Visual Neuroscience (1993), **10**, 1157–1169. Printed in the USA.

Responses of neurons in macaque MT to stochastic motion signals

KENNETH H. BRITTEN,[1] MICHAEL N. SHADLEN,[1]
WILLIAM T. NEWSOME,[1] AND J. ANTHONY MOVSHON[2]

[1] Department of Neurobiology, Stanford University School of Medicine, Stanford
[2] Howard Hughes Medical Institute, Center for Neural Science, and Department of Psychology, New York University, New York

(RECEIVED October 26, 1992; ACCEPTED May 19, 1993)

Abstract

Dynamic random-dot stimuli have been widely used to explore central mechanisms of motion processing. We have measured the responses of neurons in area MT of the alert monkey while we varied the strength and direction of the motion signal in such displays. The strength of motion is controlled by the proportion of spatiotemporally correlated dots, which we term the *correlation* of the stimulus. For many MT cells, responses varied approximately linearly with stimulus correlation. When they occurred, nonlinearities were equally likely to be either positively or negatively accelerated. We also explored the relationship between response magnitude and response variance for these cells and found, in general agreement with other investigators, that this relationship conforms to a power law with an exponent slightly greater than 1. The variance of the cells' discharge is little influenced by the trial-to-trial fluctuations inherent in our stochastic display, and is therefore likely to be of neural origin. Linear responses to these stochastic motion stimuli are predicted by simple, low-level motion models incorporating sensors having relatively broad spatial and temporal frequency tuning.

Keywords: Visual cortex, Motion processing, Motion models, MT, Direction selectivity, Direction discrimination

Introduction

It is now widely recognized that a subset of areas in primate visual cortex is specialized for the analysis of visual motion. This "motion pathway" has become a particularly fruitful model system for understanding how sensory information is processed and used at higher levels of the cerebral cortex. The pathway was first discovered by Zeki and colleagues who showed that a large majority of neurons in the depths of the superior temporal sulcus are selective for the direction of motion of visual stimuli (Dubner & Zeki, 1971; Zeki, 1974). More recently, the various areas that contribute to the motion pathway have been identified and mapped, and the anatomical connections between the areas described (e.g. Maunsell & Van Essen, 1983*a*; Desimone & Ungerleider, 1986). In parallel with these organizational studies, several physiological analyses addressed the issues of representation and information processing in the motion pathway, yielding novel insights into the mechanisms that encode the motion of complex objects, optic flowfields, and figure-ground relationships (for reviews, see Nakayama, 1985; Maunsell & Newsome, 1987).

Reprint requests to: Kenneth H. Britten, Department of Neurobiology, Stanford University School of Medicine, Stanford, CA 94305, USA.

For several years, our own efforts have been directed towards establishing an empirical link between neuronal activity in the motion pathway and the perception of visual motion. Thus far, we have concentrated on the middle temporal visual area (MT, or V5) because this is the first area on the pathway in which a large majority of neurons is directionally selective. Our basic strategy was to record and manipulate the activity of neurons in MT while trained rhesus monkeys performed a near-threshold direction discrimination task. These experiments revealed that MT neurons carry directional signals of sufficient precision to account for the psychophysical sensitivity of behaving monkeys (Britten et al., 1992), and that electrical microstimulation of MT can influence choices in the discrimination task in a directionally specific manner (Salzman et al., 1992). Additional experiments showed that lesions of MT, at least initially, impaired performance on the direction discrimination task while leaving contrast thresholds unchanged (Newsome & Pare, 1988). Taken together, these findings strongly support the idea that directional signals in the motion pathway, and specifically in area MT, contribute directly to the perception of motion.

The visual stimuli employed in these studies were stochastic random-dot patterns in which the experimenter controlled the strength of the motion signal by specifying the percentage of dots undergoing coherent rather than random motion. These stimuli, which are related to but distinct from previously used

random-dot stimuli (Morgan & Ward, 1980; Williams & Sekuler, 1984), have also proven useful for psychophysical studies of motion vision in humans (Downing & Movshon, 1989; Hiris & Blake, 1992) and for clinical analyses of neurologically impaired patients (Hess et al., 1989; Vaina et al., 1990; Braun et al., 1991). To understand how these stimuli are represented within the central nervous system, we have examined the influence of stochastic random-dot patterns on the activity of visual cortical neurons. We provided a partial description of these results in our prior comparison of neuronal and behavioral sensitivity to the stochastic motion signals (Newsome et al., 1989; Britten et al., 1992), but we now provide a more thorough analysis by describing quantitatively the response functions relating neuronal activity to the strength of the motion signal. The analysis reveals that most MT neurons respond approximately linearly to changes in the strength of the motion signal and that the variance of the responses depends approximately linearly on their magnitude. We also report the results of simulations suggesting that these response properties are to be expected from relatively simple motion filters like those which have been employed in recent models of early motion processing (e.g. Adelson & Bergen, 1985; Watson & Ahumada, 1985). Thus, the responses of MT cells to more complex stimuli can be well characterized by surprisingly simple models.

Methods

The methods used for these experiments have been described in detail elsewhere (Britten et al., 1992). Three rhesus monkeys (*Macaca mulatta*; two male and one female) were implanted with stainless-steel head holders and scleral search coils for monitoring eye movements (Judge et al., 1980), and then trained for several months on a two-choice direction-discrimination task. Immediately prior to recording, a stainless-steel cylinder was implanted over occipital cortex, allowing a posterior approach to area MT. All surgical and animal care methods conformed to NIH guidelines for the care and use of laboratory animals.

Visual stimuli

We employed a stochastic, variable-strength motion display (Morgan & Ward, 1980; Britten et al., 1992) in which random dots were continuously plotted at 6.67 kHz on the face of a large CRT (P4 phosphor), resulting in an average density of 16.7 dots$\cdot s^{-1} \cdot deg^{-2}$. The dots were approximately 0.1 deg in diameter and were of high contrast, but their low space-time density produced an average luminance of only 0.2 cd/m^2, against a background luminance of 0.01 cd/m^2. A certain proportion of these dots carried a specified motion signal, because dots were replotted with a fixed spatial offset after a fixed temporal interval. This proportion, which we express as a percentage and term the *correlation* of the stimulus, controls the strength of the motion signal without affecting the luminance, contrast, or average spatial and temporal structure of the stimulus. Dots not so replotted were randomly repositioned, producing dynamic masking noise. Thus, the correlation determines the signal-to-noise ratio between the correlated and random components of the motion in the stimulus. The speed and direction of the correlated motion were controlled by adjusting the magnitude and direction of the spatial offset of the signal dots; in these experiments, the temporal interval was held fixed at 45 ms, a value shown previously to be near optimal for psychophysical performance in the range of eccentricities where these data were gathered (Newsome & Pare, 1988). We have described these stimuli in detail in a previous publication (Britten et al., 1992).

Recording

MT was located and mapped using transdural, glass-coated platinum–iridium electrodes. A guide tube support grid with holes spaced 1 mm apart (Crist et al., 1988) was then secured inside the recording cylinder, and stainless-steel guide tubes were inserted to within 5 mm of area MT. We inserted tungsten microelectrodes through these guide tubes and recorded the responses of single neurons using standard extracellular methods. Eye-movement monitoring, behavioral reinforcement, and unit data collection were carried out by a PDP 11/73 computer running software developed at NIH for these purposes (Hays et al., 1982). A second PDP 11/73, under control of the first, generated and presented the dynamic random-dot stimuli.

In each daily session, we searched for units while the monkey maintained fixation on a spot of light from a laser or projection LED. We slowly advanced the electrode and isolated action potentials from single units while activating the local background activity with effective search stimuli. Once an MT cell was isolated, its receptive field (RF) was mapped using moving bars, and a circular aperture was fitted in front of the screen so that the random-dot stimulus just filled the RF. In cases where the initial receptive-field map was uncertain because of poor responses to the moving bar stimuli, we adjusted the size and location of the aperture to maximize the neuron's response to random-dot stimuli. The neuron's preferred direction and speed of motion were then determined using random-dot stimuli. If the neuron was reasonably responsive and directional, the preferred-null axis was determined by listening to the neuronal discharge over an audiomonitor, but if the neuron was less responsive or less directional, this assignment was made on the basis of a computer-controlled direction series (see below). The neuron's optimal speed was also estimated "by ear," and ranged from 0.4–28 deg/s.

The raw data used in this paper were the same as those employed for the analysis of neuronal sensitivities in a prior publication. We recorded from 216 MT neurons that responded to random-dot patterns in a directionally selective manner (Britten et al., 1992). The primary data were obtained from runs in which the stimuli were presented in pseudo-random sequence at the preferred speed in the preferred and null directions over a range of correlation levels spanning neuronal and behavioral threshold (Britten et al., 1992, "combined threshold series"). Most experiments contained 5–7 different correlation levels, differing from one another by a factor of 2. In all but a few cases, 0% correlation trials (pure noise) were also interleaved. The monkey was required to make a direction judgement subsequent to the stimulus presentation for each trial in a combined threshold series. Repeat blocks of data were collected as long as isolation of the cell could be maintained. Accordingly, the number of trials per condition ranged from 7–120, with an average of about 30. For 168 of the 216 cells, we also presented a direction series while the monkey was fixating to measure quantitatively the direction tuning to random-dot fields of 100% correlation. For these experiments, we presented stimuli at the optimal speed in eight directions of motion differing by 45 deg, once again randomly interleaved.

Histology

At the time of data collection, MT was identified on the basis of physiological landmarks; histological confirmation of the location of recordings was obtained for two of the three monkeys following the conclusion of experiments. The third monkey is being used in an ongoing study. Two monkeys were killed with an overdose of barbiturate, and then perfused transcardially with 0.9% saline followed by 10% buffered formalin. Each brain was removed and allowed to sink in 30% sucrose, blocked, and the relevant regions of cortex were cut frozen in the coronal plane at 48-μm thickness. Sections were mounted on slides, dried, and alternate series (each consisting of every tenth section) were stained for cell bodies with cresyl violet or for myelin with a reduced silver method (Gallyas, 1979). Although individual penetrations were impossible to recover, the region from which we recorded was evident from guide tube scars posterior to the superior temporal sulcus (STS) and from electrode track damage approaching the posterior bank of the STS. In both monkeys, the region thus identified corresponded well to the heavily myelinated region on the posterior bank of the STS, and therefore we can assert with confidence that the recordings were made from area MT in both of these animals. This procedure also serves to validate our physiological identification of area MT, so it is likely that the same will be true for the third animal as well.

Results

Response as a function of correlation

Fig. 1 illustrates the responses of a typical neuron to the stochastic stimuli employed in this study. The left-hand column of raster and spike-density plots illustrates responses to increasing stimulus correlation for motion in the neuron's preferred direction. The right-hand column shows responses to motion in the null direction. At 6.4% correlation, the neuron responded equivalently to the two directions. As stimulus correlation increased, responses to preferred direction motion increased monotonically while responses to null direction motion decreased. At the highest correlation, therefore, the neuron was strongly direction selective. Inspection of these data also reveals that the firing rate of the neuron, following a brief transient burst, remained approximately constant during the 2-s stimulus presentation interval. This response pattern was typical of the neurons in our sample, and few neurons showed any other temporal structure to their responses. This suggests that the simple spike count provides a good estimate of the visual signal encoded by these neurons, and we therefore employ integrated spike counts as our measure of neuronal response.

A more complete description of neuronal responses to our stimulus set is provided by the "correlation-response functions" shown in Fig. 2 for six MT neurons. For each neuron, the mean response to multiple trials is shown as a function of stimulus strength in the neuron's preferred (filled symbols) and null (open symbols) directions. The average maintained activity measured during the intertrial period is indicated by the horizontal dotted line. The data confirm the observation from Fig. 1 that responses increased monotonically with stimulus correlation in the preferred direction, and decreased with increasing correlation in the null direction. The figure shows, however, that considerable variability exists in the shape of the correlation-

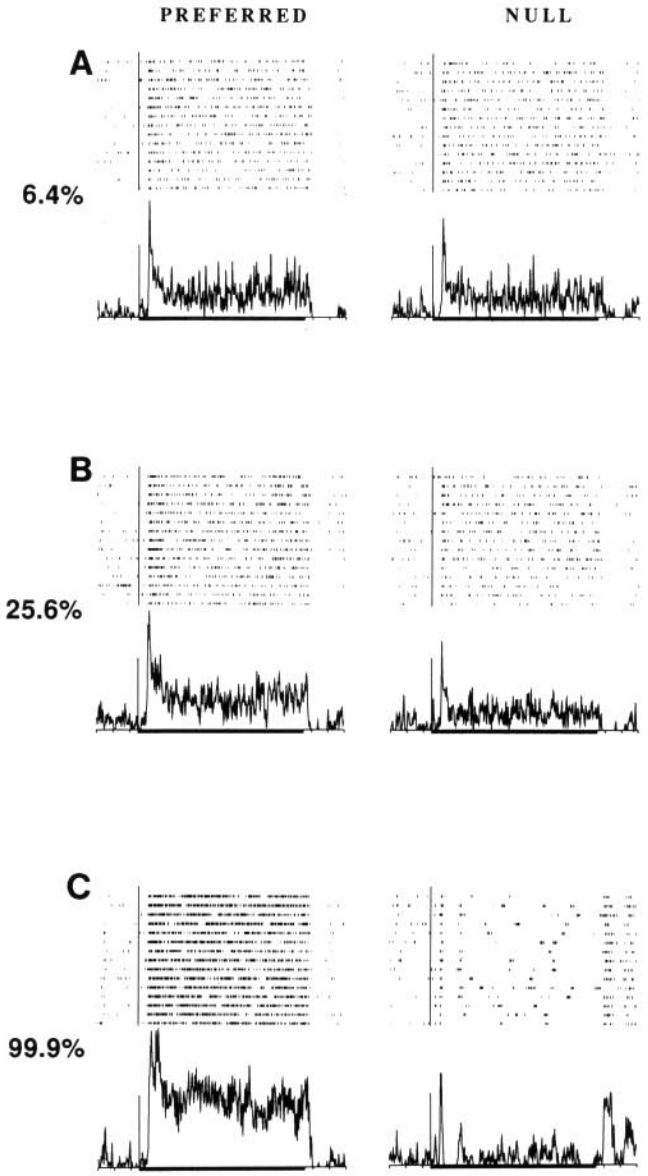

Fig. 1. Raster diagrams and spike-density profiles for the responses of a single MT neuron to stimuli presented in the preferred and null directions at three correlation levels. Each panel represents the responses to 15 presentations (a typical value for a single block of trials), which are individually plotted on the rasters forming the upper part of each panel. The spike-density plots result from averaging all of the trials in the corresponding raster plots. The vertical calibration bar at stimulus onset corresponds to 100 spikes/s; and the solid bar beneath each axis, corresponding to the entire duration of the visual stimulus, is 2 s in duration.

response functions. In Figs. 2A and 2B, for example, the relationship between firing rate and correlation is nearly linear, while the relationship is highly nonlinear in Figs. 2E and 2F. The correlation-response functions in Figs. 2C and 2D are intermediate, having relatively small nonlinear components.

To assess the linearity of these response functions, we compared maximum-likelihood fits of first- and second-order polynomials for each of the 216 neurons in our sample. The first-order polynomial fit was the best-fitting straight line; the second-order fit incorporated a quadratic term. The best-fitting

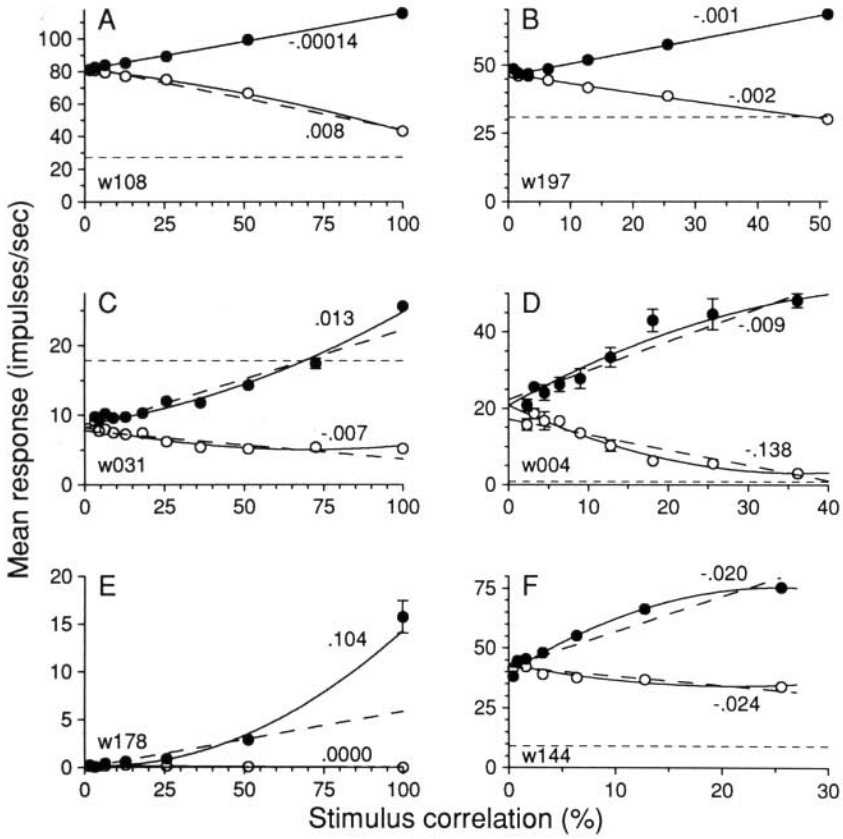

Fig. 2. Response functions indicating mean response as a function of correlation for six MT cells. Solid circles show responses to preferred direction stimuli; and open circles show responses to null direction stimuli. Error bars indicate standard errors of the mean. The curves through each set of points are the best-fitting linear (dashed) and quadratic (solid) functions for each. The horizontal dotted line in each panel represents the average maintained activity during the intertrial interval. The number next to each response function is the ratio of the quadratic term to the linear term for the best-fit quadratic function.

linear and quadratic functions are illustrated in Fig. 2 as the dashed and solid lines, respectively. For the linear responses shown in Figs. 2A and 2B, the two functions are almost exactly superimposed. Progressively larger second-order contributions to the fits are observed in Figs. 2C–2F. For the highly nonlinear response functions in Figs. 2E and 2F, the quadratic fits depart considerably from the linear fits, and provide a much better account of the data. To compare the contributions of linear and quadratic terms across the entire set of response functions, we calculated the ratio of the quadratic coefficient to the linear coefficient for each second-order fit. The absolute value of this ratio indicates the relative contribution of the quadratic term, and the sign of the ratio indicates whether the nonlinearity is positively or negatively accelerating with respect to the linear component. Thus, for example, the positively accelerating nonlinearity in the preferred direction response function in Fig. 2C yields a positive ratio because both the linear and the quadratic-term coefficients are positive, while the positively accelerated nonlinearity in the same cell's null direction response function yields a *negative* ratio because the linear-term's coefficient is negative. The distribution of this ratio for the 216 cells in our sample is shown in Fig. 3. Separate distributions are illustrated for preferred direction (Fig. 3A) and null direction (Fig. 3B) response functions. An intuitive feel for the degree of nonlinearity represented within these distributions may be obtained by examining the "q/l" ratio values provided next to each response function in Fig. 2. It is apparent that roughly 2/3 of the neurons in our sample yielded correlation-response functions that were at least as linear as those in the middle row of Fig. 2. The distributions are roughly symmetrical about a ratio of zero, indicating that nonlinearities of either sign occurred with

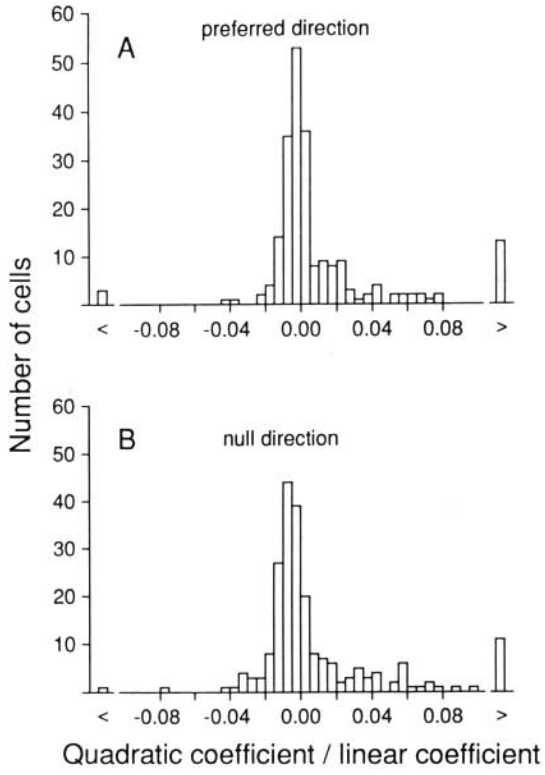

Fig. 3. Distributions of the ratio of the magnitudes of the quadratic term to the linear term for the best-fitting quadratic functions (solid curves in Fig. 2) for all 216 cells in our sample, separately plotted for preferred and null direction response functions.

approximately equal frequency. Negatively accelerating nonlinearities (the sign was expressed *relative* to the sign of the linear term) were somewhat more common for the null direction response functions, probably because of a "floor" effect resulting from complete inhibition by null direction motion (e.g. Fig. 2D).

Although the nonlinear contributions to MT correlation-response functions do not appear large on average, we wished to determine the proportion of neurons for which the quadratic functions provided a significantly better description of the data than the corresponding linear functions. To accomplish this, we conducted a nested hypothesis test (Hoel et al., 1971) of the improvement in fit resulting from inclusion of the quadratic term. The quadratic term failed to improve the fit for 76 of the 216 (35%) preferred direction response functions in our data set; the linear fit generally provided an excellent description of these data. For an additional 38 neurons (18%), inclusion of the quadratic term led to a significantly better fit ($0.001 < P < 0.05$), but the improvement in the fit was very modest. The median r^2 value for these fits, which indicates the fraction of the total variance in neuronal response accounted for, increased from 0.946–0.972 when the quadratic term was included, an improvement of only 2.7% (median values were used because these distributions were highly skewed). For these response functions the quadratic term, while significant, contributed only a minor modification to a predominantly linear response function. For the remaining 102 neurons (47%), the contribution of the quadratic term was highly significant ($P < 0.001$) and more substantial in size. The median value of r^2 for these response functions increased by 5.8% following inclusion of the quadratic term. For the null direction response functions, incorporating the quadratic term resulted in a highly significant improvement ($P < 0.001$) for 93 of the 216 cases (43%). Therefore, the nonlinear contributions to the MT correlation-response functions appear substantial for less than half of our data sets. This conclusion is consistent with the impression gained from the examples in Fig. 2 and from the distributions of q/l ratios in Fig. 3.

Response variance

The information signaling capacity of a neuron is influenced not only by the magnitude, but also by the variability of its responses. The variability of cortical responses is typically related to the mean response by a function of the form

$$\text{Variance} = k\ (\text{mean})^b$$

(Dean, 1981; Tolhurst et al., 1983; Vogels et al., 1989; Snowden et al., 1992). We used a maximum-likelihood technique that takes account of the variability of the estimates of both mean and variance to fit this relationship to our data. The average value of the exponent b for 190 cells (for 26 cells, the model II regression routine we employed was unstable) in our sample was 1.20, a value similar to the reported mean values of 1.16 (Dean, 1981) and 1.11 (Tolhurst et al., 1983; Vogels et al., 1989) obtained in studies of striate cortical neurons, and also to the value of 1.10 reported for MT neurons (Snowden et al., 1992).

It was also of interest to determine whether the power law relating variance to mean discharge would be the same for the entire sample of cells taken together as it was for the individual cells. This is not a foregone conclusion, since the range of response values covered by the population is much greater than

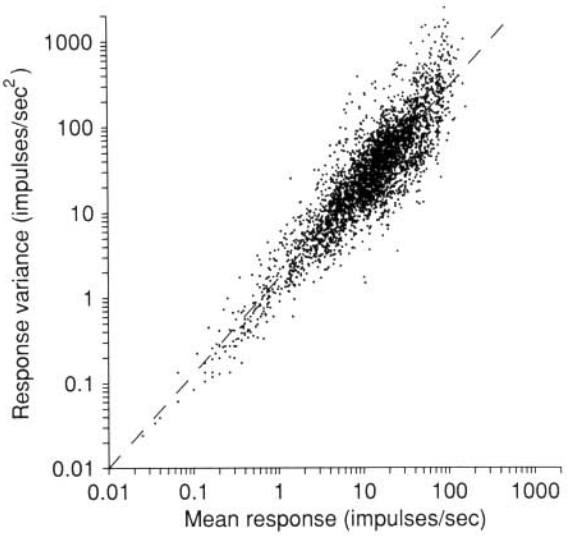

Fig. 4. Relationship between mean response level and response variance for each condition presented to all of the 216 cells in our sample. The dashed line represents the average slope of the relationship estimated using maximum-likelihood fitting to individual cells' data.

the dynamic range of each cell's responses. Fig. 4 illustrates this relationship for the entire pool of observations from our 216 MT cells. Superimposed on this scatterplot is a dashed line illustrating the *average* of the single-cell relationships as previously described. This line passes directly through the cloud of data, indicating that the relationship observed for individual neurons correctly describes the population response as well.

For most experiments in this study, the visual stimuli were stochastic motion displays in which the exact spatio-temporal pattern of signal and noise dots was unique on each trial. In other words, variance was present within the set of visual stimuli for a particular test condition (strength and direction of the motion signal), and it was therefore possible that some portion of the response variance represented in Fig. 4 originated in the visual stimulus itself rather than in the central nervous system. Alternatively, noise sources within the central visual pathway might be sufficient to obscure any contribution of variance in the visual stimulus itself. To determine whether stimulus variance contributed to the response variance of MT neurons, 56 of the 216 experiments were conducted such that the pattern of random dots was identical on all trials of a given stimulus condition. The average ratio of response variance to response mean was calculated for each neuron in both sets of experiments: those with stimulus variance and those lacking stimulus variance. The resulting distributions (log-transformed so as to be approximately normal) are illustrated in Fig. 5. These distributions are not significantly different (t-test, $t = 1.13$, $P = 0.26$), suggesting that the observed response variance of MT neurons did not result from variance in the visual stimulus. However, the two distributions are sufficiently broad that differences between cells might obscure small shifts of the means. We therefore made measurements of response variance under both conditions on a group of 31 cells, with the no-stimulus-variance trials randomly interleaved among ordinary trials. The ratios of variance to mean were then calculated for each type of trial and compared, and again there was no difference resulting from the presence of stimulus variance (paired t-test, $t = 0.27$, $P =$

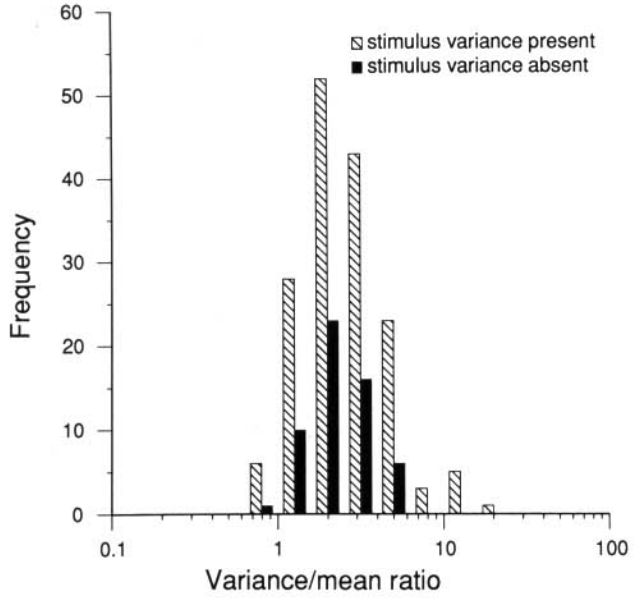

Fig. 5. Distribution of the ratio of response variance to response mean for all of the cells in our sample. Each cell provides one ratio value, which is the average of log(variance)/log(mean) across all conditions presented to the cell. Hatched bars show the values for the 161 cells in which the individual trials were novel random patterns, and the solid bars show the values for the 56 cells in which repeated trials of a given type were identical (no stimulus variance).

0.79). In both of these experiments, variance in eye position or movement may have contributed to the residual variance as well. We believe such effects to be minor, however, since the dimensions of the fixation window were small with respect to the dimensions of MT receptive fields. Unfortunately, the magnitude of this contribution is impossible to estimate because analog eye-movement data were not stored in these experiments. From these control experiments, though, we can conclude that neuronal response variance observed in our experiments does not arise from variance in the stochastic visual stimuli themselves.

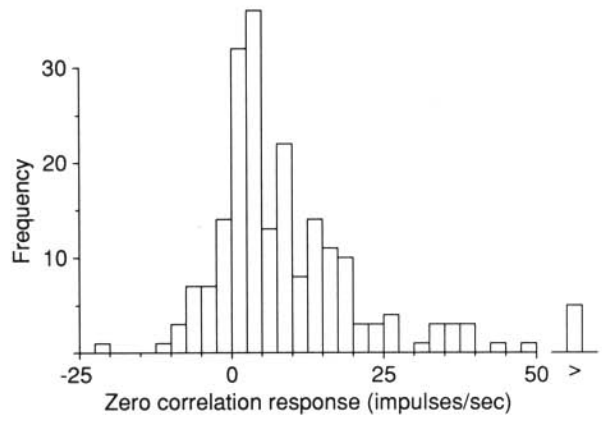

Fig. 6. Distribution of the 0% correlation responses for the 209 MT cells which were presented 0% correlation stimuli. Each was corrected for maintained activity, so that positive values indicate excitation and negative values indicate inhibition.

Responses to 0% correlation stimuli

An unexpected feature of our data was a substantial heterogeneity in the responses of MT neurons to the 0% correlation stimulus. Across the entire data set, responses to the 0% correlation condition ranged from zero to 112 impulses/s. For example, the cells in Figs. 2A and 2B were highly excited by the 0% correlation stimuli, while the cell in Fig. 2C was quite strongly inhibited. In Fig. 6, we show the distribution of these responses for the entire sample of 207 neurons for which we obtained such data (a few of the early experiments contained no 0% correlation trials). These responses have been corrected for maintained activity, so negative values show inhibition and positive values indicate excitation. Most of the cells were modestly excited by these stimuli, since responses to 0% stimuli averaged 20% of the maximum observed.

Since the 0% correlation stimulus contains motion signals that are randomly distributed in direction and speed, it seemed likely that responses to this stimulus reflect a combination of the excitatory and inhibitory inputs arising from the various motion signals in the 0% stimulus. Inhibition from null direction stimulation is a widespread feature in MT (Maunsell & Van Essen, 1983*b*; Snowden et al., 1992), and about half of the cells in our experiments (110/216) were significantly inhibited by the highest correlation null direction stimuli. The only cells which would be expected not to respond to the 0% stimulus are those with precisely balanced opponent inputs. To test this idea, we obtained an independent estimate of the net balance of excitation and inhibition for each neuron by averaging the responses of the neuron to 100% correlation stimuli moving in eight equally separated directions (see Methods). Normalized responses to individual directions were computed relative to the maintained firing rate; excitation thus yielded positive response values while inhibition yielded negative values. This average can be thought of as the predicted response from a linear cell presented with the superposition of all eight individually presented directions. Fig. 7 shows that this average response across directions is positively correlated with the similarly normalized 0% correlation response ($r = 0.596$, $P < 0.0001$). Thus, the 0% correlation response is predicted reasonably well by the linear superposition of responses to various directions of motion presented individually.

The estimate of net response provided by the average across directions is imperfect for at least two reasons: (1) inhibition is poorly measured if the maintained discharge rate of a neuron is low, and (2) the direction tuning curves were measured at a constant stimulus speed (the preferred speed of the neuron) thus omitting possible contributions from other speeds. Were we able to surmount these limitations, the relationship illustrated in Fig. 7 might be more precise. Nevertheless, this relationship confirms the basic intuition that responses to the 0% correlation stimulus reflect the net balance of excitatory and inhibitory inputs elicited by the uniformly distributed motion signals in the stimulus.

Relationship to measures of neuronal threshold

We have previously presented an analysis of these data in which we calculated thresholds for the discrimination of direction of motion from the neuronal responses (Britten et al., 1992). We used methods based on the theory of signal detection to transform neuronal firing rates into discrimination probabilities.

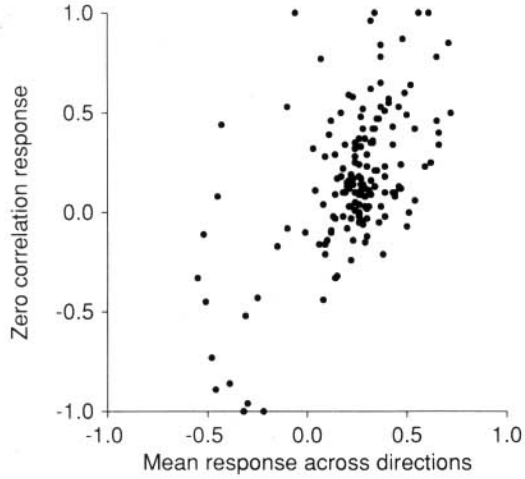

Fig. 7. The relationship between directionality and the response to 0% correlation trials. The *Y* axis shows the response to 0% correlation, and the *X* axis shows the average of the responses to stimuli presented at 100% correlation in eight directions 45 deg apart. The responses were normalized to the maximum observed excursion from maintained activity for each cell, so overall differences in cell responsiveness cannot contribute to the observed correlation. All 163 cells to which we presented direction series and 0% correlation stimuli are included.

These probabilities defined "neurometric functions" that relate the probability of correct discrimination of motion to stimulus correlation. Examples of such functions for the six neurons whose data were shown in Fig. 2 are shown in Fig. 8. The smooth curves represent fits to standard sigmoid functions (Quick, 1974) that yield an estimate of neuronal threshold, defined as the value of the stimulus correlation supporting 82% correct neuronal performance. Thresholds measured in this way vary substantially across cells, and it is natural to wonder how the factors we have described in this paper contribute to these neuronal thresholds.

For the neurons illustrated in Fig. 2, the preferred and null direction response functions diverge significantly at varying points between 0% and 50% correlation. Thus, each neuron can discriminate the direction of motion of the stimulus over a range of correlations. These curves can be used to gain a more intuitive grasp of the physiological characteristics that distinguished sensitive from insensitive neurons in the prior study. Three obvious properties can influence neuronal thresholds for discriminating direction of motion: the intrinsic variability in neuronal responsiveness as analyzed above, and the slopes (or gains) of the preferred and null direction correlation-response functions. To characterize these slopes quantitatively, we interpolated linearly between two points on the best-fitting quadratic correlation-response function: 0% and 20% correlation. We chose these two points because the thresholds of most neurons fell within this range and because data were collected up to 20% correlation for all neurons. The distributions of these slopes for the preferred and null direction response functions are shown in Figs. 9A and 9B, respectively. The distributions are broad, ranging across two orders of magnitude for the preferred direction and slightly less for the null direction. On average, the slopes are almost four times higher for the preferred direction than for the null direction (means of 39.2 and −11.1, respectively). The preferred and null direction slopes were inversely correlated;

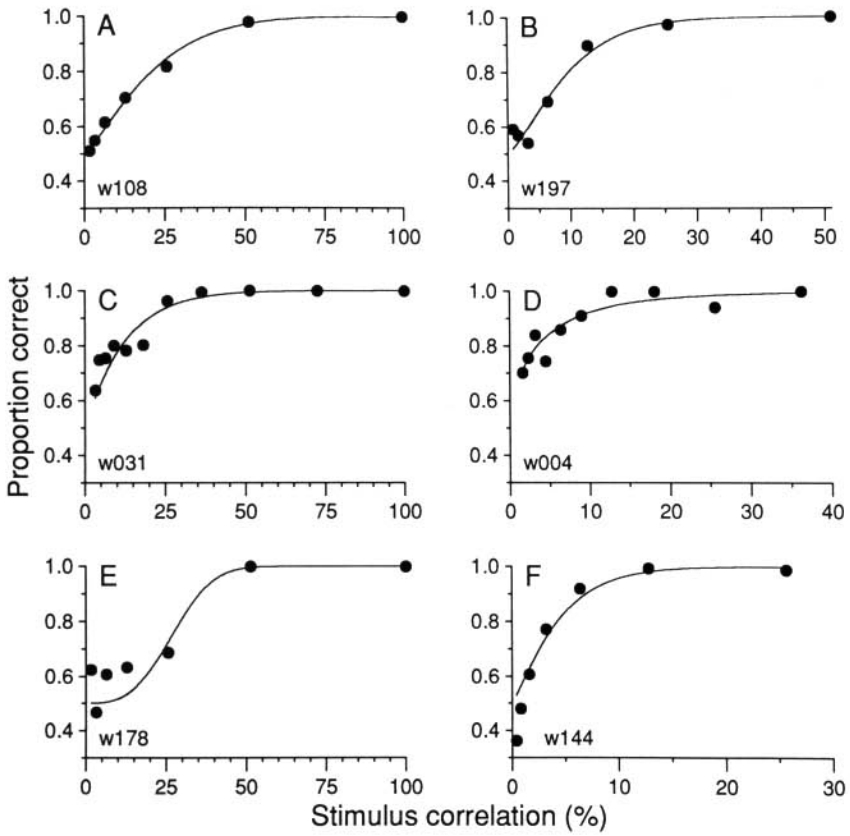

Fig. 8. Neurometric functions relating the discriminative capability of single neurons' signals to the strength of the motion stimulus. The height of each point was derived from the integrated area beneath the receiver operating characteristic curve (ROC) for that correlation level [(Green & Swets, 1966), see (Britten et al., 1992) for details]. Each indicates how well an ideal observer could discriminate the direction of motion were she to base her judgements only on the spike counts we recorded from the neuron. Each panel (A–F) is from the same neuron whose correlation-response function is shown in the corresponding panel of Fig. 2.

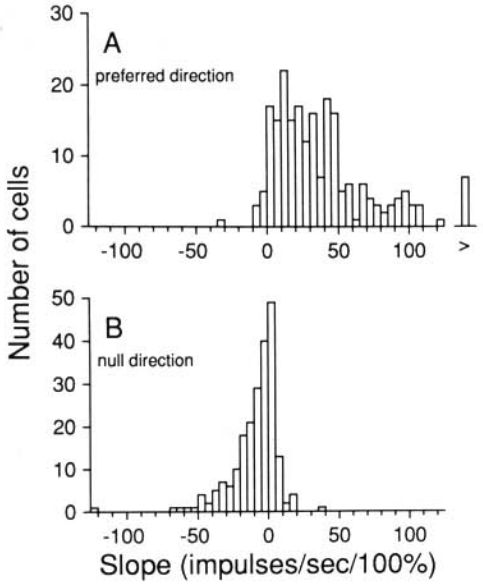

Fig. 9. Distribution of a slope measure of sensitivity for the 216 cells in our sample. The slope is derived from the best-fit quadratic function, evaluated at 0% and 20% correlation. A: Slopes of preferred direction response functions. B: Slopes of null direction response functions.

cells with large, positive slopes in the preferred direction tended to have large, negative slopes in the null direction ($r = -0.409$, $P < 0.001$).

Since most MT neurons obey a simple rule relating the variance and the mean of firing rate (Fig. 4 and associated text), we suspected that the range of response function slopes illustrated in Fig. 9 is primarily responsible for the range of neuronal thresholds alluded to above. To confirm this, we conducted a stepwise multiple-regression analysis which explored the dependence of neuronal threshold on four factors: the slopes of the preferred and null direction response functions, the average ratio of variance to mean, and the 0% correlation response. Each of these factors influenced neuronal threshold significantly ($P < 0.05$), but the slope of the preferred direction response function was by far the most important, capturing approximately ten times as much of the variance in neuronal threshold as any of the other three variables. Thus, preferred direction gain is the primary determinant of the neuronal thresholds reported in our prior study.

Discussion

We have explored the dependence of MT neurons' responses upon the strength of the motion signal contained in a stochastic random-dot display. We found that many MT cells have an approximately linear dependence of mean firing rate on stimulus correlation, although a substantial minority show strong nonlinearities. About half of the cells showing nonlinearities had negatively accelerating "compressive" nonlinearities, similar to those seen in cortical contrast-response functions (e.g. Albrecht & Hamilton, 1982; Sclar et al., 1990), but an equally numerous group showed positively accelerating correlation-response functions. Thus, the modal MT cell provides a rather linear representation of the strength of the motion signal in our display. We were struck by the frequency of linearity, and sought to explore what aspects of the stimulus also varied linearly with stimulus correlation, and thus were likely to be encoded linearly by MT neurons. By employing Fourier analysis of the stimulus set (see below), we discovered that the spectral power in the stimulus associated with the specified direction of motion varies linearly with the correlation of the stimulus. This property of the stimulus allows us to relate the neuronal responses observed in MT to recently formulated computational models of motion processing.

Spectral analysis of stimuli

The stimuli employed in these experiments consist of randomly positioned dots which can be either correlated or uncorrelated in space and time. The proportion of dots which are so correlated is the principal dimension along which we vary the stimulus. What this implies is that any mechanism which is capable of counting dots, and which responds differentially to correlated and uncorrelated dots, will modulate its output linearly with correlation. It is unlikely, however, that MT cells can be described by such a simple mechanism, since they show highly nonlinear responses to increasing numbers of dots moving in the preferred direction (Snowden et al., 1992). Alternatively, a number of recent theoretical treatments of motion signaling have made clear the utility of frequency-domain approaches to motion analysis (Fahle & Poggio, 1981; Watson & Ahumada, 1983; Adelson & Bergen, 1985; Watson & Ahumada, 1985; Emerson et al., 1992; Heeger, 1992). The application of this approach to our stimuli is illustrated in Fig. 10. For ease of presentation, we will consider only one spatial dimension of our display, thus restricting analysis to motion components parallel to the x axis. The left column of panels shows space-time portrayals of a two-dimensional simulation of our display at four different correlation levels, 0%, 30%, 60%, and 95%.[1] The horizontal position of each dot represents its position along the axis of correlated motion and the vertical position represents its time of occurrence. At 0% correlation the distribution of dots in the space-time portrayals is isotropic, corresponding to the uniform distribution of velocities in the display. As stimulus correlation increases, the longer dot lifetimes create a progressively stronger structure whose orientation corresponds to the velocity of the specified motion signal; as more and more dots join the oriented elements in the display, proportionally fewer remain in the isotropic randomly moving cloud of "noise" dots.

The task of detecting the space-time orientation in this display can usefully be viewed as that of detecting particular inhomogeneities in its spatio-temporal frequency content, because orientation in space-time produces obliquely distributed power in the frequency domain as well. The right column of panels in Fig. 10 shows two-dimensional power spectra of the corresponding space-time stimulus analogues in the left column. The two-dimensional space-time images in the left column contain one spatial and one temporal dimension (x and t), so the spectra in the right column contain one axis each of spatial frequency and temporal frequency (ω_x and ω_t). The directional content is captured in these spectra by the relative magnitudes of components in different quadrants of the display. Components whose spatial and temporal frequencies are of the same sign (quad-

[1] A more complete simulation of our display with two spatial and one temporal dimensions yielded indistinguishable results.

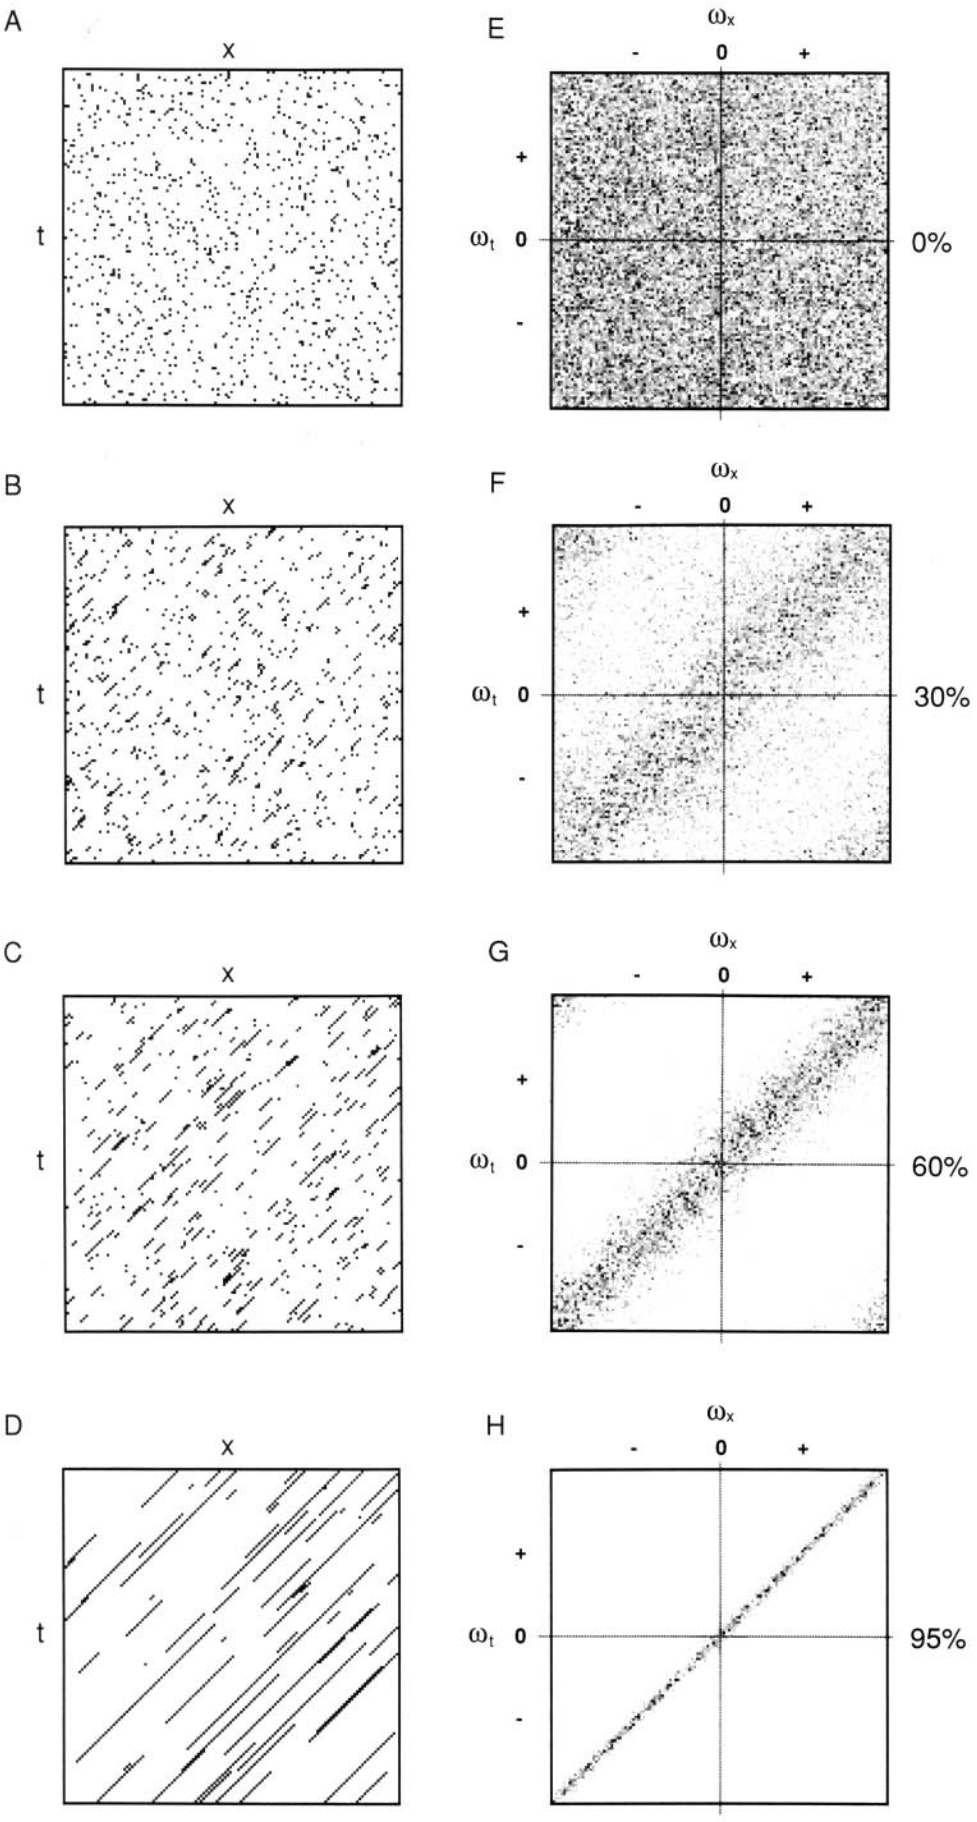

Fig. 10. A–D: Schematic representation of the random-dot stimulus. Each panel shows a space-time representation of a two-dimensional analog to our stimulus at one of four different correlation levels. Each row shows subsequent points in time, and the orientation seen at higher correlation levels thus corresponds to motion to the left over time. This version differs from our stimulus in its fine structure; in this simulation the dots are simultaneously presented (as in a frame-based display), while in the actual stimulus dots were plotted asynchronously. Note the increasingly visible orientation with higher correlation, reflecting the strength of the motion signal. Note also that the spatial and temporal parameters of the motion (embodied by the slant of the oriented segments) and the contrast (darkness of the lines) are invariant with correlation. A complete three-dimensional portrayal of our stimulus would be very similar, and would fill a volume defined by two spatial dimensions and one temporal one. Orientation with respect to the X and Y (spatial) dimensions would correspond to the direction of stimulus motion and orientation with respect to the t axis would correspond to speed. E–H: Two-dimensional Fourier power spectra of the stimuli illustrated in A–D. These have one spatial-frequency dimension (ω_x) and one temporal-frequency dimension (ω_t). Four correlation levels are illustrated here, ranging from 0% correlation (spectrally flat, or "white") to 95% correlation. Note the diagonal band corresponding to the motion energy of the specified speed. The width of this band decreases with increasing correlation, indicating that spectral power concentrates onto a narrower range of spatial and temporal frequencies. The integrated power for all correlation levels is the same, something that is not evident from these normalized images. The DC level ($\omega_x = 0$, $\omega_t = 0$) was zeroed for the purposes of illustration. Otherwise, it would dominate and the distribution of nonzero power would be concealed, especially at low correlation.

rants I and III) move in one direction, while components whose frequencies are of opposite sign (quadrants II and IV) move in the opposite direction. Thus, all components off the cardinal axes in the right column of panels in Fig. 10 correspond to motion in the stimulus. The spectral power of each component (the square of its amplitude) is represented by the darkness of each point, with darker values representing higher power (note that the grey scale used is not absolute, but is normalized for each spectrum). For low stimulus correlations, the power in the spectrum is roughly uniformly distributed in all quadrants. As stimulus correlation is increased, power concentrates in quadrants I and III (corresponding here to "forward" motion, which is the specified direction of the correlated signal dots), and is withdrawn from quadrants II and IV ("backward" motion). Of course, because these are spectra of two-dimensional slices through the full three-dimensional stimulus, only these two directions of motion are represented. In spectra such as those in Fig. 10, motion of a spatially broadband stimulus at a single speed appears as a line, on which spatial and temporal frequency are proportional. Thus, the coherent motion of the high-correlation stimulus seen in the lower left panel produces a diagonal line of elevated power in the corresponding lower right panel. Coherent motions at different speeds would form lines of different orientation. Were we able to portray the corresponding spectra for a three-dimensional case, the locus of points containing the elevated power would form a plane rather than a line.

We can now ask how the distribution of spectral power in the stimulus changes as a function of stimulus correlation. We integrated the total power in a quadrant (octant in the three-dimensional cases) corresponding to the specified direction (one of the ones containing the diagonal band of elevated power in the right column of Fig. 10) and also in the adjacent quadrant (opposite direction) at a number of correlation levels; the results are plotted in Fig. 11. The relationship between integrated spectral power and stimulus correlation is quite linear. Power accumulates in the quadrant corresponding to the specified direction and is removed from the adjacent quadrant, while the total power remains constant (as it must, since the total number of dots is not changing). Therefore, if a direction-selective cell responded to the pooled spectral power in its "preferred" quadrant, we would expect its output also to rise linearly as a function of correlation.

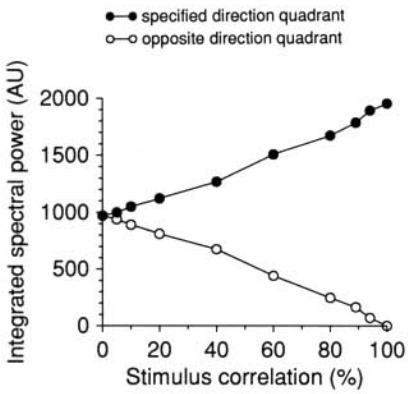

Fig. 11. The total spectral power (Fourier series amplitudes were squared, then integrated) within each of two quadrants of spatio-temporal frequency space for our stimuli as a function of correlation level. Note the approximately linear change in total power with increasing correlation for both quadrants.

This linear relationship is observed if we integrate across entire quadrants, but we know that MT cells do not respond to an entire quadrant of frequency space. Like most cells in visual cortex, their spatio-temporal frequency tuning is typically bandpass (Newsome et al., 1983; Movshon et al., 1988, and unpublished observations). We have explored the consequences of this bandpass filtering by creating Gaussian filters whose spatial and temporal bandwidths approximate those of MT neurons (Movshon, unpublished observations). These Gaussian filters, which are in fact frequency-domain implementations of "motion energy" sensors (Adelson & Bergen, 1985), limit the integration of spectral power to a restricted region of spatio-temporal frequency space. We measured the output of these filters as a function of stimulus correlation. Only the filters with the narrowest frequency bandwidths (standard deviations of 0.2 cycle/deg spatially and 1 Hz temporally) showed strikingly nonlinear, positively accelerated responses. Most were quite linear, as were the responses of typical MT cells in our sample. Therefore, the observed linearity of correlation-response functions may simply be a consequence of the breadth of the frequency tuning of MT cells, combined with a linear encoding of the spectral power within their receptive fields. Thus, a relatively simple model based on spectral power distributions can account for the main features of MT cell responses, and indeed might explain some of the nonlinearities as well. For example, we would expect that MT cells with narrow passbands would be the ones to show the most positively accelerated nonlinear correlation-response functions. This prediction remains to be tested.

We may draw two related conclusions from this analysis: (1) spectral power in our random-dot stimuli changes linearly as a function of correlation when one considers any reasonably broad region of frequency space; and (2) motion filtering models which detect spectral power will also respond linearly as a function of increasing correlation, provided that the spectral bandwidths of the filters are sufficiently broad. The first conclusion makes no assumption about mechanism; it merely draws attention to a relationship between dot correlation in our display and a quantity of more general interest. The second is of more biological relevance, and makes predictions about neuronal responses to other classes of stimuli.

Comparison with contrast-response functions

The analysis that we have just described explores the consequences of increasing the amount of motion energy within the RF of a spatially and temporally tuned neuron. In some ways, this is analogous to increasing the contrast of a grating within the cell's RF, and we reasoned that there might be a relationship between correlation-response functions and contrast-response functions. The contrast-response functions of MT neurons are typically nonlinear, positively accelerating at low contrast and saturating at medium or high contrast (Sclar et al., 1990). How can we reconcile the predominantly linear nature of our correlation-response functions with the profoundly nonlinear contrast-response functions of MT neurons? First, it is important to remember that the contrast of our random-dot displays is constant across correlation levels. Thus, it is perfectly reasonable for single neurons to yield both a linear correlation-response curve and a nonlinear contrast-response curve. More

generally, linear correlation-response functions at the level of MT can arise from the pooled outputs of lower level neurons whose transducer (i.e. contrast-response) functions are nonlinear provided that the pooling operation itself is linear. Nonlinearities in the pooling process would generate corresponding nonlinearities in the correlation-response functions of MT neurons irrespective of the shape of the transducer functions of the individual inputs. This analysis is, of course, consistent with the simulations presented in the preceding section which yielded consistently linear correlation-response functions despite the inclusion of an explicitly nonlinear contrast-response relationship (spectral power is the square of the contrast). Thus, our data are consistent with the nonlinear contrast-response functions of MT and V1 neurons as long as the pooling that forms the RF of an MT cell is approximately linear.

This reasoning might account for the lack of a positively accelerating nonlinearity at low correlations, but does not explain the lack of a saturating nonlinearity at high correlations. Saturating nonlinearities are nearly universal in MT contrast-response functions, but in our data are no more frequently seen than accelerating nonlinearities. One possible explanation for this discrepancy is the range of correlation values over which we recorded—we did not take each cell up to the maximum achievable correlation of 100%. Two analyses render this explanation unlikely. First, the cells which we did stimulate with a full range of correlations were no more likely to saturate (i.e. have negative q/l ratios) than those which were stimulated with only lower values. Second, we divided each response function into two approximately equal parts; the upper parts showed saturation no more pronounced than the lower parts. Lower q/l ratio values would be expected if the cell were nearing saturation at the high end of the stimulus range. Another possible explanation for the lack of saturation results from the relatively low spectral energy density (and thus effective contrast in any particular frequency band) of our stimuli. Thus, despite the high local peak-to-peak contrast of our dots, our stimuli might not invade the saturating portion of the cells' response functions.

Varying dot correlation in our display differs in one important way from varying the contrast of moving gratings or the number of coherently moving dots: for all but the strongest motion stimulus (100% correlated dots), our display contains opposing motion signals due to the presence of uncorrelated random dots. These "noise" dots carry a mixed motion signal and stimulate most MT neurons (Fig. 7). It is reasonable to conclude that our stimuli effect a mixture of excitation and inhibition to the MT neuron. We believe that this property may account for the linear behavior we have observed with our stimuli and the nonlinear responses to gratings and coherent dot displays (Snowden et al., 1992). Berman et al. (1992) have suggested that inhibition may reduce the effective gain of excitatory inputs by the action of cortical feedback circuits, producing a strongly nonlinear response to dominantly excitatory input and a graded linear response to more balanced stimuli.

Lastly, differences between correlation-response functions and contrast-response functions, both of exponent and with regard to saturation, might be a consequence of the action of cortical contrast-gain control mechanisms (e.g. Albrecht et al., 1984; Ohzawa et al., 1985; Bonds, 1991). Since the contrast of our stimuli is constant at all correlation levels, they would not modulate the gain-setting activity of these circuits. On the other hand, responses to stimuli whose contrast varies will be affected to some degree by changes in cortical gain.

Relationship to models of motion analysis

It is useful to consider how our data compare with predictions that arise from computational models of motion analysis (e.g. Adelson & Bergen, 1985; van Santen & Sperling, 1985; Watson & Ahumada, 1985; Heeger, 1987). These models, while different in important details, all extract local estimates of image velocity using filters localized in space-time and in frequency space. We wish to explore how far such simple, low-level motion detectors can go towards predicting the responses of MT cells to more complex stimuli, such as those employed in the present experiments. Broadly, we can say on the basis of the simulations described above that the principal result of these experiments—responses approximately linear with stimulus correlation—is consistent with the predictions of a low-level motion filter model, assuming that the filters are sufficiently broadly tuned, or that a summation step accomplishes the same result. The model we employed was, in fact, very similar to the nonopponent stage of an Adelson-Bergen motion energy model, although we have explored related models lacking a squaring nonlinearity which produce very similar results. The similarity of the predictions of the linear and energy models makes the point that the present data do not serve well to distinguish between different specific models. Most local motion filters would probably be consistent with the present results. The same is true with regard to motion opponency, another specific point of divergence between different early motion models. The predominantly negative slopes of null direction response functions (Fig. 9) are consistent with either opponent or nonopponent mechanisms, since increasing motion strength in the null direction withdraws power from the components of the noise in the preferred direction quadrant. However, we do have some evidence for opponency in that about half of the cells show overt inhibition to strong null direction stimuli. Opponency is also strongly supported by the results of Snowden et al. (1992) showing suppression of preferred direction responses by the addition of superimposed null direction motion.

Although broadly consistent with the predictions of early motion models, one aspect of the present results remains unexplained. A substantial minority of cells showed negatively accelerating preferred direction response functions, while the simulations produced only positively accelerating nonlinearities. Although this might be interpreted legitimately as a fundamental discrepancy, relatively minor modifications to early motion models could predict this form of response function as well.

Given that simple motion filters predict fairly well the responses of MT cells to these stimuli, can we make the general statement that MT cells are analogous to such filters? To answer this, we considered other experimental results as well. It is clear from the relationship between the spatial dimensions and the spatial tuning of MT cells (Movshon et al., 1988) that the simplest form of local motion operators are inadequate: some spatial pooling is required. In addition, because the spatial- and temporal-frequency bandwidths of MT cells are larger than those in V1, there is probably summation in the frequency domain as well. Appropriate pooling of local motion filter inputs, then, might be sufficient to account for known MT response properties. Linear pooling of local inputs would explain many of the response functions reported in this paper. The shape of the local detectors' transducer functions does not matter; if their outputs are linearly summed, linear correlation-response functions will result. Linear summation of local inputs is also con-

sistent with the results of the superposition test illustrated in Fig. 7. Such pooling across inputs will also predict some MT cells' responses to plaid stimuli, since "component" direction selectivity (Movshon et al., 1985) is direction tuning predicted by the linear superposition of responses to the individual stimulus components. It is perhaps coincidental but nonetheless intriguing that in both those experiments and the present ones, approximately half of the cells sampled showed significant departures from the simple linear prediction. A natural question to ask is whether the same cells showing "pattern" direction selectivity would show nonlinear responses to the stimuli employed in the present experiments.

Overall, then, with the addition of a relatively simple summation step, we can answer the question posed above in the affirmative: most MT cells can be well described by spatially summed local motion operators. Where the simple linear summation model fails, we might glean important clues to biologically significant nonlinear operations. The broadband stimuli employed in the present experiments are not ideal for exploring the nature of nonlinearities when these are present, but the present experiments certainly help to guide future experiments. Careful investigation of the summation mechanisms that serve to build the spatially and spectrally broadly tuned MT cell receptive field would be timely and potentially very rewarding.

Acknowledgments

We thank Daniel Salzman and David Heeger for useful comments and suggestions, and greatly appreciate the expert technical assistance of Judy Stein. Supported by National Eye Institute (EY-5603 and EY-2017) and by a McKnight Development Award to W.T. Newsome. K.H. Britten is supported by NIH Training Grant NS 07158-11, and M.N. Shadlen is supported by a Howard Hughes Medical Institute postdoctoral research fellowship.

References

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

Albrecht, D. & Hamilton, D.B. (1982). Striate cortex of monkey and cat: Contrast response function. *The Journal of Neurophysiology* **48**, 217–237.

Berman, N.J., Douglas, R.J. & Martin, K.A.C. (1992). GABA-mediated inhibition in the neural networks of visual cortex. In *Progress in Brain Research*, ed. Mize, R.R., Marc, R.E. & Sillito, A.M., pp. 443–476. Amsterdam: Elsevier.

Bonds, A.B. (1991). Temporal dynamics of contrast gain in single cells of the cat striate cortex. *Visual Neuroscience* **2**, 239–255.

Braun, D., Boman, D. & Hotson, J. (1991). Asymmetries in smooth pursuit do not predict asymmetries in motion perception. *Investigative Ophthalmology and Visual Science* (Suppl.) **32**, 897.

Britten, K.H., Shadlen, M.N., Newsome, W.T. & Movshon, J.A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *Journal of Neuroscience* **12**, 4745–4765.

Crist, C.F., Yamasaki, D.S.G., Komatsu, H. & Wurtz, R.H. (1988). A grid system and a microsyringe for single-cell recording. *Journal of Neuroscience Methods* **26**, 117–122.

Dean, A.F. (1981). The variability of discharge of simple cells in cat striate cortex. *Experimental Brain Research* **44**, 437–440.

Desimone, R. & Ungerleider, L.G. (1986). Multiple visual areas in the caudal superior temporal sulcas of the macaque. *Journal of Comparative Neurology* **248**, 164–189.

Downing, C.J. & Movshon, J.A. (1989). Spatial and temporal summation in the detection of motion in stochastic random-dot displays. *Investigative Ophthalmology and Visual Science* (Suppl.) **30**, 72.

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

Emerson, R.C., Bergen, J.R. & Adelson, E.H. (1992). Directionally selective complex cells and the computation of motion energy in cat visual cortex. *Vision Research* **32**, 203–218.

Fahle, M. & Poggio, T. (1981). Visual hyperacuity: Spatio-temporal interpolation in human vision. *Proceedings of the Royal Society B* (London) **213**, 451–477.

Gallyas, F. (1979). Silver staining of myelin by means of physical development. *Neurological Research* **1**, 203–209.

Green, D.M. & Swets, J.A. (1966). *Signal Detection Theory and Psychophysics*. New York: John Wiley and Sons, Inc.

Hays, A.V., Richmond, B.J. & Optican, L.M. (1982). A UNIX-based multiple process system for real-time data acquisition and control. *WESCON Conference Proceedings* **2**, 1–10.

Heeger, D.J. (1987). Model for the extraction of image flow. *Journal of the Optical Society of America* **4**, 1455–1471.

Heeger, D.J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience* **9**, 181–198.

Hess, R.H., Baker, C.L. & Zihl, J. (1989). The "motion blind" patient: Low-level spatial and temporal filters. *Journal of Neuroscience* **9**, 1628–1640.

Hiris, E. & Blake, R. (1992). A new perspective on an old phenomenon, the visual motion afteraffect. *Investigative Ophthalmology and Visual Science* (Suppl.) **33**, 1139.

Hoel, P., Port, S. & Stone, C. (1971). *Introduction to Statistical Theory*. Boston, Massachusetts: Houghton Mifflin Co.

Judge, S.J., Richmond, B.J. & Chu, F.C. (1980). Implantation of magnetic search coils for measurement of eye position: An improved method. *Vision Research* **20**, 535–538.

Maunsell, J.H.R. & Newsome, W.T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience* **10**, 363–401.

Maunsell, J.H.R. & Van Essen, D.C. (1983*a*). The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *Journal of Neuroscience* **3**, 2563–2586.

Maunsell, J.H.R. & Van Essen, D.C. (1983*b*). Functional properties of neurons in the middle temporal visual area (MT) of the macaque monkey: I. Selectivity for stimulus direction, speed, and orientation. *Journal of Neurophysiology* **49**, 1127–1147.

Morgan, M.J. & Ward, R. (1980). Conditions for motion flow in dynamic visual noise. *Vision Research* **20**, 431–435.

Movshon, J.A., Adelson, E.H., Gizzi, M.S. & Newsome, W.T. (1985). The analysis of moving visual patterns. In *Pattern Recognition Mechanisms*, ed. Chagas, C., Gattass, R. & Gross, C., pp. 117–151. New York: Springer-Verlag.

Movshon, J.A., Newsome, W.T., Gizzi, M.S. & Levitt, J.B. (1988). Spatio-temporal tuning and speed sensitivity in macaque visual cortical neurons. *Investigative Ophthalmology and Visual Science* (Suppl.) **29**, 327.

Nakayama, K. (1985). Biological image motion processing: A review. *Vision Research* **25**, 625–660.

Newsome, W.T., Britten, K.H. & Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. *Nature* **341**, 52–54.

Newsome, W.T., Gizzi, M.S. & Movshon, J.A. (1983). Spatial and temporal properties of neurons in the macaque MT. *Investigative Ophthalmology and Visual Science* (Suppl.) **24**, 106.

Newsome, W.T. & Pare, E.B. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *Journal of Neuroscience* **8**, 2201–2211.

Ohzawa, I., Sclar, G. & Freeman, R.D. (1985). Contrast gain control in the cat's visual system. *Journal of Neurophysiology* **54**, 651–667.

Quick, R.F. (1974). A vector magnitude model of contrast detection. *Kybernetik* **16**, 65–67.

Salzman, C.D., Murasugi, C.M., Britten, K.H. & Newsome, W.T. (1992). Microstimulation in visual area MT: Effects on direction discrimination performance. *Journal of Neuroscience* **12**, 2331–2355.

Sclar, G., Maunsell, J.H.R. & Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research* **30**, 1–10.

Snowden, R.J., Treue, S. & Andersen, R.A. (1992). The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random-dot patterns. *Experimental Brain Research* **88**, 389–400.

Tolhurst, D.J., Movshon, J.A. & Dean, A.F. (1983). The statistical

reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research* **23**, 775–785.

VAINA, L.M., LEMAY, M., BIENFANG, D.C., CHOY, A.Y. & NAKAYAMA, K. (1990). Intact "biological motion" and "structure from motion" perception in a patient with impaired motion mechanisms: A case study. *Visual Neuroscience* **5**, 363–369.

VAN SANTEN, J.P.H. & SPERLING, G. (1985). Elaborated Reichardt detectors. *Journal of the Optical Society of America A* **2**, 300–321.

VOGELS, R., SPILEERS, W. & ORBAN, G.A. (1989). The response variability of striate cortical neurons in the behaving monkey. *Experimental Brain Research* **77**, 432–436.

WATSON, A.B. & AHUMADA, A.J., JR. (1983). A look at motion in the frequency domain. In *Motion: Perception and Representation*, ed. TSOTSOS, J.K., pp. 1–10. New York: Association for Computing Machinery.

WATSON, A.B. & AHUMADA, A.J., JR. (1985). Model of human visual-motion sensing. *Journal of the Optical Society of America* **2**, 322–341.

WILLIAMS, D.W. & SEKULER, R. (1984). Coherent global motion percepts from stochastic local motions. *Vision Research* **24**, 55–62.

ZEKI, S.M. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *Journal of Physiology* **236**, 549–573.