From: Motion Vision - Computational, Neural, and Ecological Constraints
Edited by Johannes M. Zanker and Jochen Zeil
Springer Verlag, Berlin Heidelberg New York, 2001

The Organization of Global Motion and Transparency

Oliver Braddick[1] and Ning Qian[2]

[1]Department of Psychology, University College London, UK; [2]Center for Neurobiology and Behavior, Columbia University, New York, USA

Contents

1. Abstract

The visual system has the task of computing global motions associated with objects and surfaces. This task strongly involves extrastriate brain areas, particularly V5/MT. Motion transparency provides a particular challenge for understanding how global motions are computed and represented in the brain. Psychophysical experiments show that, for a single region, multiple motions can be quantitatively represented. However, at the most local scale, motion signals have a suppressive interaction so that only a single motion can be represented. Neurophysiological experiments show that this suppression is a property of MT, not of V1, reflecting a subunit structure within MT receptive fields and showing that transparency perception is related to MT rather than V1 activity. A full understanding of transparency perception and other global motion phenomena will require us to understand how perceived motions are related to the distribution of activity across a population of directionally selective neurones, and how the brain implements the representation of motions assigned to extended objects rather than to specific retinotopic locations.

2. Global representations of motion

Much research on visual motion processing has been concerned with the processes that measure image velocity in a local region of the image. The computational goal of motion vision is sometimes described as "computing the velocity field" i.e. obtaining valid measurements of this kind at each point in the field of view. However, it is not adequate to consider visual motion processing at a purely local level, for several reasons.

First, it is well known that purely local operations are not adequate to obtain a valid velocity field. The "aperture problem" means that, in regions of the image where variation is mostly along one dimension (i.e. contours) local operations can only measure the component of motion at right angles to the contour (Adelson and Movshon 1982). Directionally selective mammalian V1 neurones with small, oriented receptive fields will generally make ambiguous motion measurements of this kind. To derive the true velocity vector for such locations, information must be integrated from other locations on the same moving object. This process is believed to depend on "pattern-motion-selective neurones" of the kind found in the area known as V5 or MT in the primate brain (Movshon et al. 1986). Second, spatial integration is important to smooth noisy or sparse motion signals (Braddick 1993). Processes that disambiguate and smooth motion signals require operations over an extended area, but still, their goal can be regarded as achieving a representation of velocity for each point in the field.

However, the purpose of vision is not to compute fields but to create a representation that is useful for understanding, and acting on, the environment. Generally, the significant entities in the environment are large-scale structures, objects and surfaces. The visual organism needs to assign to these structures their motions, or properties of the structures which are derived from image motion, such as depth ordering or relation to the perceiver's trajectory. Thus the visual system has to compute global motions associated with objects and surfaces, not just the local motions associated with locations in the field.

The difference between local and global motion is most obvious when local motions are not uniform. For example, the turbulent flow of water in a stream, gravel being dumped from a truck, or a flock of birds taking to the air, each give the viewer a sense of overall or global direction despite containing a wide and disorderly range of local motions. Neural mechanisms must exist that can represent this global direction.

3. Experimental approaches to global motion

3.1 Motion coherence

Sensitivity to motion coherence has become an established measure of the performance of global motion mechanisms. It is tested with a random dot display in which some fraction of the dots are "signal" dots which share a common "coherent" motion, while the remaining "noise" dots are displaced at random. The noise can be defined in a number of different ways, but these do not appear to make much difference to the measurement of sensitivity (Scase et al. 1996). The direction of the signal dots can be detected when the coherence (proportion of signal dots) is as low as 5-7%. This direction cannot be derived by inspecting individual dots, but must depend on the integration of motion information over a large area of the display (Downing and Movshon 1989).

Several lines of evidence support the idea that this integration depends on processing in extrastriate cortical areas. First, lesions in MT of the macaque greatly impair coherence thresholds (Newsome and Pare 1988), and the presence of noise dots also has a very deleterious effect on motion discrimination in a patient with a lesion including the human homologue of this area (Baker et al. 1991). Secondly, the response of single neurones in macaque MT shows a systematic dependence on motion coherence (Britten et al. 1992). The estimated threshold coherence for such neurones agrees well with the animal's behavioural coherence threshold (Britten et al. 1993).

Third, a recent fMRI study (Braddick et al. 1998) has looked for cortical areas that are differentially activated by coherent motion compared to random noise (Fig. 1). Such activation is not found in V1, but is found in human V5 and some other extrastriate areas, particularly the area identified by Tootell et al. (1997) as V3A. In fact, V1 showed stronger responses to noise than to coherent motion. These findings confirm that in the human visual system, the integrative processes that underlie sensitivity to coherence occur beyond the initial stage of cortical processing in V1. They are also consistent with single unit findings in macaque, discussed below, that V1 cells respond much more strongly to dynamic noise than do cells in MT (Qian and Andersen 1994).

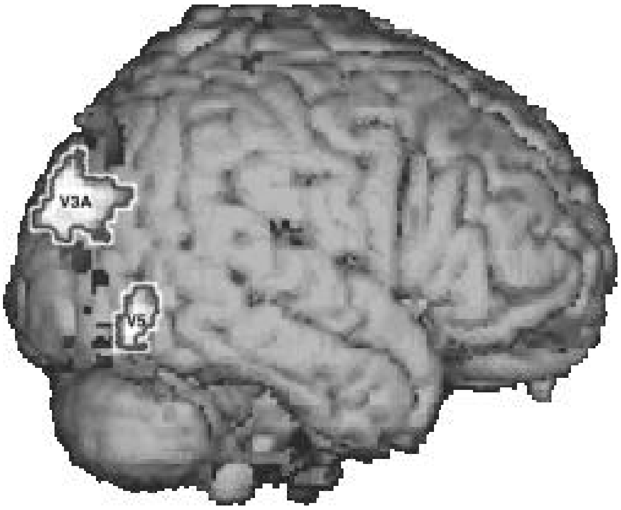

Fig. 1 Lateral view showing brain areas that are activated significantly more by coherent motion than by dynamic noise, in an individual subject. Highlighted areas are voxels within 3 cm of the cerebral surface with $z > 3.09$ ($p<0.001$) for this contrast in Statistical Parametric Mapping (Friston et al. 1995). The two major foci of activation marked correspond to areas identified as V5 and V3A. Other areas of differential activation were found on the ventral occipital surface, in the intraparietal sulcus and the superior temporal sulcus.

3.2 Directional distributions

An alternative approach to global motion is to use, rather than displays containing random noise, displays in which there is a broad distribution of elementary motions (Williams and Sekuler 1984). Human subjects' ability to extract global motion is exemplified by results from Watamaniuk et al. (1989), who showed that the mean direction of a random dot display could be judged with an accuracy of 2º, even when the motion vectors of the individual dots were distributed over a range of up to 120º. However, this performance should not be taken to imply that the motion is perceived as unitary. As in the real-world examples of the flock of birds or the tumbling gravel, subjects are aware that diverse local motions exist within the global flow.

In fact, global direction processing appears to be a flexible process that is capable of parsing the overall direction distribution differently for different purposes. This became apparent in experiments by Zohary et al. (1996). They tested subjects with skewed distributions of local motion vectors, in an attempt to determine whether global direction judgments depended on a mechanism using the peak or the mean of the velocity distribution. The result was that neither model alone gave an adequate account; rather subjects could freely switch between these two kinds of performance, as if they had access to the whole distribution of local motion vectors, and could perform different optional operations upon it to derive different global results. We return to this question below in discussion of bimodal and unimodal distributions of directional signals.

3.3 Multiple global motions: transparency

A particularly significant case where a distribution of directions does not automatically yield a single global outcome is the case of motion transparency. When there is a sufficiently large gap in the distribution of local directions (Smith et al. 1999), the distribution can be parsed into two separate global motions which are seen as spatially co-extensive and as superimposed in a transparent manner. This is not simply a special laboratory situation. It reflects the need to deal with visual situations where elements belonging to entities with different motions are interleaved or superimposed. Real-world examples are a moving object partially occluded by lacy foliage, or a moving shadow cast on a stationary, or differently moving, surface.

The phenomenon of transparency is an important test case for the nature of motion computation and representation. This is because it is incompatible with the simple goal of a single-valued velocity field. If perceived transparency is taken at its face value, it implies that the brain can represent two different motions at the same location. To make this point, it is important to test the nature of this multiple representation.

4. The analysis of motion transparency

4.1 Performance-based evaluation of motion transparency

Most earlier work on motion transparency has rested on simple subjective reports; does the display appear to contain two (or more) superimposed motions? Such reports cannot provide a very strong basis for understanding the underlying neural representation. In particular, they cannot clearly distinguish whether there is a true representation of two motions, or whether one motion is registered as segmented from its background. For instance, when signal dots are detected amid noise, they can be represented as a coherent entity that is distinct from the noise background of noise dots; this does not imply any necessary representation of the motion properties of the background, beyond its presence. Subjects given the choice between reporting one motion and two might well describe an asymmetrical representation of this kind as "two motions".

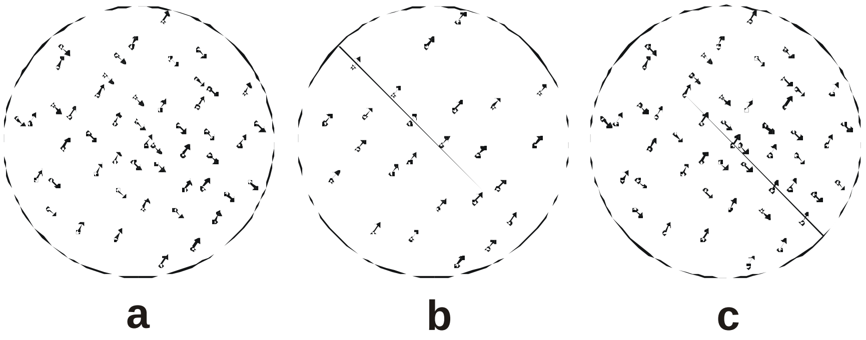

Fig. 2 a "Motion vs motion" display used in performance-based evaluation of transparency. The motions illustrated are 13º away from orthogonal, corresponding to the precision of subjects' judgments (s.d. of the psychometric function). **b** "Motion vs line" display serves as a control for the judgment illustrated in (a). The single motion illustrated is 5º away from being orthogonal to the line, corresponding to the precision of subjects' judgments (s.d. of the psychometric function). **c** "Transparent motions vs line" display. One motion is parallel to the line; the subject's task is to judge whether the other is orthogonal to the line.

If there are equivalent representations of the two velocities in a transparent display, subjects should have access to information about both of them. This can be tested by a perceptual task that necessarily depends jointly on both motions (Braddick 1997; Wishart and Braddick 1997a). Subjects are required to judge whether the angle between the directions of two superimposed streams of moving dots is greater or less than 90º (Fig. 2a). The orientation at which the pair of motions is presented is randomized from trial to trial, to ensure that there is no way to make the discrimination using one motion alone. An adaptive psychophysical method (Watt and Andrews 1981) is used to estimate the psychometric function for this discrimination. The standard deviation (s.d.) of this function is a measure of the precision with which the subject can use joint information about the two directions.

The results show that the angle between the motions can be judged with a precision of around 13º. Quantitative directional information is available from both dot streams jointly. By this criterion, the perception of transparency in this situation corresponds to a genuine multi-valued representation of velocity, not simply to a segregation of one motion from a background.

It is worth noting, however, that the precision, as reflected in the s.d. of these judgments, is somewhat impaired compared to that for a single motion direction relative to a line, which is around 5º (Fig. 2b). Some increase might be expected on a simple model of independent errors arising in the estimate of each of the pair of directions. This would predict a standard deviation about 7º, i.e. about half what is actually observed, so there must be some additional penalty associated with the transparent case. Such a penalty might arise from interference between the different local motions of dots in the two streams, or from an interaction at the level of the two global representations. In the latter case, it might be associated specifically with transparency, or it might be a more general penalty incurred by computing and comparing two directions at the same time.

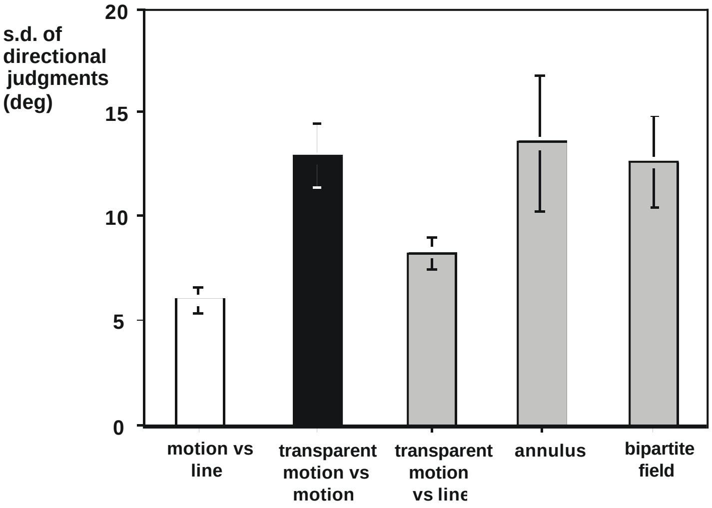

Fig. 3 Performance on directional judgments for the three conditions illustrated in figure 2a, b, c (columns 1-3) and for conditions where tow dot streams are separated into an annulus and central region (column 4) or a bipartite field (column 5). Mean and s.e. of results for 4 subjects are shown.

This issue has been tested by three further conditions. In one condition, the two motions are superimposed, but the direction judgment (relative to a line) has to be made for only one of them (Fig. 2c). This shows only a small loss of precision (s.d. around 8º) relative to a single motion, implying that local interference is not the major penalty in the dual motion judgment. In the second and third conditions, the required judgment is of the angle between two spatially separated motions. In one case these are a central region and a surrounding annulus; in the other, two halves of a bipartite field with a randomly oriented boundary. In these cases no interference will arise from superimposed local motions. Nonetheless, the loss of precision compared to judgment of a single stream is comparable to that for the transparent case; figure 3 shows the standard deviations obtained in each of the conditions. Thus the penalty in direction judgements is associated with the need to compute and compare two global representations of motions, whether or not these are spatially co-localized.

In summary, a measure based on visual performance shows that the visual system can represent two transparent motions in the same region of the field and operate with these represented motions as readily as when they are in adjoining regions. It should be noted, though, that this result does not tell us about the spatial properties of these representations. In some sense, two globally represented velocities occupy “the same place”, but this leaves open whether at a local level, there is a multi-valued representation of velocity associated with each location. The study of locally paired displays, discussed below, throws some light on this question.

4.2 Transparent motions at the single-unit level

Along with the psychophysical investigation of transparency in random-dot patterns, there has also been an increasing number of physiological investigations of this topic in recent years (Snowden et al. 1991; Qian and Andersen 1994; Qian and Andersen 1995). Single unit studies have also tested plaid patterns (Movshon et al. 1986; Rodman and Albright 1989; Stoner and Albright 1992), where variations in stimulus parameters can change the perceived effect from that of two gratings sliding over each other transparently in different directions, to a single coherently moving plaid.

4.3 Directional suppression in V5/MT

Snowden et al. (1991) were among the first to record neurones from visual cortical areas V1 and MT of awake monkeys using transparent random dot patterns. MT (also known as V5) is a later stage than V1 along the well established motion processing pathway in primates (Andersen 1997). Intuitively, one might expect that neurones at the higher level would be able to represent the transparent motions as two independent global entities. However, the most prominent finding of the study was that for most MT cells, stimulated with a transparent display containing one set of dots moving in the cell's preferred direction and another in the antipreferred directions, responses were significantly reduced compared to a single set of dots moving in the preferred direction alone. This result indicates strong suppression in MT between the preferred and the antipreferred directions of motion, consistent with other physiological studies of this cortical area (Mikami et al. 1986a; Rodman and Albright 1987; Britten et al. 1992, 1993). In this review, we will use the terms "suppression" and "inhibition" interchangeably, both referring generically to any reduction in the preferred-direction response. We will not try to distinguish various specific mechanisms for the response reduction such as normalization, subtraction or averaging, although these might have quite different synaptic bases.

Snowden et al. (1991) found that this kind of suppression was much weaker in area V1. Since MT is usually considered to be a major site for motion analysis, the finding raises the question of how MT cells can represent transparent motions given the directional suppression in MT. Indeed, the subpopulation of directionally selective V1 cells whose responses to transparent displays are not much suppressed, would appear better able to support the perception of motion transparency. However, as Snowden et al. (1991) point out, motion perception is not unaffected by a superimposed transparent motion; psychophysical detection threshold is higher in the transparent condition than for unidirectional motion (Snowden 1989; Verstraten et al. 1996). They argued that this effect could be a correlate of the suppressive interaction seen in MT cells.

5. "Locally balanced" motions

5.1 The perceptual effects of locally balanced motions

Interactions between directions at different levels of the motion pathway can be examined by using a modification of the transparent random-dot display. In the psychophysical and physiological experiments discussed so far, dots belonging to the two streams are randomly distributed. Qian et al. (1994a, b) and Qian and Andersen (1994, 1995) noted that, in patterns of two components moving in opposite directions, the distribution of the dots had a critical effect upon the appearance of transparency. They designed "paired" dot patterns, composed of many (typically 100) randomly located pairs of dots. The two dots in each pair move across each other in opposite directions over a short distance (about 0.4°) and then jump to a new random location to repeat the process. This display can be compared with the "unpaired" case where the dots moving in opposite directions are positioned independently. Qian et al. observed that, although both types of stimuli contain two sets of dots moving in opposite directions, only the latter gives a perception of two transparent surfaces sliding across each other. The paired display appeared like flickering noise, with no percept of opposing motions.

Wishart and Braddick (1997b) and Braddick (1997) have shown that the effects of local pairing are reflected in directional judgments based on the two motions. They used the task of figure 2a with paired dots in short (0.14º) trajectories, varying the distance between centres of trajectories in a pair. Precision in judging orthogonal directions began to decline when the separation was reduced below 0.5º, and for separations of 0.2-0.1º, subjects found it impossible to make the judgments at all. A similar variation in performance occurred when paired trajectories had a common centre, and their length was manipulated. Testing at different speeds confirmed that the spatial separation of the dots was the critical variable (Fig. 4). It should be noted that in the case of orthogonal motions, the alternative to transparency may be coherent motion rather than flicker, an issue we discuss later.

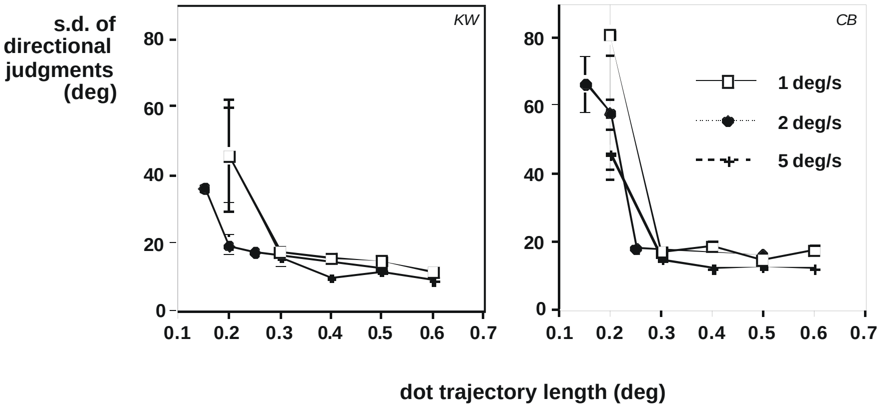

Fig. 4 Performance on directional judgements for locally paired dots moving in directions 90º apart, at 1-5º/s, with varying trajectory length. Panels show results of two subjects. At all speeds, performance radically deteriorates and variability (shown by s.e. bars) in-creases for trajectories shorter than 0.25º, indicating the abolition of motion transparency.

The effect of local pairing suggests that within any small area of the field, local motion detectors tuned to different directions have a strong interaction, such that only a single directional signal can survive at that location. In the case of opposed motions, the signals are locally balanced and the result is an abolition of any perceived motion. The co-representation of different motions in transparency can occur only at a scale coarser than this local interaction. In the case of unpaired dot patterns, local fluctuations of dot density can lead to locally unbalanced motion signals in opposite directions. For example, in one local area there might be three dots moving to the left and only one to the right; in a nearby location the opposite balance might occur. Transparency would then depend on the integration of two sets of spatially mixed local signals, into two global motions.

5.2 Neural responses to transparent and locally balanced motion stimuli

The local interaction which leads to the representation of motion being locally single-valued has been investigated neurophysiologically, through single unit recordings from behaving monkeys using paired and the unpaired dot patterns. Qian et al. (1994a,b) suggested that the interaction determining the presence or absence of transparency might correspond to the directional suppression observed in area MT. Two predictions can be made from this hypothesis. Firstly, if one set of dots move in the preferred direction and the other in the antipreferred direction for an MT cell, its response to both the paired and the unpaired dot patterns should be reduced compared to a single set of dots moving in the preferred direction alone. This is simply because the presence of motion in the antipreferred direction should always generate some suppression. Secondly, and more importantly, the amount of suppression should be stronger (i.e., the response smaller) for the paired dot patterns than for the unpaired ones, because the balanced opposing motion signals in the paired dot patterns should cancel each other more completely.

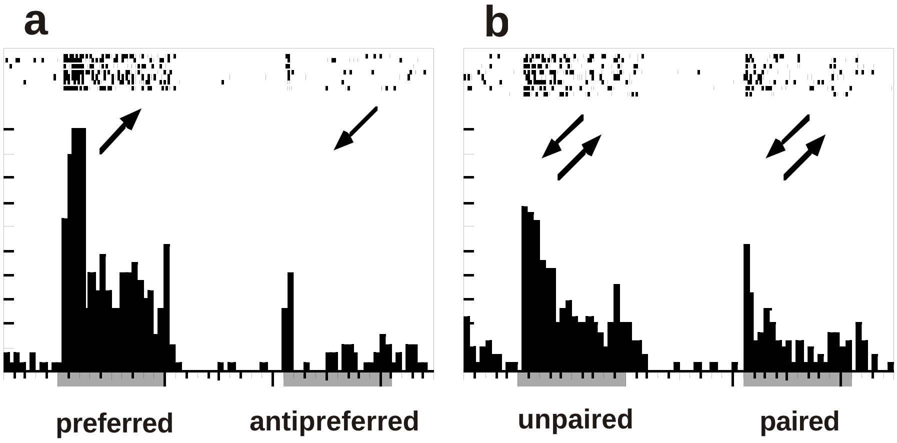

Fig. 5 An MT cell with significantly stronger response to the unpaired dot pattern than to the paired dot pattern. **a** Responses to single sets of random dots moving in its preferred and the antipreferred directions. **b** Responses to the paired and the unpaired dot patterns with one set of dots moving in the preferred direction and the other the antipreferred direction. The raster on the top of each diagram represent the spike records from several repeated trials. Each small dot in the raster represents the occurrence of a spike. The response histograms compiled from the rasters are shown at the bottom of the diagrams. The arrows below the rasters indicate the directions of motion. Each small division in the horizontal axis represent 10 ms. The one-second periods during which the stimuli are presented were marked by the thick black lines under the histograms. One small vertical division represents 7.7 spikes/s.

Qian and Andersen (1994) found that many MT cells did indeed behave as predicted. An example is shown in figure 5. Figure 5a shows the cell's responses to single sets of dots moving in its preferred and antipreferred directions respectively. It is evident that the cell is highly directionally selective, which is typical for MT cells. Figure 5b shows the cell's responses to the paired and the unpaired dot patterns. Both responses are weaker than the preferred direction response and the response to the paired dot pattern is significantly weaker than that to the corresponding unpaired dot patterns.

To see if the overall behaviour of MT cells is consistent with the predictions, the population results from a total of 91 recorded MT cells are summarized in figure 6. To characterize the cells quantitatively, two suppression indices for each cell were computed, one for the paired dot pattern and the other for the unpaired dot pattern along the preferred--antipreferred axis of motion. They are defined as:

$$SI_p = 1 - (\text{paired response})/(\text{preferred response}),$$

$$SI_{up} = 1 - (\text{unpaired response})/(\text{preferred response}).$$

The background firing rate was subtracted from all responses before calculation. These indices represent the percent reductions of a cell's responses to the paired and the unpaired dot patterns, respectively, in comparison with its preferred direction response. They therefore measure the degrees of suppression between preferred and antipreferred directions of motion. An index near zero indicates no suppression, a large value indicates strong suppression, and a negative value means that enhancement instead of suppression has occurred.

Figure 6 plots, for each cell, the unpaired suppression index, SI_{up}, against the paired suppression index, SI_p. First note that the suppression indices for almost all MT cells, whether measured with the paired or the unpaired dot patterns, are positive. This can be seen by projecting the dots in figure 6 along either coordinate axis. The result agrees with the previous finding by Snowden et al. (1991) that MT cells show significant suppression between different directions.

To examine how cells responded differently to the transparent and the non-transparent patterns, note that cells with similar responses (thus equal degrees of suppression) to the paired and the unpaired dot patterns lie near the diagonal line. The responses of these cells do not distinguish the transparent patterns from the non-transparent ones. Those falling well below the diagonal line show stronger suppression (or less response) for the paired dot patterns than for the unpaired ones. Finally, cells well above the diagonal line have the opposite behaviour: they show stronger suppression for the unpaired dot patterns than for the paired ones. There are significantly more cells below the diagonal line than above in figure 6 (Wilcoxon signed-rank test, $p < 0.0001$), consistent with the prediction.

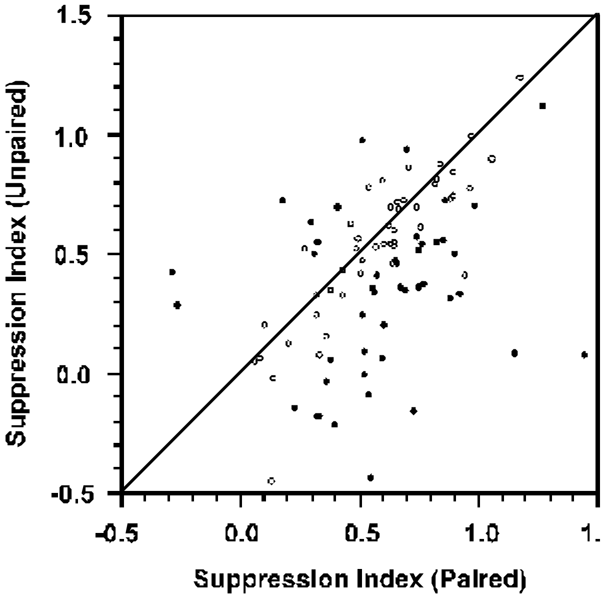

Fig. 6 MT population data. The suppression index for the paired dot pattern of each MT cell is plotted against its suppression index for the unpaired dot pattern. The cells with significantly different responses to the two types of patterns based on a two-tailed t-test are shown as solid dots. The rest are shown as open dots.

Each cell can also be considered individually. A two-tailed t-test was performed for each cell to see if its mean response to the unpaired dot pattern over repeated trials was significantly different from that to the paired dot pattern. The cells that did show a significant difference between the two patterns are plotted as solid dots in figure 6 while those did not are plotted as open dots. It was found that 40% of cells responded significantly more strongly to the unpaired dot patterns than to the paired dot patterns, while only 10% displayed the opposite behaviour. A closer examination indicated that the cells with the opposite behaviour are more responsive to flicker noise and therefore less selective to motion. The remaining 50% of cells did not show significant differences between displays of the two types. For these cells, the suppression was already very strong for the unpaired dot patterns, and there was perhaps no scope for further suppression when the paired dot patterns were used. When all recorded MT cells were averaged together, the whole population still showed significantly stronger responses to the transparent, non-paired, dot patterns than to the non-transparent paired ones. Therefore, measuring the residual responses of the directional suppression mechanism in the population of MT cells as a whole could reliably distinguish the two types of pattern.

5.3 MT suppression as a function of separation

If suppressive interaction within MT is the basis of the psychophysical effect of local pairing, the suppression should be reduced, like the psychophysical effect, when the local balance is reduced by spatial separation of differently moving dots. This effect was indeed observed by Snowden et al. (1991). In that study, displays were divided into either two or six adjacent but non-overlapping bands. Dots in alternate bands moved in opposite directions. Individual MT cells were recorded using these less balanced patterns as well as the overlapping transparent patterns. The results indicate that the suppression under the less balanced condition is indeed weaker than under the overlapping condition (shown by more points appearing below the diagonal lines in figure 9 of Snowden et al. (1991).

6. MT and V1 organization and transparency

6.1 Possible subunit structure of MT receptive fields

Since the physical difference between the paired and the unpaired dot patterns occurs at the scale of about 0.4º, these displays would best be distinguished by units operating at about this scale. The receptive fields of the MT cells recorded, however, were much larger (6-10º in diameter); a single receptive field could contain the whole of the paired or unpaired patterns. For MT cells to respond differentially to the paired and the unpaired dot patterns, it has to be assumed that the receptive field is composed of many small subunits of size about 0.4º, and that directional suppression occurs within each subunit. If it is further assumed that the response of an MT cell is determined by the sum of the thresholded outputs of all its subunits, the cell will then respond to the two types of dot patterns differently. The assumption of subunit structure in MT receptive fields is consistent with a study by Shadlen et al. (1993) who used stimuli containing local motion in one direction and global apparent motion in the opposite direction. They found that MT responses to such stimuli were largely determined by the local motion, suggesting that MT cells summate inputs to the subunits in their receptive fields.

However, there is evidence which makes it unlikely that inhibitory interactions occur only within subunits and not between them. First, the spatial range of directional interaction is three times as large in MT as in V1 (Mikami et al. 1986b). If V1 inputs provide the subunits, MT directionality must arise at least in part from mechanisms on a larger scale than the subunit fields. Second, Recanzone et al. (1997) recently recorded MT cells’ responses to two small objects moving along different paths in the receptive field, and found suppressive interactions between the objects similar to those in the random dot case. Since the two objects were spatially close only briefly along their trajectories, the observed suppression must partially result from longer-range inhibitory interactions between different subunits in a receptive field. Clearly, more direct experimental evidence is needed to elucidate the structure and interactions of MT subunits. However, the difference between paired and unpaired patterns, both psychophysically and neurophysiologically, makes it clear that there is a large difference between the inhibitory interactions within a subunit field and any that go beyond it.

6.2 V1 responses to transparent and non-transparent motion stimuli

It is plausible that MT subunits should correspond to the inputs from individual V1 cells that build up the larger receptive fields of MT. If so, can the suppressive interactions that occur within subunit fields be found in V1 cells? To test this point, Qian and Andersen (1994) also recorded from V1 cells using the paired and the unpaired dot patterns. However, only 17% of V1 cells responded significantly more strongly to the unpaired dot patterns than to the paired ones, while 8% showed the opposite behaviour. The remaining 75% showed no significant difference of response to the two types of patterns. Moreover, unlike area MT, measurements of the average V1 responses did not reliably distinguish between the paired and the unpaired stimuli. V1 cells also showed overall much weaker suppression between different directions of motion, and stronger responses to the flicker noise whose effect in MT are discussed below.

The inhibitory interaction found within subunits of MT receptive fields cannot therefore be present in their inputs from V1. Rather, the interaction must be a feature of the connectivity that builds up the MT receptive field. For example, each subunit might receive mutually inhibitory inputs from a set of V1 cells tuned to the same spatial location but to different directions.

Because of the relatively weak suppressive interaction between different directions of motion in V1, cells in this area behave rather like unidirectional motion energy detectors (Emerson et al. 1992) that signal the presence of moving components in a pattern (Movshon et al. 1986; Snowden et al. 1991; Qian and Andersen 1994, 1995) regardless of the presence of other components. This means that the pattern of responses in MT cells correlates better than that in V1 cells with our perception of motion transparency.

6.3 Why directional suppression in MT?

It is natural to ask why there should be strong directional suppression in the organization of MT cells. The effect on transparency must be a side-effect, since it is hard to see any functional reason why paired dot patterns should not be perceived as transparent. One possible explanation is that suppression may be useful for enhancing the directional selectivity of neurones (Barlow and Levick 1965; Mikami et al. 1986a; Rodman and Albright 1987). Indeed, a positive correlation was found between the cells' directional index and their suppression index. Another, and perhaps more important, functional role of the observed directional suppression could be motion noise reduction. Noise, for the motion system, is any change in the stimulus light intensity distribution that is not generated by coherent motion. For example, motion noise is generated when the overall intensity of the scene changes as when the light is turned on or off, or when an object or part of an object, appears or disappears. Such situations are quite common in a natural environment. When leaves and branches move in a forest, they also generate flicker noise due to the changes of reflectance with the orientation of their surfaces, and randomly occlusions or removals of occlusions. Such spatially and temporally uncorrelated noise has a relatively uniform spatiotemporal Fourier spectrum and, from the point of view of motion energy detectors (Adelson and Bergen 1985; Watson and Ahumada 1985), it contains equal amounts of "motion" signals in all directions. Mutual suppression among detectors tuned to different directions of motion can greatly reduce the responses of these detectors to noise.

Qian and Andersen (1994) tested this hypothesis by recording from MT cells using two types of noise pattern. These were identical to the paired and unpaired dot patterns, except that instead of moving, each dot stayed in a fixed random location during its lifetime (7 frames or 117 ms) and then was replotted at a new location. Thus, in the "paired" noise, the two dots of a pair are plotted stationary but transient in nearby locations. If directional suppression is important for noise reduction, then cells with stronger suppression should respond less to the noise patterns. An example of an MT cell's responses to the noise patterns is shown in figure 7, together with its responses to preferred, antipreferred, paired and unpaired motions. This cell shows very strong directional inhibition, and also the small noise responses that would be expected if the inhibition serves to suppress noise.

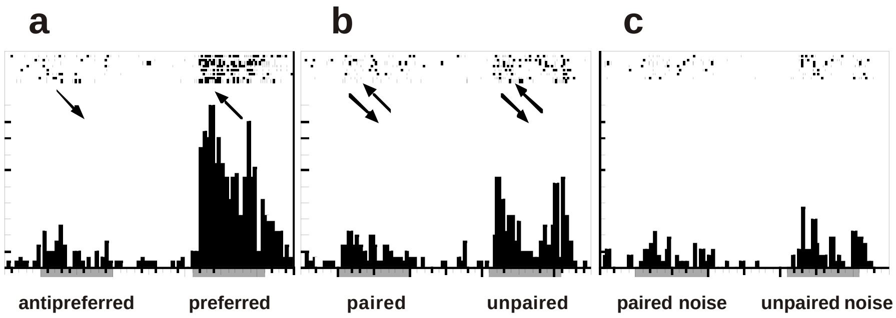

Fig. 7 Noise response of an MT cell. **a** Responses to single sets of random dots moving in its preferred and the antipreferred directions. **b** Responses to the paired and the unpaired dot patterns with one set of dots moving in the preferred direction and the other the antipreferred direction. **c** Responses to the paired and the unpaired noise patterns of flickering dots. One small vertical division represents 7.8 spikes/s. See figure 5 for an explanation of other details.

This relationship can be quantified over population in terms of two noise indices:

NI_p = (paired noise response)/(preferred response),

NI_{up} = (unpaired noise response)/(preferred response).

The background rate was again subtracted before calculation. Figure 8a shows that cells with a large index (i.e. strong noise response) tend to show a low suppression index (Spearman rank correlation = -0.75, $p < 0.0001$) in the case of the paired noise and motion patterns. The cells' responses to the paired dots and the paired noise patterns were similar. The correlation is weaker but still highly significant (Spearman rank correlation = -0.52, $p < 0.0001$) for the unpaired case (Fig. 8b). The stronger relationship in the former case is presumably because the paired dot patterns provide a better probe for the cells' local suppressive mechanisms.

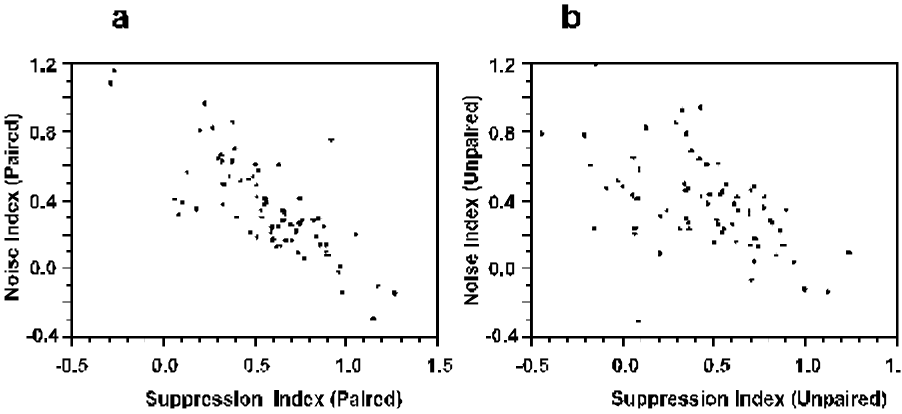

Fig. 8 Noise index verses suppression index for MT cells. **a** The paired case; **b** the unpaired case.

This discussion suggests that there is a conflict between noise reduction and the representation of multiple motions; local directional suppression in MT reduces responses to noise but makes it impossible for MT cells to represent more than one motion vector in a small area. Fortunately, there are other cues present in the real world that helps to minimize the conflict. For example, different objects in the natural environment will often be separated in depth and hence binocular disparity. Psychophysical experiments and computer simulations (Qian et al. 1994a, b) have demonstrated the effect of the disparity cue; paired dot patterns can become perceptually transparent when binocular disparity introduced between the two dots in each pair. These findings suggest that the directional suppression in MT should be strongest when the two different directions of motion are in the same disparity plane, and decrease with increasing separation in disparity (Qian and Andersen 1994). This prediction has been subsequently verified to be consistent with MT physiology (Bradley et al. 1995).

7. MT responses and perception

7.1 Attentional modulation

Subjectively, we can shift our attention between the two moving components in a transparent motion stimulus. Lankheet and Verstraten (1995) have shown that such switches of attention can affect the induction of motion aftereffects (MAEs); attention to one component significantly increased the MAE in the opposite direction. If area MT is involved in transparent motion perception, can MT activity be modulated by attention in an analogous way? Such modulation has indeed been demonstrated recently by Treue and Anderson (1996). They placed two dots, moving in opposite directions back and forth, in MT cells' receptive fields. At any time, one dot moved along the preferred direction of a cell and the other along the anti-preferred direction. They found that MT cells' activities were enhanced when the monkeys paid attention to the dot in the preferred direction, and were reduced when they attended the anti-preferred dot, even though under both conditions, the sensory inputs in the receptive fields were identical. The median attentional modulation in MT was found to be as large as 86%. Although these stimuli were quite different from transparent random dot patterns, it is plausible that the enhanced activity for the attended stimulus could be the basis for the increased adaptation found in Lankheet and Verstraten's study.

These effects raise the intriguing possibility that attention might act to reduce the suppression between different directions of motion, and consequently facilitate transparent motion perception, in the same way as does the visual cue of binocular disparity. In locally paired displays, it does not seem to be possible to perceive one of the superimposed motions by directing attention to it. Nonetheless, the hypothesis should be tested for MT cells with transparent random dot patterns in Treue and Maunsell's experimental paradigm.

7.2 Directional interactions at different angles

The physiological data discussed above was mostly concerned with random dot components moving in directions 180º apart. Snowden et al. (1991) also investigated interactions at smaller relative angles. They measured tuning curves of the suppressive effect on responses of MT cells to dots moving in the preferred direction, from another set of dots moving at different directions on different trials. These suppression tuning curves had a width of about 90°, similar to the width of ordinary directional tuning curves. A simple interpretation of the finding is that

MT cells show strong suppression for any two directions differing by more than 90º, and less suppression when the angular difference is smaller. Alternatively, Simoncelli and Heeger (1998) demonstrated that the result can be explained by a directionally isotropic mechanism of divisive inhibition (normalization). According to this view, with smaller angular differences the second set of dots excites the preferred direction more strongly, thus cancelling out a larger proportion of the uniform directional suppression and showing a reduced net suppressive effect.

These physiological data can be compared with psychophysical observations on locally paired stimuli with various directional differences. Qian et al. (1994a) noted informally that paired dot patterns with directional differences of 90º or 135º between the two dots in each pair were appeared neither transparent nor coherent, just as in the 180º-case. However, with a direction difference of only 45º, they observed a single coherent motion along the average direction of the two moving components. In all these cases, motion transparency was observed when the dots were not paired.

More recently, Curran and Braddick (1999, 2000) have tested subjects' ability to extract directional information from locally paired dot displays using a method similar to that illustrated in figure 1c. With direction differences of 60, 90, and 120º, they found that subjects could reliably judge a global direction of motion, although performance (s.d. in the range 10-18º) was somewhat impaired relative to directional judgment of a single set of dots. The judged direction followed closely the average of the two component directions. Measurements when the paired sets of dots had different speeds, and judgments of speed as well as direction, confirmed that the perceived velocity corresponded to the vector average of the components under a range of direction/speed conditions. These results imply that at no angle up to at least 120º is there a wholly suppressive interaction between locally paired directions. Of course, with equal speeds differing in direction by 180º, the vector average of the two velocities would be zero, a result indistinguishable from suppression of directional motion signals. It remains to be explored whether any effective combination of directions occurs in the range between 120-180º, although it should be realized that in this range the vector average velocity would become progressively smaller. The results are not incompatible with physiological models involving an interplay between summation and inhibitory effects, but suggest that a more detailed quantitative exploration of these models will be necessary to test them against psychophysical data.

In an earlier section on motion coherence, we mentioned the response of MT cells to coherently moving "signal" dots embedded in randomly moving noise (Newsome et al. 1989; Britten et al. 1992, 1993). The detection of the coherent signal in such displays must be influenced, and possibly may be aided, by the suppressive interactions between different directions of motion in MT. Simoncelli and Heeger (1998) found that the experimental dependence of MT responses on signal-to-noise ratio can be modelled by their directionally uniform, divisive inhibitory mechanism.

7.3 Bimodal versus unimodal population activity distribution

When a transparent random dot pattern containing two independent sets of dots moving in opposite directions is presented, the population activity of MT cells as a function of their preferred directions would be expected to form a bimodal distribution, with the two peaks approximately centred around the two directions in the stimulus. Although it is appealing to suppose that perception of transparent motion might correspond to a bimodal neuronal activity distribution in MT, such a relation cannot be expected to hold in general. Since the typical direction tuning width of MT cells is about 90º, transparent patterns containing two directions at an acute angle would most likely produce a merged unimodal distribution. This is indeed what Treue et al. (2000) found physiologically. They also found that the merging of two distributions in the transparent motion condition produced a wider distribution of activity in MT than in the unidirectional condition, and suggested that the brain could use the width of the population distribution as an indicator of the presence of transparent motion. Treue et al. further noted that a population distribution of a certain width can be generated either by two motion components, or by an appropriate mixture of more than two components. They then demonstrated through psychophysical experiments that observers could not tell the stimuli with different mixtures apart, and reported the minimum number of directions consistent with the population activity.

The experiment of Treue et al. illustrates the problem of parsing a distribution of activation into the components arising from different moving entities in the visual environment. This parsing operation appears to be flexible and may be influenced by the visual task the subject is performing. So, for example, in the experiments of Watamanuik et al. (1989) subjects could use broad directional distributions to define a single global direction with high precision. However, Zohary et al. (1996) showed that with asymmetrical distributions, the same distribution could either yield a global direction judgement or alternatively be parsed into a dominant direction corresponding to the peak plus a separate “background”. Another aspect of the parsing operation is illustrated by the results of Smith et al. (1999). They tested judgements based on motion transparency for pairs of broad directional distributions with a gap between them. The size of gap that yielded effective transparency was not fixed, but depended on the overall range of directions represented. They concluded that parsing the internal representation of directions into two motions did not depend only on local features such as peaks and gaps, in the directional distribution. Rather, the parsing operation took account of the distribution as a whole. The overall range of directions of distributions shown in Smith et al. (1999) was wider than that in the experiment of Treue et al. (2000). Given this difference, it is not yet clear whether the parsing rules proposed by Treue et al., deriving transparency from a unimodal distribution, are compatible with those suggested by Smith et al., which imply bimodal distributions. It is also possible that the way the distributions were parsed in each case depend on the specific visual tasks which subjects were performing.

8. Transparency in plaid patterns

8.1 MT responses to transparent and coherent plaids

A plaid pattern consists of two superimposed gratings with different orientations and moving in different directions. The perceptual problem posed by such displays was first analysed by Hans Wallach (Wuerger et al. 1996), and following the work of Adelson and Movshon (1982) they have been widely used for psychophysical and physiological investigations of how the brain combines local motion measurements. Depending on a variety of factors (such as the differences in contrasts, spatial frequencies, and binocular disparities, and the relative angle between the two gratings), the two component gratings in a plaid can either appear to slide across each other transparently, or move together coherently in a single direction. An early physiological study of plaid coherence was provided by Movshon et al. (1986) who demonstrated that V1 cells responded to the individual motion components, but that a subpopulation of MT cells appeared to combine these components and respond to the single coherent direction, as perceived.

More recently, Stoner et al. (1990) introduced a clever manipulation of plaid coherence versus transparency, which they subsequently employed to study the neural correlates of both coherence and transparency in area MT (Stoner and Albright 1992). Instead of adding two moving sinusoidal gratings as done by Adelson and Movshon (1982), Stoner et al. constructed plaids by superimposing two drifting square wave gratings, and discovered that the perception depended critically on the luminance of the diamond-shaped intersections between the two gratings. Specifically, they found that when the intersection luminance relative to that of the gratings satisfied the rule of static transparency, a surface segmentation cue, the pattern tended to appear as two independent gratings sliding across each other transparently. Otherwise, a single coherently moving plaid was more likely to be perceived. An alternative explanation of these observations could come from the Fourier energy in the coherent pattern motion direction provided by the intersections (Simoncelli and Heeger 1998; Wilson and Kim 1994). However, Stoner and Albright (1996) later provided evidence against the Fourier interpretation by introducing a novel plaid stimuli in which surface segmentation cues could be varied independently of changes in the Fourier energy distribution.

Stoner and Albright (1992) recorded responses of MT cells to plaids whose intersection luminance was manipulated so as to generate either transparent or coherent percepts. They compared the cells' directional tuning curves under the two conditions, and found that neurones were more tuned to the component directions in the condition yielding transparency, and to the single combined pattern direction under the coherent condition. These results support those with random dot patterns, in suggesting a neural substrate for perceptual transparency and coherence in area MT.

8.2 Random dot stimuli compared with plaid patterns

Although both random dot stimuli and plaid patterns have been used in studies of motion transparency, they have some quite different properties. They can be viewed, respectively, as highly simplified versions of two different real world situations of motion transparency: (1) partial and patchy occlusions between moving objects (e.g., a tiger moving behind waving bushes), and (2) overlapping motions involving semi-transparent surfaces (e.g., an object seen through the window of a moving car). They pose different problems for the visual system; the ambiguities which arise in plaids from the coupling of motion and orientation are absent in a dot pattern, where the modulation of luminance is fully two-dimensional.

There are different reasons for using each kind of stimulus in electrophysiological and psychophysical research. Random dot patterns have a broad spatiotemporal frequency spectrum and therefore excite most V1 and MT cells effectively. Also the small sizes of dots make them suitable for investigating the spatial ranges of interaction through the pairing manipulation. On the other hand, the narrow spatiotemporal spectra of plaids render them less convenient in physiological experiments because most cortical cells are frequency tuned, and cells do not respond well unless the grating frequency matches their preferred frequencies. In addition, cells in different areas often have different tuning characteristics, which makes comparison across different areas, for example V1 and MT, somewhat difficult. However, there are important questions that cannot be investigated with random dot patterns. Since the dots moving in different directions rarely overlap, the luminance of intersections has negligible effect. In contrast, plaid patterns have large intersection regions whose luminance can be readily manipulated to study the influence of static transparency cues. In addition, the directional ambiguity of each component grating in plaids makes them suitable for investigating how the brain solves the “aperture problem”. However, when these issues are not the focus of investigation, the presence of intersection regions and corners in plaids may complicate the interpretation of results. Thus the two types of patterns appear to complement each well other in research on multiple motion directions and transparency.

Despite the differences between random dot stimuli and plaid patterns, experimental investigations with both types of patterns point to some common properties and mechanisms for motion transparency. Psychophysical experiments indicate that surface segmentation cues such as binocular disparity, spatial frequency, and static transparency enhance transparent motion perception. Physiological recordings on both types of patterns agree that MT activities are well correlated with our perception of motion transparency, or the lack of it (Movshon et al. 1986; Snowden et al. 1991; Stoner and Albright 1992; Qian and Andersen 1994), while V1 cells, in contrast, appear to be more concerned with representing individual components in a motion pattern through motion energy computations, regardless of the pattern’s perceptual appearance (Movshon et al. 1986; Snowden et al. 1991; Qian and Andersen 1994, 1995).

9. Computational considerations

9.1 The role of suppressive interactions

Qian et al. (1994b) demonstrated through computer simulations that a motion energy computation, followed by disparity and spatial frequency specific suppression among different directions of motion can indeed explain the perceptual difference of the transparent and non-transparent displays used in their psychophysical experiments (Qian et al. 1994a). According to the discussion above, these two stages approximately correspond to physiological responses seen in areas V1 and MT respectively. Qian et al. (1994b) found that the non-transparent displays, such as paired dot patterns, generate relatively weak responses in the simulations (opponent or normalized energies) at the suppression stage due to strong cancellations of locally balanced motion signals. In fact, these responses are not significantly higher than those generated by flicker patterns. On the other hand, the perceptually transparent displays, such as unpaired dot patterns, or paired dots with disparity, generate much stronger responses along more than one direction of motion at the suppression stage due to the presence of unbalanced motion signals. These responses in different directions are located either in different but mixed small areas, as in the unpaired dot pattern, or in different disparity or spatial frequency channels over the same spatial regions, as in the paired dot pattern with binocular disparity. A later stage could integrate these responses in different directions separately to form two overlapping transparent surfaces. Simulation examples for the unpaired, paired, and flicker patterns, and for the paired dots pattern with disparity cue are shown in figure 9.

The idea that transparent motion perception can be viewed as detection of locally unbalanced motion signals in different spatial locations, or in different disparity (or other) channels, is also consistent with a recent selection model by Nowlan and Sejnowski (1995) who proposed that motion computation should be based on those regions in the visual field where the velocity estimates are most reliable. For a transparent motion display containing two directions of motion, the most reliable regions for motion estimation are clearly those with strong unbalanced motion signals in one direction or the other. For the non-transparent paired dot patterns (without disparity cues), such reliable regions do not exist.

The suppression among different directions of motion makes it impossible for a stimulus to generate strong responses along more than one direction of motion in each small spatial area at the opponent stage when there are no other cues in the stimulus, such as disparity or spatial frequency. The minimum size of the small area is determined by the size of the front-end filters. In this regard, the suppression stage is rather like the pooling or regularization step commonly used in machine vision systems (Horn and Schunk 1981; Lucas and Kanade 1981; Hildreth 1984; Heeger 1987; Poggio et al. 1988; Grzywacz and Yuille 1990). Such a step provides a means to solve the aperture problem and to average out noise in the initial measurements. At the same time, it prevents those models from having more than one velocity estimate over each area covered by the pooling operator. In this connection, it is interesting to note that some versions of the pooling procedures for combining local gradient constraints are equivalent to a mixture of subtractive and divisive types of suppression (Simoncelli 1993). The agreement between our simulations and the psychophysical observations implies that machine vision systems can be made more consistent with transparent motion perception if the pooling operation in these systems is restricted to small areas and to each frequency and disparity channel. We suggest that the difficulty most machine vision systems have with motion transparency can be partly attributed to the fact that these systems typically apply pooling operations over a relatively large region, and that they usually do not explore other cues such as disparity and spatial frequency to restrict the scope of pooling.

We identified the suppression stage with the pooling operation in machine vision systems above. Since the function of the pooling operation is to solve the aperture problem and to average out noise, we suggest that directional suppression in MT has similar functions. Indeed unidirectional energy detectors like V1 cells suffer the aperture problem (Movshon et al. 1986). Also, V1 cells are very responsive to dynamic noise patterns made of flickering dots (Qian and Andersen 1994). Therefore, directional V1 cells are not adequate for representing transparent motion, and in fact, as a population, they could not reliably distinguish transparent patterns from non-transparent ones. The suppression in MT could help to solve these problems, just as the pooling operations do in machine vision systems. The noise response of MT cells is indeed much lower than that of V1 cells (Qian and Andersen 1994). There is also evidence that the human visual system may solve the aperture problem by averaging local motion measurements (Ferrera and Wilson 1990, 1991; Rubin and Hochstein 1993). We have discussed above how the outcome of the suppressive interaction appears to be a vector averaging, at least for a certain range of stimuli. A negative effect of the suppression is the reduced sensitivity to transparent motion. This problem is minimized, however, by applying suppression locally and by restricting it within each disparity and spatial frequency channel, since multiple motions in the real world are usually not precisely balanced in each local area, and different objects tend to have different disparity and spatial frequency distributions. While the inhibition among the cells within each disparity and spatial frequency channel could help to combine V1 outputs into a single motion signal at each location in order to solve the aperture problem and to reduce noise, different disparity and spatial frequency channels make it possible to represent multiple motions at the same spatial location.

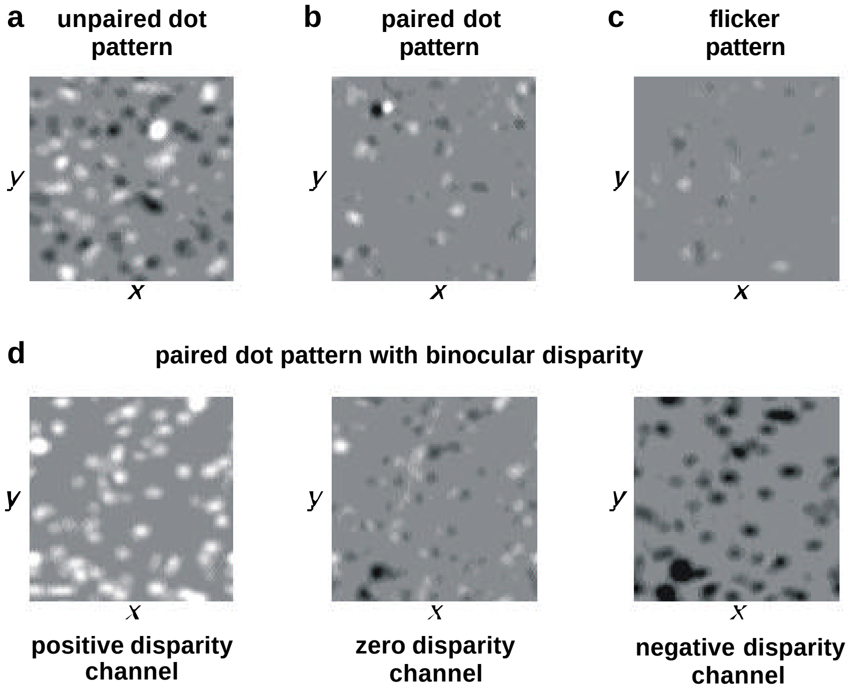

Fig. 9 Computer simulations of opponent motion energies at the motion suppression stage for **a** unpaired dot pattern, **b** paired dot pattern, **c** flicker noise pattern, and **d** paired dot pattern with binocular disparity. White and black colors code for rightward and leftward opponent energies, respectively, and grey color indicates little motion energy. In (d), opponent energies in three different disparity channels are shown.

9.2 Representation of motion, objects, and scenes

Since local suppressive mechanisms mean that the representation of motion is locally single valued, the perception of transparent motions must depend on linking spatially interdigitated signals into global entities. This can be considered as a special case of the wider process of perceptual grouping. The global entities which results from a grouping process arises from signals that have their own locations specified in a retinotopic matrix of signals. However, the global entities themselves do not have a location that is readily specified in this way, especially in the case of transparency which implies that two objects occupy overlapping locations. Therefore, the representation of transparent objects, and indeed, objects in general, cannot easily be understood as based on the kind of retinotopic array that is familiar from the early stages of cortical processing. The representation, rather than being indexed by spatial locations, must be indexed by objects – perhaps the kind of representation that has been referred to as an “object file” (Triesman 1988). Nonetheless, the information in such a representation must allow it to be associated with (distributed) spatial locations. The way in which such an representation can be implemented in a neural system is a challenge for future accounts of neural computation. Motion transparency may provide a useful route to attack this important general problem.

Acknowledgments

The work from OJB’s laboratory described here was supported by project grant G9509331 and programme grant G7908507 from the Medical Research Council. NQ was supported by a Sloan Research Fellowship and an NIH grant (MH54125) during the preparation of this chapter.

References

Adelson EH, Bergen JR (1985) Spatiotemporal energy models for the perception of motion. J Opt Soc Am A 2: 284-299

Adelson EH, Movshon JA (1982) Phenomenal coherence of moving visual patterns. Nature 300: 523-525

Andersen RA (1997) Neural mechanisms of visual motion perception in primates. Neuron 18: 865-872

Baker CL, Hess RF, Zihl J (1991) Residual motion perception in a ‘motion-blind’ patient, assessed with limited-lifetime random dot stimuli. J Neurosci, 11: 454-481

Barlow HB, Levick WR, (1965) The mechanism of directionally selective units in the rabbit's retina. J Physiol 178: 477-504

Braddick O (1993) Segmentation versus integration in visual motion processing. Trends in Neurosci 16: 263-268

Braddick O (1997) Local and global representations of velocity: transparency, opponency, and global direction perception. Perception 26: 995-1010

Braddick OJ, Hartley T, O’Brien J, Atkinson J, Wattam-Bell J, Turner R (1998) Brain areas differentially activated by coherent visual motion and dynamic noise. NeuroImage 7: S322

Bradley DC, Qian N, Andersen RA (1995) Integration of motion and stereopsis in middle temporal cortical area of macaques. Nature 373: 609-611

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci 12: 4745-4765

Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. Visual Neurosci 10: 1157-1169

Curran W, Braddick OJ (1999) Perceived motion direction and speed of locally balanced stimuli. Perception 28 (suppl): 49

Curran W, Braddick OJ (2000) Speed and direction of locally-paired dot patterns. Vision Res: in press

Downing CJ, Movshon AJ (1989) Spatial and temporal summation in the detection of motion in stochastic random dot displays. Invest Ophthalmol Vis Sci (Suppl) 30: 72

Emerson RC, Bergen JR, Adelson EH (1992) Directionally selective complex cells and the computation of motion energy in cat visual cortex. Vision Res 32: 203-218

Ferrera VP, Wilson HR (1990) Perceived direction of moving two-dimensional patterns. Vision Res 30: 273-287

Ferrera VP, Wilson HR (1991) Perceived speed of moving two-dimensional patterns. Vision Research 31: 877-894

Friston KJ, Holmes AP, Worsley KJ, Poline JB, Frith CD, Frackowiak RSJ (1995) Statistical parametric maps in functional imaging: a general approach. Human Brain Mapp 2: 189-210

Grzywacz NM, Yuille NL (1990) A model for the estimate of local image velocity by cells in the visual cortex. Proc R Soc Lond A 239: 129-161

Heeger DJ (1987) Model for the extraction of image flow. J Opt Soc Am A 4: 1455-1471

Hildreth EC (1984) Computations underlying the measurement of visual motion. Art Intell 23: 309-355

Horn BKP, Schunck B (1981) Determining optical flow. Art Intell 17: 185-203

Lankheet MJM, Verstraten FJ (1995) Attentional modulation of adaptation to two-component transparent motion. Vision Res 35: 1401-1412

Lucas BD, Kanade T (1981) An iterative image registration technique with an application to stereo vision. Proc 7th Internat Joint Conf Art Intell, Vancouver: 674-679

Mikami A, Newsome WT, Wurtz RH (1986a) Motion selectivity in macaque visual cortex. I. Mechanisms of direction and speed selectivity in extrastriate area MT. J Neurophysiol 55: 1308-1327

Mikami A, Newsome WT, Wurtz RH (1986b) Motion selectivity in macaque visual cortex. II. Spatiotemporal range of directional interactions in MT and V1. J Neurophysiol 55: 1328-1339

Movshon JA, Adelson EH, Gizzi MS, Newsome WT (1986) The analysis of moving visual patterns. In: Chagas C, Gattass R, Gross C (eds) Experimental Brain Research Supplementum II: Pattern recognition mechanisms. Springer, New York, pp 117-151

Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. Nature 341: 52-54

Newsome WT, Paré EB (1988) A selective impairment of motion processing following lesions of the middle temporal area (MT). J Neurosci 8: 2201-2211

Nowlan SJ, Sejnowski TJ (1995) A selection model for motion processing in area MT of primates. J Neurosci 15: 1195-1214

Poggio T, Torre V, Koch C (1988) Computational vision and regularization theory. Nature 317: 314-319

Qian N, Andersen RA (1994) Transparent motion perception as detection of unbalanced motion signals II: Physiology. J Neurosci 14: 7367-7380

Qian N, Andersen RA (1995) V1 responses to transparent and non-transparent motions. Exp Brain Res 103: 41-50

Qian N, Andersen RA, Adelson EH (1994a) Transparent motion perception as detection of unbalanced motion signals I: Psychophysics. J Neurosci 14: 7357-7366

Qian N, Andersen RA, Adelson EH (1994b) Transparent motion perception as detection of unbalanced motion signals III: Modeling. J Neurosci 14: 7381-7392

Recanzone GH, Wurtz RH, Schwarz UC (1997) Responses of MT and MST neurons to one and two moving objects in the receptive field. J Neurophysiol 78: 2904-2915

Rodman HR, Albright TD (1987) Coding of visual stimulus velocity in area MT of the macaque. Vision Res 27: 2035-2048

Rodman HR, Albright TD (1989) Single-unit analysis of pattern-motion selective properties in the middle temporal visual area (MT). Exp Brain Res 75: 53-64

Rubin N, Hochstein S (1993) Isolating the effect of one-dimensional motion signals on the perceived direction of moving two-dimensional objects. Vision Res 33: 1385-1396

Scase MO, Braddick OJ, Raymond JE (1996) What is noise for the motion system? Vision Res 16: 2579-2586

Shadlen MN, Newsome WT, Zohary E, Britten KH (1993) Integration of local motion signals in area MT. Soc Neurosci Abstract 19: 1282

Simoncelli EP (1993) Distributed analysis and representation of visual motion. Ph.D. Thesis, MIT, Cambridge

Simoncelli EP, Heeger DJ (1998) A model of neuronal responses in visual area MT. Vision Res 38: 743-761

Smith AT, Curran W, Braddick OJ (1999) What motion distributions yield global transparency and spatial segmentation? Vision Res 39: 1121-1132

Snowden RJ (1989) Motions in orthogonal directions are mutually suppressive. J Opt Soc Am A 7: 1096-1101

Snowden RJ, Treue S, Erickson RE, Andersen RA (1991) The response of area MT and V1 neurons to transparent motion. J Neurosci 11: 2768-2785

Stoner GR, Albright TD (1992) Neural correlates of perceptual motion coherence. Nature 358: 412-414

Stoner GR, Albright TD (1996) The interpretation of visual motion: Evidence for surface segmentation mechanisms. Vision Res 36: 1291-1310

Stoner GR, Albright TD, Ramachandran S (1990) Transparency and coherence in human motion perception. Nature 344: 153-155

Tootell RBH, Mendola JD, Hadjikhani NK, Ledden PJ, Liu AK, Reppas JB, Sereno MI, Dale AM (1997) Functional analysis of V3A and related areas in human visual cortex. J Neurosci 17: 7060-7078

Treue S, Maunsell JHR (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. Nature 382: 539-541

Treue S, Hol K, Rauber HJ (2000) Seeing multiple directions of motion – physiology and psychophysics. Nature Neurosci 3: 270-276

Triesman, A (1988) Features and objects. Quart J Exp Psychol A 40: 201-237

Verstraten FA, Fredericksen RE, van Wesel RJ, Boulton JC, van de Grind WA (1996) Directional motion sensitivity under transparent motion conditions. Vision Res 36: 2333-2336

Watamaniuk SNJ, Sekuler R, Williams DW (1989) Direction perception in complex dynamic displays - the integration of direction information. Vision Res 29: 47-59

Watson AB, Ahumada AJ (1985) Model of human visual-motion sensing. J Opt Soc Am A 2: 322-342

Watt RJ, Andrews DP (1981) APE: Adaptive probit estimation of psychometric functions. Curr Psychol Rev 1: 205-214

Williams DW, Sekuler R (1984) Coherent global motion percepts from stochastic local motions. Vision Res 24: 55-62

Wilson HR, Kim J (1994) A model for motion coherence and transparency. Visual Neurosci 11: 1205-1220

Wishart KA, Braddick O (1997a) Performance-based measures of motion transparency. Invest Opthalmol Vis Sci 38: S75

Wishart KA, Braddick OJ (1997b) Performance based measures of transparency in locally-balanced motions. Perception 26 (Suppl): 86

Wuerger S, Shapley R, Rubin N (1996) “On the visually perceived direction of motion” by Hans Wallach: 60 years later. Perception. 25: 1317-1368

Zohary E, Scase MO, Braddick OJ (1996) Integration across directions in dynamic random-dot displays: vector summation or winner-take-all? Vision Res 16: 2321-2331