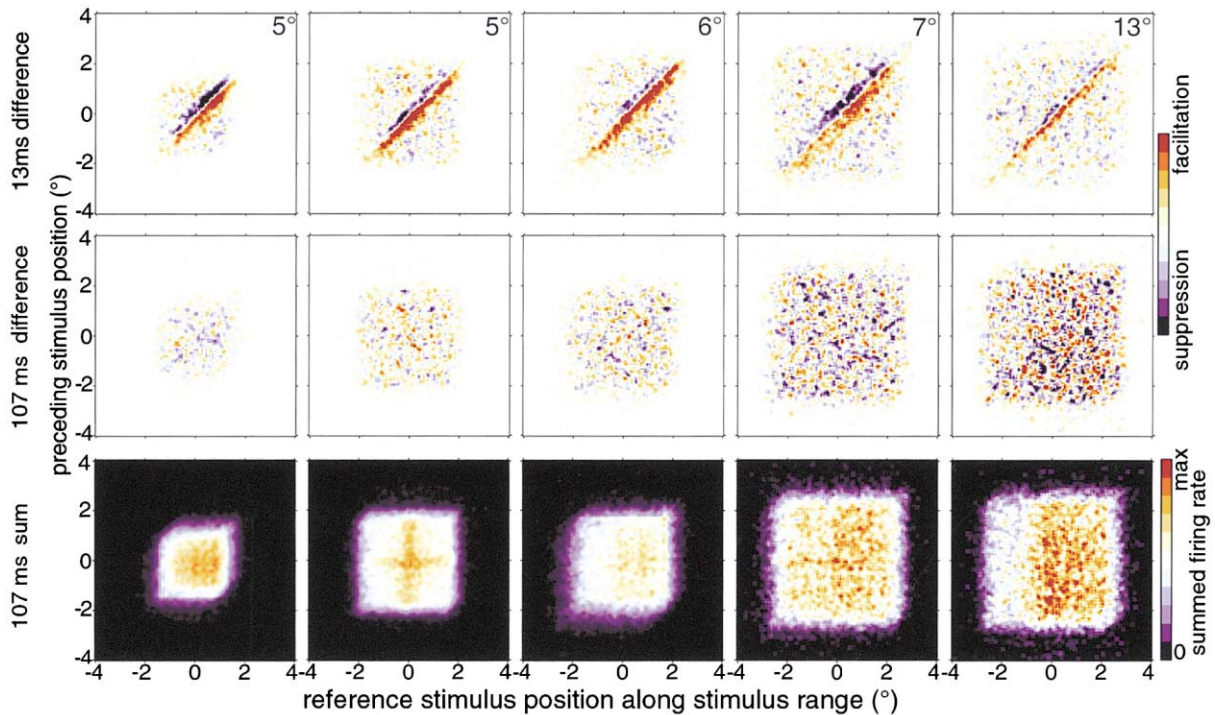
Figure 2. 1D Sequential Interaction Maps for Five Additional MT Cells

Each column represents one cell. Top row, 13 ms stimulus intervals, difference maps (same-contrast sequences minus inverted-contrast sequences). Middle row, 107 ms stimulus intervals, difference maps. Bottom row, 107 ms stimulus intervals, summation maps (same-contrast sequences plus inverted-contrast sequences). The eccentricity of each cell is indicated in the upper right corner of each top panel. The stimulus range was centered on each cell's receptive field; it was 3° long for the first cell, 4° for the second and third cells, and 5° for the fourth and fifth cells. For all cells the positive direction on each axis corresponds to the preferred direction of stimulus motion. Reverse correlation delay was between 45 and 55 ms. Color scale for upper and middle rows ranges from maximal facilitatory interaction (+1) to maximal suppressive interaction (−1); scale for bottom row represents summed firing rate. For each cell, the color scale for the top and middle maps represents the same percentage facilitation or suppression. For the first cell, the magnitude of the same-contrast facilitation was 32% larger than the response to the two stimuli presented independently; for the second cell the same-contrast facilitation was 45% larger than the response to the two stimuli presented independently; for the third cell it was 31% larger; for the fourth cell it was 42% larger; for the fifth cell it was 77% larger.

tional interactions using both short (13 ms) and long (up to 150 ms) interstimulus intervals, but we failed to see long-range or contrast-independent interactions in any of the cells we studied. In the second and third rows of Figure 2 we show examples, from the same five cells, of our attempts to look for longer-lasting and/or contrast-independent directional interactions. The second row of Figure 2 shows *difference* interactions between stimuli eight frames (107 ms) apart to look for longer-lasting, contrast-inverting interactions. These difference maps do not show any directionality. However, because the difference maps sum same-contrast sequences and subtract different-contrast sequences, any interaction that did not reverse with inverting-contrast sequences would cancel out in these difference maps. Because psychophysical experiments suggest that long-range motion perception does not reverse with inverting-contrast stimuli (Braddick, 1974), we therefore also looked for long-lasting interactions that did not reverse with inverting-contrast stimuli by *summing* same- and different-contrast sequences at 107 ms intervals. The third row in Figure 2 shows *summation* interactions between stimuli eight frames (107 ms) apart to look for interactions that did not reverse with inverting contrast. Again,

there is no evidence of directionality. And there was no evidence for directional interactions in any of the individual same- or different-contrast maps at 107 ms intervals.

The first four columns of rows two and three of Figure 2 show interactions when intervening stimuli were present (i.e., the same spike trains were used as for generating the upper maps), and the last column was derived from a stimulus sequence when stimuli were presented only every eighth frame (with background/blank intervening frames). In these, as well as in other cells in which we used both kinds of stimulus sequences, we saw no evidence for contrast-dependent or contrast-independent directional interactions at intervals between 100 and 150 ms. Long-lasting directional interactions should be manifest as some kind of difference across the same-position diagonal, as in the short-duration interactions in the top row, or some heterogeneity in the maps favoring one side of the same-position diagonal. The vertical striping in some of the maps reflects inhomogeneities in the receptive field, and not directional interactions. That is, there is relatively higher firing whenever a stimulus appears in one part of the receptive field, compared to the rest. Because the receptive field is bigger than
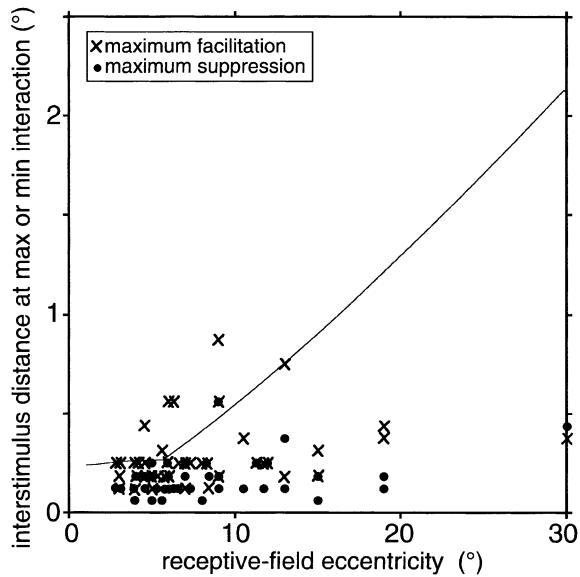
Figure 3. Interstimulus Distance Giving Maximum Facilitatory Inter-actions (x's) and Maximum Suppressive Interactions (Filled Circles) as a Function of MT Receptive-Field Eccentricity

Data were derived from graphs generated as in last panels of Figure 1 for 48 MT cells. Solid line shows average V1 receptive-field diameters (from Van Essen et al., 1984).

the stimulus range, the striping means that one part of the receptive field is more responsive than nearby regions.

We explore interactions as a function of interstimulus interval in more detail in the following section.

Two-Dimensional Interactions

The 1D maps show that MT cells' directional interactions act across distances much shorter than their receptive-field diameter, in confirmation of previous results (Mikami et al., 1986), and consistent with the idea that MT cells inherit their direction selectivity directly from V1 inputs (Movshon and Newsome, 1996). If the directional interactions in MT are indeed largely derived from V1, then directional interactions should be short-range not only in the dimension parallel to the preferred-motion axis, as demonstrated, but also perpendicular to it.

To explore the 2D structure of the directional interactions in MT, we mapped 28 cells using a two-dimensional sparse noise stimulus (Szulborski and Palmer, 1990) consisting of pairs of small squares (0.19° or 0.25° across), one black, one white, on a gray background, flashed at 75 Hz at random positions within a 2.5° square stimulus range (Figure 4, top). Spikes were then reverse correlated with the *relative position* in 2D visual space of pairs of sequential stimuli. Maps from five cells are illustrated in Figure 4; each row represents one cell. The panels in the lower part of Figure 4 show 2D interaction maps for each contrast combination and *difference* maps (same minus inverted sequences) of average spike activity as a function of the position of one stimulus relative to an immediately subsequent (13 ms later) reference stimulus in visual space. Each contrast stimulus (black or white) of each stimulus presentation is consid-

ered as one reference stimulus, and the position of this stimulus, even though it could have been anywhere within the stimulus range, is assigned to position (0,0) on the interaction map. Neural activity is then plotted as a function of the position of the immediately preceding stimulus (of each contrast) minus the reference position, with a baseline linear map subtracted (see Experimental Procedures). The maps are oriented so that activity is plotted as a function of the position (in visual space) of the first stimulus minus the second (reference) stimulus. For all five cells, maps for each sequential contrast combination are shown in the first four panels, and in the fifth panel are shown *difference* maps (same contrast maps minus inverting-contrast maps).

For all of the cells mapped, the interactions were local in two dimensions, and were consistent with the cell's preferred direction of motion (indicated by arrows in fifth panels). For example, the first cell illustrated in Figure 4 preferred rightward motion, and the 2D interaction maps show facilitation of a reference stimulus response whenever an immediately preceding same-contrast stimulus was anywhere within a small, slightly elongated region to the left of the reference stimulus; the cell's activity was suppressed when the preceding same-contrast stimulus was to the right of the reference stimulus. The reverse was true, though to a lesser degree, for inverting-contrast sequences. The distance from the reference stimulus, both perpendicular to and parallel to the preferred-direction axis, within which a preceding stimulus can influence the reference-stimulus response, was only a few tenths of a degree, even though the receptive field was roughly 10° in diameter. Similarly, the other cells shown in Figure 4 were typical in showing facilitation for same-contrast sequences in the preferred direction, within a small, elongated region, suppression for same-contrast sequences in the null direction, and a reversal of the directional interactions for inverting-contrast sequences.

Crescent-Shaped 2D Interactions

The most unexpected finding from these 2D mapping studies is that about half the cells were like those shown in the first three rows of Figure 4, with *difference* maps in which the preferred-direction positive interaction region showed an inflection around the reference position, in some cases forming a crescent curving toward the more localized null-direction negative region. This curvature is clear in many cells, and when it is clear, the curvature is always with the concavity facing the preferred direction. This finding may bear on the much debated and much studied question of how direction selectivity and orientation selectivity are integrated (Movshon et al., 1985). As far as we know, no models for direction selectivity predict such a crescent shape, so we will examine this phenomenon and possible explanations for it in some detail.

Though the ultimate explanation for the asymmetry in the 2D difference maps certainly could lie at the level of V1 receptive fields, we can still ask what aspects of the MT 2D difference maps correlate with the crescent shape, and these analyses may be revealing because they probably reflect the average behavior of V1 directional cells. To simply describe the crescent shapes, and
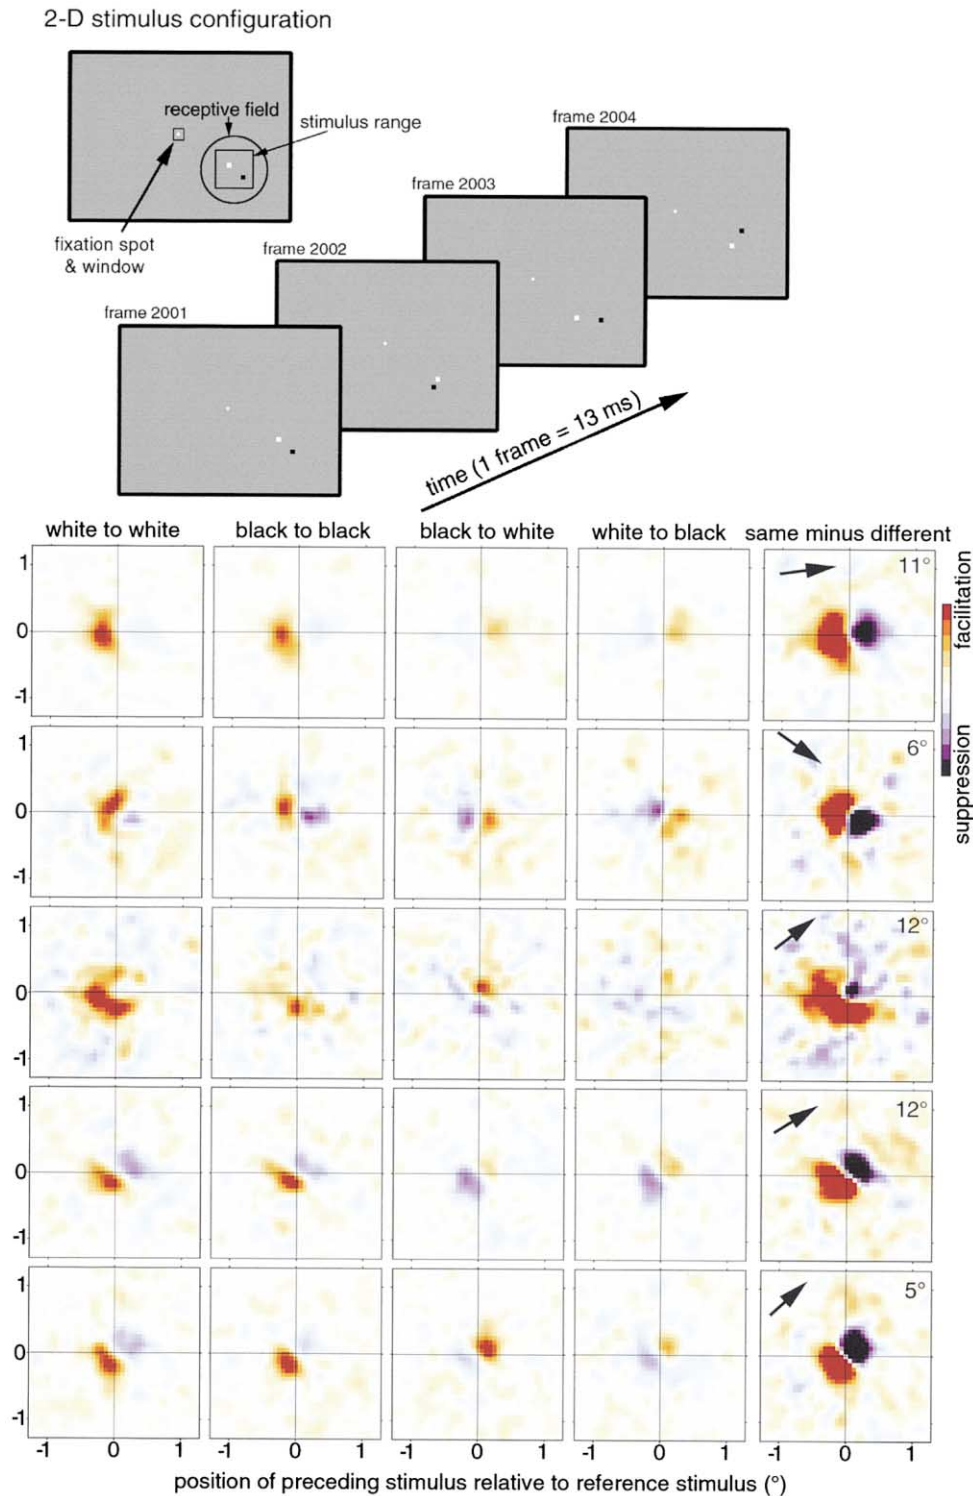
Figure 4. Two-Dimensional Directional Interactions for Five MT Cells

(Top) Diagram of the stimulus configuration for 2D mapping. (Bottom) 2D interactions for each sequential contrast combination (first four columns) and *difference* maps (last column) for five MT cells. Each row represents one cell. For each cell, the preferred direction of motion (arrows) and receptive-field eccentricity are indicated on the last panel. Color scale ranges from maximal facilitatory interaction (+1) to maximal suppressive interaction (−1) and is consistent within each cell. For the first cell, the magnitude of the same-contrast facilitation was 85% larger than the response to the two stimuli presented independently; for the second cell the same-contrast facilitation was 49% larger than the response to the two stimuli presented independently; for the third cell it was 58% larger; for the fourth cell it was 41% larger; for the fifth cell it was 57% larger. For the first four maps for each cell (the individual paired-contrast maps) a baseline first order map was subtracted (see Experimental Procedures).
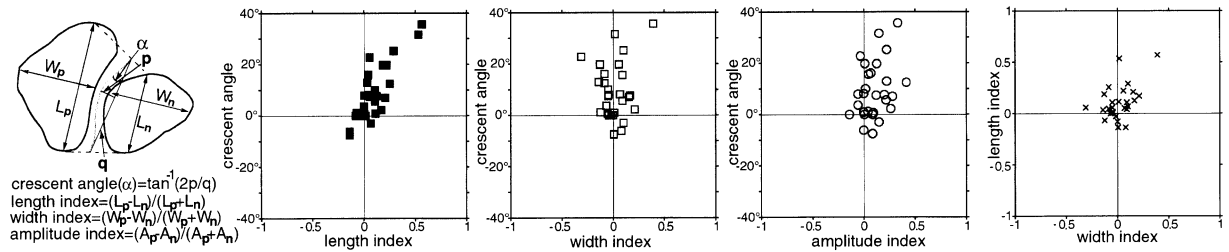
Figure 5. Scatter Plots of Length, Width, Amplitude, and Crescent Angle for 2D Maps from 28 Cells

(First panel) Diagram of how parameters of 2D interaction maps were determined. Solid black contours are $+20\%$ and -20% contours of the *difference* map of the second cell in Figure 4. L_p and L_n are measured perpendicular to the axis of preferred motion, and W_p and W_n perpendicular to that. The gray contour is the zero contour (between positive interactions and negative interactions) terminated at the lines joining the ends of L_p and L_n (panels 2–5).

to explore possible explanations for them, we measured various parameters in the 2D difference maps, as diagramed in Figure 5 (first panel). For each 2D difference map we drew contours at $+20\%$ and -20% of the peak positive interaction or peak negative interaction (whichever was larger in magnitude) and at zero. From the contour around the positive region ($+20\%$ contour), we calculated the longest dimension perpendicular to the cell's preferred motion axis, L_p, and the longest dimension perpendicular to that, W_p. The length and width of the negative region, L_n and W_n, were similarly calculated from the -20% contour. To quantify the degree of curvature of the border between positive and negative interactions, we calculated a "crescent angle." This angle was evaluated from the zero contour (the boundary between positive and negative interactions, indicated in gray) lying between lines connecting the ends of L_p and L_n (dotted lines). The concavity of this contour was estimated as the $\tan^{-1}(2p/q)$ where q is the distance between the ends of the zero contour and p is the longest perpendicular from q to the zero contour. This measurement represents the average of the angles (α) the zero contour forms with a straight line (q). Positive angles indicate curvature concave toward the negative (null-direction) region and negative angles the reverse. The distribution of crescent angles for the 28 cells in our population is shown as the vertical scatter of points in the first three graphs in Figure 5 (as a function of other metrics described later). The distribution is clearly skewed toward the positive (mean $= 9°$; this is significantly greater than zero; one-tailed t test $p < 0.0002$), corresponding to our impression that the preferred-direction positive-interaction region is commonly crescent shaped.

It would seem that a crescent-shaped positive-interaction region, or a curved boundary between preferred-direction and null-direction interactions, implies an asymmetry either between positive and negative interactions or between preferred-direction and null-direction interactions. If the difference is between preferred-direction and null-direction interactions, it could imply either that the preferred-direction-interaction region was actually crescent shaped (i.e., more broadly tuned than the null-direction-interaction region) or that it was simply longer. On the other hand, the difference could be between positive and negative interactions, arising from the fact that facilitation might be larger in magnitude than inhibition because of rectification, but for this

latter explanation to make positive interactions larger than negative, there would also have to be a difference between same-contrast interactions and different.

To clarify this issue, we asked, for the 28 cells for which we had 2D interaction maps, whether the crescent angle correlated with the relative lengths of the positive and negative-interaction regions, with their relative widths, or with their relative magnitudes. To do this, we calculated a length index to reflect the relationship between the length of the positive-interaction region and the length of the negative-interaction region (Figure 5, first panel). This value could vary between -1 and 1, with a value of 0 when the two interaction regions had the same length, and positive values when the positive-interaction region was longer than the negative-interaction region. Similarly, a width index indicated whether the positive-interaction region was wider than the negative. The amplitude index reflected the relative magnitudes (A) of the peak positive and negative interactions.

The second panel in Figure 5 shows the crescent angle plotted as a function of the length index. First, it is clear from the scatter of the crescent angles alone that the curvature of the boundary between the positive region and the negative region *does* tend to curve with the concave side toward the null-direction-negative-interaction region, that is, the angles are mostly positive. The second panel also shows that the crescent angle does significantly correlate with the length index (linear regression, $r^2 = 0.71$, $p < 0.00003$, F test). The third panel of Figure 5 shows that the width index does not show any correlation with the crescent angle (linear regression, $r^2 = 0.022$, $p < 0.46$ F test), and the fourth panel shows that the amplitude index is weakly correlated with the crescent angle (linear regression, $r^2 = 0.13$, $p < 0.053$ F test).

The fact that the crescent angle is correlated with the length index but not the width index suggests that there is something special and independent about the relative lengths of the preferred-direction and null-direction-interaction regions. We cannot rule out the possibility that the longer positive interactions are simply due to their having a larger magnitude than the negative interactions, though if the positive-interaction regions were relatively longer than the negative simply because they were larger in magnitude, we would expect the widths to also be larger. The last panel of Figure 5 shows that the length index tends to be larger than 0, but the width
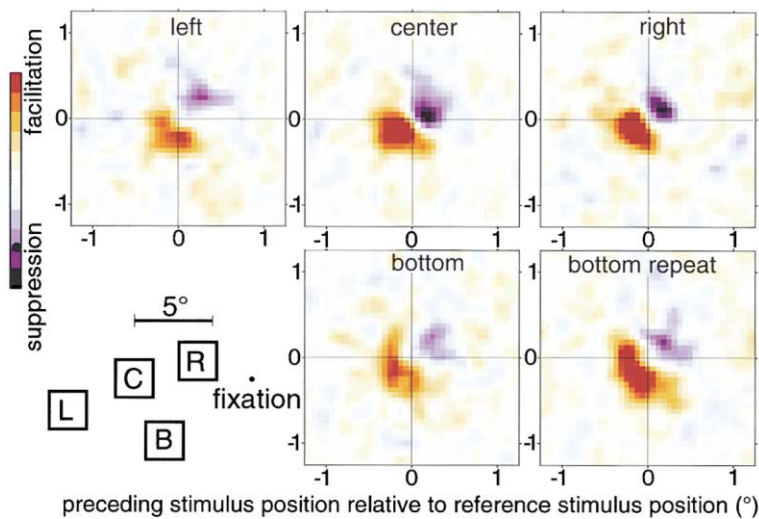
Figure 6. 2D Interaction Maps at Various Positions within the Receptive Field for One MT Cell (Same Cell as the Fourth Cell in Figure 7)

Interaction maps were generated by reverse correlating spike trains (at a delay of 50 to 60 ms) with the difference in position of two sequential stimuli. The stimuli were pairs of small squares 0.19° across presented at random positions at 75 Hz within a 2.5° stimulus range. The maps represent same-contrast stimulus sequences minus inverted-contrast stimulus sequences. This cell's receptive field was approximately 10° in diameter, centered 8° to the left of the fovea. The four mapping locations are indicated in the lower left of the figure with respect to the fovea. Maps for each of the locations are shown. The bottom location (B) was mapped twice, once at the beginning of the series and again at the end, an hour later. Color scale ranges from maximal facilitatory interaction (+1) to maximal suppressive interaction (−1) and the magnitude of the scale is the same for all panels. The magnitude of the same-contrast facilitation was 53% larger than the response to the two stimuli presented independently for the center location.

index does not, and the two are not correlated. The mean length index is 0.1011, which is significantly different from zero (one-tailed t test p < 0.0017), but the width index is 0.012, which is not significantly different from zero (two-tailed t test, p < 0.6394). That is, the preferred-direction positive-interaction regions are on average longer than the null-direction negative-interaction regions, but not wider. This observed difference in length of the two regions could explain the crescent shape; our results are less consistent with a simple difference in magnitude because the correlation is weaker, and a difference in magnitude should affect the relative width as well as the length.

2D Interactions in Different Parts of the Receptive Field

To determine the consistency of these directional interactions across individual receptive fields, we examined the receptive-field 2D substructure at various nonoverlapping subregions within the larger receptive fields of 5 MT cells. In the representative example shown in Figure 6, the spatial organization of the positive and negative interactions of the substructure is roughly similar at all four locations and correlates with the overall directionality of the MT cell, which was up and to the right. We did not see any striking or reproducible differences between different parts of the receptive field in our sample of MT cells. The maps were also reproducible over time, as shown by repeating the map for the same subregion after an interval of 1 hr.

Time Course of 2D Interactions

We looked at the time course of the 2D directional interactions by reverse correlating the spike train with the difference in position between each reference stimulus and preceding stimuli at different temporal intervals; i.e., with the immediately preceding stimulus, the stimulus

before that, etc.. Figure 7 shows the 2D interaction *difference* maps at various interstimulus intervals for five representative MT cells (each row represents one cell). For each cell, the first column shows a polar plot of the direction preference determined using a field of moving dots. By inspection, this direction preference is most strongly reflected in the spatial arrangement of positive and negative interactions at the smallest interstimulus interval tested (13 ms).

Figure 7 further shows that the receptive-field substructure organization shifts spatially with different interstimulus intervals, but the shift is not linear with time. Interactions between simultaneously presented stimuli show weak positive interactions along a region centered on the reference stimulus. (Note that because only one dark and one light stimulus were present in each frame, at time = 0 only opposite-contrast interactions exist for these maps.) At 13 ms intervals, regions of positive and negative interactions abruptly become much stronger and sharply segregated to either side of the reference location. Over the next 40 ms the positive interaction fades away faster than the negative interaction, and the negative interaction moves toward the reference location. In some cases, row three for example, the negative interaction eventually moves to occupy the location previously occupied by the initial positive interaction, perhaps representing a rebound effect. Thus directional interactions comprise a spatially and temporally complex interplay of facilitation and suppression. There is a caveat to the interpretation of the changing interaction magnitude over time, because for the longer intervals there are increasing numbers of intervening stimuli falling in the cells' receptive field.

Discussion

Our results show that MT directional interactions act only over very short distances both perpendicular and
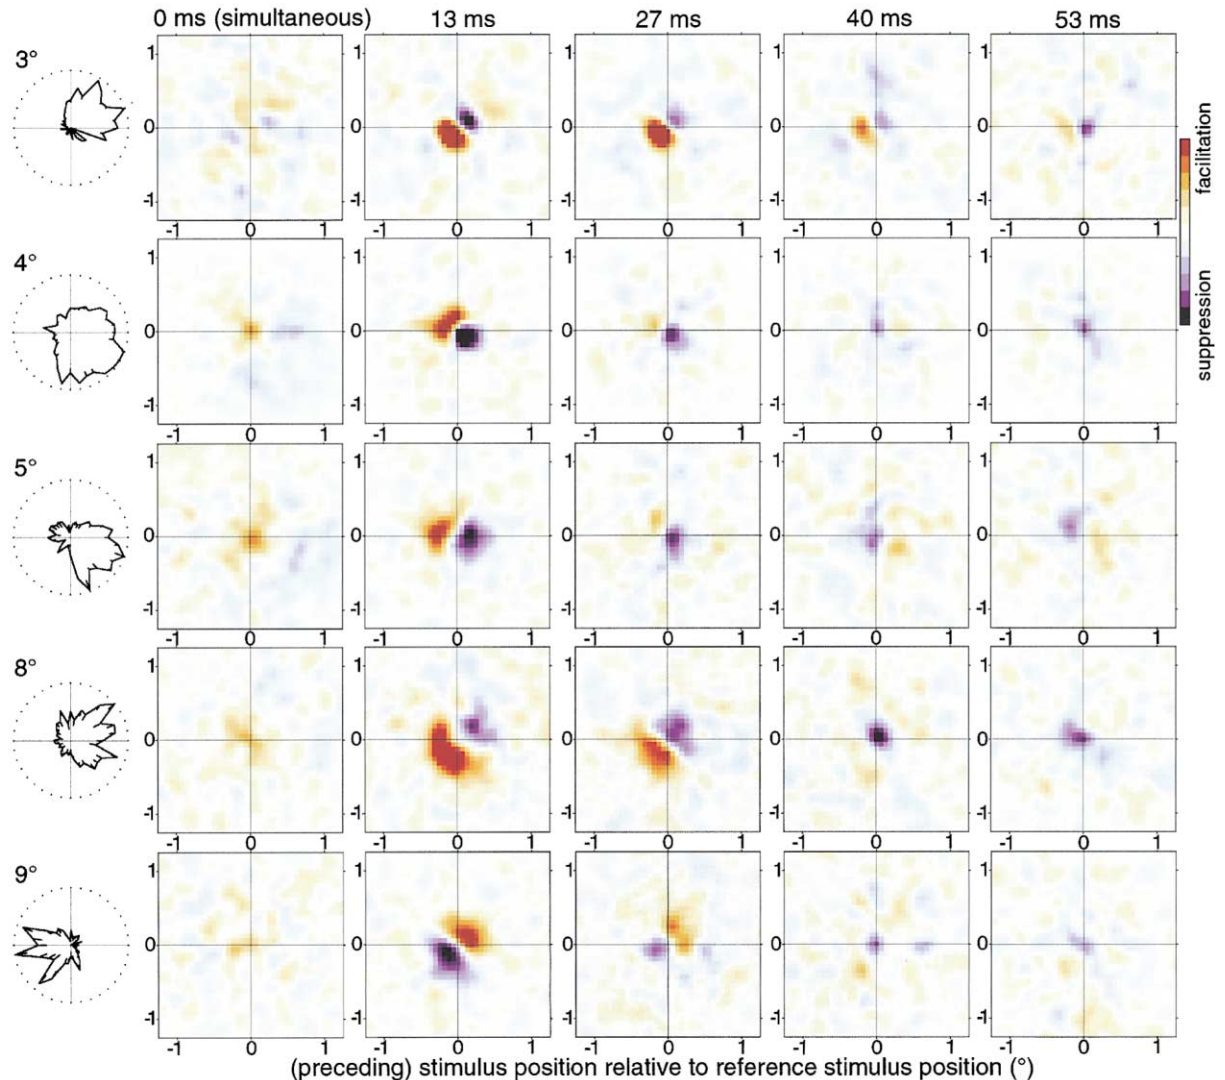
Figure 7. Direction Tuning Plots and 2D Sequential Interaction Maps at Various Interstimulus Intervals for Five MT Cells

The eccentricity of each cell is indicated in the upper left of each row. Direction tuning was generated using small fields of moving white dots. Interaction maps were generated as described for Figure 4. The maps represent same-contrast stimulus sequences minus inverted-contrast stimulus sequences. All five maps for each cell were derived from the same spike data, but using different interstimulus intervals. The 0 ms maps were derived from interactions between white and black stimuli presented simultaneously (there were no same-contrast conditions at 0 ms). Color scale ranges from maximal facilitatory interaction (+1) to maximal suppressive interaction (−1) and is consistent within each cell. For the first cell, for the 13 ms sequential interaction, the magnitude of the same-contrast facilitation was 42% larger than the response to the two stimuli presented independently; for the second cell the same-contrast facilitation was 11% larger than the response to the two stimuli presented independently; for the third cell it was 52% larger; for the fourth cell it was 53% larger; for the fifth cell it was 58% larger.

parallel to the axis of preferred motion, and that these interactions change in magnitude and position over time. The results do not contradict any well-established dogma, yet we found the precision, the size, and the shape of the interactions novel. No model we are aware of predicts the crescent-shaped interactions we observed in both MT and V1.

The short distances over which directional interactions occur are consistent with the idea that MT directionality is derived from antecedent cells with smaller receptive fields, as previously proposed (Movshon and Newsome, 1996). These interactions need not, however, represent the receptive-field structure of an immediately

antecedent input cell, as has been assumed in some previous studies of second-order interactions (Barlow and Levick, 1965; Movshon et al., 1978; Szulborski and Palmer, 1990). We know only that the 2D maps represent nonlinear interactions, and that these interactions could arise at any stage in the hierarchy. They could represent any combination of nonlinearities from several stages: nonlinear behavior of V1 simple cells (Heeger, 1991), nonlinearities in the integration properties of V1 complex cells (Emerson et al., 1987; Livingstone, 1998), or nonlinearities in the integration properties of MT cells (Britten and Heuer, 1999).

We cannot distinguish whether the directional interac-

tions we measure arise from linear or nonlinear mechanisms. Linear mechanisms alone (slanted space/time responses [Adelson and Bergen, 1985]) cannot give direction-selective responses to spots, bars, or any other non-repetitive stimulus (Poggio and Reichardt, 1976). Linear mechanisms alone can generate direction selectivity only for the amplitude of response modulation in response to a moving grating stimulus, but cannot affect the total number of spikes. Linear mechanisms can give a direction-selective response to bars or spots if they are followed by a nonlinearity like a threshold operation or an expansive non-linearity like squaring. Our method cannot distinguish between a linear mechanism followed by a nonlinearity versus a nonlinear mechanism like asymmetric or shunting inhibition.

The directional interactions are contrast-selective and exceedingly precise. Even at eccentricities of 20 and 30 degrees, directionality is seen for sequential presentations of stimuli less than a tenth of a degree apart (one pixel), a distance much smaller than the average V1 receptive-field size, and smaller still than the average size of V1 direction-selective receptive fields. If MT directional interactions arise in MT as a result of interactions *between MT cells*, they would be expected to be about the same scale as MT receptive fields, and they would not be expected to reverse with inverting-contrast stimuli (unless we postulate the existence of non-directional simple cells in MT). If MT directional interactions arise in MT as a result of interactions *between V1 inputs to MT cells* (from some kind of order-dependent dendritic calculation), they would be expected to be about the scale of V1 receptive fields. If MT directional interactions arise in V1 as a result of *interactions between V1 cells*, they would also be expected to be about the scale V1 receptive fields. The fact that MT directional interactions are smaller than V1 receptive fields suggests that they arise from interactions between subregions of V1 receptive fields, or that they arise from interactions between V1 cells with smaller-than-average receptive fields.

The exquisite precision of the directional interactions, and the fact that these interactions change in relative magnitude and spatial organization over time, with suppressive interactions outlasting facilitatory interactions, constrain models of direction selectivity. The time-dependent shifts between preferred-direction positive interactions and null-direction negative interactions perhaps explain why there has been variability in attributing directionality to facilitatory or inhibitory mechanisms (Ferster, 1994).

Reversed directionality to inverting-contrast stimuli indicates that the initial stages of generating motion selectivity must occur in or between cells that themselves have inverted *responses* to opposite-contrast stimuli (Anstis and Rogers, 1975; Adelson and Bergen, 1985). This may in fact be a general feature of neural mechanisms of direction selectivity, as similar reversed directional responses to inverting-contrast stimuli have previously been observed in primate V1 directional cells (Livingstone et al., 2000), as well as in directional cells in rabbit retina (Barlow and Levick, 1965), cat striate cortex (Emerson et al., 1987), and fly lobula plate (Egelhaaf and Borst, 1992). Simple cells (and their antecedents—geniculate cells, retinal ganglion cells) show in-

verted responses to opposite-contrast, but complex cells do not. Therefore any two-stimulus interaction generated at the simple-cell stage should invert when the two stimuli are of opposite contrast; any interaction generated between complex cells should not invert. The fact that direction selectivity inverts for inverting-contrast stimuli implies that it is largely generated within or between simple cells, geniculate inputs, or within dendritic compartments that behave like simple cells. This reversal of the directional response with contrast-inverting stimuli has a perceptual correlate in humans: reversed motion is seen in apparent motion demonstrations in which the sequential stimuli are of opposite contrast (Anstis and Rogers, 1975). For a nice demonstration of this phenomenon, see http://www.biols.susx.ac.uk/home/George_Mather/Motion/Harley.HTML.

Our results raise several theoretical questions. The first concerns the crescent-like shape of the 2D interaction maps: What is the significance of the inflection, or curvature at the reference location in our 2D maps? It could result from convergence of V1 inputs having different preferred directions and orientations (Simoncelli and Heeger, 1998), but our results in V1 indicate that the inputs themselves can exhibit the same curvature (M.S.L., B.R. Conway, and D.Y. Tsao, unpublished data). This does not eliminate the question of how the curvature arises, but it does push it back to V1. At the top of Figure 8 is diagrammed the simplest question of whether the crescent shape represents centripetally biased directionality, or simply a difference in length between positive and negative interactions. If the left diagram is correct, and the interactions are really crescent shaped, then the curvature must reflect some kind of convergence of directional inputs. Figure 8A is a diagram of the idea that excitatory inputs to a directional cell might be more broadly tuned for direction than its inhibitory inputs. This is a modification of the idea that directional cells are interconnected, with opposite direction preferences inhibiting each other, and cells with the same preference exciting each other (Adelson and Bergen, 1985; Qian et al., 1994; Simoncelli and Heeger, 1998). The biggest problem with this explanation is that such interactions would have to occur at an early enough stage that the interactions would still be extremely local.

If we entertain the idea that the crescent shape arises from the positive interactions being more elongated than the negative ones (Figure 8, top right), we need to spell out what the positive and negative interactions represent. One might think that an asymmetry in magnitude between *facilitation* and *suppression* could arise trivially from rectification: if excitatory inputs summate linearly, but inhibitory inputs do not (because a cell cannot exhibit a negative firing rate), then facilitation could become more powerful than suppression. However, the crescent shape cannot be a simple consequence of inhibition being weaker than excitation (or rectification) because the negative (blue) regions in the difference maps do not represent suppressive interactions alone, but are the combination of same-contrast inhibition and opposite-contrast excitation. Why should opposite-contrast excitation be weaker than same-contrast excitation? In the underlying, original simple cell (or center/surround cell) where these interactions must originate, stimulating with a light bar in an ON region and a dark
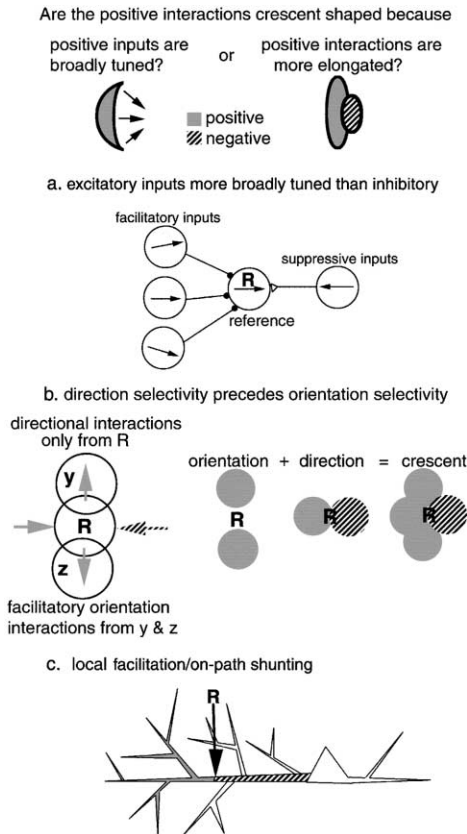
Figure 8. Three Kinds of Mechanisms that Could Lead to Crescent-Shaped 2D Directional Interactions

In diagram (a), a directional cell at the reference location (R) receives excitatory and inhibitory inputs from other directional cells; the excitatory interactions cover a broader range of directions than the inhibitory interactions.

In diagram (b), an orientation-selective, direction-selective cell is shown as being made up from three aligned cells that are direction-selective but not orientation-selective. The reference location "R" would be facilitated by locations directly above and below, but directional interactions would be more local. Why having direction selectivity precede orientation selectivity would give a crescent-shaped interaction is diagramed to the right.

Diagram (c) is a model in which excitatory and inhibitory inputs are distributed all along the dendritic tree, which is postulated to sample a retinotopic input map. For any given point in the dendritic tree (reference location, R) preceding distal inputs are facilitatory (gray regions) and preceding proximal inputs are suppressive (hatched regions).

bar in an OFF region is usually at least as good a stimulus as two light bars in the ON region (Hubel and Wiesel, 1962). So we would expect opposite-contrast excitation to be as strong as same-contrast excitation, but it apparently is not. There is a weak tendency in Figure 5 for the crescent angle to be correlated with the amplitude index, so some kind of asymmetry between preferred-direction and null-direction interactions could explain our results, though a difference in amplitude should result in a difference in width index as well as length index, which we do not see. It is nevertheless still possible that some difference in the magnitudes of the preferred-direction and the null-direction interactions might explain our results, but the explanation cannot be as simple

as an asymmetry between excitatory and inhibitory interactions.

Another possible explanation for the crescent shape is that direction selectivity arises before orientation selectivity (Figure 8B) (if we consider orientation selectivity to be equivalent to receptive-field elongation). If directional interactions occur between orientation-selective simple cells, we would expect the interactions to form parallel bands, not crescents. The crescent shape could represent a combination of nondirectional orientational facilitation (that would not be displaced from alignment with 0,0), plus a more localized directional interaction. The inflection at 0,0 implies that the directional interactions are more local than the interactions that are elongated along the preferred orientation. Another possible explanation for the crescent shape is that direction selectivity arises before orientation selectivity (Figure 8b) (if we consider orientation selectivity to be equivalent to receptive-field elongation). If directional interactions occur between orientation-selective simple cells, we would expect the interactions to form *elongated*, *parallel* bands of facilitation and suppression, not crescents. The crescent shape could represent a combination of nondirectional, orientationed facilitation (along the orientation axis) plus a more localized directional interaction (along the direction axis). The inflection at 0,0 implies that the directional interactions are more local than the interactions that are elongated along the preferred orientation, which would be consistent with direction selectivity arising before orientation selectivity.

Figure 8C shows a third possible mechanism to explain the crescent-shaped interactions, based partially on a model for direction selectivity proposed by Koch et al, 1982 and Koch and Poggio, 1987): that inhibitory inputs exert a more powerful suppressive effect on distal excitatory inputs than on more proximal excitatory inputs because they produce a shunting conductance that can block more distal epsp's. Direction selectivity would naturally arise in cells with excitation and inhibition distributed along the dendrites, if the dendrites were asymmetrically distributed about the cell body. If we combine the idea that an excitatory input will be blocked (or suppressed) by any inhibitory inputs that lies between it and the cell body (on-path) with an expansive nonlinearity, such as a local (nearby on the dendrite) facilitation (Mel, 1993), we obtain a model in which any point in the dendritic field would be facilitated by preceding distal inputs, and suppressed by preceding more proximal inputs. In a cell with a branching asymmetric dendritic field that samples a retinotopic input field (Livingstone et al., 2000), this would result in preferred-direction interactions that are longer (in the dimension perpendicular to the direction preference) than the null-direction interactions.

Which, if any, of these models is correct should clearly also be addressed at the level of V1 receptive fields, and it will be interesting to see how V1 interactions compare with those in MT.

Time Course
We did not look at speed tuning in these cells, and therefore did not try to compare the time course of directional interactions with speed tuning. But we infer

that any kind of long-lasting null-direction suppression should be pretty broadly tuned for velocity (see Livingstone, 1998). Moreover, Pack and Born have found that the earliest directional responses of MT cells are generally not well tuned for stimulus speed (Pack and Born, unpublished data), and because our stimuli were presented at 75 Hz, we were probably in the time-range when cells are poorly tuned for speed.

We did not find any evidence for longer-range, longer-lasting, or contrast-independent directional interactions using both 1D and 2D stimulus configurations using both short (13 ms) and long (up to 150 ms) interstimulus intervals. This is consistent with a previous study (Shadlen et al., 1993) that found local interactions to nearly completely dominate MT neuronal responses, when they were opposed in direction to long-distance interactions. Some previous studies have reported relatively long-distance directional interactions in MT (Mikami et al., 1986; Mikami, 1991), but in these studies the long-distance interactions were not constrained to be distinct (e.g., opposite in direction or contrast-invariant) from the short-range interactions, so they could represent the upper limit of the short-range system. Britten and Heuer, 1999 looked at summation properties of MT cells using a motion energy pulse designed to stimulate directional subunits of the kind we describe here. They did find interactions between directional subunits, in the form of nonlinear summation properties, but they did not analyze these interactions in terms of directionality.

It has been suggested that certain kinds of long-range motion perception represent a more cognitive positional tracking of identified objects (Lu and Sperling, 1995). If long-range motion detection does require object identification, attentional processes, or feedback from higher levels, our noisy stimulus might not have engaged this kind of motion processing (Casco et al., 1989), and the stimulus intervals we used (up to 150 ms) may not have been long enough. The ability of MT cells to signal the correct direction of motion for plaids (Movshon et al., 1985) and for long bars moving obliquely to their preferred-orientation axis (Pack and Born, 2001) suggests that some kind of long-range interactions do occur in MT. It is not clear, however, whether these interactions represent correlates of psychophysically characterized long-range (or second or third order) motion perception or something else.

Experimental Procedures

We recorded from 48 single units in MT of two alert rhesus macaque monkeys while they performed a simple fixation task. The monkeys were rewarded for maintaining fixation within $2°$ of a small fixation spot.

For each unit studied, we first determined the cell's preferred direction of motion using fields of dots or moving bars. We then presented the two-bar (1D) stimulus centered on the cell's receptive field, while the monkey fixated a small spot in the center of the monitor. The 1D mapping stimulus consisted of pairs of narrow bars, one white, one black on an intermediate gray background. At each monitor refresh (every 13 ms) one black and one white bar were presented at random positions along a one-dimensional stimulus range parallel to the preferred motion axis. The orientation of the bars was perpendicular to the cell's preferred axis of motion. The 2D mapping stimulus consisted of pairs of small squares (one black one white on a gray background) presented at random positions (fully random, not in a grid) at 75 Hz within a square stimulus range.

The white and black stimuli were 19 cd/m^2 above and below the mean background gray luminance of 20 cd/m^2. When black and white stimuli overlapped, the resultant stimulus was the same gray as background.

A computer recorded the evoked spike train, at 1 ms resolution, each stimulus position, and the monkey's eye position, at 4 ms resolution. For each map, between 5,000 and 30,000 spikes were collected over a 5 to 20 min period. The stimulus is "white" in the sense that the spatial and temporal autocorrelation functions are flat. It is "sparse" in that the stimulus at each position is the same as the background luminance most of the time. Because the probability distribution of luminances used in our stimulus is not Gaussian, it would not fit the criteria for Wiener kernel analysis. However, Emerson et al., 1987 developed a modified Wiener kernel analysis for use with ternary white noise, and our difference maps are equivalent to their second-order Wiener-like kernel calculation. For the 1D maps, pairs of bars, one white, one black, were presented at random positions along a one-dimensional stimulus range that was perpendicular to the cell's preferred orientation. Spikes were reverse correlated with the positions of sequential pairs of bars at a delay corresponding to the peak of the response to the second stimulus; bar position was corrected for eye position at stimulus onset. For the 1D maps, the x axis represents the second bar position along the stimulus range, with rightward corresponding to the preferred direction of motion; the y axis represents the preceding bar position along the stimulus range, with upward corresponding to the preferred direction. For the 2D maps, pairs of small squares, one white, one black, were presented at 75 Hz within a square (two-dimensional) stimulus range. Spikes were reverse correlated with the difference in position between sequential pairs of spots at a delay corresponding to the time to peak response to the second stimulus. For the 2D maps, the x axis corresponds to the horizontal separation, in visual space, of the two stimuli, and the y axis corresponds to the vertical separation.

In order to examine the portions of the maps that depend on the relative positions of pairs of stimuli, we must eliminate those aspects of the response that do not depend on the relative positions of two stimuli. The first of these is the linear contribution of each individual stimulus, and one way to eliminate them is by subtracting the different-contrast maps from the same-contrast maps [(*white*-to-<u>white</u> + *black*-to-<u>black</u>) − (*white*-to-<u>black</u> + *black*-to-<u>white</u>)]. For both the 1D and the 2D maps we use this technique (which is equivalent to Emerson et al.'s (1987) calculation of second-order Wiener-like kernels) to calculate *difference* maps by subtracting different-contrast maps from same-contrast maps.

Our difference maps give a picture of the interactions between stimuli, independent of the contrast of the stimuli. To look at interactions between specific pairs of contrast stimuli in the 2D maps, we had to use a different means to eliminate the linear component of the response. In these 2D maps, we want to study directional interactions as a modulation of firing rate independent of spatial position within the receptive field. This requires subtracting out some measure of the average number of spikes at each position in the interaction map. However, this average is affected both by variations in sensitivity across the receptive field and by the fact that we did not sample directional interactions perfectly uniformly. It is impossible to sample stimulus interaction space uniformly and still be able to examine various intervals from the same spike train. Using a fixed stimulus range results in edge effects at the border of the stimulus range: for example, a reference stimulus anywhere within the leftmost column of the stimulus range cannot be preceded by a stimulus further to the left. The magnitude of these contributions can be determined by reverse correlating the spike train with pairs of stimuli separated by long temporal intervals. (We used 250 ms because our 1D mapping experiments had shown that there were no interactions between stimuli separated by such a long interval. Selecting this early, "meaningless" stimulus as the probe stimulus for the map has the effect of randomizing the positions of the stimuli immediately preceding the reference stimulus.) For large numbers of stimulus presentations, this is equivalent to convolving the stimulus probability distribution with the receptive field sensitivity profile. The resulting map averages over motion direction while preserving the mean spatial distribution of spikes. Because this value depends only on the sensitivity of the cell and the number of stimulus presen-

tations at each spatial position, we subtracted the baseline map from the 2D maps, leaving only the modulation due to paired-stimulus interactions.

The monkeys were prepared for chronic recording from MT (Born et al., 2000). Following surgery each animal underwent a magnetic resonance imaging scan to locate MT within the coordinates of a plastic grid inserted in the recording cylinder (Crist). The same grid was used to guide insertion of the microelectrode after which MT was identified based on its depth, prevalence of direction-selective neurons, receptive field size, and visual topography. MT was easily distinguished from MST based on the ratio of receptive field size to eccentricity, which is much lower in MT than in MST (Van Essen et al., 1981). All procedures were approved by the Harvard Medical Area Standing Committee on Animals.

References

Adelson, E.H., and Bergen, J.R. (1985). Spatiotemporal energy models for the perception of motion. J. Opt. Soc. Am. A-Optics & Image Science *2*, 284–299.

Anstis, S.M., and Rogers, B.J. (1975). Illusory reversal of visual depth and movement during changes of contrast. Vision Res. *15*, 957–961.

Barlow, H.B., and Levick, W.R. (1965). The mechanism of directionally selective units in rabbit's retina. J. Physiol. *178*, 477–504.

Born, R.T., Groh, J.M., Zhao, R., and Lukasewycz, S.J. (2000). Segregation of object and background motion in visual area MT: effects of microstimulation on eye movements. Neuron *26*, 725–734.

Braddick, O. (1974). A short-range process in apparent motion. Vision Res. *14*, 519–527.

Britten, K.H., and Heuer, H.W. (1999). Spatial summation in the receptive fields of MT neurons. J. Neurosci. *19*, 5074–5084.

Casco, C., Morgan, M.J., and Ward, R.M. (1989). Spatial properties of mechanisms for detection of moving dot targets in dynamic visual noise. Perception *18*, 285–291.

Cavanagh, P., Tyler, C.W., and Favreau, O.E. (1984). Perceived velocity of moving chromatic gratings. J. Opt. Soc. Am. A-Optics & Image Science *1*, 893–899.

deBoer, E.K.P. (1968). Triggered correlation. IEEE Transactions on Biomedical Engineering *153*, 169–179.

Dubner, R., and Zeki, S.M. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. Brain Res. *35*, 528–532.

Egelhaaf, M., and Borst, A. (1992). Are there separate ON and OFF channels in fly motion vision? Vis. Neurosci. *8*, 151–164.

Eggermont, J.J., Johannesma, P.M., and Aertsen, A.M. (1983). Reverse-correlation methods in auditory research. Q. Rev. Biophys. *16*, 341–414.

Emerson, R.C., Citron, M.C., Vaughn, W.J., and Klein, S.A. (1987). Nonlinear directionally selective subunits in complex cells of cat striate cortex. J. Neurophysiol. *58*, 33–65.

Ferster, D. (1994). Linearity of synaptic interactions in the assembly of receptive fields in cat visual cortex. Curr. Opin. Neurobiol. *4*, 563–568.

Heeger, D.J. (1991). Half-squaring in responses of cat striate cells. Vis. Neurosci. *9*, 427–443.

Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J. Physiol. *160*, 106–154.

Jones, J.P., and Palmer, L.A. (1987). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. J. Neurophysiol. *58*, 1187–1211.

Koch, C., Poggio, T., and Torre, V. (1982). Retinal ganglion cells: a functional interpretation of dendritic morphology. Philos. Trans. R. Soc. Lond. B. Biol. Sci. - Series B. Biological Sciences *298*, 227–263.

Koch, C., and Poggio, T. (1987). Biophysics of Computational Systems: Neurons, Synapses, and Membranes. In Synaptic Function, G.M. Edelman, W.E. Gall, and W.M. Cowan, eds. (New York: John Wiley and Sons) pp. 637–697.

Livingstone, M.S. (1998). Mechanisms of direction selectivity in macaque V1. Neuron *20*, 509–526.

Livingstone, M.S., Tsao, D.Y., and Conway, B.R. (2000). What happens if it changes contrast when it moves? Soc. Neurosci. Abst. *26*, 162.6.

Lu, Z.L., and Sperling, G. (1995). The functional architecture of human visual motion perception. Vision Res. *35*, 2697–2722.

Mel, B.W. (1993). Synaptic integration in an excitable dendritic tree. J. Neurophysiol. *70*, 1086–1101.

Mikami, A. (1991). Direction selective neurons respond to short-range and long-range apparent motion stimuli in macaque visual area MT. Int. J. Neurosci. *61*, 101–112.

Mikami, A., Newsome, W.T., and Wurtz, R.H. (1986). Motion selectivity in macaque visual cortex. II. Spatiotemporal range of directional interactions in MT and V1. J. Neurophysiol. *55*, 1328–1339.

Movshon, J.A., Adelson, E.H., Gizzi, M.S., and Newsome, W.T. (1985). The analysis of moving visual patterns. In Pattern recognition mechanisms. C. Chagas, R. Gattas and C.G. Gross, eds. (New York:Springer). 117–151.

Movshon, J.A., and Newsome, W.T. (1996). Visual response properties of striate cortical neurons projecting to area MT in macaque monkeys. J. Neurosci. *16*, 7733–7741.

Movshon, J.A., Thomson, I.D., and Tolhurst, D.J. (1978). Spatial summation in the receptive fields of simple cells in the cat's striate cortex. J. Physiol. *283*, 53–77.

Pack, C.C., and Born, R.T. (2001). Temporal dynamics of a neural solution to the aperture problem in visual area MT of macaque brain. Nature *409*, 1040–1042.

Poggio, T., and Reichardt, W. (1976). Nonlinear Interactions Underlying Visual Orientation Behavior of the Fly. Cold Spring Harbor Symp. on Quant. Biol. *XL*, 635–645.

Qian, N., Andersen, R.A., and Adelson, E.H. (1994). Transparent motion perception as detection of unbalanced motion signals. III. Modeling. J. Neurosci. *14*, 7381–7392.

Ramachandran, V.S., and Gregory, R.L. (1978). Does colour provide an input to human motion perception? Nature *275*, 55–56.

Shadlen, M.N., Newsome, W.T., Zohary, E., and Britten, K.H. (1993). Integration of local motion signals in area MT. Soc. Neurosci. Abst. *19*, 1282.

Simoncelli, E.P., and Heeger, D.J. (1998). A model of neuronal responses in visual area MT. Vision Res. *38*, 743–761.

Szulborski, R.G., and Palmer, L.A. (1990). The two-dimensional spatial structure of nonlinear subunits in the receptive fields of complex cells. Vision Res. *30*, 249–254.

Van Essen, D.C., Maunsell, J.H., and Bixby, J.L. (1981). The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization. J. Comp. Neurol. *199*, 293–326.

Van Essen, D.C., Newsome, W.T., and Maunsell, J.H. (1984). The visual field representation in striate cortex of the macaque monkey: asymmetries, anisotropies, and individual variability. Vision Res. *24*, 429–448.