Compound Stimuli Reveal the Structure of Visual Motion Selectivity in Macaque MT Neurons

Andrew D. Zaharia,^{1,2} Robbe L. T. Goris,¹ J. Anthony Movshon,¹ and Eero P. Simoncelli^{1,2,3}

<https://doi.org/10.1523/ENEURO.0258-19.2019>

¹Center for Neural Science, New York University, New York, NY 10003, ²Howard Hughes Medical Institute, New York, NY 10003, and ³Courant Institute of Mathematical Sciences and Department of Psychology, New York University, New York, NY 10003

Abstract

Motion selectivity in primary visual cortex (V1) is approximately separable in orientation, spatial frequency, and temporal frequency (“frequency-separable”). Models for area MT neurons posit that their selectivity arises by combining direction-selective V1 afferents whose tuning is organized around a tilted plane in the frequency domain, specifying a particular direction and speed (“velocity-separable”). This construction explains “pattern direction-selective” MT neurons, which are velocity-selective but relatively invariant to spatial structure, including spatial frequency, texture and shape. We designed a set of experiments to distinguish frequency-separable and velocity-separable models and executed them with single-unit recordings in macaque V1 and MT. Surprisingly, when tested with single drifting gratings, most MT neurons’ responses are fit equally well by models with either form of separability. However, responses to plaids (sums of two moving gratings) tend to be better described as velocity-separable, especially for pattern neurons. We conclude that direction selectivity in MT is primarily computed by summing V1 afferents, but pattern-invariant velocity tuning for complex stimuli may arise from local, recurrent interactions.

Key words: extrastriate visual cortex; macaques; visual motion processing; receptive fields; normalization

Significance Statement

How do sensory systems build representations of complex features from simpler ones? Visual motion representation in cortex is a well-studied example: the direction and speed of moving objects, regardless of shape or texture, is computed from the local motion of oriented edges. Here, we quantify tuning properties based on single-unit recordings in primate area MT, then fit a novel, generalized model of motion computation. The model reveals two core properties of MT neurons, speed tuning and invariance to local edge orientation, result from a single organizing principle: each MT neuron combines afferents that represent edge motions consistent with a common velocity, much as primary visual cortex (V1) simple cells combine thalamic inputs consistent with a common orientation.

Introduction

Most neurons in extrastriate area MT (V5) are tuned for the speed and direction of visual motion (Dubner and Zeki, 1971; Van Essen et al., 1981; Maunsell and Van

Essen, 1983), and many of them are selective for the coherent motion of complex patterns (Movshon et al., 1985). Such tuning is absent from the earliest stages of visual processing in primates, the retina and lateral genic-

Received July 3, 2019; accepted August 22, 2019; First published October 11, 2019.

The authors declare no competing financial interests.

Author contributions: A.D.Z. and R.L.T.G. performed research; A.D.Z. and R.L.T.G. analyzed data; A.D.Z., J.A.M., and E.P.S. wrote the paper; R.L.T.G., J.A.M., and E.P.S. designed research.

This work was supported by the Howard Hughes Medical Institute and the National Institutes of Health Grant R01 EY04440.

A. D. Zaharia’s present address: Zuckerman Mind Brain Behavior Institute, 3227 Broadway, New York, NY 10027.

R. L. T. Goris’s present address: Center for Perceptual Systems, University of Texas at Austin, Austin, TX, 78712.

ulate nucleus. There, incoming visual signals are filtered without regard to direction, and are approximately separable in space and time (Enroth-Cugell et al., 1983; Derdikson and Lennie, 1984). Motion-selective simple cells in primary visual cortex (V1) are tuned for motion in a manner that treats spatial and temporal frequency roughly separably (Tolhurst and Movshon, 1975), while a quarter of V1 complex cells treat them jointly (Priebe et al., 2006), consistent with speed tuning. V1 neurons provide input to MT, where neurons also tend to be speed tuned (Perrone and Thiele, 2001; Priebe et al., 2003).

Motion-selective V1 neurons are also orientation-selective, and their responses confound the direction of motion and the orientation of moving stimuli. In particular, they respond independently to each oriented component rather than to the pattern as a whole (Movshon et al., 1985). Under many conditions, humans perceive such complex patterns as moving coherently in a single direction (Wallach, 1935; Adelson and Movshon, 1982). Similarly, MT neurons signal coherent pattern motion, with some neurons being completely invariant to component orientation (Movshon et al., 1985). The degree to which MT neurons respond to the motion of individual components or the whole pattern lies on a continuum, quantified by a “pattern index” (Fig. 1A–C; Materials and Methods; Movshon et al., 1985).

Speed tuning and pattern motion selectivity in MT were typically studied separately. Furthermore, previous studies in MT were performed in at most two of three dimensions: spatial and temporal frequency (Perrone and Thiele, 2001; Priebe et al., 2003, 2006), or direction and speed (Rodman and Albright, 1987). Recently, Nishimoto and Gallant (2011) and Inagaki et al. (2016) quantified MT selectivity in all three dimensions simultaneously, but did not relate their findings to pattern motion selectivity.

The Simoncelli and Heeger (1998) model of MT motion computation proposes that speed tuning and pattern motion selectivity both emerge from selective weighting of V1 afferents, parameterized in all three frequency dimensions. The model posits that MT neurons sum responses of V1 neurons whose preferred stimuli are consistent with a common velocity. MT neurons could, however, sum V1 afferents whose preferences share a common temporal frequency.

Here, we unify previous theory and experimental data in a coherent framework, by modifying the Simoncelli and Heeger (1998) model to allow direct fitting to electrophysiological recordings. Specifically, we compared the two hypotheses of MT computation above in their ability to explain the responses of neurons in areas V1 and MT of

anesthetized and awake macaques to a large collection of sinusoidal gratings and plaids (superimposed gratings with different orientations and temporal frequencies). We fit these responses with a linear-nonlinear model of MT computation, in which the MT receptive field was constructed by summing velocity-specific or temporal frequency-specific combinations of V1 afferents. We refer to the former model variant, in which selectivity to spatial and temporal frequency varies jointly, as the velocity-separable model, and the latter model as the frequency-separable model.

Nearly all V1 neurons were better described by the frequency-separable model. When probed with drifting sinusoidal gratings, MT responses were equally well-described by both models. However, when probed with plaid stimuli, the velocity-separable model systematically outperformed the frequency-separable model for pattern-selective neurons. This is the first direct evidence establishing speed tuning and pattern motion selectivity in area MT as consequences of a single organizing principle: selectivity organized along a preferred velocity plane.

Materials and Methods

Anesthetized recording procedures

We recorded from seven anesthetized, paralyzed, adult male macaque monkeys (*Macaca fascicularis*) and one adult female macaque (*Macaca mulatta*) using standard procedures for surgical preparation and single-unit recording, as described previously (Cavanaugh et al., 2002). We maintained anesthesia and paralysis by intravenously infusing sufentanil citrate ($6\text{--}30\ \mu\text{g kg}^{-1}\ \text{h}^{-1}$), and vecuronium bromide (Norcuron, $0.1\ \text{mg kg}^{-1}\ \text{h}^{-1}$), respectively, in isotonic dextrose-Normosol solution ($4\text{--}10\ \text{ml kg}^{-1}\ \text{h}^{-1}$). Vital signs [heart rate, lung pressure, electroencephalogram (EEG), electrocardiogram (ECG), body temperature, urine flow and osmolality, and end-tidal CO_2 partial pressure (pCO_2)] were continuously monitored and maintained within appropriate physiological ranges. Atropine was applied topically to dilate the pupils. Gas-permeable contact lenses protected the eyes, which were refracted with supplementary lenses chosen by direct ophthalmoscopy. Experiments typically lasted 5–7 d at the end of which the monkey was killed with an overdose of sodium pentobarbital. We conducted all experiments in compliance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals and with the approval of the New York University Animal Welfare Committee.

The monkey was positioned so their eyes were 57–114 cm from the display. Grating and plaid stimuli each lasted for 1000 ms and were presented in randomly interleaved blocks. We used $0.5\text{--}3\ \text{M}\Omega$ impedance quartz-platinum-tungsten microelectrodes (Thomas Recording) to make extracellular recordings in the brain through a craniotomy and small durotomy. Electrolytic lesions were made at the end of each recording track for histological confirmation of MT recording sites. For each isolated unit, we determined eye dominance and occluded the non-preferred eye. While isolating neurons in V1 for recording, we selected those with strong direction tuning.

Acknowledgements: We thank Romesh Kumbhani, Michael Gorman, Najib Majaj, and other members of the Movshon Laboratory for help with physiological experiments.

Correspondence should be addressed to Andrew D. Zaharia at zaharia@cns.nyu.edu or Eero P. Simoncelli at eero.simoncelli@nyu.edu.

<https://doi.org/10.1523/ENEURO.0258-19.2019>

Copyright © 2019 Zaharia et al.

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International license, which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

Figure 1. Pattern index, frequency-separable and velocity-separable hypotheses, and their predicted tuning. **A**, Tuning curve of an idealized direction-selective neuron, responding to drifting gratings. **B**, An ideal pattern-selective neuron exhibits a unimodal tuning curve for drifting plaids (red), while an ideal component neuron shows a bimodal tuning curve (blue). The peaks of the component neuron tuning curve correspond to the directions of the two gratings that comprise the plaid. **C**, The pattern index captures the degree to which a given neuron is pattern or component selective. Each point represents the correlation of a given neuron's measured tuning curve with the ideal component and pattern tuning (abscissa and ordinate, respectively; see Materials and Methods for details), as predicted from its actual grating responses. Open and filled points correspond to neurons featured in this article in V1 ($n = 21$) and MT ($n = 112$), respectively. **D**, 3D frequency domain representation of moving images, with two spatial frequency axes and one temporal frequency axis. These coordinates can alternatively be expressed as orientation, spatial frequency, and temporal frequency. A single point in the frequency domain represents a single drifting sinusoidal grating. **E**, The motion of a rigidly translating pattern (e.g., a field of dots moving with the same velocity) contains frequency components that lie on a plane through the origin. **F**, **G**, Two possible hypotheses for MT selectivity in the frequency domain. In a velocity-separable receptive field (**F**), spatial and temporal frequency tuning are concentrated along a tilted, preferred velocity plane. In the frequency-separable prediction (**G**), spatial and temporal

continued

frequency tuning are independent. Note the velocity-separable hypothesis depicted shears along the vertical (temporal frequency) dimension, rather than along the vector orthogonal to the preferred velocity plane (for details, see Materials and Methods). **H, I**, Contour plots of slices through the two selectivity volumes from **F, G** at the optimal direction, superimposed with stimuli for two “classical” tuning experiments (black lines) containing the optimal stimulus (black ball) and suboptimal stimuli (dark gray): spatial frequency tuning at optimal and low temporal frequencies (**H**) and temporal frequency tuning at optimal and low spatial frequencies (**I**). **J, K**, Temporal and spatial frequency tuning for the two models is the same for classical stimuli (red-blue dashed lines), but different for non-optimal stimuli (red and blue solid lines). The velocity-separable (light red) spatial and temporal frequency tuning curves are shifted away from the tuning curves observed for optimal stimuli.

Awake recording procedures

To verify that the observations we made in the anesthetized preparation were not affected by anesthesia, we also recorded from two awake, actively fixating, adult male macaques (one *M. mulatta* and one *Macaca nemestrina*). No differences were observed between awake and anesthetized data. A headpost was surgically implanted for head stabilization using the design and methods described previously (Adams et al., 2007). In a second surgical procedure, a chamber was implanted for chronic electrode recording over the superior temporal sulcus (STS) of the left hemisphere, using the techniques and a variant of the design described in (Adams et al., 2011). Before surgery, we used structural magnetic resonance imaging (MRI) and Brainsight software (Rogue Research) to design a chamber with legs matched to the curvature of the monkey’s skull (Johnston et al., 2016) above the STS.

We acclimated each monkey to his recording chair and experimental surroundings. After this initial period, he was head-restrained and rewarded for looking at the fixation target with dilute juice or water. Meanwhile, we used an infrared eye tracker (EyeLink 1000; SR Research) to monitor eye position at 1000 Hz via reflections of infrared light on the cornea and pupil. The monkey sat 57 cm from the display.

The monkey initiated a trial by fixating on a small white spot (diameter 0.1°), after which he was required to maintain fixation for a random time interval between 2350 and 4350 ms. A grating or plaid stimulus would appear 100 ms after fixation began and last for 250 ms. Stimulus conditions were presented in randomly interleaved blocks. The monkey was rewarded if he maintained fixation within 1–1.75° from the fixation point for the entire duration of the stimulus. No stimuli were presented during the 300–600 ms in which the reward was being delivered. If the monkey broke fixation prematurely, the trial was aborted, a timeout of 2000 ms occurred, and no reward was given.

For extracellular recordings, we advanced 0.5–3 M Ω impedance tungsten microelectrodes with epoxylite insulation (FHC) through a 23 Gauge stainless steel dura-penetrating guide tube. We identified area MT from gray matter-white matter transitions and isolated neurons’ brisk, direction-selective responses.

Analysis of neuronal response

We used Expo software (<http://corevision.cns.nyu.edu>) on an Apple Mac Pro to process signals and sort single units. Signals were amplified 1000 $\times$, bandpass filtered (300 Hz to 10 kHz), fed into a multiple-window time-amplitude discriminator, and time-stamped with 100 μ s

resolution. During each single-unit isolation attempt, discriminator windows and thresholds were manually set, online, to most unambiguously and stably distinguish a single spike wave form from noise and other units. If spike waveforms from multiple neurons overlapped to the extent to which they could not be separated, the discriminator allowed us to detect this based on subrefractory period interspike intervals. If this type of multiunit activity was detected, we would then either manually refine the discrimination windows to effectively isolate the unit, move the electrode a small distance, or abandon that neuron for experimentation and move on to the next unit that could be fully isolated. During the experiments, each unit’s waveform was continuously monitored for isolation stability. If necessary, minor, manual adjustments detailed above were made to preserve isolation quality. Spike waveform consistency and interspike interval distributions over the entire duration of the experiments were verified *post hoc*. While running experiments, Expo automatically generated and displayed tuning curves, from which half-maximum spiking response rates, and their corresponding stimulus values, could be directly computed and displayed through the graphical interface.

Visual stimulation

We presented visual stimuli on a gamma-corrected CRT monitor (an Eizo T966 during anesthetized experiments, and an HP P1230 during awake experiments; mean luminance, 33 cd/m²) at a resolution of 1280 $\times$ 960 with a refresh rate of 120 Hz. Stimuli were generated and presented on the same Mac Pro using Expo.

For each isolated unit, we presented windowed sinusoidal grating stimuli to determine, by hand, initial estimates of each cell’s receptive field and preferred size. We used a standard sequence of tuning experiments to make precise estimates of the cell’s tuning preferences. Each tuning curve measured in this sequence featured 100% contrast single gratings varying along a single stimulus dimension, beginning with size tuning and followed by direction, spatial frequency, and temporal frequency tuning. After each of these individual tuning experiments finished, we determined the preferred stimulus value of the dimension tested and used it in subsequent experiments. Next, we measured pattern direction selectivity, at optimal spatial and temporal frequencies, with interleaved drifting gratings and plaids. In the rare cases in which this experiment yielded a different grating direction preference from the previously determined value, we repeated the full sequence of tuning curve measurements, to make sure optimal values would be used for the single component

and planar plaid experiments which followed. The receptive fields of all recorded neurons were centered between 2° and 30° from the fovea.

Next, we ran the planar plaid study, which required no further stimulus optimization. In the planar plaid study, stimuli were chosen to span four different direction tuning curves at the optimal spatial frequency (Fig. 7A,B). The first two were based on single gratings at 50% contrast, one with temporal frequency held constant at the optimal value (“constant frequency”), presented in all directions in 30° intervals, and the other with constant, optimal velocity (“constant velocity”) from −90° to 90° relative to the preferred direction, in 15° intervals. Since a given velocity corresponds to spectral content lying on a tilted plane in frequency space, constant velocity gratings had a temporal frequency that varied with the cosine of their direction.

The last two tuning curves consisted of 120° “plaids” (sums of two gratings with orientations 120° apart). The component gratings had the same temporal frequency, and were presented at 50% contrast each. The “pattern direction” of motion (direction consistent with rigid translation, equal to the average direction of the two gratings) was sampled the same set of directions used for single grating tuning curves.

Following the planar plaid study, we ran the single component study. It included presentation of 225 drifting grating stimuli, each at 100% contrast. Stimuli were arranged to widely sample the three dimensions of spatiotemporal frequencies near a given neuron’s tuning preferences, using multiple tuning curves, each varying along a single dimension: direction, spatial frequency, or temporal frequency. These tuning curves were measured at optimal and suboptimal values, the latter of which were determined by reading out (or if necessary, linearly interpolating) the stimulus values which elicited a response at half the neuron’s maximum spike rate (computed directly in Expo, see above, Analysis of neuronal response) in the preceding standard tuning curve experiments (for details, see Extended Data Table 4-3). By sampling spatiotemporal frequencies in this way, we could efficiently concentrate stimuli to reveal subtle changes of each neuron’s selectivity in a manner that does not assume a particular shape of selectivity or manner of tuning specific to either V1 or MT.

Note that the first two direction tuning curves of the single component study differ from the two grating tuning curves in the planar plaid study in that in the planar plaid study: (1) gratings were at 50% contrast instead of 100%, and (2) constant frequency gratings spanned the whole range of directions rather than just the semicircle of directions centered at the preferred one. Although gratings were only presented at 50% contrast in the planar plaid study, its implications for 3D frequency selectivity shape should generalize to the 100% contrast case, since differences in response strength for these contrast levels are negligible (Sclar et al., 1990; Carandini et al., 1997).

Frequency-separable and velocity-separable models

The MT linear weighting functions for both the frequency-based and velocity-based models are defined

as a separable product of tuning functions over direction $w_d(d)$, spatial frequency $w_s(s)$, and temporal frequency $w_t(t)$.

Specifically, the frequency-separable linear weighting for a grating is defined as follows:

$$F(d, s, t) = w_d(d) \cdot w_s(s) \cdot w_t(t) . \quad (1)$$

Direction tuning above is represented by a von Mises function:

$$w_d(d) = \frac{e^{\sigma_d \cos(d - \mu_d)}}{2\pi I_0(\sigma_d)} , \quad (2)$$

where μ_d and σ_d represent the direction preference and bandwidth, respectively, and $I_0()$ is the modified Bessel function of order 0 (which normalizes the integral of the numerator). Spatial frequency is represented by a log-Normal function, parameterized by spatial frequency preference μ_s and bandwidth σ_s :

$$w_s(s) = \frac{1}{\sigma_s \sqrt{2\pi}} e^{-(\log_2(s) - \log_2(\mu_s))^2 / 2\sigma_s^2} . \quad (3)$$

Finally, temporal frequency is represented by a Gaussian in coordinates which are linear at low frequencies and logarithmic at higher ones:

$$w_t(t) = \frac{1}{\sigma_t \sqrt{2\pi}} e^{-(g(t) - g(\mu_t))^2 / 2\sigma_t^2} , \quad (4)$$

where $g(t)$ is

$$g(t) = \text{sgn}(t) \log_2 \left(\frac{|t|}{\tau} + 1 \right) . \quad (5)$$

Using this functional form for temporal frequency tuning allows $w_t(t)$ to be logarithmic at high temporal frequencies, but also be zero-valued and continuous at zero temporal frequency. The parameter τ determines the temporal frequency at which the function transitions from linear to logarithmic, and μ_t and σ_t are the temporal frequency preference and bandwidth, respectively.

The velocity-separable linear weighting function is defined as follows:

$$V(d, s, t) = w_d(d) \cdot w_s(s) \cdot v_t(t; d, s) , \quad (6)$$

where the velocity-separable temporal frequency function, v_t , is defined as a Gaussian, again linear at low frequencies and logarithmic at higher ones:

$$v_t(t; d, s) = \frac{1}{\sigma_t \sqrt{2\pi}} e^{-(g(t) - g(P(d, s)))^2 / 2\sigma_t^2} \quad (7)$$

The only difference between $w_t(t)$ (Eq. 4) and $v_t(t; d, s)$ (Eq. 7) is that in the latter, temporal frequency tuning is centered on the preferred speed plane $P(d, s)$:

$$P(d, s) = s \frac{\mu_t}{\mu_s} \cos(d - \mu_d) . \quad (8)$$

Note that a consequence of Equations 6–8 is that the velocity-separable model “shears” vertically in the direction of the temporal frequency axis, rather than along the vector orthogonal to the preferred velocity plane (Fig. 1). This was done deliberately, to account for the broad temporal frequency tuning MT neurons exhibit near the preferred direction and spatial frequency (Figs. 4, 8).

Previous models included a V1 normalization stage, either explicitly or implicitly simulated, at this part of the computation (Simoncelli and Heeger, 1998; Rust et al., 2006; Nishimoto and Gallant, 2011). Normalization at the V1 stage has been previously shown to be an important contributor to MT tuning properties. While we could have made it an explicit piece of the model here, Rust et al. (2006) showed that, in some cases, it can be combined with the MT normalization stage to yield a single normalization computation. Moreover, V1 contrast normalization is engaged only when contrast varies widely, and cross-orientation suppression is strongest for components close in orientation. Since the grating components in the plaid stimuli in our experiments are always 50% contrast and 120° apart, we assume such cross-orientation and contrast normalization effects in V1 are negligible. Thus, we assume the next model stage consists of MT neurons summing the responses to each plaid component.

Finally, the full MT model response is computed by raising the linear responses to a power β , and then normalizing them:

$$R_f(d, s, t, t_{max}) = \alpha_0 + \frac{\alpha_1 n^{(1-2\beta)/3} \left(\sum_i^n F_i(d, s, t) \right)^\beta}{\alpha_2 + \sum_i^n N_i(t, t_{max})} \quad (9)$$

$$R_v(d, s, t, t_{max}) = \alpha_0 + \frac{\alpha_1 n^{(1-2\beta)/3} \left(\sum_i^n V_i(d, s, t) \right)^\beta}{\alpha_2 + \sum_i^n N_i(t, t_{max})}$$

where the sums are over the components of the stimulus ($n = 2$ for plaids, $n = 1$ for gratings), and the α_0 and α_1 parameters represent the spontaneous and maximum discharge rates of the cell. The relative gains of responses to grating and plaid are controlled by the $n^{(1-2\beta)/3}$ term in the numerator in Equation 9.

The normalization signal, $N_i(t, t_{max})$, is meant to approximate the effects of tuned normalization. In the original Simoncelli and Heeger cascade model, MT normalization signals were computed by summing over a simulated population of MT neurons, but this construction would be computationally prohibitive in the context of fitting the model to spiking data. We parameterize the tuning as follows:

$$N_i(t, t_{max}) = (1 - \gamma_0) \left(1 - \frac{t_i}{t_{max}} \right)^{\gamma_1} + \gamma_0. \quad (10)$$

This function is maximally active at zero temporal frequency (with a value of 1) and minimally active with a

value of γ_0 at t_{max} . t_{max} is the highest temporal frequency simulated and experimentally presented. We used this form of normalization for fitting tractability and because it has useful properties, namely, it: (1) ensures there is no suppression at the preferred temporal frequency, (2) can be completely disabled by setting $\gamma_0 = 1$, and (3) can be sublinear, linear, or super-linear.

The rationale for this particular approximation of MT tuned normalization is based on the following thought experiment. We start with the assumption that neurons in the MT normalization pool fall into component-, intermediate-, and pattern-selective subpopulations (Fig. 1C). Next, we assume all neurons' selectivities in each of these subpopulations evenly tile the space of spatiotemporal frequencies. Now we will consider the consequences of summing together the spatiotemporal frequency selectivities of all the neurons in each subpopulation. Specifically, the manner in which these selectivities sum together and overlap each other in frequency space determine if and how normalization is tuned. Any systematic biases in tuning overlap, as a function of spatiotemporal frequency, yield tuned normalization.

Component neurons from both model variants have narrow tuning, overlapping only between neurons with adjacent spatiotemporal tuning preferences. As a consequence, the summed responses of the population of component neurons evenly tile frequency space, producing an untuned normalization signal.

Frequency-separable pattern-selective neurons have broad direction tuning, so their overlap will occur most strongly in direction. The overlap, however, will be separable in spatial and temporal frequency, so for any subpopulation with the same spatial and temporal frequency tuning at all directions, the overlap will be confined to a donut-shaped region centered on those spatial and temporal frequencies. Since we assume the population of frequency-separable pattern selective neurons are evenly distributed across all preferred spatial and temporal frequencies, the tuning overlap will also be evenly distributed, yielding an untuned normalization signal.

Finally, velocity-separable pattern-selective neurons, which are organized along tilted planes which pass through the origin, will have strong overlap at zero and low temporal frequencies regardless of their preferred direction. As such, a pool of velocity-separable pattern neurons tiling frequency space generate a tuned normalization signal which strongly emphasizes low/zero temporal frequency. The same conclusions can be drawn for intermediate neurons in each separable model, although their selectivities will overlap less, yielding a similar, but weaker, tuned normalization signal.

Estimating model parameters for individual cells

In total, the model has nine free parameters for the single-grating study and 10 for the planar plaid study. For the former, they are: the direction preference and bandwidth (μ_d and σ_d), spatial frequency preference and bandwidth (μ_s and σ_s), temporal frequency preference, bandwidth, and log-linear transition (μ_t , σ_t , and τ), and the spontaneous and maximum firing rates (α_0 and α_1). For

the latter experiment, μ_s , σ_s , and μ_t are unconstrained by the data and are therefore held fixed at experimentally determined values, but the exponent (β), semi-saturation constant (α_2), and normalization parameters (γ_0 and γ_1) are free. To avoid model fits producing spuriously wide temporal frequency tuning, we included temporal frequency tuning data collected immediately prior in the fitting of the planar plaid dataset. That temporal frequency tuning data, along with the planar plaid stimuli which sample different directions, constrain μ_d , σ_d , and σ_t . In each study, the frequency-separable and velocity-separable models have the same parameters, and only differ in the parameterization of their temporal frequency linear weighting functions, w_t and v_t (Eqs. 4, 7).

For each cell, we optimized the model parameters by minimizing the negative log-likelihood (NLL) over the observed data, assuming spike counts arise from a modulated Poisson model. An additional parameter, σ_G , describes across-trial fluctuations in neural response gain (Goris et al., 2014) and was optimized to the data independently from the frequency-separable and velocity-separable models and held constant during model fitting. We performed the optimization in successive steps, using optimal values from one step as initialization values for the next. First, we fit τ , then added the rest of the MT linear weighting parameters, and then in the case of the planar plaid experiment, the MT parameters controlling the MT nonlinearity.

Experimental design and statistical analysis

In the single component study, we recorded single-unit responses of 13 V1 neurons and 39 MT neurons from seven anesthetized, paralyzed, adult male macaque monkeys (*M. fascicularis*) and one adult female macaque (*M. mulatta*). From those same eight monkeys, we recorded 21 V1 neurons and 53 MT neurons for the planar plaid study. These monkeys were also used for other visual system experiments not reported here, 12 of the 53 MT neurons reported had another experiment run after the single component and planar plaid studies were run. Of the 53 MT neurons in the planar plaid study, 23 were also in the single component study. All of the 13 V1 neurons from the single component study are in the set of 21 V1 neurons in the planar plaid study. We additionally recorded 58 MT neurons from two awake, actively fixating, adult male macaques for the planar plaid study (31 from *M. mulatta* “Albert” and 27 from *M. nemestrina* “LW”). In all studies, neurons were only excluded from analysis if spike isolation degraded during the experiment, or if spike rates were too low (e.g., always below 10 spikes/s) or variable to reliably predict direction tuning. No additional neurons were rejected from any subsequent analyses.

Following stimulus onset, we counted spikes within a 1000 ms window (anesthetized experiments) or a 250 ms window (awake experiments). For each cell, latency of these windows (relative to stimulus onset) were chosen by maximizing the sum of response variances computed for each stimulus condition (Smith et al., 2005). Error bars on tuning curve responses indicate ± 1 SD.

We used standard methods to compute each cell’s pattern index (Movshon et al., 1985; Smith et al., 2005). First, we computed partial correlations between the actual (constant temporal frequency) plaid responses and idealized predictions of pattern and component direction selectivity (r_p and r_c , respectively). We then converted these values to Z-scores to stabilize the variances of the correlations (Z_p and Z_c). Finally, the pattern index is the difference of these two quantities: $Z_p - Z_c$. Cells were classified as pattern selective if $Z_p - Z_c > 1.28$, or component-selective if $Z_c - Z_p > 1.28$. Both thresholds correspond to a significance of $p = 0.90$. Confidence intervals on pattern index were computed from the 95th percentile of 100 bootstrapped estimates (Efron and Tibshirani, 1993; Rust et al., 2006).

For optimal and non-optimal spatial and temporal frequency tuning curves in the single component study in Figure 2B–D, we fit a difference of log2-Gaussians (Hawken et al., 1996). For each neuron, the stimulus value corresponding to the peak of this fitted difference of log2-Gaussians function was used as the fitted preferred stimulus in the figure. To test the robustness of these tuning curve fits, we ran a bootstrap analysis in which trials from each tuning curve were pseudorandomly resampled 1000 times, with replacement, with the restriction that no stimulus condition had zero trials sampled. The error bars in Figure 2B–D represent the 95% confidence intervals of these bootstrapped fitted peak stimulus values. Some tuning curves had flat tops, yielding unreliable tuning preference estimates. We therefore excluded neurons (7 MT, 0 V1) from all analyses in Figure 2B–D, which had a confidence interval exceeding 1.5 decades in any of the three. The conclusions are the same with or without these neurons.

For the constrained parameter search during model fitting, we used a simplex algorithm (the MATLAB function ‘fmincon’). To avoid overfitting and obtain estimates of parameter stability (i.e., the error bars in Figs. 9A,B, 10), we fit the model on 100 bootstrap resamplings of the data. Bootstrapping was done on a per stimulus-condition basis; that is, trials within each stimulus condition were sampled with replacement, ensuring that there were no stimulus conditions without data. Error bars on model fits indicate ± 1 SE.

To compare model fits to a given neuron, we computed “velocity superiority”, the difference of the normalized NLLs of the velocity-separable and frequency-separable models. The NLLs were normalized by their corresponding “null” and “oracle” models, which serve as lower and upper bounds, at 0 and 1, respectively. The null model assumes the cell has two possible response rates: one when a stimulus is present and another when there is no stimulus. These are fixed to the measured mean spontaneous and maximal stimulus-driven response rates, respectively. The oracle model serves as an upper bound for the models’ performance, computed by using the measured mean responses to each stimulus condition to predict the neuron’s response to any individual presentation of that stimulus. We used the Wilcoxon signed-rank test to test velocity superiority significance and Pearson’s

r to assess correlation between velocity superiority and other quantities, such as pattern index.

Results

Joint and independent representations of motion in the frequency domain

Any image sequence can be decomposed (using a 3D Fourier transform) into a sum of sinusoidal gratings of differing orientation, and spatial and temporal frequency. A single point in this 3D frequency domain corresponds to a drifting sinusoidal grating with a unique orientation, spatial frequency, and temporal frequency (Fig. 1D). More complex spatial patterns contain mixtures of gratings of different orientations and spatial frequencies. If these patterns are rigidly translating over time, their frequency domain constituents lie on a tilted plane through the origin, the slope of which is equal to the object's speed (Fig. 1E; Watson and Ahumada, 1983, 1985).

How do V1 and MT neurons represent visual motion? Most V1 neurons are selective for a relatively narrow range of orientations, and spatial and temporal frequencies, corresponding to a ball in the frequency domain (Goris et al., 2015). If MT neurons are specialized for analyzing rigid motion, their receptive fields should be organized along just such a plane with slope equal to a preferred speed (Fig. 1F, "velocity-separable"; Simoncelli and Heeger, 1998). While there is some direct physiological evidence for velocity-separable organization (Rodman and Albright, 1987; Perrone and Thiele, 2001; Priebe et al., 2003; Nishimoto and Gallant, 2011), as well as perceptual evidence (Adelson and Movshon, 1982; Schrater et al., 2000), this is not the only kind of receptive field organization consistent with known MT properties.

However, almost all experimental measurements of grating direction selectivity use stimuli that lie along a horizontal plane of constant temporal frequency. By treating spatial and temporal frequency independently, they implicitly assume that MT direction selectivity is organized along these planes ("frequency-separable"; Fig. 1G). Evidence exists for this alternative possibility (Perrone and Thiele, 2001; Priebe et al., 2003), an MT neuron with this type of organization would still be direction-selective, but in a manner that is more strongly influenced by variations in spatial pattern.

These two model structures make different, testable predictions about how MT tuning should change in response to preferred and non-preferred stimuli. Experimenters typically assess a neuron's spatiotemporal frequency tuning preferences by presenting gratings varying along one of the three dimensions of the frequency domain, while keeping the values in the other two dimensions fixed at the best estimate of the neuron's preferences (Fig. 1H,I, black lines, for spatial frequency, temporal frequency, and direction, respectively). For a simulated neuron, the tuning curves (Fig. 1J,K, red and blue dashed lines) generated from these optimized stimuli have their peaks at the neuron's preferred spatiotemporal frequency (Fig. 1H,I, black points).

The two predictions differ most for tuning curves measured at non-preferred frequencies (Fig. 1H,I, dark gray lines and points), most notably in the tuning curves' peak

locations. The frequency-separable hypothesis predicts tuning in response to stimuli of non-preferred spatial and temporal frequency that is lower in amplitude but with a peak at the same frequency (Fig. 1J,K, blue lines). However, the velocity-separable hypothesis predicts that the tuning curve will shift (Fig. 1J,K, red lines), such that if the non-preferred tuning experiment is run at a frequency below preferred, the peak will also be at a lower frequency, and vice versa.

To test these hypotheses in V1 and MT, we measured tuning curves at optimal and suboptimal spatial and temporal frequencies and asked whether or not there was a shift in their peak location. For "suboptimal" frequencies, we used the stimulus values corresponding to the half-maximum responses when measured at optimal frequencies (for details, see Materials and Methods). Many cells (Fig. 2A), exhibited a peak spatial frequency tuning that increased with increases in grating temporal frequency, consistent with the velocity-separable hypothesis. To quantify this shift, and compare across neurons, we computed the peak spatial frequency and plotted it as a function of the relative temporal frequency at which it was measured (Fig. 2B). The degree to which the neuron is velocity tuned can be captured by the slope of the line through the data (0 for no speed tuning, 1 for ideal speed tuning, 0.37 for the neuron in Fig. 2A). V1 neurons show, on average, no slope (0.08 ± 0.22 SEM, $n = 13$; Fig. 2B, blue), while MT neurons have a significantly positive slope (0.50 ± 0.07 SEM, $n = 39$; Fig. 2B, red). Performing the same analysis for changes in temporal frequency preferences as a function of stimulus spatial frequency (Fig. 2C) yields similar slopes in MT (0.36 ± 0.03 SEM) and V1 (0.05 ± 0.04 SEM).

These measurements, all performed at the neuron's preferred direction, support previous findings that V1 tends to be frequency-separable and MT velocity-separable (Simoncelli and Heeger, 1998; Perrone and Thiele, 2001; Priebe et al., 2003, 2006; Nishimoto and Gallant, 2011). Since our goal was to characterize tuning in all three dimensions, we also assessed peak temporal frequency changes when measured at different directions (Fig. 2D), which should either remain constant or decrease (for the frequency-separable and velocity-separable hypotheses, respectively). When averaged across the populations, slopes were flat [Fig. 2D, V1 (blue triangles) mean -0.001 ± 0.004 SEM; MT (red triangles) mean -0.0004 ± 0.0007 SEM]; however, on a neuron-by-neuron basis, tuning at non-preferred directions was inconsistent. To probe the 3D selectivity more finely, we presented stimuli at many more spatiotemporal frequencies, and fit velocity-separable and frequency-separable models directly to the responses.

The velocity-separable and frequency-separable models

To examine MT receptive field organization in the frequency domain, we fit two modified versions of the Simoncelli and Heeger (1998) model of MT direction selectivity to the responses of individual neurons. Both models have the same structure: two stages, each with a linear weight-

Figure 2. At the preferred direction, V1 is frequency-separable and MT is velocity-separable. **A**, Spatial frequency tuning curve data from an example MT cell, measured at three temporal frequencies (error bars denote ± 1 SD). The light gray shaded area denotes the spontaneous firing rate, ± 1 SD. The fitted SF tuning preferences for the three curves are shown above as triangles. **B**, The fitted SF preferences from each cell are plotted against the TFs at which they were presented. Both axes are on a normalized scale, representing the ratio of the non-optimal frequencies, relative to the optimal frequency. Each line is the best fit line to the data for one cell. The data along the ordinate axis are aligned to the offset of each best fit line. Lines and points are shaded by the pattern index corresponding to each individual cell. Red corresponds to MT neurons, with darker shades corresponding to higher pattern index, and blue corresponds to V1 neurons, with darker shades corresponding to lower pattern index. The blue and red triangles indicate the mean slopes for all V1 and MT neurons, respectively. Error bars indicate the 95% confidence intervals of 1000 bootstrapped fitted peak stimulus values (for details, see Materials and Methods). **C**, Same as **B** but based on TF tuning curves measured at optimal and suboptimal SFs. **D**, Same as **C**, but based on TF tuning curves at optimal and suboptimal directions. Here, the points are aligned to the origin, which represents the preferred TF at the preferred direction. Mean slopes in **D** are computed separately for the different suboptimal directions.

ing followed by a point nonlinearity and normalization (Fig. 3). The first (V1) stage consists of narrowly-tuned direction-selective complex cells, simulated with a linear weighting of a narrow band of frequencies, followed by squaring. The second (MT) stage computes a weighted linear combination of its V1 inputs, followed by another point nonlinearity and normalization.

Linear weighting in the MT stage is the primary determinant of the MT neuron's tuning properties, including pattern motion selectivity. We constrain it to be a separable product of three tuning curves. The first two (direction and spatial frequency tuning) are common to both models. In the frequency-separable model, the third separable function is temporal frequency tuning, independent of the other two dimensions. In the velocity-separable model, temporal frequency tuning co-varies with direction tuning such that the peak lies on a tilted plane whose slope is the preferred velocity of the neuron. This temporal frequency tuning parameterization is the only difference between the two models.

The MT stage nonlinearity controls interactions between multiple spatiotemporal frequencies simultaneously present in the stimulus, and thus plays an important role in establishing pattern motion responses. In the full models, the MT nonlinearity is composed of a point-wise power function, followed by divisive normalization. The divisive normalization operates on a uniform population of pattern and component cells which, taken in aggregate, are assumed to uniformly cover direction and spatial frequency, while exhibiting tuning for temporal frequency (for details, see Materials and Methods). Single grating stimuli do not constrain this model component, and thus for the single grating study presented below, the exponent is fixed to a value of two.

Single grating responses do not differentiate the models

How can we distinguish the two models? We designed a study in which we measured 17 tuning curves, chosen on a neuron-by-neuron basis, to sample the frequency domain where the predictions of the models should devi-

Figure 3. The separable models. A stimulus is passed through a narrowly tuned V1 linear weighting, then squared and normalized. V1 output is then passed to the MT neuron, which applies either a frequency-separable or velocity-separable linear weighting, then raises the output to a super-linear power, and undergoes another stage of normalization. Finally, a modulated Poisson process determines spike variability.

Figure 4. Comparison of actual and model-predicted responses to single gratings for four example MT neurons. **A, B**, Two example component neurons, one better fit by the frequency-separable model (**A**) and one better fit by the velocity-separable model (**B**). **C, D**, Two example pattern neurons, one better fit by the frequency model (**C**) and one better fit by the velocity model (**D**). Measured spike rate mean and SD are shown in black. Velocity model predicted spike rates are shown in red, frequency model predictions in blue. All subsequent figures follow this color convention. Means are indicated by the dark lines, ± 1 SD by the lighter shaded areas. The first four columns from the left show four of the most informative tuning curves for distinguishing the models (for data and model predictions for all tuning curves measured for the neuron in **D**, see Extended Data Fig. 4-2; for details on all the tuning curves measured, see Extended Data Figure 4-1, Extended Data Table 4-3). In the scatter plots on the right, each point represents how well the frequency and velocity models predict the mean firing rate for one spatiotemporal frequency among the 225 presented across all experiments. Goodness of fit is expressed in terms of log likelihood under the modulated Poisson process, where values closer to zero indicate a better fit. The log likelihoods are normalized to a scale between 0 and 1, which represent the null and oracle model prediction log likelihoods, respectively (for details, see Materials and Methods). Each point is colored on a Fisher transformed scale (i.e., in units of SD). The difference between the velocity and model predictions for each neuron are summarized as a single value ($\Delta NLL = NLL_v - NLL_f$). The frequency and velocity model linear weightings for each example neuron are rendered in the rightmost column. All four neurons were recorded under anesthesia.

ate the most. The stimuli were full contrast sinusoidal gratings; five tuning experiments included a grating at the optimal spatiotemporal frequency, while 12 suboptimal tuning experiments did not (see Materials and Methods; for details, see Extended Data Fig. 4-1, Extended Data Table 4-3). By comparing how responses fall off as stimuli deviate from the preferred spatiotemporal frequency, a picture of 3D tuning should emerge in support of one model or the other.

For each cell, we fit the frequency and velocity models to data from all 17 tuning experiments simultaneously. Figure 4A–D shows four of the 17 tuning curves of the optimized model, fit to data from two example MT component neurons (Fig. 4A,B) and two MT pattern neurons

(Fig. 4C,D; for all 17 tuning curves, see Extended Data Fig. 4-2). As expected, the models make substantially different predictions for spatial and temporal frequency tuning (Fig. 4, first three columns), but not direction tuning (Fig. 4, fourth column). In the first two columns, for example, the velocity model predicts tuning peak shifts, whereas the frequency model does not.

Most tuning curves from each neuron are well fit by one of the two models (frequency model for Fig. 4A,C and velocity model for Fig. 4B,D), including changes in relative gain across tuning experiments. Relative model performance for each stimulus condition from all 17 tuning experiments (Fig. 4, rightmost column, points in the scatter plots) show that while some spatiotemporal frequen-

Figure 5. Single grating stimuli do not distinguish the two models. Fit quality, expressed as the normalized log likelihood of the velocity and frequency models, is plotted for each neuron on Fisher transformed axes. V1 neurons ($n = 13$) are shown with open circles, MT ($n = 39$) with closed circles. Blue, black, and red colors indicate whether a neuron is classed as component, intermediate, or pattern selective, respectively. Error bars denote SEM. On average, the two models are equally good at explaining the single grating MT data for any class of cells. The frequency model explains the V1 single grating data better.

cies strongly distinguish the two models, most do not. This reflects the fact that some tuning curves are well-described by both models (e.g., the constant-velocity grating direction tuning; Fig. 4, fourth column).

This range of behavior was observed across the population. We assessed overall fit quality on a cell-by-cell basis by normalizing the log likelihoods of the models to null and oracle models. The null model assumes the cell has two possible response rates: one when a stimulus is present and another when there is no stimulus. These are fixed to the measured mean spontaneous and maximal stimulus-driven response rates, respectively. The oracle model serves as an upper bound for the models' performance, computed by using the measured mean responses to each stimulus condition to predict the neuron's response to any individual presentation of that stimulus. "Velocity superiority" is the difference of the normalized log likelihoods of the velocity-separable model and the frequency-separable model (Fig. 5).

In general, V1 cells were better fit by the frequency model (average velocity superiority of -0.03 , $p = 0.0046$ Wilcoxon signed-rank test; Fig. 5, open circles). Most MT neurons were clearly better fit by one model or the other, but overall, neither model was significantly better (-0.005 mean difference, $p = 0.12$ Wilcoxon signed-rank test; Fig. 5, filled circles), regardless of pattern index.

Model fits to single grating responses provided further evidence that V1 neurons are frequency-separable, but were inconclusive for MT neurons. MT neurons tend to be velocity-separable for stimuli at the preferred direction (Fig. 2B,C) but have inconsistent tuning at off-directions (Fig. 2D). In theory, comparing direction tuning curves measured at either a given neuron's optimal velocity or

optimal temporal frequency should distinguish the models: they predict direction tuning bandwidth to be wider when the stimulus and model type match (e.g., velocity-separable model direction tuning for constant-velocity gratings should have wider bandwidth than tuning for constant frequency gratings). In fact, measured tuning to these two stimuli are nearly identical for component neurons, and slightly broader, on average, for constant velocity gratings for intermediate and pattern neurons (Fig. 6). These data provide more evidence that MT neurons are likely velocity-separable, using different tuning measurements at non-preferred directions.

Taking these observations into account, we concluded that single grating stimuli are not rich enough to fully distinguish the models. In particular, since only one spatiotemporal frequency is presented at a time, single gratings do not constrain the MT nonlinearity which underlies pattern computation. We therefore sought to use more complex stimuli and focused on sampling the frequency domain at non-preferred directions, those spatiotemporal frequencies which were the most informative in distinguishing the two models (both in theory and in practice).

Compound stimuli reveal velocity-separable organization in MT

Selectivity for pattern motion is a defining property of MT neurons. Since single gratings alone are not rich enough to characterize MT, we ran a second study in which direction tuning curves were measured for gratings and 120° plaids, presented either at a given neuron's optimal velocity or optimal temporal frequency (Fig. 7A,B). All stimuli were fixed at the neuron's optimal spatial frequency. These stimuli can be equivalently described as gratings and plaids drifting either in the preferred direction of the cell or along the direction normal to the mean orientation (constant velocity or constant frequency, respectively; Fig. 7A,B).

The two models make dramatically different predictions (Fig. 7C) for pattern-selective neurons: the frequency model predicts tuning for constant-velocity plaids to have a trough at the preferred pattern motion direction, and peak 90° from preferred. The velocity model predicts a near-constant, elevated response to all constant-velocity plaids.

Responses to this new stimulus family were complex enough to fully constrain the models' MT nonlinearity (a power function and divisive normalization; for details, see Materials and Methods). There were three free parameters in the nonlinearity that were fixed in the previous study. Since the spatial frequency preference and bandwidth and temporal frequency preference parameters were unconstrained by this dataset, they were fixed to values determined in preceding tuning measurements. In total, there was one additional free parameter fit compared to the single-grating study.

Qualitatively, the models predict that direction tuning should be flatter when the coordinate system of the model matches that of the stimuli. Two features of the measured responses stand out. First, constant frequency and constant velocity direction tuning curves to gratings are again

Figure 6. Velocity grating direction tuning tends to be slightly wider. **A**, Direction bandwidth (in degrees) for each neuron was calculated separately for constant- frequency and velocity grating tuning curves. is plotted for each neuron. Blue indicates a component neuron, black intermediate, and red a pattern neuron. V1 neurons ($n = 21$) are shown with open circles, MT ($n = 112$) with closed circles. This figure includes isomorphic data recorded from the next experiment, described in the Compound stimuli reveal velocity-separable organization in MT section. **B**, Differences of velocity and frequency grating bandwidth, by pattern classification. Proportions are expressed within each classification type, but including both V1 and MT neurons (open and filled stacked bars, respectively). Pattern and intermediate neurons have wider velocity grating direction bandwidth, significant below $p < 0.0005$ (Wilcoxon signed-rank test).

clearly indistinguishable for all cells (Fig. 8, two leftmost columns). Second, the pattern selective MT neuron (Fig. 8D) exhibits much wider direction tuning bandwidth for constant velocity plaids as opposed to constant frequency plaids (Fig. 8, fourth and third columns from the left), while the other cells show more similar tuning band-

width for the two plaid types. For all cells, both models capture grating responses well. However, the frequency model cannot account for the pattern selective neuron's responses to both types of plaids simultaneously (Fig. 8D, blue). The best it can do is pick a compromise direction tuning bandwidth that is too wide for constant frequency

Figure 7. Two-component “planar plaid” experiment design and predictions. Constant-velocity and constant-frequency direction tuning experiments were done with gratings and plaids. Constant-velocity plaids (**A**) were constructed by superimposing two gratings 120° apart and drifting at a temporal frequency determined by the optimal velocity plane. Constant-frequency plaids (**B**) were two gratings 120° apart superimposed and drifting at the optimal temporal frequency. The example plaids shown contain the same orientations, but have different component drift rates, and thus different perceived drift directions. **C**, For the two models matched in constant-frequency plaid direction tuning (red and blue dashed line), the velocity model (red) predicts a high response rate to all constant-velocity plaids. The frequency model (blue) is more narrowly tuned.

Figure 8. Comparison of actual and model-predicted responses to gratings and plaids for four example neurons. First five columns show data (points) and tuning curves predicted by the frequency-separable (blue) and velocity-separable (red) models. The first four columns are responses to gratings and plaids with constant frequency and velocity. The fifth column is temporal frequency data collected in a separate session, but included in the model fits. **A**, V1 component-selective neuron. **B**, MT component neuron. **C**, MT intermediate neuron. **D**, MT pattern-selective neuron. The right-most column shows goodness-of-fit across all stimulus conditions, next to renderings of the fitted models. See Figure 4 caption for details. Differences between the two model predictions become more apparent with increasing pattern selectivity. Neurons in **A–C** are from recordings done under anesthesia; **D** is from an awake recording.

plaids and too narrow for constant velocity plaids. The velocity model, on the other hand, is able to account for all the data simultaneously, including the different plaid tuning bandwidths. This pattern cell is the only one of the four example cells that has substantial differences in the frequency-separable and velocity-separable model predictions. The increasingly divergent model predictions as pattern index increases is further illustrated by the increasingly different MT linear weighting functions (Fig. 8, last column). The models make nearly identical predictions for narrowly tuned neurons, yielding velocity superiority indices at or near 0 (Fig. 8, scatter plots on right).

The relationship between velocity superiority and pattern index holds across cell populations. For the single grating data set, there is overall no significant correlation [for all cells (Pearson's $r = 0.05$, $p = 0.70$) or MT alone ($r = -0.01$, $p = 0.95$); Fig. 9A]. There was, however, a

significant negative correlation for V1 ($r = -0.70$, $p = 0.007$). Furthermore, there was no significant relationship between pattern index and the number of tuning curves per cell better fit by one model or the other (Pearson's $r = -0.22$, $p = 0.13$). In contrast, responses to plaid stimuli indicate a significant correlation between velocity superiority and pattern index, and (Pearson's $r = 0.34$, $p = 5.6e-5$; Fig. 9B); 86% of all pattern cells were better fit by the velocity model ($p = 1.4e-4$, Wilcoxon signed-rank test).

As a more direct test, we compared pattern selectivity of model predictions against the measured pattern selectivity of the cells. The velocity model accounts for the full range of pattern selectivity across the population (Pearson's $r = 0.60$; Fig. 9C). The frequency model, however, fails to produce any cells with pattern tuning (Pearson's $r = 0.50$; Fig. 9D), due to the compromises it must make

Figure 9. Compound stimuli reveal velocity-separable organization for pattern cells. **A, B**, Velocity superiority, or the difference of normalized log likelihoods between the velocity and frequency models, per cell as a function of pattern index. V1 cells appear as open circles, MT closed. Example cells featured in [Figures 4, 8](#) are highlighted in gray. Light and dark lines indicate the running mean, with a window of $\pm 1/3$ of cells in each population. Error bars indicate ± 1 SD, calculated from model fits to bootstrapped data (note most are smaller than the plotted points). **A**, On average, for the single grating dataset, neither model better explains the single grating MT data (**A**) for any class of cells ($n = 39$). The frequency model explains the V1 single grating data better ($n = 13$). **B**, Pattern cell responses to the compound stimulus dataset (V1: $n = 21$, MT: $n = 112$) are clearly better explained by the velocity model. Error bars indicate ± 1 SE. **C, D**, Observed and predicted pattern indices for each cell, derived from the compound stimulus dataset, for the velocity model (**C**) and frequency model (**D**). The velocity model can account for pattern index across all cell types, whereas the frequency model fails to predict the pattern selectivity of neurons classified as pattern-selective based on measured responses. Error bars indicate 95th percentiles, generated from pattern indices calculated by bootstrapping measured and predicted spike trains.

when fitting both constant frequency and constant velocity plaid responses simultaneously.

Model validation

Which characteristics are needed to describe the motion selectivity of a given neuron in MT? They are, in order of increasing complexity: (1) its preferred direction and speed, (2) the degree to which responses fall off as stimuli deviate from the preferred stimulus, and (3) the extent to which multiple overlapping motion components are treated independently or as a single, coherently moving pattern.

Experimentally, these attributes are established by identifying the stimulus that evokes the neuron's maximum response, the shape of tuning curves for direction, spatial frequency, and temporal frequency, and calculating the pattern index based on correlating (constant frequency) grating and plaid direction tuning. The recorded direction tuning curves we report ([Fig. 8](#)), with identical bandwidths for constant frequency and velocity gratings (all cells) and wider bandwidth tuning to constant velocity

plaids (pattern cells), provide a novel, fourth criterion for describing MT motion selectivity.

The velocity model accurately captures the first, second, and fourth attributes (direction tuning, response fall-off, and matched/unmatched frequency bandwidths) by accurately reproducing tuning curves ([Figs. 4, 8](#)). We verified that the velocity model also accounts for the third attribute (pattern motion selectivity), and accounts for the full range of pattern selectivity across the population ([Fig. 9C](#)). The frequency model, in addition to failing on the second and fourth criteria ([Fig. 8](#)), also fails on the third by failing to predict any pattern motion selective cells ([Fig. 9D](#)).

How does the velocity model provide a full account of motion selectivity? Capturing selectivity in all three frequency dimensions accounts for the first two attributes (direction tuning and response fall-off). Pattern selectivity, the third attribute, is controlled by increasing direction tuning bandwidth (Pearson's $r = 0.73$; [Fig. 10A](#)) and the exponent in the nonlinearity (Pearson's $r = 0.60$; [Fig. 10B](#)). Divisive normalization, which allows the model to adjust

Figure 10. Relationship between velocity-separable model parameters and pattern index. Pattern index is strongly correlated with direction tuning bandwidth (**A**) and the log of the MT nonlinearity's exponent (**B**). **C**, Pattern index is negatively correlated with the log2 of the semi-saturation, or uniform divisive normalization parameter. This means that for neurons with higher pattern index, the temporal frequency-dependent suppression is stronger.

the fourth attribute (matched/unmatched frequency bandwidths), does so by producing grating direction tuning curves with more similar bandwidths than would be predicted in its absence. The semi-saturation, or “uniform” divisive normalization parameter, is very weakly correlated with pattern index (on a log2 scale, Pearson's $r = -0.16$, $p = 0.07$; Fig. 10C; for details, see Materials and Methods). This means that for neurons with higher pattern index, the temporal frequency-dependent suppression is stronger.

Taken together, the two datasets and associated model fits reveal important aspects of MT computation. First, sinusoidal grating stimuli drifting in a neuron's preferred direction can reveal a frequency-separable receptive field organization in V1, and a velocity-separable one in MT. These stimuli, however, are not sufficient to reveal the nonlinear behaviors that distinguish direction selectivity observed in MT from that observed in V1. Second, compound stimuli, which constrain nonlinear receptive field behavior in MT, reveal receptive fields that are organized along a neuron's preferred velocity plane.

Discussion

To date, attempts to characterize MT motion selectivity have generally followed two distinct strategies. They focused either on how multiple superimposed spatiotemporal frequencies are integrated into a single, coherent drifting pattern, or on how tuning varies across multiple dimensions of the spatiotemporal frequency domain. Here, we present a model that unifies these two approaches in a common framework, and for the first time, generalizes previous findings to all three dimensions of the spatiotemporal frequency domain.

We recorded responses of a large population of neurons in both MT and V1 to simple stimuli specifically designed to extensively quantify tuning in the spatiotemporal frequency domain as well as tuning for pattern motion. We fit two compact two-stage models to each individual neuron in the population. We found temporal frequency tuning in MT to be much broader than that instantiated in the Simoncelli and Heeger (1998) model. Furthermore, by comparing two models' performance, we provide model-based evidence that MT neurons' selectivity is best described by a tilted, constant velocity plane

in the spatiotemporal domain. Finally, compound stimuli were necessary to reveal this organization; single sinusoidal gratings were not sufficient.

Motion computation in the velocity-separable model

The separable models we developed and tested are generalizations of previous models (Simoncelli and Heeger, 1998; Rust et al., 2006). The Simoncelli and Heeger (1998) model was constructed using populations of V1 and MT neurons, each having their own rectifying nonlinearities and divisive normalization. The second (MT) stage of the model linearly weighted the afferent signals from V1 along a tilted, constant velocity plane in the frequency domain. But this model was not explicitly fit to single cell data, and comparisons of predicted and measured tuning curves were qualitative. The Rust et al. (2006) paper used a simplified model variant that predicted (and was fit to) responses to gratings and plaids at a single temporal frequency. The paper showed that pattern selectivity could be explained by incorporating opponent suppression and direction-tuned normalization. By fitting to a more diverse set of stimuli, and a model that includes a full range of temporal frequencies, we find that selectivity for both speed and direction of moving patterns can be captured in a single model. Note that we have incorporated temporal frequency-dependent normalization in the MT stage [as opposed to the direction-tuned normalization of the Rust et al. (2006) model]. For parameter values optimized to neurons in the compound stimulus dataset, this tends to sharpen direction tuning for constant velocity gratings.

To characterize MT receptive field structure in all three dimensions of the frequency domain, it was not feasible to simulate entire populations of V1 and MT neurons. By restricting our stimuli to gratings and plaids which would not be affected by normalization in V1, we could avoid explicitly simulating the V1 stage. Rather, the model evaluated tuning directly based on the separable product of tuning curves along three dimensions in the frequency domain. Since all three tuning curves are exponential functions, the separable tuning volume and exponent approximately accounts for both the linear weighting stages and power function nonlinearities of V1 and MT. As a result, three computational elements, all implemented in

Figure 11. Effects of removing model elements for one example neuron and the population. **A–C**, Plots and renderings are from fits to the same neuron shown in Figure 8D. In **A**, **B**, **D**, three nested model fits are shown. Lighter shades denote fits with nonlinear elements removed from the full model. LN-N is the full model, which includes linear weighting (L), a point-wise power function nonlinearity (N), and “temporal frequency-dependent” normalization (-N). LN has no normalization at all, and L has no nonlinearity at the MT stage. Direction tuning data are shown as black points and lines for constant-frequency plaids (**A**) and constant-velocity plaids (**B**). The red and blue shaded curves show the different model fits to those data, with red and blue corresponding to the velocity and frequency models. The darkest traces are for the full model, the lighter ones for the model with normalization removed, and the lightest for the model with no MT nonlinearity. Each (nested) model was optimized separately. (**C**) The leftmost plots show the strength of normalization as a function of temporal frequency, for the velocity (red, top row) and frequency (blue, second row from top) models. The middle two renderings show the linear weights (at one level set) as a function of spatial and temporal frequency, at two different viewing angles. The temporal frequency scale in these renderings match that of the normalization plots on the left. The renderings on the right are a “birds-eye” view, showing the same weights as function of the two spatial frequency dimensions. **D**, Fit quality, expressed as the normalized log likelihood of the velocity and frequency models, is plotted for all component, intermediate, and pattern neurons (in blue, black, and red, respectively), for the three nested models (lighter shades indicate nonlinear model elements removed, see above).

the MT stage, determine how a given MT neuron responds to moving stimuli: linear weights, a point-wise power nonlinearity, and normalization.

The linear weights capture the first-order aspects of a given MT neuron’s tuning: its tuning preferences and a coarse estimate of tuning breadth. In fact, both frequency-separable and velocity-separable (“linear”) models can capture single grating tuning curve shape well. This is why the two separable models are indistinguishable, on average, when fit to the single grating dataset (Fig. 5). The model captures the continuum of selectivity for pattern motion in MT (characterized by a unimodal constant frequency plaid direction tuning curve; Figs. 8D, 11B) by simply increasing the direction tuning bandwidth in the linear weights (Simioncelli and Heeger, 1998). Direction tuning bandwidth (calculated from measured tuning curves) is correlated

with pattern index, as observed in our data (Pearson’s $r = 0.27$, $p = 0.0027$, $n = 112$) and previous studies (Pearson’s $r = 0.35$, $p < 0.0002$, $n = 788$; Wang and Movshon, 2016). If the velocity plaid responses are not included during model fitting, linear weighting alone in the velocity-separable model is sufficient to capture the unimodality of constant frequency plaid (pattern) direction tuning, but not in the frequency-separable model (unpublished data). However, when all data are included, linear weighting alone is not sufficient for either model to capture this unimodality (Fig. 11A, lightest red and blue traces). Further, the frequency-separable model severely underestimates constant velocity plaid responses (Fig. 11B). It is important to note that in order for the two separable models to make realistic and distinguishable predictions, temporal frequency tuning data were also included during

fitting; all models capture this tuning well (Fig. 8). Ultimately, the linear model fails by overestimating the tuning bandwidth to frequency plaids (Fig. 11A).

The original [Simoncelli and Heeger \(1998\)](#) model solved this problem by applying an expansive, point-wise nonlinearity in the MT stage. The nonlinearity, when fit to data, only changes tuning to compound stimuli. By adding a point-wise power function, both separable “linear-nonlinear” (LN) models achieve better, sharper constant frequency plaid tuning (Fig. 11A, medium blue and red traces). Frequency-separable predictions to constant velocity plaids (Fig. 11B, medium blue) however, worsen, since tuning for all mixture stimuli are sharpened by the power function. This model parameter is strongly correlated with measured pattern index (Fig. 10), indicating its role in pattern motion computation.

If the goal was simply to correctly reproduce direction tuning for constant frequency gratings and plaids, a separable model including linear weighting and a power function nonlinearity alone would be sufficient. However, the unexpected, nonlinear property of MT selectivity that we discovered is that direction tuning bandwidth is wider for constant velocity gratings than for constant frequency gratings (Fig. 6), but by much less than expected, and by much less than is the case for measured constant velocity plaid tuning (Fig. 8D). To account for this small change in grating tuning and large change in plaid tuning, the separable models require normalization at the MT stage.

We used a closed-form approximation of MT normalization, rather than simulating an entire population of MT neurons [as was done in the original [Simoncelli and Heeger \(1998\)](#) model; for details, see Materials and Methods], to make model fitting tractable. Despite being an approximation, it is a functionally interpretable one. Its primary effect is to suppress responses at low temporal frequencies; its tuning relative to the linear weights is plotted in Figure 11C. Suppression for low temporal frequencies has been observed experimentally ([Maunsell and Van Essen, 1983](#)). Incorporating normalization improves contrast gain control (darkest blue and red tuning curves are better scaled to the data in Fig. 11A,B). This normalization also sharpens tuning to both single gratings and conjunctions of gratings by concentrating suppression at low temporal frequencies. In the case of pattern neuron responses, these conjunctions are components consistent with a preferred velocity, so velocity-separable model predictions for constant velocity plaids are appropriately widely tuned (Fig. 11B) and frequency-separable model predictions for constant frequency plaids are too widely tuned (Fig. 11A).

For each nested version of the models, namely, the linear, linear-nonlinear, and the full model, the velocity-separable version performs better on pattern neurons as a group (Fig. 11D). For component and intermediate neurons, the two types of model are indistinguishable regardless of model version.

Relationship to previous work

[Perrone and Thiele \(2002\)](#) observed broader temporal frequency tuning than predicted by the [Simoncelli and](#)

[Heeger \(1998\)](#) model. Their weighted intersection mechanism (WIM) model was able to capture joint spatial and temporal frequency tuning in MT. It employed a weighting function on V1 inputs organized along a common speed, and only responded when two types of V1 inputs, “sustained” and “transient,” had equal response levels. With particular choices of parameters, the WIM model can also simulate pattern direction selectivity ([Perrone, 2004](#)). The authors stated that a model with velocity-based tuning at the MT stage, such as in [Simoncelli and Heeger \(1998\)](#), would not be capable of producing realistic spatiotemporal tuning. Here we fit just such a velocity-based separable model directly to pattern motion data and to data simultaneously spanning all three dimensions of frequency space. It achieved similar realism in reproducing tuning in both temporal frequency and pattern motion direction, each recorded from heterogeneous populations of neurons, without the need for the two specific V1 neuron types.

[Priebe et al. \(2003\)](#) investigated joint tuning for spatial and temporal frequency and pattern motion selectivity in MT. Consistent with our findings, they reported stronger evidence for speed tuning with compound stimuli such as plaids or square-wave gratings, as compared to single sinusoidal gratings. Additionally, speed tuning for single sinusoidal gratings and degree of pattern selectivity were independent. They concluded that speed tuning arises in MT only when multiple spatial frequencies are present. Our findings are consistent with these conclusions, and arise in our velocity-separable model simulations as a result of the MT normalization. Our model additionally predicts that pattern tuning will be correlated with speed tuning for square wave gratings and random dots ([Kumano and Uka, 2013](#); [McDonald et al., 2014](#); [Xiao and Huang, 2015](#)).

More recent studies ([Nishimoto and Gallant, 2011](#); [Inagaki et al., 2016](#)) have explored MT selectivity in all three dimensions of the frequency domain. [Nishimoto and Gallant \(2011\)](#) used “motion-enhanced” natural movies to visualize 3D spectral receptive fields of individual MT neurons for the first time. These weights followed a simulated V1 population, which performed linear filtering, a compressive nonlinearity, and divisive normalization. They reported weights with excitation organized along a partial ring on the plane, with a gap in the ring occurring at low temporal frequencies. Suppression also appeared as partial rings off the preferred velocity plane, much like the opponent suppression reported in [Rust et al. \(2006\)](#). [Inagaki et al. \(2016\)](#) performed linear regression directly on the frequencies of their stimulus, which was comprised of multiple gratings superimposed spatially and partially overlapping in time. They observed broadly tuned receptive fields at mid- and high temporal frequencies in two pattern cells and observed diffuse suppression off the preferred velocity plane. The absence of excitation at low temporal frequencies observed in both studies provides indirect support for our use of suppression there.

Neither [Nishimoto and Gallant \(2011\)](#) nor [Inagaki et al. \(2016\)](#) directly confirmed that their models could produce pattern tuning, making the connection between the recep-

tive field structure they observed and pattern selectivity harder to interpret. The velocity separable model is able to reproduce pattern tuning in both frequency-separable and velocity-separable coordinates, while making slightly different predictions of receptive field structure. Pattern cells have excitation on a full ring on the preferred velocity plane, with partially overlapping suppression at low temporal frequencies (Fig. 11*F*).

Including normalization only at the V1 stage (Rust et al., 2006; Nishimoto and Gallant, 2011) or using a purely subtractive form of suppression in MT (as in all three aforementioned studies) in the separable models was not sufficient to simultaneously account for the broad tuning observed for constant velocity plaids.

Conclusions drawn from the separable model

Selectivity to moving patterns is a hallmark of MT response. How does this selectivity arise? Orientation selectivity in V1 provides a useful analogy. There, first-order properties of selectivity to simple stimuli, such as simple/complex classification, can be attributed to linearly weighting of LGN afferents (Reid and Alonso, 1995; Goris et al., 2015). Responses to compound stimuli, however, are likely a result of additional (possibly recurrent) computation within V1. Likewise, basic MT direction selectivity along the component/pattern continuum can be constructed by appropriate summing of V1 inputs (on a constant velocity plane) and shaped on a per-neuron basis by their own point-wise nonlinearities. Further nonlinear tuning behaviors, such as the different tuning bandwidths for constant velocity gratings and plaids, is likely shaped by recurrent computation within MT.

There is some evidence of recurrent computation shaping pattern motion signals in MT. Using drifting fields of bars, Pack and Born (2001) showed that pattern motion tuning emerges later in a pattern neuron's response, ~70 ms after its earliest response to stimulus onset, a result later replicated with sinusoidal gratings and plaids (Smith et al., 2005; Solomon et al., 2011). Additional experiments could be done to verify this recurrent computation prediction. If feasible, imaging a population of MT neurons and fitting a population-level model could reveal these recurrent computations, as has been done in V1 (Cossell et al., 2015; Antolík et al., 2016; Klindt et al., 2017). Examining dynamics of tuning to compound stimuli, possibly with whole-cell recording techniques, could also provide empirical evidence regarding the nature of suppression in MT.

While the velocity separable model unifies data and theory regarding tuning for pattern direction and velocity, there is much work to be done to further incorporate other aspects of MT selectivity into the model. The velocity separable model includes rudimentary gain control, and we used stimuli which only had two different contrast values. However, accounting for gain control in MT, and its interactions with pattern motion selectivity, motion opponency, and stimulus size (Britten and Heuer, 1999; Heuer and Britten, 2002), will likely require more experiments, and perhaps inclusion of a full normalization pool at the MT stage.

A strict interpretation of the Simoncelli and Heeger (1998) model predicts broad direction tuning to both constant velocity gratings and plaids, yet we only observed broad tuning to the latter (Figs. 6, 8*D*). This is consistent with later findings (Priebe et al., 2003, 2006) suggesting that some speed tuning in MT is inherited from V1, but that full form-independent speed computation occurs within MT, and is evident only when multiple spatial frequencies are present. Our results further suggest that individual MT pattern neurons always signal motion direction, but only signal speed when it is uniquely specified (i.e., when multiple orientations or spatial frequencies are present).

Finally, our findings change our understanding of the degree to which MT corresponds with motion perception. Consider a single drifting contour or grating, viewed through an aperture. The true direction of motion is inherently ambiguous: any drift direction $\pm 90^\circ$ from normal is a valid interpretation. Perceptually, however, this so-called "aperture problem" is unambiguously solved: the grating is perceived to be drifting in the direction normal to its orientation (Stumpf, 1911; Wohlgenuth, 1911; Wallach, 1935; Marr and Ullman, 1981; Adelson and Movshon, 1982; Todorović, 1996). Previously, it was thought that pattern selective neurons in MT, as a population, would signal a single grating's drift direction ambiguously (Movshon et al., 1985; Simoncelli and Heeger, 1998). Our findings show MT pattern neurons can unambiguously signal such motion, and that such a population can represent the translational motion of a stimulus regardless of whether it contains a mixture of orientations or a single one. The representation of motion in MT may thus be even closer to perception than previously thought.

References

- Adams DL, Economides JR, Jocson CM, Horton JC (2007) A bio-compatible titanium headpost for stabilizing behaving monkeys. *J Neurophysiol* 98:993–1001.
- Adams DL, Economides JR, Jocson CM, Parker JM, Horton JC (2011) A watertight acrylic-free titanium recording chamber for electrophysiology in behaving monkeys. *J Neurophysiol* 106:1581–1590.
- Adelson EH, Movshon JA (1982) Phenomenal coherence of moving visual patterns. *Nature* 300:523–525.
- Antolík J, Hofer SB, Bednar JA, Msríc-Flogel TD (2016) Model constrained by visual hierarchy improves prediction of neural responses to natural scenes. *PLoS Comput Biol* 12:1–22.
- Britten KH, Heuer HW (1999) Spatial summation in the receptive fields of MT neurons. *J Neurosci* 19:5074–5084.
- Carandini M, Heeger DJ, Movshon JA (1997) Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci* 17:8621–8644.
- Cavanaugh JR, Bair W, Movshon JA (2002) Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *J Neurophysiol* 88:2530–2546.
- Cossell L, Iacarus MF, Muir DR, Houlton R, Sader EN, Ko H, Hofer SB, Msríc-Flogel TD (2015) Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* 518:399–403.
- Derrington A, Lennie P (1984) Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. *J Physiol* 357:219–240.
- Dubner R, Zeki SM (1971) Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Res* 35:528–532.

- Efron B, Tibshirani R (1993) An introduction to the bootstrap. Boca Raton, FL: Chapman & Hall/CRC.
- Enroth-Cugell C, Robson JG, Schweitzer-Tong DE, Watson AB (1983) Spatio-temporal interactions in cat retinal ganglion cells showing linear spatial summation. *J Physiol* 341:279–307.
- Goris RLT, Movshon JA, Simoncelli EP (2014) Partitioning neuronal variability. *Nat Neurosci* 17:858–865.
- Goris RLT, Simoncelli EP, Movshon JA (2015) Origin and function of tuning diversity in macaque visual cortex. *Neuron* 88:1–13.
- Hawken MJ, Shapley RM, Gross DH (1996) Temporal-frequency selectivity in monkey visual cortex. *Vis Neurosci* 13:477–492.
- Heuer HW, Britten KH (2002) Contrast dependence of response normalization in area MT of the rhesus macaque. *J Neurophysiol* 88:3398–3408.
- Inagaki M, Sasaki KS, Hashimoto H, Ohzawa I (2016) Subspace mapping of the three-dimensional spectral receptive field of macaque MT neurons. *J Neurophysiol* 116:784–795.
- Johnston JM, Cohen YE, Shirley H, Tsunada J, Bennur S, Christison-Lagay K, Veeder CL (2016) Recent refinements to cranial implants for rhesus macaques (*Macaca mulatta*). *Lab Animal* 45:180–186.
- Klindt DA, Ecker AS, Euler T, Bethge M (2017) Neural system identification for large populations separating “what” and “where.” *Adv Neural Inf Process Syst* 31:4–6.
- Kumano H, Uka T (2013) Responses to random dot motion reveal prevalence of pattern-motion selectivity in area MT. *J Neurosci* 33:15161–15170.
- Marr D, Ullman S (1981) Directional selectivity and its use in early visual processing. *Proc R Soc Lond B Biol Sci* 211:151–180.
- Maunsell JHR, Van Essen DC (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *J Neurophysiol* 49:1127–1147.
- McDonald JS, Clifford CW, Solomon SS, Chen SC, Solomon SG (2014) Integration and segregation of multiple motion signals by neurons in area MT of primate. *J Neurophysiol* 111:369–378.
- Movshon JA, Adelson EH, Gizzi MS, Newsome WT (1985) The analysis of moving visual patterns. In: *Pattern recognition mechanisms (Pontificiae Academiae Scientiarum Scripta Varia)* (Chagas C, Gattass R, Gross CG, eds), Vol 54, pp 117–151. Rome: Vatican Press.
- Nishimoto S, Gallant JL (2011) A three-dimensional spatiotemporal receptive field model explains responses of area MT neurons to naturalistic movies. *J Neurosci* 31:14551–14564.
- Pack CC, Born RT (2001) Temporal dynamics of a neural solution to the aperture problem in visual area MT of macaque brain. *Nature* 409:1040–1042.
- Perrone JA (2004) A visual motion sensor based on the properties of V1 and MT neurons. *Vision Res* 44:1733–1755.
- Perrone JA, Thiele A (2001) Speed skills: measuring the visual speed analyzing properties of primate MT neurons. *Nat Neurosci* 4:526–532.
- Perrone JA, Thiele A (2002) A model of speed tuning in MT neurons. *Vision Res* 42:1035–1051.
- Priebe NJ, Cassanello CR, Lisberger SG (2003) The neural representation of speed in macaque area MT/V5. *J Neurosci* 23:5650–5661.
- Priebe NJ, Lisberger SG, Movshon JA (2006) Tuning for spatiotemporal frequency and speed in directionally selective neurons of macaque striate cortex. *J Neurosci* 26:2941–2950.
- Reid R, Alonso J (1995) Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* 378:281–283.
- Rodman HR, Albright TD (1987) Coding of visual stimulus velocity in area MT of the macaque. *Vision Res* 27:2035–2048.
- Rust NC, Mante V, Simoncelli EP, Movshon JA (2006) How MT cells analyze the motion of visual patterns. *Nat Neurosci* 9:1421–1431.
- Schreiner PR, Knill DC, Simoncelli EP (2000) Mechanisms of visual motion detection. *Nat Neurosci* 3:64–68.
- Sclar G, Maunsell JHR, Lennie P (1990) Coding of image contrast in central visual pathways of the macaque monkey. *Vision Res* 30:1–10.
- Simoncelli EP, Heeger D (1998) A model of neuronal responses in visual area MT. *Vision Res* 38:743–761.
- Smith MA, Majaj NJ, Movshon JA (2005) Dynamics of motion signaling by neurons in macaque area MT. *Nat Neurosci* 8:220–228.
- Solomon SS, Tailby C, Gharaei S, Camp AJ, Bourne JA, Solomon SG (2011) Visual motion integration by neurons in the middle temporal area of a New World monkey, the marmoset. *J Physiol* 589:5741–5758.
- Stumpf P (1911) Über die Abhängigkeit der visuellen Bewegungsrichtung und negativen Nachbildes von den Reizvorgängen auf der Netzhaut. *Z Psychol* 59:321–330.
- Todorović D (1996) A gem from the past: Pleikart Stumpf’s (1911) anticipation of the aperture problem, Reichardt detectors, and perceived motion loss at equiluminance. *Perception* 25:1235–1242.
- Tolhurst DJ, Movshon JA (1975) Spatial and temporal contrast sensitivity of striate cortical neurones. *Nature* 257:674–675.
- Van Essen D, Maunsell J, Bixby J (1981) The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization. *J Comp Neurol* 199:293–326.
- Wallach H (1935) Über visuell wahrgenommene Bewegungsrichtung. *Psychol Forsch* 20:325–380.
- Wang HX, Movshon JA (2016) Properties of pattern and component direction-selective cells in area MT of the macaque. *J Neurophysiol* 115:2705–2720.
- Watson AB, Ahumada AJ (1983) A look at motion in the frequency domain. In: *Motion: representation and perception*, pp 1–10. Baltimore: ACM.
- Watson AB, Ahumada AJ (1985) Model of human visual-motion sensing. *J Opt Soc Am A* 2:322–342.
- Wohlgemuth A (1911) On the after-effect of seen movement. *Br J Psychol* 1:1–117.
- Xiao J, Huang X (2015) Distributed and dynamic neural encoding of multiple motion directions of transparently moving stimuli in cortical area MT. *J Neurosci* 35:16180–16198.