Spatial integration by MT pattern neurons: A closer look at pattern-to-component effects and the role of speed tuning

John A. Perrone

The University of Waikato, Hamilton, New Zealand

Richard J. Krauzlis

The Salk Institute for Biological Studies, La Jolla, CA, USA

The primate visual system faces a difficult problem whenever it encounters the motion of an object moving over a patch of the retina. Objects typically contain a number of edges at different orientations and so a range of image velocities are generated within the receptive field of a neuron processing the object movement. It is still a mystery as to how these different velocities are combined into one unified and correct velocity. Neurons in area MT (V5) are considered to be the neural substrate for this motion integration process. Some MT neurons (pattern type) respond selectively to the correct global motion of an object, whereas others respond primarily to the individual components making up the pattern (component type). Recent findings from MT pattern cells tested with small patches of motion (N. J. Majaj, M. Carandini, & J. A. Movshon, 2007) have put further constraints on the possible mechanisms underlying MT pattern motion integration. We tested and refined an existing model of MT pattern neurons (J. A. Perrone, 2004) using these same small patch stimuli and found that it can accommodate these new findings. We also discovered that the speed of the test stimuli may have had an impact on the N. J. Majaj et al. (2007) results and that MT direction and speed tuning may be more closely linked than previously thought.

Keywords: visual motion, extrastriate, MT, V5, pattern, component, speed tuning

Citation: Perrone, J. A., & Krauzlis, R. J. (2008). Spatial integration by MT pattern neurons: A closer look at pattern-to-component effects and the role of speed tuning. *Journal of Vision, 8*(9):1, 1–14, http://journalofvision.org/8/9/1/, doi:10.1167/8.9.1.

Introduction

A long-standing puzzle in the field of motion perception concerns how the visual system is able to correctly estimate the direction of a moving pattern containing multiple edges and orientations (Adelson & Movshon, 1982; Albright, 1984; Hildreth, 1990; Movshon, Adelson, Gizzi, & Newsome, 1985; Nakayama & Silverman, 1988; Wallach, 1976). Each edge making up the pattern generates its own "component" velocity and yet we tend to perceive a coherent rigid object moving in the correct overall direction. Somehow the individual velocity estimates generated by the pattern are combined into one unified percept that mirrors the correct motion of the pattern. The neural locus for this integration stage has long been considered to be a sub-class of neurons ("pattern neurons") in the middle temporal (MT/V5) region of the primate brain (Albright, 1984; Born & Bradley, 2005; Movshon et al., 1985).

While a number of suggestions have been made as to the mechanism underlying this MT pattern integration (e.g., Adelson & Movshon, 1982; Albright, 1984; Bowns, 2002; Grzywacz & Yuille, 1990; Johnston, McOwan, & Buxton, 1992; Movshon et al., 1985; Nowlan & Sejnowski, 1995; Perrone, 2004; Qian, Andersen, & Adelson, 1994b; Rust, Mante, Simoncelli, & Movshon, 2006; Simoncelli & Heeger, 1998; Snowden, Treue, Erickson, & Andersen, 1991), the exact details of the neural processes underlying pattern motion detection remain elusive. A recent finding regarding the way MT neurons respond to small patches of motion (Majaj et al., 2007) has raised doubts about many of the theories concerning pattern motion estimation in primates.

Pattern-to-component effects in MT pattern neurons

Figure 1 shows the experimental paradigm adopted by Majaj et al. (2007) and some of their (replotted) direction tuning data. Small patches of motion stimuli were presented to non-overlapping sub-regions of the receptive field of the MT neuron being tested. For a neuron tuned to upwards motion (90°), the stimuli were either single gratings presented in the top part (Figure 1A) or bottom (Figure 1B) part of the receptive field; plaids (made up of the sum of two gratings separated by 120°, Figures 1C

doi: 10.1167/8.9.1 Received December 20, 2007; published July 2, 2008 ISSN 1534-7362

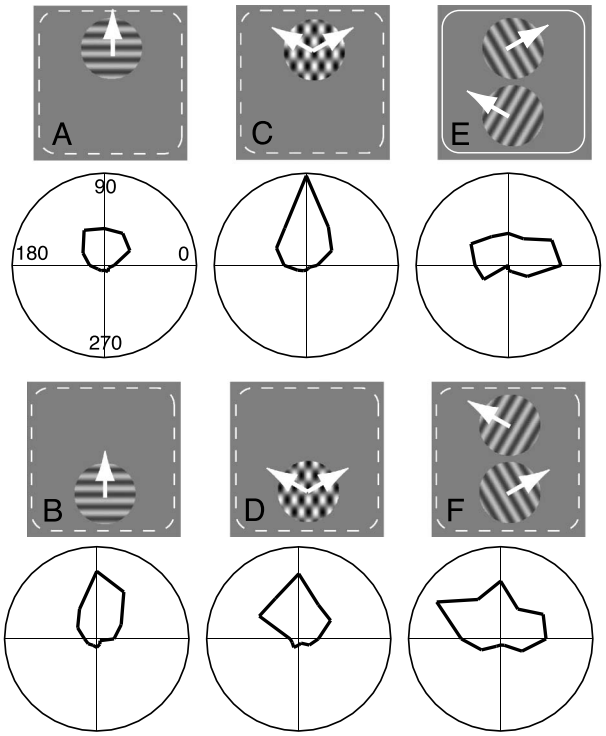

Figure 1. Stimuli used by Majaj et al. (2007) and their re-plotted data. The top part of each panel shows a representation of an MT neuron's receptive field (dashed line) along with the small patch stimuli used by Majaj et al. The bottom part of the panel shows (in polar plot form) the direction tuning curve that was produced by the stimulus configuration shown at the top. Note that the Majaj et al. direction tuning curves have been rotated 60° counter-clockwise to bring them in line with the cardinal directions.

and 1D); or "pseudoplaids" where the two gratings making up a plaid were spatially separated and presented in two different regions of the rf (Figures 1E and 1F). The pseudoplaids appeared in two different polarities depending on which grating orientation occupied the top position.

MT cells vary in their responses to grating and plaid stimuli (Albright, 1984; Movshon et al., 1985; Pack & Born, 2001; Rodman & Albright, 1987; Rust et al., 2006). A class of MT neurons in primates have been shown to act as pattern motion detectors (also referred to as Type II neurons; Albright, 1984). When a "pattern" type neuron was tested by Majaj et al. (2007) both grating and plaid patches produced similar direction tuning curves (Figures 1A, 1B, 1C, and 1D) with the maximum response occurring in the preferred direction of the cell (90° for Figure 1). In contrast, another category of MT neurons ("component" or Type I) produce bi-lobed direction tuning plots (similar to that shown in Figures 1E and 1F) when tested with plaids. The pattern types somehow integrate the motion signals generated by the individual gratings making up the plaid, whereas the component type appear to respond to the two gratings separately.

The novel result found by Majaj et al. (2007) for "pattern" type neurons is shown in Figures 1E and 1F. When the two gratings making up the plaid pattern were spatially separated to produce a "pseudoplaid" stimulus, the neuron changed from a pattern type (Figures 1C and 1D) into a component type (Figures 1E and 1F). A large proportion of their neurons (50%) showed this shift from pattern to component behavior and the remainder shifted from pattern to "unclassified" (a classification region between component and pattern types).

The Majaj et al. (2007) data and models of pattern motion estimation

Majaj et al. concluded that their data were problematical for "simple" models that assume that motion signals are pooled across the entire MT receptive field. This eliminates many models of pattern motion estimation (Britten & Heuer, 1999; Grzywacz & Yuille, 1990; Movshon et al., 1985; Rust et al., 2006; Simoncelli & Heeger, 1998; Snowden et al., 1991). Instead, Majaj et al. suggested that their data indicate that motion opponent suppressive mechanisms must be localized within small sub-regions of the receptive field of an MT pattern neuron, consistent with earlier results that also pointed to local suppression effects (Qian & Andersen, 1994; Qian, Andersen, & Adelson, 1994a; Qian et al., 1994b). Qian et al. (1994b) proposed an opponent motion model to account for the local suppression. This model used motion energy filters (Adelson & Bergen, 1985; Watson & Ahumada, 1985) as the MT pattern neuron sub-units. Unfortunately standard spatiotemporal energy filters lack the properties required for pattern motion detection (Perrone, 2004); they are broadly tuned for temporal frequency and cannot account for the speed tuning properties of V1 and MT neurons (Perrone & Thiele, 2001; Priebe, Lisberger, & Movshon, 2006).

Many other "pattern motion" models can also be eliminated as an explanation for the Majaj et al. (2007) data simply because they are not consistent with MT neuron properties. Any model that requires a velocity signal prior to the MT integration stage cannot qualify as an "MT model" because MT neurons exhibit speed-tuning—their activity is not related to velocity per se, but instead decreases for speeds both higher and lower than their preferred speed (Lagae, Raiguel, & Orban, 1993; Maunsell & Van Essen, 1983; Perrone & Thiele, 2001). Consequently, MT neurons lack the velocity signal required by schemes that use "gradient" mechanisms (Johnston et al., 1992), incorporate the "intersection of constraints" (IOC) rule (Adelson & Movshon, 1982), vector averaging (Wilson, Ferrera, & Yo, 1992), or feature tracking (Alais, Wenderoth, & Burke, 1994; Bowns, 2002). Some of these latter models may account for the overall detection of pattern movement at later stages of the visual system (e.g., MST) and may help explain some human perceptual effects but they do not qualify as models of MT pattern-motion processing.

Another MT pattern model that starts with motion energy filters and which considers spatially distributed inputs was proposed by Nowlan and Sejnowski (1995). This model uses a "selection stage" to isolate regions of the visual field where velocity estimates are most reliable. Although this model has had some success in accounting for the perception of complex stimuli such as those containing transparent motion, it is not clear how it would perform with the Majaj et al. (2007) pseudoplaid stimuli. The location of the selection units in the model was never specifically stated and Nowlan and Sejnowski indicated that some of the functions of these units could be distributed outside of MT (MST and parietal cortex). If this is the case, then their pattern model cannot be directly compared to the MT pattern neurons used to generate the Majaj et al. data. Another problem with the Nowlan and Sejnowski selection model is that the details of the mechanisms between the V1 (spatiotemporal energy) and MT stages remain unknown. The properties of the units in the model evolved through a training procedure and so the specific functionality of the V1-MT stage is unspecified.

In contrast, an MT pattern model proposed by Perrone (2004) has an explicit V1-MT stage that is closely tied to the properties of V1 neurons (Perrone & Thiele, 2002). It has been shown to provide a close match to many of the speed and direction tuning properties of MT neurons (Perrone, 2004) but it remains unknown whether this model is also able to account for the Majaj et al. (2007) small patch data.

The Perrone (2004) model

Figure 2A is a representation of an object moving vertically upwards with velocity $\overline{V} = (v, \theta)$, where v is the speed and θ is the direction (90° in the figure). Because the object is made up from edges with different orientations and because only the component of motion orthogonal to an edge can be detected (the "aperture problem;" Adelson & Movshon, 1982; Hildreth, 1990; Movshon et al., 1985; Nakayama & Silverman, 1988; Wallach, 1976), the motion present in the receptive field of an MT neuron (dashed line) will contain many different velocities. Specifically, if each edge facet has an orientation, α, then the velocity at edge location (x_i, y_i) is given by $\overline{V}_i = [v\cos(\theta - \alpha_i + 90°), \alpha_i - 90°]$. In order to register the range of possible image velocities, Perrone suggested that an MT pattern neuron should have the structure shown in Figure 2B. A number of sub-unit motion sensors tuned to a 360° range of directions in 30° steps are used to sample any motion occurring in the area enclosed by the receptive field of the model neuron. Each sub-unit motion sensor is represented as a single arrow in Figure 2B. The direction of the arrow indicates the direction tuning of the sensor and the length of the arrow is proportional to the speed tuning. The speed tuning of these sub-unit sensors varies as a cosine function of the difference between the overall direction tuning of the model neuron and the direction tuning of the sensor $(\theta - \alpha)$. We will refer to a directional set of these speed tuned sub-units as a "cluster." A total of nine (3×3) of these directional clusters is mapped out over the receptive field of the model neuron (hereafter referred to as a "pattern motion detector" or PMD). In the Perrone MT model, the design of the sub-unit sensors is based on a model of V1 motion processing (weighted intersection mechanism or WIM model; see below) (Perrone, 2005; Perrone & Thiele, 2002).

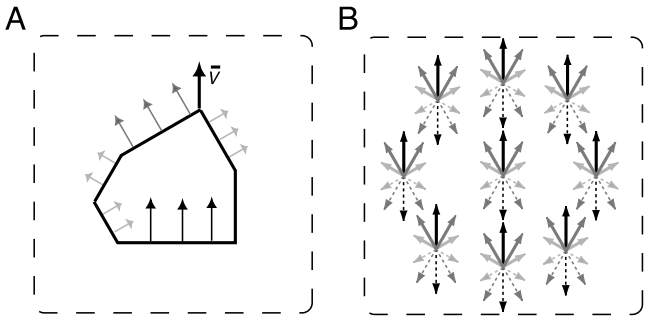

Figure 2. Detecting the overall image motion of a moving pattern. (A) Object containing multiple edge orientations moving vertically upwards. Each edge generates its own velocity depending on the orientation. (B) Pattern motion detector from Perrone (2004). The arrows represent speed- and direction-tuned V1 neurons. The detector is tuned to respond maximally to a pattern such as that shown in panel A.

Note carefully that the arrows in Figure 2B do not represent a velocity signal proportional to the image speed. These units are speed *tuned* and this speed tuning stage is what distinguishes the Perrone (2004) pattern motion model sensors from most other pattern motion models. Unlike many other pattern models, the estimation of pattern direction in the Perrone model is not based on the Intersection of Constraints rule (Adelson & Movshon, 1982). This latter rule requires that an actual speed estimate (e.g., 2 deg/s) is obtained for a number of different component directions. The pattern sensor sub-units in the Perrone model respond optimally to a particular speed but do not output a signal proportional to the component speed. The pattern motion sensors in the Perrone model are based on a "template" or "matched filter" approach (e.g., Perrone, 1992; Perrone & Stone, 1994) and do not solve the pattern estimation problem on their own; the final estimation of pattern velocity is left to later stages (e.g., MST) and the velocity is derived from a population of these MT units (Perrone & Krauzlis, 2007).

The overall plan of the model leading up to the PMD description of MT neurons is shown in Figure 3. A unique aspect of our model is the development of tight speed tuning at the stage prior to the pattern motion detection computation. Directly under the "V1 speed tuned neurons" label, the figure shows a typical spatiotemporal frequency sensitivity map obtained using the WIM model. Such maps closely match those recently found in V1

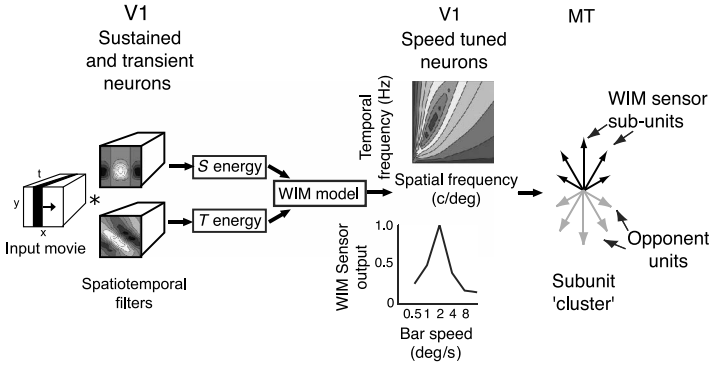

Figure 3. The early stages of the Perrone (2004) model. Nine of the MT sub-unit clusters shown on the right are incorporated into the final pattern motion detector (PMD) design.

speed tuned neurons (Perrone, 2006; Priebe et al., 2006), leading us to view the WIM sub-units as analogues of the speed-tuned, directional, complex V1 neurons described by Priebe et al. (2006). The speed tuned WIM sub-units are particular to our model. Other MT pattern models (e.g., Simoncelli & Heeger, 1998) attempt to detect the pattern motion using spatiotemporal filters that do not possess the type of speed tuning found in typical complex V1 neurons and/or which have unrealistic temporal frequency tuning properties (see Perrone, 2004). In spatiotemporal frequency space, the WIM sub-units in our model can be represented as oriented "flying saucer" like disks that are narrow in the temporal frequency dimension. The Simoncelli and Heeger (1998) model uses round "blobs" (see their Figure 3) that are not as well suited for detecting the spectral plane generated by a moving pattern in frequency space (Watson & Ahumada, 1983).

The speed tuning curve below the map shows the tightness of the tuning when tested with moving bars. This tight speed tuning is an important part of the PMD design and it enables the sub-units in the clusters to respond selectively to particular directions and speeds. The assignment of different optimum speeds to the WIM sensors in the cluster sub-units can be carried out in a very economical fashion and many different speeds can be generated using a small number of V1 neurons (Perrone, 2005).

Our PMDs represent an integration stage in which the outputs from a number of WIM sub-units tuned to a range of different speeds, directions, and possessing different spatial locations are summed. WIM sub-units tuned to directions opposite to the preferred overall direction of the PMD act in an inhibitory (opponent) way. We therefore propose that the motion opponency occurs at the stage where the inputs for individual MT pattern neuron are integrated rather than within V1 itself or between MT neurons.

The basic design of the PMDs has been tested using gratings and plaids that cover the entire model MT pattern detector receptive field and it can account for many of the properties of MT pattern neurons (Perrone, 2004). However, we do not know how it will perform with small patches of stimuli such as those used by Majaj et al. (2007). We therefore decided to test the model with the Majaj et al. (2007) stimuli.

The sub-units in the Perrone (2004) pattern motion model are sensitive to the speed of the image motion and stimulus speed has been shown to have an impact on MT neuron direction tuning (Albright, 1984; Okamoto et al., 1999). We therefore also investigated what role, if any, the speed of the test stimuli may have had on the pattern-to-component effects discovered by Majaj et al. (2007).

Methods

The WIM model

The sub-unit sensors making up the PMDs described in this paper are based on the Perrone and Thiele (2002) WIM model. The main details of this model can be found in earlier publications (Perrone, 2004, 2005, 2006; Perrone & Thiele, 2002). Features of the model used in the simulations reported in this paper that differ from the previously published versions are mainly described here. We have shown that the oriented spatiotemporal frequency (STF) response surfaces found in complex directional V1 neurons (Perrone, 2006; Priebe et al., 2006) and MT neurons (Perrone & Thiele, 2001; Priebe, Cassanello, & Lisberger, 2003) can be generated from two V1 neurons; one with low pass temporal frequency tuning (sustained, S) and another with band-pass temporal frequency tuning (transient, T). In the time domain, the S type has a unimodal temporal response profile that extends for the duration of the stimulus (sustained) and the T type has a biphasic profile with the response primarily at stimulus onset and offset (transient). The spatiotemporal energy (Adelson & Bergen, 1985; Watson & Ahumada, 1985) outputs from the S and T neurons are combined using the following equation:

$$\mathrm{WIM}(\mathrm{sf},\mathrm{tf}) = \frac{\log(\phi T + S + \alpha)}{|\log\phi T - \log S| + \delta}. \quad (1)$$

The overall S and T responses are determined by the multiplicative combination of their separate temporal and spatial frequency sensitivity functions (see below). The α and δ parameters are constants that control the overall tuning of the WIM sensor (Perrone & Thiele, 2002). Alpha can be used to control the range of spatial frequencies that the sensor will respond to and the δ parameter is used to set the gain and the speed tuning bandwidth of the sensor. The optimum speed tuning for the sensor can be controlled using ϕ (Perrone, 2005).

Perrone and Thiele (2002) never specified the neural locus of the sensors in their WIM model (other than suggesting that it preceded MT) but it is now apparent that these sensors could be analogues of the direction- and speed-tuned complex V1 neurons recently discovered by Priebe et al. (2006). We have shown that the WIM model can replicate key properties of these complex V1 neurons (Perrone, 2006). In this paper, we refer to each of these WIM motion sensors as cluster "sub-units" to distinguish them from the more "global" pattern motion detector they feed into.

Temporal frequency tuning

The low-pass, S-neuron temporal frequency tuning used in the WIM model stage was based on a Gaussian function. In the frequency domain the equation used was

$$\begin{aligned}\tilde{f}_{\text{sust}}(\text{tf}) = &\exp(-0.5\text{tf}^2\sigma^2)\cos(2\pi\text{tf}\vartheta)\\ &-\exp(-0.5\text{tf}^2\sigma^2)\sin(2\pi\text{tf}\vartheta)i,\end{aligned} \quad (2)$$

where tf is the temporal frequency measured in Hz and $i = \sqrt{-1}$. The ϑ term (phase) controls the temporal delay (lag) of the response and σ controls the spread of the Gaussian. The T-neuron temporal frequency tuning function is bandpass in shape and is given by

$$\tilde{f}_{\text{trans}}(\text{tf}) = k\tilde{f}_{\text{sust}}(\text{tf})\text{tf}i. \quad (3)$$

The magnitude parts of both of these functions are good matches to the temporal frequency tuning functions often observed in V1 neurons (Perrone, 2005). For all of the simulations reported here, $\sigma = 0.06$, $\vartheta = 0.07$, and $k = 0.25$.

Spatial frequency tuning

The spatial frequency tuning functions used in the WIM model are based on the difference of difference of Gaussians with separation (dDOGs) function used by Hawken and Parker (1987) to fit their V1 spatial frequency tuning data. The T-neuron sf tuning function, $\tilde{u}_{\text{trans}}(\text{sf})$ differs from the S-neuron sf function $\tilde{u}_{\text{sust}}(\text{sf})$ in a special way so that, when they are combined with the S and T temporal frequency tuning functions, a WIM sensor is generated which has an oriented (inseparable) spatiotemporal frequency response surface (Perrone & Thiele, 2002). For all of the model simulations in this paper, the sustained and transient spatiotemporal energy (S and T in Equation 1) was determined from the combined magnitudes of the spatial and temporal frequency functions, i.e.,

$$\begin{aligned}S(\text{sf},\text{tf}) &= |\tilde{u}_{\text{sust}}(\text{sf})|\cdot|\tilde{f}_{\text{sust}}(\text{tf})| \text{ and } T(\text{sf},\text{tf})\\ &= |\tilde{u}_{\text{trans}}(\text{sf})|\cdot|\tilde{f}_{\text{trans}}(\text{tf})|,\end{aligned} \quad (4)$$

and the values of α, δ in Equation 1 of the WIM model were set to 1.0, 0.6. The level of ϕ was dependent upon the required speed tuning of the sensor. The peak spatial frequency tuning of the S neuron determines the peak spatial frequency tuning of the PMD unit. It was set at 2 c/deg for all of the simulations.

Basic pattern motion detector design

Each cluster within a PMD is made up from 7 "positive" WIM sub-units and 5 inhibitory (opponent) sub-units (Figure 2B). Their direction tuning ranges from 0° to 330° in 30° steps. The speed tuning (set using ϕ in Equation 1) is a cosine function of the difference between the direction tuning value and the optimum overall direction tuning (θ) for the PMD. Therefore, if the overall velocity tuning of the PMD is $\overline{V}_P = (v_P, \theta_P)$, then the speed tuning of the cluster sub-units making up the detector in the Perrone (2004) model is given by $s_i = v_p\cos(\theta_P - \beta_i)$, where β_i ranges from 0° to 330° in 30° steps (see Figure 2B). The sub-units for ±90° are not shown in the figure because their speed value is 0°/s. This set of β_i values was designed to sample the range of possible edge orientations that could be present in the receptive field of the PMD and it sets up the speed tuning of each cluster sub-unit to match the expected speed of the different possible edge configurations (see Figure 2A). The location and number of possible edges present in the receptive field of a PMD varies from object to object and the edges are often spatially separated (Figure 2A). The distribution of edges cannot be predicted ahead of time and so the pattern detectors in the Perrone model have spatially separated direction-clusters in a 3 × 3 hexagonal array (Figure 2B). The separation distance between clusters was set to be 22 pixels for the 2 deg/s PMDs.

The outputs from the different sub-units in the cluster array are summed to gather evidence for a particular overall object velocity. The PMD is considered to be acting as a "template" for a particular object velocity. An object or pattern moving with velocity $\overline{V}_P = (v_P, \theta_P)$ and with multiple edge orientations will maximally stimulate this pattern detector.

There is a question as to whether or not all directions should be weighted equally in the determination of the overall pattern motion direction. The output weights affect the shape of the PMDs direction tuning curve and we selected values that provided the best match to "average" pattern-motion direction tuning curves (Albright, 1984). We weighted the outputs from the primary direction units (θ_p) by 1.0, the ±30° off-axis units were weighted by 0.87 and the ±60° units by 0.5. The weights for the ±90° units were set to 0.0 in this paper but they can be used to detect "static" features oriented 90° to the direction of motion (Perrone, 2004). Sub-units that are tuned to directions within the range $\theta_P - 180° \pm 60°$ (dashed lines

in Figure 1B) contribute in an "opponent" fashion and their output is subtracted from the total activity generated in the direction clusters; their output was weighted by 1.0. In the original model, the total positive activity across the whole receptive field was summed and the total negative activity was subtracted from it. In this paper, we have changed this feature and the output from the opponent units in a patch is subtracted from the positive activity in the same patch. The net local output (cluster positive activity − cluster negative activity) from all of the 9 clusters in the receptive field is half-wave rectified, then summed. This total activity represents the output of the MT model pattern detector and is considered to be the equivalent of the average firing rate (spikes/sec) generated by primate MT neurons. The model was implemented using digital filters and was tested using 2-dimensional image sequences that matched those used in the electrophysiological studies. Consequently, the output of the model changes over time (12 frames). We report the response at the particular frame (number 8), which typically corresponded to the peak response of the MT model units.

Normalization stage

Our model does not currently include an explicit normalization stage, although contrast normalization has been shown to have an important impact on MT neurons properties (Rust et al., 2006). We have found that the logarithmic transformations in the WIM stage (see Equation 1) act very much like a normalization process (e.g., division of the S and T energy outputs by a signal equal to the sum of total S and T energy across all 12 directions at a particular patch location). Both normalization and the log transformation produce WIM output response curves that follows the typical saturating contrast sensitivity curves found in MT neurons (Sclar, Maunsell, & Lennie, 1990). However, the log transformations are easier and faster to implement computationally and so we retained this version of the WIM model for the simulations in this paper.

Test stimuli

We tested the model with stimuli that were as similar as possible to those used by Majaj et al. (2007; see Figure 1). The grating, plaid, and pseudoplaid test image sequences used in the simulations were in digital form (128 × 128 pixels × 12 frames). The first 6 frames were set to the mean intensity level in order to minimize temporal onset effects. The relationship between pixels/frame and deg/s is somewhat arbitrary (it depends on the assumed sampling rate and the assumed field of view covered by the image) but for convenience, pixels/frame and deg/s were assumed to be equivalent in this paper. The stimulus patches were 32 pixels in diameter and their centers were separated vertically by 16 pixels. For the simulations shown in Figure 4, the test pattern was not centered exactly in the middle of the PMD but was displaced 2 pixels vertically. This did not affect the overall behavior of the PMD but it did provide a closer match to some of the features present in the Majaj et al. tuning curves. The contrast of the gratings was set at 0.5. The components of the plaids also had a contrast of 0.5, consistent with the stimuli used by Majaj et al. Stimulus test directions ranged from 0° to 330° in 30° steps. The speed of grating patches (v) varied depending upon the simulation. For the speed tests (Figure 8), the speed of the components and the overall speed of the plaids/pseudoplaids were set at: 1, 1.5, 2, 2.5, and 3 deg/s. The spatial frequency of the gratings was set to be equal to the 2 c/deg so as to match the peak spatial frequency tuning of the WIM sub-units making up the PMDs.

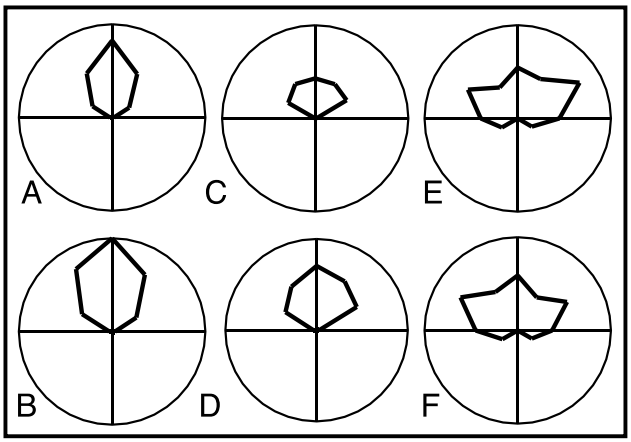

Figure 4. MT model output. Each panel shows the direction tuning curve generated by the model when tested with the stimuli depicted in the top sections of Figure 1. The top and bottom rows correspond to patches presented in the upper and lower parts of the receptive field respectively.

Quantifying the degree of pattern versus component selectivity

For a particular PMD, we measured the direction tuning curve using a grating stimulus. This represents the pattern model (Movshon et al., 1985; Smith, Majaj, & Movshon, 2005). Two of these grating responses were shifted by ±60° and summed to produce the component model. We then followed the same procedure used by Majaj et al. (2007) and calculated the partial correlation coefficients (transformed into Z scores) between these two models and the PMD output in response to plaids or pseudoplaids. If the absolute value of the difference between the two Z scores exceeded 1.28 (the $p = 0.9$ value), then the PMD

was classified as either a pattern or component type (depending on the sign of the difference). If the difference was less than 1.28, the PMD was categorized as unclassed (U). The different classification zones are marked by the two solid lines in Figure 8.

Results

The results from a slightly modified version (see Methods) of the Perrone (2004) model are shown in Figures 4A–4F. The model is able to replicate the pattern-to-component effects discovered by Majaj et al. (2007) in many of their MT pattern neurons (compare Figure 4 with the tuning curves in Figure 1). Since the upper and lower rows of the model plots are almost identical, we arbitrarily selected the bottom row for a detailed analysis of the tuning curve shapes. Using standard methods for assessing the pattern or component nature of the tuning curves (Movshon et al., 1985; Smith et al., 2005), we verified that the curve for the plaid condition in Figure 4D was clearly of the pattern type (Fisher's *r*-to-*Z* transformation score difference: $Z_p - Z_c = 2.58 > 1.28$, $p = 0.9$), and for the pseudoplaid condition (Figure 4F) it was definitely of the component type ($Z_c - Z_p = 3.1$). The model PMD behaved as an MT pattern neuron when presented with the plaid stimuli in either the top or bottom of their receptive field but acted like MT component neurons when the gratings making up the plaid were spatially separated and placed in different parts of the receptive field. We next sought to discover which features of the model were responsible for generating these pattern-to-component effects.

Identifying the source of the pattern-to-component effects

Figure 5 shows an animated movie sequence depicting the components of a plaid stimulus (red arrows) located over the bottom cluster of the PMD as the overall direction of the plaid varies through 360° in 30° steps.

On the right of the PMD receptive field is the direction tuning curve in standard polar plot form. For plaid directions close to the primary direction of the PMD (90°), the two plaid components stimulate the ±60° off-axis WIM sensors and quite a high response is generated (solid line in polar plot). As the direction of the plaid moves more than 30° away from the primary direction, one of the plaid components begins to stimulate the cluster sub-units that form part of the opponent mechanism (gray lines). The positive activity generated by the upward pointing component is inhibited and lessened by the negative activity generated in the downward pointing opponent sub-units. It become apparent from the movie

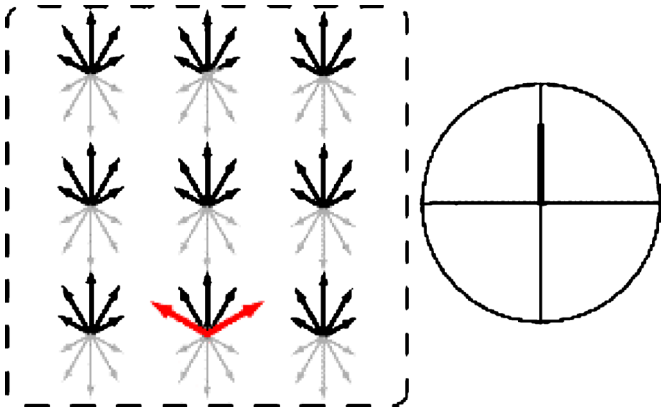

Figure 5. Animated movie sequence showing the two components of a plaid stimulus (red arrows) located in the bottom part of the receptive field as the overall plaid direction shifts clockwise from 90° to 120° in 30° steps. The pattern motion detector (PMD) is tuned to vertical motion (90°) at 2 deg/s. The polar plot on the right depicts the output of the detector as solid black lines drawn out from the origin. Use the slider bar on the QuickTime viewer to see the frame-by-frame responses.

that the amount of inhibition will be affected by the angle between the plaid components and that the plaid tuning curve is being shaped by the opponent units; the tuning curve would be much less directional if the opponency mechanism was not in place.

This movie demonstrates that the actual shape of the tuning curve is a complex mixture of a number of variables including the plaid angle and the speed of the plaid. To better understand the mechanisms affecting the tuning curves shapes, we carried out a more detailed analysis of the activity in each of the sub-units.

Figure 6 shows two maps that summarize the activity present in each of the different WIM sensor sub-units that make up the direction clusters of a PMD. In the case of the grating (Figure 6A) and plaid (Figure 6B) stimuli, the motion patches are located over a single sub-unit cluster (either at the top or bottom position of the hexagonal array) and so only a single map needs to be considered (the bottom location has been chosen arbitrarily). The horizontal axis of each map shows the direction tuning of the cluster sub-units. These are equivalent to the WIM sub-units depicted by the arrows in Figure 2B but the speed tuning is not represented. The vertical axis of the map (left hand side) shows the direction of motion of the particular test stimulus patch under study (gratings in Figure 6A, plaids in Figure 6B). Only test directions that lie ±90° from the PMD optimum direction tuning are shown for economy and clarity. The remaining directions produce very small responses (see Figures 4A–4D).

The optimum overall speed tuning for the PMD shown in Figure 6A was 2°/s and the preferred direction was 90° (upwards). The spatial frequency (sf) of the grating (2 c/deg) matched the overall peak sf of the PMD (see Methods) and the speed of all of the gratings matched the

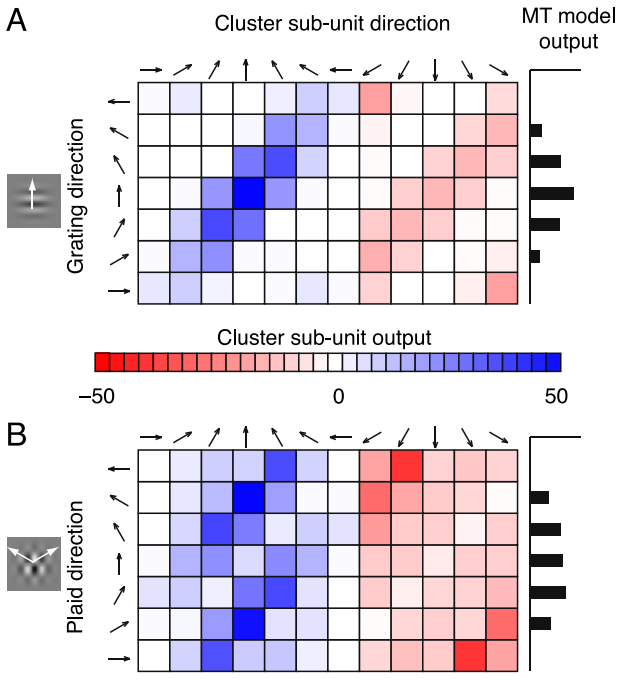

Figure 6. Activity maps from direction clusters making up a PMD tested with a grating and a plaid. Each square in the maps represents the output from a WIM sensor sub-unit tuned to the direction indicated by the arrows at the top of the map and in response to the stimuli moving in the direction indicated at the left of the map. The histogram on the right of the plot is obtained by summing the activity along the rows and represents the overall response of the PMD for each stimulus test direction. (A) Grating map. (B) Plaid map. The red squares represent inhibitory (suppressive) responses that block the individual component responses.

overall PMD speed tuning (2 deg/s). When the grating moves directly upwards (middle horizontal row in the map), the sub-unit tuned to 90° (4th column from the left) responds maximally (approximately 50 units of output) because the sf, direction and speed of the grating closely match the optimum speed, sf and direction tuning of the sub-unit at that location. A 90° direction grating will also produce some output in the 60° and 120° sub-units on either side of the 90° sub-unit. It will be less than that produced in the 90° unit however because of direction and speed discrepancies and because of the lower output weights assigned to the off-axis units (see Methods). Test directions that lie ±90° from the optimum direction for the PMD (top and bottom rows) generate a small amount of inhibitory activity (pink squares) because they partially activate some opponent type sub-units (210° and 330°).

For localized stimulus patches (as used by Majaj et al., 2007), the output from the other clusters in the receptive field array is negligible and can be ignored. The total activity from the PMD can therefore be represented in one map and is depicted in this type of plot by the net sum of the activity along each row of the map (see histogram on the right). For gratings, the output peaks in the direction that matches the optimum direction tuning of the PMD (θ_p = 90° in Figure 6A). This activity profile is consistent with the direction tuning results (see the polar plot in Figure 4B) for the same PMD and patch location. The map provides a powerful alternative means of visualizing the separate mechanisms at work inside the PMDs. This becomes particularly apparent when the plaid stimuli are treated to the same activity map analysis (Figure 6B).

In the plaid stimulus map, the component gratings making up the plaid have directions that are ±60° (see white arrows in stimulus subplots) from the test direction shown on the right hand side of the map. Therefore, the activity map is more complex because multiple sub-units are activated at the same time. However, the key feature to notice in the plaid map is the presence of strong inhibitory (negative) activity (red squares). When the plaid is moving at 30° (row 2 from bottom) one of the grating components will be moving in the 90° direction and generates a large response in the 90° sub-unit (4th column). However, the other grating component moves in the 330° direction. It therefore stimulates the 330° sub-unit in the PMD (row 2, last column on the right of the map). This 330° sub-unit is part of the opponent (inhibitory) mechanism of the PMD (see Figure 2B) and its output suppresses the positive activity from the 90° sub-unit. This is why the 30° plaid test direction produces little output from the PMD (tuned to 90°). A similar scenario exists for the 150° test direction (2nd row from top). For plaid directions greater than or equal to 90° from the optimum for the PMD, the inhibition is even stronger (red squares, top and bottom rows) and the output is totally suppressed. The map shows in more detail, the effect that was apparent in the animated movie (Figure 5).

As can be seen from the map, the only test direction relatively immune to the inhibitory effects is 90° (upwards). It stimulates the 30°, 60°, 90°, 120°, and 150° sub-units but has very little impact on the inhibitory units. The actual level of activity generated in each of the "positive" units varies and is not maximal because of direction and/or speed tuning discrepancies. Despite the modest levels of output from these sub-units, they have "strength in numbers" and the total response for the 90° test direction is the high and the PMD produces a distribution of activity corresponding to a tuning curve centered around 90° (see histogram on the right and Figure 4D).

This analysis with the plaid stimulus condition reveals that the inhibitory sub-units play an important part in shaping the direction tuning curves of the PMDs. It also provides a valuable clue for why the pseudoplaid stimuli are special and why the pattern-to-component effects occur.

Figure 7 shows the results of applying the activity map analysis to the pseudoplaid stimuli used by Majaj et al.

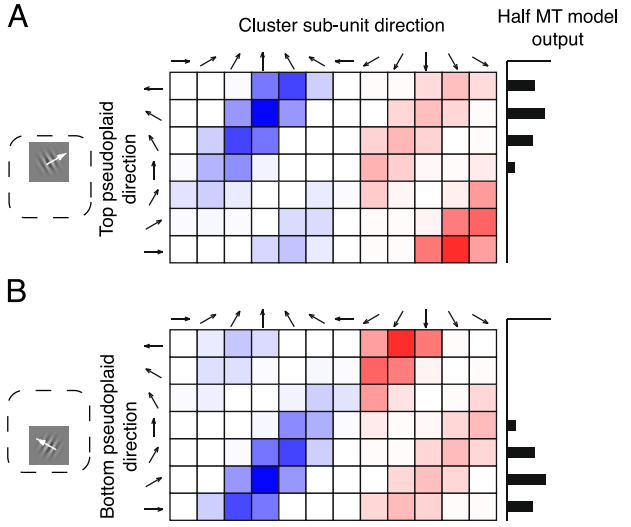

Figure 7. Activity maps for a PMD tested with pseudoplaids. The inhibition does not block the component responses in this case. (A) Map for a pseudoplaid component in the top part of the receptive field of the PMD. (B) Map for bottom part of the receptive field. The overall output of the PMD is found from the sum of the two histograms at the right of the maps.

(2007). These stimuli contain identical component gratings to the ones in the plaids but each grating is spatially separated and presented in different parts of the PMD receptive field. They therefore activate different sub-unit clusters in our model (see Figure 2B) and so each cluster must be analyzed separately. The pseudoplaid polarity shown in Figure 1E (top grating moving at 30°) was arbitrarily chosen for the analysis. The Figure 7 maps show the activity distribution for the center top sub-unit cluster (Figure 7A) and for the center bottom cluster (Figure 7B). The maps for the pseudoplaid component are very similar to the map obtained for a single grating (Figure 6A) but the pseudoplaid components have directions that are rotated ±60° from the test direction (see white arrows). Therefore, a 150° test direction for the top pseudoplaid patch now maximally stimulates the 90° sub-unit (row 6 from the bottom, column 4) instead of the 90° test for the grating. Similarly the 90° sub-unit is maximally stimulated by a 30° pseudoplaid test at the bottom position (Figure 7B, 2nd row from bottom).

Notice that unlike the plaid map (Figure 6B), there is no inhibition suppressing the activity being generated in the 90° sub-unit by the 30° and 150° tests. This is because there is no second grating present at the particular sub-unit cluster location (top or bottom) to generate inhibitory signals. Therefore, the maximum net response for the top cluster is generated by a 150° test and for the bottom cluster it occurs for the 30° test (see histograms to the right of the maps). The total activity for the PMD is determined by the activity from all of the sub-unit clusters in the 9 cluster array. Each histogram in Figure 7 only shows the activity for a single sub-unit cluster (top or bottom). The total activity is found from the sum of the two histograms and it is apparent that the distribution is going to be bimodal. This is verified by examining the direction tuning polar plot for this pseudoplaid shown in Figure 4E. The PMD now displays the tuning of a component detector rather than a pattern detector.

We have established that localized inhibition within a sub-unit cluster is the main cause of the pattern-to-component behavior we observed in our model MT neurons. It is also a possible explanation for the effects noted by Majaj et al. (2007) in their MT neurons and is consistent with earlier suggestions that local suppression is a feature of MT pattern neuron behavior (Majaj et al., 2007; Qian et al., 1994b). In order to test the robustness of our result, we varied some of the parameters of the model such as the spacing between the clusters to see what impact they had on the basic pattern-to-component results. We found that the main result persisted in spite of variations in many of the model settings. One notable exception however was the speed tuning of the PMD in relation to the speed of the test stimuli.

The role of stimulus speed in the pattern-to-component effect

In the above tests, the speed of the gratings in the plaid stimuli (v) was set to be equal to the speeds in the grating and pseudoplaid stimuli consistent with the procedure adopted by Majaj et al. (2007). This means that the overall speed of the plaids was in fact equal to $v/\cos(60°) = 2v$ and was not optimum for the speed tuning of the PMD under study (v). We decided to test the effect that this had on the behavior of our pattern motion detectors. In one series of speed tests, the component speed of the plaid and pseudoplaid was changed over a range that included the optimum tuning for the PMD (2 deg/s). In a second series of simulations, we varied the overall speed of the plaid or pseudoplaid across the same range. We then determined the Z-correlation coefficients (see Methods) for the different direction tuning curves that were generated and these coefficients are shown plotted in Figure 8.

Figure 8A shows the pattern-component Z-correlation values obtained for a PMD tuned to 2 deg/s when the plaid and pseudoplaid *components* had a speed given by the numbers in the plot. The blue numbers refer to plaid stimuli and the red numbers are for pseudoplaid stimuli. When the component speed matched the tuning of the PMD (bold "2"), the standard Majaj et al. (2007) pattern-to-component effect is seen. The plaid stimuli produce Z-correlation values that definitely qualifies the PMD as a pattern unit (upper left region of the plot). When a pseudoplaid was used as the stimulus (red "2"), the PMD definitely falls into the component region (lower right). The direction tuning curves for the 2 deg/s plaid

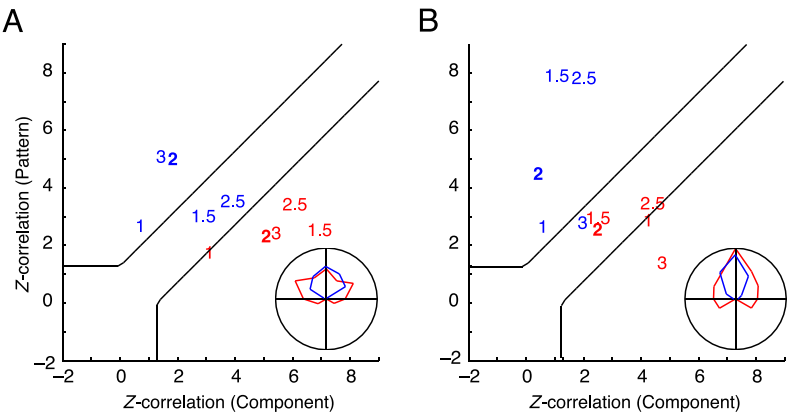

Figure 8. Plots showing the degree to which a single pattern motion detector is classified as a pattern, unclassed or component unit depending upon the test speed of the stimuli. The plots are based on a standard classification technique (see Methods) developed by Smith et al. (2005). Each location in the plots shows the *Z*-transformed partial correlations for pattern and component models derived from the direction tuning curves. (A) The PMD was tuned to 2 deg/s and was tested with plaids (blue numbers) or pseudoplaids (red numbers) in which the components moved at the speeds depicted by the numbers in the plot. The bold number indicates when the component speed matched the overall speed tuning of the PMD. The inset shows the direction tuning curve for the case in which a pseudoplaid was used and the component speed matched the speed tuning of the PMD (solid red 2) (B). The overall speed of the plaids or pseudoplaids was varied and is given by the numbers in the plot. Many fewer of the pseudoplaid tests (red numbers) are in the component region of the plot compared to panel A.

(blue) and pseudoplaid (red) tests are shown as an inset in Figure 8A and they match the typical curves found by Majaj et al. for their MT neurons (Figures 1D and 1F).

Test speeds that did not exactly match the speed tuning of the PMD (e.g., 1.5 and 2.5 deg/s) had the effect of moving the plaid tests into the "unclassed" region (central part of plot). Thus, the overall trend is for PMD units to lose their "pattern" classification when non-preferred component speeds are used. The details of how the classification changes with speed depend on the properties of the individual sub-units. Small changes to the speed of the stimulus can emphasize certain directions more than others, depending on which off-axis WIM sub-units (±30° or ±60°) are being activated the most (see Figures 6 and 7). The 30° units are tuned to 1.73°/s and the 60° units are tuned to 1°/s so small changes in stimulus speed can bias one response direction over the other. The impact on the *Z*-correlation values is predictable but non-linear and small speed alterations can result in both decrements and increments of the *Z*-correlation value (see 1.5 and 3 deg/s values in Figure 8A). Overall, although the different tests speeds moved the *Z*-correlation values around, the pseudoplaid stimuli values tended to remain in the component region of the plots and the overall trend (pattern-to-component) followed that discovered by Majaj et al. (2007) for the MT neuron tests.

This was not the case when the overall speed of the plaids and the pseudoplaids was set to be equal to the speed tuning of the PMD (Figure 8B). When the overall speed was set to 2 deg/s (bold "2s"), the pseudoplaid stimulus resulted in an "unclassed" rating for the PMD and the tuning curve for the pseudoplaid (red tuning curve in inset plot) has a peak in the overall pattern direction rather than two peaks in the component directions. This is because when it moves in the 90° direction, the speed of the components (1 deg/s) exactly matches the speed tuning of the ±60° cluster sub-units. The tuning curve differs from that generated by a plaid stimulus (blue tuning curve) but the *Z*-correlation scores are not sufficient to rate it as clearly "component."

Variations in the overall speed sometimes increased the pattern rating for the plaids (see 1.5 and 2.5 values in Figure 8B) but the pseudoplaid scores tended to remain in the "unclassed" region of the plots. The only exception was for an overall speed of 3 deg/s where the pseudoplaid produced a definite component rating. However, note that this speed did not produce a pattern rating for the plaid stimulus and so this would be an unclassed-to-component shift, not a pattern-to-component shift.

These simulations indicate that the pattern-to-component effects we noted are very much dependent upon the construction of the plaids/pseudoplaids and how they are matched to the speed tuning of the PMDs; when the speed of the components making up the plaids/pseudoplaids is set to match the overall speed tuning of the PMD, the pattern-to-component effects are observed. When the overall speed of the plaids/pseudoplaids was matched to the speed tuning of the PMD, the pattern-to-component effect (as judged by *Z*-correlation scores) disappeared. If MT pattern neurons incorporate the same basic mechanism as that of the model PMDs, then the spatial integration effects noted by Majaj et al. (2007) could be partly caused by the particular speeds chosen to test their neurons. The overall speed of the moving pattern (plaid or pseudoplaid) in their study was most likely too high, given the speed preference of the neuron.

Discussion

For a long time now, MT neurons have been considered to be the locus for some form of motion integration which enables the visual system to combine the separate local velocities generated by a moving object into one correct "pattern velocity" estimate (Adelson & Movshon, 1982; Albright, 1984; Movshon et al., 1985). Majaj et al. (2007) tested MT neurons with small patches of motion and found differences in the directional responses of the neurons when the two components of a plaid pattern were located in one location of the receptive field compared to when the components were presented in separate regions of the field (see Figures 1C–1F). Based on these results, Majaj et al. (2007) argued that "These results do not support a simple model in which motion signals are pooled across an entire MT neuron." Their data create problems for a number of popular MT models that do not take spatial effects into consideration (e.g., Movshon et al., 1985; Rust et al., 2006; Simoncelli & Heeger, 1998).

We have demonstrated that an existing model of MT pattern neurons (Perrone, 2004) can replicate the Majaj et al. (2007) results. Two changes had to be made to the model and testing procedure in order to successfully simulate the Majaj et al. (2007) data: (1) the inhibition used in the model needs to be applied locally (within a sub-unit cluster) prior to integration across the receptive field of the model pattern motion detectors, and (2) the speed of the components making up the plaid and pseudoplaid stimuli had to match the overall speed tuning of the PMD. Previous direction tuning tests of the Perrone (2004) model mainly used plaid stimuli for which the overall speed (not the component speed) matched the speed tuning of the PMD. In order to obtain good matches to the Majaj et al, data the components had to move at the preferred speed of the PMD, consistent with the paradigm adopted by Majaj et al. (2007).

The first change (local inhibition) represents a minor (but critical modification) to the original Perrone (2004) pattern neuron design and is consistent with Majaj et al's own interpretation of their results. They suggested that "local opponency" is required to explain their results in agreement with an earlier similar suggestion by Qian et al. (Qian & Andersen, 1994; Qian et al., 1994a, 1994b). Majaj et al. (2007) never specified how this local opponency could be implemented and never alluded to the type of sub-unit structure outlined in the Perrone (2004) model. We have now shown that this multiple "cluster" layout can explain their data and have provided a detailed analysis of how the local opponency acts to shape the direction tuning curves.

The second "change" to the model (stimulus speed matches) was not anticipated and is a novel result. The speed of the test stimuli was found to be critical to the outcome of the simulations. On one level, this is not surprising because tight speed tuning in the WIM sensors making up the sub-unit clusters is a distinctive feature of the Perrone (2004) pattern detector model (Figure 3). However, this speed sensitivity means that the speed of the plaid or pseudoplaid stimuli has a big impact on the resulting direction tuning curves (Figure 8). When the Majaj et al. (2007) paradigm was adopted (component speeds matched the PMD speed), the pattern-to-component trend discovered by Majaj et al. was replicated as the stimuli shifted from plaid patches to pseudoplaid patches (Figure 4). However, when the overall speed of the plaids and pseudoplaids was matched to the overall speed tuning of the PMD, the pattern-to-component effects were no longer detectable using the standard Majaj et al. statistical tests (Movshon et al., 1985; Smith et al., 2005). In other words, the spatial integration effects cannot be interpreted independently from speed tuning effects.

In addition, we discovered that PMDs were more likely to be classified as "pattern" types rather than "unclassed" if the overall speed of the plaid stimuli was matched to the overall speed of the PMD (see Figure 8). Studies of MT neurons using plaid stimuli typically match the speed of the components, not the speed of the plaid, to the overall speed tuning of the MT neurons under study (J. A. Movshon, 2007, personal communication); this choice promotes the situation depicted in Figure 8A where a "pattern" classification for plaids is less likely. An accurate estimate of the speed tuning of the MT neuron is difficult to obtain and rarely carried out. Therefore, the likelihood of an exact match between the MT speed tuning and the stimulus speed is low. In our simulations, a deviation of just 0.5 deg/s was sufficient to alter the *Z*-correlations values significantly (see blue 1.5 and 2.5 data points in Figure 8A). We wonder if this speed "mismatch" can explain the large percentage (usually around 35%) of neurons that have typically received an "unclassed" classification in previous MT studies (Majaj et al., 2007; Movshon et al., 1985; Smith et al., 2005).

The fact that our model is image-based and that it uses input sequences that match the stimuli used in physiological studies of MT pattern cells means that it could be experimentally tested. For example, an MT pattern cell could be first tested with a single grating and the model parameters (e.g., optimum speed tuning) set to provide the best fit to the grating data. Model tests with plaids of different speeds could then be used to generate a set of component and pattern classification predictions similar to those shown in Figure 8. The neuron could then be tested with the same speeds to see if the predicted patterns evolve. Based on our own simulations results (Figure 8), we predict that the "pattern" classification of the MT neuron will be reduced when non-preferred speeds are used for the components.

Relationship to other MT models

The majority of other MT models (e.g., Grzywacz & Yuille, 1990; Simoncelli & Heeger, 1998) do not sample motion across space and so cannot account for the Majaj et al. (2007) small patch data. They also have severe limitations in the physiological validity of their V1 input stage (see Introduction and Perrone, 2004). One of the most extensive attempts to model the visual pattern processing that occurs in MT neurons was recently presented by Rust et al. (2006). Their "L–N" (linear–non-linear) model was based on an earlier model of MT neurons (Simoncelli & Heeger, 1998) and it acts as a fitting function for optimizing the fit to a number of MT neuron properties.

The L-N model does not sample over spatially separated locations and so it cannot account for the Majaj et al. (2007) data. However, it is interesting to note that two of the factors that Rust et al. (2006) discovered to be crucial for the creation of pattern-direction selectivity in MT neurons are features of the Perrone (2004) MT pattern neuron model: a broad pooling of excitatory inputs from cells with varied preferred directions and a strong motion-opponent inhibition (see Figure 2B).

The biggest difference between our model and other approaches is the development of tight speed tuning in the component motion sensors prior to their inclusion in the MT pattern motion structure (see Figure 3). As Perrone (2004) has argued, the cosinusoidal speed tuning required for the pattern detector mechanism to work necessitates tight speed tuning in the sub-unit motion sensors. The sub-unit sensors need to be able to discriminate speeds differing by a very small amount (~13%) and this requires the type of speed tuning only offered by a WIM-like process (Perrone & Thiele, 2002). Other models lack this intermediate speed-tuned stage, yet our simulations have shown that speed tuning is a critical aspect of the directional tuning properties of MT pattern neurons.

The link between speed tuning and direction tuning has been noted previously in physiological studies (Albright, 1984; Okamoto et al., 1999). Simoncelli and Heeger (1998) also noted the impact that stimulus velocity had on their model MT pattern neuron direction tuning curves. It is not clear however if their sensors would exhibit the same degree of sensitivity to speed demonstrated in Figure 8. As mentioned above, the speed tuning of the sensors in the Simoncelli and Heeger model is derived from unrealistic V1 temporal frequency tuning properties (Perrone, 2004) and is not as sharp as the tuning in our sensors. A systematic comparison of the spatiotemporal frequency response properties of their model MT neurons against V1 and MT neuron data has never been carried out. In addition to being able to account for the Majaj et al. (2007) data, our model is currently the only one that can generate both the tight, concave-flanked speed tuning curves found in MT neurons (Lagae et al., 1993; Maunsell & Van Essen, 1983) and the spatiotemporal frequency tuning maps found in V1 and MT neurons (Perrone, 2006; Perrone & Thiele, 2001; Priebe et al., 2003; Priebe et al., 2006).

Velocity coding

The creation of a sensor which is able to respond maximally to the correct direction of motion of an object is still only a preliminary step in the extraction of the object velocity. MT pattern neurons (and our model PMDs) are still only speed and direction tuned. They have a peak output when the object moves at their optimum speed and direction but they still respond (albeit at a lower level) when the object moves at a different speed and/or direction (see tuning curves in Figures 1 and 4). Therefore, a single MT neuron cannot code for the velocity of the object, yet we know that at some stage of the visual system the brain can discriminate between objects moving at different speeds and directions. Without this ability, we would not be able to extract structure-from-motion or sense changes in our rate of movement through the environment. Some sort of population velocity code must be in place at a higher cortical level and one of our next major challenges is to further develop and test possible neural mechanisms that can accomplish this computation (Perrone & Krauzlis, 2007).

Acknowledgments

Supported by the Marsden Fund Council from Government funding, administered by the Royal Society of New Zealand.

Commercial relationships: none.
Corresponding author: John A. Perrone.
Email: jpnz@waikato.ac.nz.
Address: The University of Waikato, Private Bag 3105, Hamilton, New Zealand, 3240.

References

Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America A, Optics and Image Science, 2,* 284–299. [PubMed]

Adelson, E. H., & Movshon, J. A. (1982). Phenomenal coherence of moving visual patterns. *Nature, 300,* 523–525. [PubMed]

Alais, D., Wenderoth, P., & Burke, D. (1994). The contribution of one-dimensional motion mechanisms to the perceived direction of drifting plaids and their

after effects. *Vision Research, 34,* 1823–1834. [PubMed]

Albright, T. D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology, 52,* 1106–1130. [PubMed]

Born, R. T., & Bradley, D. C. (2005). Structure and function of visual area MT. *Annual Review of Neuroscience, 28,* 157–189. [PubMed]

Bowns, L. (2002). Can spatio-temporal energy models of motion predict feature motion? *Vision Research, 42,* 1671–1681. [PubMed]

Britten, K. H., & Heuer, H. W. (1999). Spatial summation in the receptive fields of MT neurons. *Journal of Neuroscience, 19,* 5074–5084. [PubMed] [Article]

Grzywacz, N. M., & Yuille, A. L. (1990). A model for the estimate of local image velocity by cells in the visual cortex. *Proceedings of the Royal Society of London B: Biological Sciences, 239,* 129–161. [PubMed]

Hawken, M. J., & Parker, A. J. (1987). Spatial properties of neurons in the monkey striate cortex. *Proceedings of the Royal Society of London B: Biological Sciences, 231,* 251–288. [PubMed]

Hildreth, E. C. (1990). The neural computation of the velocity field. In B. Cohen & I. Bodis-Wollner (Eds.), *Vision and the brain* (pp. 139–164). New York: Raven Press.

Johnston, A., McOwan, P. W., & Buxton, H. (1992). A computational model of the analysis of some first-order and second-order motion patterns by simple and complex cells. *Proceedings of the Royal Society B: Biological Sciences, 250,* 297–306. [PubMed]

Lagae, L., Raiguel, S., & Orban, G. A. (1993). Speed and direction selectivity of macaque middle temporal neurons. *Journal of Neurophysiology, 69,* 19–39. [PubMed]

Majaj, N. J., Carandini, M., & Movshon, J. A. (2007). Motion integration by neurons in Macaque MT is local, not global. *Journal of Neuroscience, 27,* 366–370. [PubMed] [Article]

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

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

Nakayama, K., & Silverman, G. H. (1988). The aperture problem. II. Spatial integration of velocity information along contours. *Vision Research, 28,* 747–753. [PubMed]

Nowlan, S. J., & Sejnowski, T. J. (1995). A selection model for motion processing in area MT of primates. *Journal of Neuroscience, 15,* 1195–1214. [PubMed] [Article]

Okamoto, H., Kawakami, S., Saito, H., Hida, E., Odajima, K., Tamanoi, D., et al. (1999). MT neurons in the macaque exhibited two types of bimodal direction tuning as predicted by a model for visual motion detection. *Vision Research, 39,* 3465–3479. [PubMed]

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

Perrone, J. A. (1992). Model for the computation of self-motion in biological systems. *Journal of the Optical Society of America A, Optics and Image Science, 9,* 177–194. [PubMed]

Perrone, J. A. (2004). A visual motion sensor based on the properties of V1 and MT neurons. *Vision Research, 44,* 1733–1755. [PubMed]

Perrone, J. A. (2005). Economy of scale: A motion sensor with variable speed tuning. *Journal of Vision, 5*(1):3, 28–33, http://journalofvision.org/5/1/3/, doi:10.1167/5.1.3. [PubMed] [Article]

Perrone, J. A. (2006). A single mechanism can explain the speed tuning properties of MT and V1 complex neurons. *Journal of Neuroscience, 24,* 11987–11991. [PubMed] [Article]

Perrone, J. A., & Krauzlis, R. J. (2007). Image velocity estimation based on vector averaging of MT neuron responses: The problem of spatial scale [Abstract]. *Journal of Vision, 7*(9):776, 776a, http://journalofvision.org/7/9/776/, doi:10.1167/7.9.776.

Perrone, J. A., & Stone, L. S. (1994). A model of self-motion estimation within primate extrastriate visual cortex. *Vision Research, 34,* 2917–2938. [PubMed]

Perrone, J. A., & Thiele, A. (2001). Speed skills: Measuring the visual speed analyzing properties of primate MT neurons. *Nature Neuroscience, 4,* 526–532. [PubMed]

Perrone, J. A., & Thiele, A. (2002). A model of speed tuning in MT neurons. *Vision Research, 42,* 1035–1051. [PubMed]

Priebe, N. J., Cassanello, C. R., & Lisberger, S. G. (2003). The neural representation of speed in macaque area MT/V5. *Journal of Neuroscience, 23,* 5650–5661. [PubMed] [Article]

Priebe, N. J., Lisberger, S. G., & Movshon, J. A. (2006). Tuning for spatiotemporal frequency and speed in directionally selective neurons of macaque striate cortex. *Journal of Neuroscience, 26,* 2941–2950. [PubMed] [Article]

Qian, N., & Andersen, R. A. (1994). Transparent motion perception as detection of unbalanced motion signals. II. Physiology. *Journal of Neuroscience, 14,* 7367–7380. [PubMed] [Article]

Qian, N., Andersen, R. A., & Adelson, E. H. (1994a). Transparent motion perception as detection of unbalanced motion signals. I. Psychophysics. *Journal of Neuroscience, 14,* 7357–7366. [PubMed] [Article]

Qian, N., Andersen, R. A., & Adelson, E. H. (1994b). Transparent motion perception as detection of unbalanced motion signals. III. Modelling. *Journal of Neuroscience, 14,* 7381–7392. [PubMed] [Article]

Rodman, H. R., & Albright, T. D. (1987). Coding of visual stimulus velocity in area MT of the macaque. *Vision Research, 27,* 2035–2048. [PubMed]

Rust, N. C., Mante, V., Simoncelli, E. P., & Movshon, J. A. (2006). How MT cells analyze the motion of visual patterns. *Nature Neuroscience, 9,* 1421–1431. [PubMed]

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

Simoncelli, E. P., & Heeger, D. J. (1998). A model of the neuronal responses in visual area MT. *Vision Research, 38,* 743–761. [PubMed]

Smith, M. A., Majaj, N. J., & Movshon, J. A. (2005). Dynamics of motion signaling by neurons in macaque area MT. *Nature Neuroscience, 8,* 220–228. [PubMed]

Snowden, R. J., Treue, S., Erickson, R. G., & Andersen, R. A. (1991). The response of area MT and V1 neurons to transparent motion. *Journal of Neuroscience, 11,* 2768–2785. [PubMed] [Article]

Wallach, H. (1976). On perceived identity: I. The direction of motion of straight lines. In H. Wallach (Ed.), *On perception,* New York: Quadrangle.

Watson, A. B., & Ahumada, A. J. (1983). A look at motion in the frequency domain. In J. K. Tsotsos (Ed.), *Motion: Perception and representation* (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 A, Optics and Image Science, 2,* 322–341. [PubMed]

Wilson, H. R., Ferrera, V. P., & Yo, C. (1992). Psychophysically motivated model for two-dimensional motion perception. *Visual Neuroscience, 9,* 79–97. [PubMed]