

available at www.sciencedirect.comwww.elsevier.com/locate/brainres
**BRAIN
RESEARCH**

Research Report

A model of encoding and decoding in V1 and MT accounts for motion perception anisotropies in the human visual system

Ariel Rokem^a, Michael A. Silver^{a,b,*}

^aHelen Wills Neuroscience Institute, University of California, Berkeley, USA

^bSchool of Optometry, University of California, Berkeley, USA

ARTICLE INFO

Article history:

Accepted 8 July 2009

Available online 30 July 2009

Keywords:

Vision

Motion perception

Psychophysics

Neural coding

Computational model

Computational cognitive

neuroscience

Adaptation

Motion aftereffect

ABSTRACT

We used the motion aftereffect (MAE) to psychophysically characterize tuning of motion perception in the human visual system. The function relating MAE strength and the range of directions present in the adapter stimulus provides information regarding the width of direction tuning of motion adaptation. We compared the directional anisotropy in MAE tuning width to the well-known oblique effect in motion direction discrimination. In agreement with previous research, we found that subjects had lower motion direction discrimination thresholds for cardinal compared to oblique directions. For each subject, we also estimated MAE tuning width for a cardinal and an oblique direction by measuring the strength of the MAE for adapter stimuli containing different directional variances. The MAE tuning width was smaller for the cardinal direction, suggesting a fundamental similarity between motion direction discrimination and tuning of the MAE. We constructed a model of encoding of motion stimuli by V1 and MT and decoding of stimulus information from the cells in area MT. The model includes an anisotropy in the representation of different directions of motion in area V1. As a consequence of the connections implemented in the model, this anisotropy propagates to cells in MT. Model simulations predicted an oblique effect for both direction discrimination thresholds and MAE tuning width, consistent with our experimental results. The model also concurs with a recent report that the magnitude of the oblique effect for direction discrimination is inversely proportional to the directional variance of the stimulus. The agreement between model predictions and empirical data was obtained only when the model employed a maximum likelihood decoding algorithm. Alternative decoding mechanisms such as vector averaging and winner-take-all failed to account for the psychophysical results.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

Performance in visual tasks is often asymmetric, depending on the location, orientation, and/or motion direction of visual

stimuli. In some cases, these differences in performance may stem from asymmetries that exist in the natural environment and can provide insight into the developmental origins of perceptual and behavioral asymmetries (Dakin et al., 2005). In

* Corresponding author. 360 Minor Hall, #2020, School of Optometry, University of California, Berkeley, CA 94720-2020, USA. Fax: +1 510 643 5109.

E-mail address: masilver@berkeley.edu (M.A. Silver).

URL: <http://argentum.ucbso.berkeley.edu> (M.A. Silver).

Abbreviations: MAE, motion aftereffect; RDK, random dot kinetogram

addition, these asymmetries may be used to illuminate the mechanisms of neural encoding and decoding underlying the performance of visual tasks. In this work, we have used anisotropies in motion perception to investigate encoding and decoding of motion stimuli by the human visual system.

Thresholds for perceptual tasks performed on moving stimuli or on oriented stimuli are often lower for stimuli with orientation or direction of motion that is parallel to the cardinal axes (up/down, left/right) than for stimuli oriented or moving along the oblique directions (the off-cardinal diagonals), a phenomenon referred to as the oblique effect (Appelle, 1972). This behavioral anisotropy probably stems from a more robust representation of cardinal orientations in the visual system. Furmanski and Engel (2000) and Furmanski et al. (2004) showed that the oblique effect in detection of low-contrast gratings (lower detection contrast threshold for cardinal than for oblique orientations) was correlated with a difference in the magnitude of primary visual cortical fMRI responses to presentation of cardinal and oblique gratings. In addition, in a large sample of cat primary visual cortical neurons, randomly sampled in many different experiments, there were more cells preferring cardinal than cells preferring oblique orientations (Li et al., 2003). In motion perception, thresholds for discriminating two similar motions of direction are higher when the stimuli are centered at oblique directions compared to cardinal directions (Ball and Sekuler, 1982; Dakin et al., 2005; Gros et al., 1998). By analogy with the oblique effect for stimulus orientation, we assume that the oblique effect for motion perception is also based on an anisotropy in the representations of different motion directions in the visual system. A significant proportion of cells in primary visual cortex is not only orientation-selective but also direction-selective (De Valois et al., 1982, 2000; Hubel and Wiesel, 1959; Peterson et al., 2004). The preferred direction and preferred orientation are always approximately orthogonal in macaque V1 cells, based on responses to moving bar stimuli (Albright, 1984). 2D motion direction information may not always be available to the cell, due to the aperture problem (Horn, 1986). However, when 2D motion direction information is available to V1 neurons, preferred direction is independent of stimulus orientation (Pack et al., 2003). Therefore, it is reasonable to assume that there are more cells in V1 that show a preference for cardinal motion directions than cells that prefer oblique directions. Moreover, the average orientation tuning width of primary visual cortical neurons tuned to cardinal orientations was smaller than the average tuning width of those tuned to oblique orientations (Li et al., 2003). Therefore, the average tuning width of motion selectivity is likely to be smaller for cells representing the cardinal directions compared to cells preferring oblique motions, though this has not yet been tested experimentally in primary visual cortex.

We used two tasks to characterize the oblique effect in motion perception. The first, a motion direction discrimination task, exhibited an oblique effect in direction discrimination threshold and was used to identify the cardinal direction associated with lowest discrimination threshold and the oblique direction associated with highest threshold in each of our subjects. We then measured the tuning width of motion adaptation for these two directions. Estimates of the tuning width were obtained by measuring the strength of adaptation

(magnitude of the motion aftereffect, or MAE) following prolonged viewing of a field of coherently moving dots in one of the two directions. Previous work has shown that the magnitude of the MAE for random dot kinetogram (RDK) adapter stimuli was greater when the adapter stimulus included a moderate range of directions compared to a single direction of motion (Hiris and Blake, 1992). Thus, the relationship between MAE strength and the range of directions in the adapter stimulus allows estimation of the width of direction tuning of motion perception.

In our experiments, the RDK adapting stimuli were generated by assigning a direction to each dot from a distribution of directions centered on either a cardinal or oblique direction. The variance of this distribution determines the directional variance of the stimulus. Our results show that like motion direction discrimination performance, the tuning width of motion adaptation also exhibited an oblique effect: direction tuning was sharper for cardinal adapter stimuli than for oblique stimuli.

We constructed a computational model of encoding and decoding of motion information by cells in areas V1 and MT that accounts for the observed oblique effects in motion direction discrimination and tuning width of motion adaptation. The model contains a set of V1 units with feedforward connections to a set of MT units. The V1 units are anisotropic in their representation of motion: V1 cells representing cardinal directions are more numerous, and their directional tuning widths are narrower than the tuning widths of V1 cells representing oblique directions. The tuning properties of MT cells are then inherited through feedforward projections from V1 cells.

Information about stimulus motion direction is then decoded from the activity in the entire population of MT cells (as in Pouget et al., 2000). The decoding method is based on a maximum likelihood procedure (Jazayeri and Movshon, 2006). Our model quantitatively accounts for the observed psychophysical results, generating oblique effects for motion discrimination and for motion adaptation tuning width. It also agrees with previous findings that the oblique effect for motion discrimination is only present for stimuli with low directional variance (Dakin et al., 2005).

Our modeling results demonstrate that oblique effects in motion perception could arise from a combination of an anisotropy in the encoding of the stimulus by the visual system and a decoding mechanism that employs a statistically optimal strategy to read out this information. This suggests that complex perceptual phenomena such as the oblique effect should be understood as a consequence of specific encoding and representation schemes as well as specific decoding strategies employed by the brain.

2. Results

2.1. The oblique effect in motion direction discrimination

To compare perceptual abilities for different directions of motion, we employed a motion direction discrimination task. Subjects viewed an annulus centered at the fixation point and containing a random dot kinetogram (RDK). For each trial, two

RDKs were presented in succession. Subjects were required to press a button to indicate whether the RDKs were moving in the same or different directions (Fig. 1A). For half of the trials, the RDKs were moving in the same direction in both intervals. For the other half of the trials, the motions were different, separated by a small angle α . The magnitude of α was adaptively adjusted based on a psychophysical staircase and according to the subject's previous performance. Discrimination thresholds were obtained for each subject for eight different directions. We found a robust and reliable oblique effect in the direction discrimination task: the mean threshold ($\sim 80\%$ performance) for direction discrimination was $12.4 \pm 9.0^\circ$ for cardinal directions and $17.9 \pm 10.9^\circ$ for oblique directions (Fig. 2). This difference was statistically significant (within-subject paired *t*-test, $n=16$, $p < 0.001$) and replicates




Fig. 1 – Task design. A: Motion direction discrimination task. Subjects observed motion in a standard direction, followed by either motion in the standard direction or motion in a direction similar but not identical to the standard direction. During the intertrial interval, subjects reported whether the two stimuli were moving in the same direction or not. **B: Motion aftereffect task.** Adapting RDK stimuli spanned a range of variances of motion directions. The directional variance was controlled by drawing the motion direction of each dot from Gaussian distributions with different widths (standard deviations). **C: Motion aftereffect task.** Subjects were initially presented with 40 s of an adapting stimulus. Then, at the beginning of each trial, there was an additional period of top-up adaptation. Subjects then made a direction judgment on a probe RDK with low motion coherence. The strength of adaptation was determined by measuring the amount of coherence that was needed in order to counteract the MAE.




Fig. 2 – The oblique effect in motion direction discrimination. Average thresholds (± 1 SEM) from the direction discrimination task (Fig. 1A) are presented. There was a robust oblique effect—mean thresholds for cardinal directions were always lower than mean thresholds for oblique directions. Numbers surrounding the plot represent angular directions of motion in the standard directions; numbers within the plot represent thresholds, expressed as the angular difference between two stimuli at threshold (in units of degrees).

previous findings of a robust oblique effect in similar tasks (Ball and Sekuler, 1982; Gros et al., 1998).

2.2. The oblique effect in direction tuning width of motion adaptation

In order to characterize width of motion direction tuning in the visual system, we used the motion aftereffect (MAE; also known as the 'waterfall effect'), whereby prolonged viewing of a moving adapter stimulus causes subjects to have a perceptual bias towards perceiving motion in the opposite direction of the adapting stimulus (Anstis et al., 1998). We manipulated the directional variance of the adapting RDK by varying the standard deviation of the distribution from which dot directions were assigned (Fig. 1B).

Initially, the adapting stimuli were presented at 100% coherence for 40 s. Each trial began with 4 s of top-up adaptation, followed by a second probe RDK with motion either in the adapting direction or in the opposite direction. Subjects discriminated the direction of motion in this probe stimulus (Fig. 1C). We insured that the discrimination was made at threshold by adjusting the proportion of coherently moving dots in the probe stimulus based on a psychophysical staircase.

When coherence of the probe stimulus was very low, it appeared to be moving in the opposite direction from the direction of the adapting stimulus due to the MAE. However, when the coherence of the physical motion present in the stimulus was increased, the MAE was eventually overcome. The proportion of coherent dots in the post-adaptation probe stimulus was adjusted for each trial according to the subject's previous responses, and the threshold ($\sim 70\%$ of responses

corresponding to perception of movement in the same direction as the adapting stimulus, in units of percent coherent dots) served as a measure of the strength of motion adaptation (Blake and Hiris, 1993). For each subject, thresholds were computed for eight different adapting stimuli that spanned a range of directional variances. Additionally, each subject performed the task for two different adapter directions: the cardinal direction in which motion direction discrimination performance was best and the oblique direction in which motion direction discrimination performance was worst. In all but two subjects, this pair of directions corresponded to the directions in which the subjects achieved their best and worst direction discrimination performance across all eight directions.

When the standard deviation of the adapting stimulus was zero (all dots in the adapting RDK moved in the same direction), there was no significant difference in motion adaptation magnitude for cardinal and oblique directions (Fig. 3A). However, when the standard deviation of the

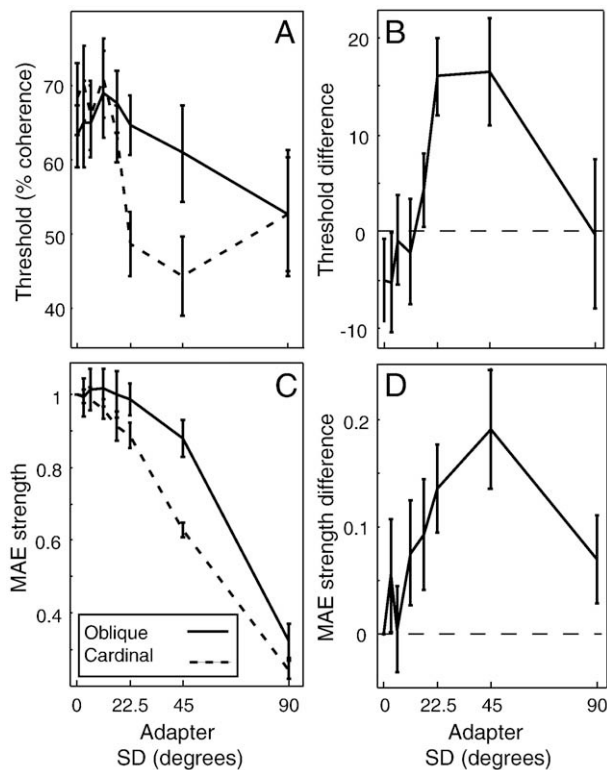


Fig. 3 – The oblique effect for motion adaptation. A: Average thresholds (± 1 SEM) for the MAE task are presented. The threshold in this task corresponded to the strength of the MAE for a given combination of adapter direction and standard deviation of motion directions. Solid line: oblique directions, dashed line: cardinal directions. B: Average difference (± 1 SEM) between MAE strength for oblique and cardinal directions. C: Model simulations of the MAE task provide an excellent fit of the experimental data. Average model thresholds (± 1 standard deviation for 10 repetitions of the simulation) are presented. D: The differences between the conditions in the model match the experimental data shown in panel B.

adapting stimulus was 22.5 or 45 degrees, a significant oblique effect was observed, with the oblique adapters resulting in stronger adaptation than cardinal adapters ($p < 0.05$). When the standard deviation was very large (90°), substantially less adaptation was observed for either adapting direction. This MAE oblique effect can be represented as the difference between oblique and cardinal MAE strength (Fig. 3B). These results indicate that the ‘optimal width’ for adaptation differs between the oblique and cardinal directions. For oblique directions, there was still significant adaptation even for adapting stimuli with widths of 22.5 and 45°, while much less adaptation was observed for these widths for cardinal adapters.

2.3. A model of encoding and decoding of motion direction in V1 and MT

To better understand the mechanisms underlying these psychophysical results, we constructed a computational model of encoding of motion stimuli in V1 and MT. The model contains two layers of units, one representing primary visual cortex (area V1), and the other representing area MT. There is a direct feedforward projection from the V1 layer to the MT layer. Each V1 unit has a profile of direction preference described by a circular Gaussian (a von Mises function, see Experimental procedures). The profile of synaptic inputs to each MT unit from a group of V1 units is also described by a circular Gaussian. The model contains an anisotropy in the numbers of V1 units representing different directions and in the widths of tuning of units representing different directions (Li et al., 2003). MT units inherit this anisotropy through the synaptic connections between V1 and MT. Specifically, the width of tuning is relatively large in MT units tuned to oblique directions and relatively small in MT units tuned to cardinal directions. The model also applies untuned divisive normalization to the output of the units in the V1 layer: the output of each unit is passed through a static nonlinearity and then normalized by the summed activity of all the V1 units before being passed as input to the MT layer.

The activity in the population of MT units is decoded using statistically optimal decoding based on a maximum likelihood algorithm (Jazayeri and Movshon, 2006). This scheme takes into account the activity of all the cells in the MT layer and selects the direction of motion most likely to be present in the stimulus, given the activity of all the MT units, their tuning widths, and their preferred directions (see Experimental procedures).

2.3.1. Modeling of motion direction discrimination

Direction discrimination relies on a comparison of the representation of motion direction in the two consecutive RDK presentations. Chance-level performance occurs when the difference between the two directions is below the resolution of the representation. In the task studied here, chance-level performance was 50%, as the task was a two alternative forced choice task (the subjects indicated whether the two stimuli had the same direction or different directions). In our computational model, we simulated motion direction discrimination by presenting the same stimulus twice. Because the spike rates produced by our model were

stochastically drawn from Poisson distributions, the direction deemed to be the most likely by the decoding mechanism was different in two subsequent presentations of the same exact stimulus. The difference between the directions estimated to be the most likely by the decoding mechanism, $\Delta\theta$, is a measure of the fidelity of the representation of motion direction. When this procedure was repeated multiple times, a distribution of the estimated $\Delta\theta$ values was obtained. In order to fulfill the 50% chance performance level requirement, any $\Delta\theta$ smaller than the median of this distribution was considered to be a trial for which the subject's response would be that there was no difference between the two directions of motion. Hence, we took the median of this distribution to be an estimate of the direction discrimination threshold of the model.

For stimuli with no directional variance, there was a reliable difference in the thresholds predicted by the model for stimuli with oblique and cardinal directions (Fig. 4A), replicating our psychophysical findings (Fig. 2). However, as the directional variance of the stimuli increased, this oblique effect diminished, until at a standard deviation of 22.5–45°, it disappeared. This pattern is strikingly similar to results reported by Dakin et al. (2005). In this study, human observers were presented with oblique and cardinal motion patterns containing varying amounts of added directional noise (variance in the direction of motion assigned to each element

in the pattern of moving stimuli). Consistent with our modeling results, Dakin et al. also observed an oblique effect in motion direction discrimination for low but not high levels of directional noise (Fig. 4B).

2.3.2. Modeling of the MAE oblique effect

We simulated the motion adaptation experiment in the model by defining the strength of the MAE as the relative likelihood of the two antagonistic directions in the probe stimulus (the direction of the adapting stimulus and the opposite direction), given the profile of activity in the units (Gold and Shadlen, 2001).

This readout of the strength of the MAE from the population activity of the MT units produced an excellent fit to our psychophysical results from the motion adaptation task (Fig. 3C). In particular, the model captured the substantial difference in adaptation strength between cardinal and oblique directions for intermediate adapter standard deviations and the minimal oblique effect for small and 90° standard deviations (Fig. 3D).

2.3.3. Comparison between different decoding mechanisms:

Decoding of the representations of stimuli in the model was performed by a maximum likelihood mechanism. There is psychophysical evidence that this is the mechanism underlying decoding of motion direction in humans (Jazayeri and Movshon, 2007). However, other mechanisms have also been suggested for decoding of motion direction in area MT, including vector averaging (Zohary et al., 1996; Nichols and Newsome, 2002) and winner-take-all (Nichols and Newsome, 2002; Salzman and Newsome, 1994; Zohary et al., 1996). We compared the abilities of these alternative decoding mechanisms to account for the psychophysical results. The encoding portion of the model was the same for all three decoding mechanisms, including the V1 directional anisotropy and the connectivity between V1 and MT. In vector averaging, each MT cell generates a vector pointing in the direction of that cell's preferred direction and proportional in length to that cell's firing rate. The direction of the average of these individual vectors is considered to be the direction coded by the population. The strength of the MAE was computed from the relative length of the component of the population vector for the direction opposite to the adapting direction.

The other decoding mechanism we considered is "winner-take-all". Here, the output of the model simply corresponds to the direction of motion preferred by the most active MT unit. The strength of the MAE was computed from the ratio between the activity in the unit which prefers the adapting direction and the activity in the unit which prefers the direction opposite to the adapting direction. Fig. 5 shows a comparison of the experimental results and the predictions of models based on the three decoding mechanisms. The maximum likelihood model best accounted for the width of tuning of motion adaptation as measured psychophysically (Fig. 5A). In contrast, the vector averaging model did not predict any difference in the MAE strength between oblique and cardinal directions except for an adapter standard deviation of 90° (Fig. 5B). Also, the winner-take-all model predicted MAE strength differences only for adapting stimuli with a standard deviation of 45° or greater (Fig. 5C). The failure

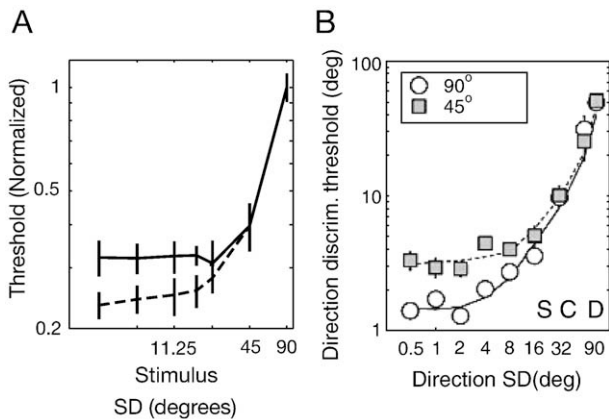


Fig. 4 – Model simulation of motion direction discrimination.

A: Motion direction discrimination thresholds predicted by our model (± 1 standard deviation for 10 repetitions of the simulation) for cardinal (dashed) and oblique (solid) directions. When the stimuli contained only a single direction of motion (zero directional variance), discrimination thresholds were lower for cardinal than for oblique directions. This matches the experimental results presented in Fig. 2. With increasing stimulus directional variance, the oblique effect diminished and eventually disappeared. B: The prediction of the model matches experimental results from a previous study (data from Dakin et al. 2005, copyright of ARVO, reproduced with permission). Thresholds in a motion discrimination task are presented for a single subject. Two directions, an oblique (grey squares) and a cardinal (white circles), are compared for different levels of direction standard deviation (SD, equivalent to the stimulus standard deviation in our model).




Fig. 5 – Comparing different decoding mechanisms. Model predictions of differences between cardinal and oblique adapting stimuli in the strength of the MAE (dashed lines) were compared to coherence threshold differences in the experimental results (solid line, same as the data presented in Fig. 3B). Three decoding mechanisms were compared: (A) the statistically optimal maximum likelihood model, (B) the vector averaging model, and (C) the winner-take-all model. The maximum likelihood model clearly provided the best fit of the experimental data.

of the vector averaging and winner-take-all models to account for the psychophysical results is probably due to the fact that both of these algorithms are necessarily invariant with regard to the tuning widths of the units in the model MT population. The population vector algorithm assigns equal weights to all the units in computing the population average, whereas the maximum likelihood algorithm weighs evidence from some units more than others, depending on their tuning width. The winner-take-all model also does not utilize the anisotropies in the encoding of oblique and cardinal directions that are present in the population of MT units.

3. Discussion

3.1. A novel directional anisotropy in motion perception

Perception of motion is not isotropic. Motion in some directions is perceived more accurately than motion in other

directions. We measured motion direction discrimination thresholds for eight directions and found lower discrimination thresholds for the cardinal directions (up/down, left/right) than for the oblique direction (off-cardinal diagonals). This result is a replication of previous findings (Ball and Sekuler, 1982; Dakin et al., 2005; Gros et al., 1998).

In addition, we have demonstrated a novel anisotropy in motion perception following motion adaptation. The adapting stimuli were RDKs containing dots moving in different directions with a distribution of directions centered at either a cardinal or an oblique direction. Directional variance of the adapter was manipulated by changing the variance of the distribution of directions of the individual dots. The strength of adaptation was measured by determining the amount of coherent motion required to null the resulting motion aftereffect.

For adapting stimuli with a small standard deviation of motion directions (0–17°), there was no difference between the magnitude of the MAE induced in cardinal and oblique directions. However, for intermediate standard deviations (22.5–45°), the MAE was significantly stronger for oblique than for cardinal directions. When the standard deviation of the adapting stimulus was very large (90°), minimal adaptation occurred for both cardinal and oblique directions.

Our results suggest that the oblique effect in the MAE and in motion direction discrimination may reflect common neural mechanisms. Specifically, there may be directional anisotropies in the encoding and decoding of stimuli in the lower levels of the visual system that produce an oblique effect for both motion discrimination and motion adaptation.

One account of the MAE posits that it stems from a temporary imbalance in the activity levels of populations of cells representing opposite directions (Barlow and Hill, 1963). Direction-selective cells in area MT are known to change their response characteristics following adaptation to a moving stimulus (Petersen et al., 1985). Among other changes, the response of these cells to moving stimuli was reduced following adaptation. Additional evidence, collected in the human brain using fMRI, also suggests that activity in area MT may be contributing to the MAE. Thus, presentation of an adapting stimulus caused direction-specific adaptation in human area MT+ and other visual cortical areas (Huk et al., 2001). However, there was no increase in the net activity measured in area MT+, suggesting that the MAE was induced not by a change in overall magnitude of activity in area MT+, but rather from differences in activity in different populations of direction-selective MT cells. Taken together, these results suggest that the sensation of motion relies on the distribution of activity within large populations of cells coding for direction, rather than an isolated change in the activity of a particular subset of direction-selective cells.

3.2. A model of encoding and decoding in V1 and MT

In order to explore possible mechanisms underlying our psychophysical results, we constructed a model of encoding of motion stimuli by populations of cells, based on the hierarchical organization of cortical areas V1 and MT. These areas contain neurons that are responsive to motion stimuli and selective for motion direction. Additionally, we

implemented a decoding scheme based on a statistically optimal maximum likelihood decoding algorithm. Our results cannot be fully explained by reference to only the encoding or decoding aspects of our model, suggesting that an explanation of complex perceptual phenomena, such as the directional anisotropy in motion perception, requires an understanding of the mechanisms underlying both encoding and decoding of stimulus information. A similar approach has been successful in accounting for anisotropies in texture perception (Cohen and Zaidi, 2007).

3.2.1. Encoding

Area V1 contains direction-selective cells, and there are direct excitatory monosynaptic projections from area V1 to cells in area MT. Therefore, previous models of encoding by cells in area MT often contained a V1 layer with feedforward projections to a second MT layer (Rust et al., 2006; Simoncelli and Heeger, 1998).

Another typical feature of these models is divisive normalization of the input to each cell in area MT by the summed V1 activity. Divisive normalization has been demonstrated physiologically in V1 (Carandini et al., 1997). Moreover, introducing divisive normalization in these models produces behaviors characteristic of MT. For example, the tuning of divisive normalization in the V1 to MT projection determines whether the MT cells integrate and average the pattern of motion of several different elements within their receptive field or whether they respond to each part of the pattern separately (Rust et al., 2006).

In order to account for the directional anisotropy we observed in our behavioral experiments, we introduced a directional anisotropy in the encoding process, based on anisotropies revealed in physiological experiments in cat primary visual cortex (Li et al., 2003). There were more V1 model units preferentially tuned to cardinal directions than units tuned to oblique directions, and the width of tuning of the units encoding cardinal directions was narrower than the width of the units encoding oblique directions. The model units representing MT cells then inherited the encoding anisotropy through the feedforward connections implemented in the model. Single-cell recordings from macaque area MT indicated no directional anisotropy in the population of recorded neurons (Churchland et al., 2003), a result that is inconsistent with our model predictions. However, stimuli moving in cardinal directions activate a larger cortical area within owl monkey MT than stimuli moving in oblique directions, as measured using intrinsic signal optical imaging (Xu et al., 2006). Optical imaging has produced inconsistent results regarding the possible existence of a directional anisotropy in area V1, possibly related to differences across studies in the portion of the visual field representation that was imaged (Xu et al., 2006, 2007).

In order to understand the origins of these anisotropies, Dakin et al. (2005) performed an analysis of the statistics of motion energy present in movies recorded in natural environments. This analysis revealed greater motion energy in cardinal than oblique directions during movement through natural environments. If the visual system is able to learn these statistical regularities, the anisotropy in motion perception may be a consequence of experience. Indeed, the oblique

effect can be partially abolished with training (Ball and Sekuler, 1982; Furmanski et al., 2004). However, comparisons between subjects from different ethnic groups, living in similar environments, indicated slight differences in the oblique effect in sensitivity to different orientations (Ross and Woodhouse, 1979; Timney and Muir, 1976). This suggests that there may be a genetic component of at least some types of oblique effect. Thus, anisotropies in visual inputs could possibly generate perceptual and neural anisotropies through natural selection as well as through experience-dependent development.

The encoding of the motion aftereffect was implemented in our model by applying activity-dependent adaptation to the MT cells. Consistent with the physiologically measured reduction in response amplitude of MT cells following exposure to motion in their preferred direction (Petersen et al., 1985), we assumed that MT cells integrate information about ongoing activity in the V1 cells through their synaptic connections. Specifically, the cells in our model that responded most vigorously to the adapting stimulus were the cells that responded the least to the post-adaptation probe stimulus.

3.2.2. Decoding

Our model implements an optimal decoding scheme based on the activity of the population of units representing MT neurons. This scheme determined the direction of motion that was most likely for each stimulus, given the population of MT neurons that responded to that stimulus. In our implementation of the model, the likelihood of a direction was weighted not only by the activity of the units representing that direction, but also by the relative reliability of the cells' responses, as reflected in the relative width of their tuning.

This kind of algorithm has been shown to be neurally plausible. A network of realistic neurons can implement maximum likelihood decoding (Deneve et al., 1999; Jazayeri and Movshon, 2006), and neurons in higher-level visual areas, reading out the information from area MT, modulate their activity pattern in a way which is consistent with this algorithm (Gold and Shadlen, 2001). Also, human observers have been found to act near-optimally when integrating information in their sensory environment (Battaglia et al., 2003; Ernst and Banks, 2002). These findings suggest that decoding of information that is represented at an intermediate level of processing, such as area MT, may proceed in a statistically optimal fashion.

We tested our model of decoding by comparing the optimal maximum likelihood strategy to other decoding algorithms. Specifically, we tested two decoding schemes which have been proposed for MT neurons: a vector averaging algorithm (Nichols and Newsome, 2002; Zohary et al., 1996) and a winner-take-all algorithm (Salzman and Newsome, 1994; Zohary et al., 1996; Nichols and Newsome, 2002). The best description of our data is clearly provided by the optimal maximum likelihood decoding algorithm. A recent study has shown that different decoding strategies may be used in solving different tasks, even if the strategies employ the same decoding algorithm. A maximum likelihood algorithm accounted for subjects' performance for both coarse and fine direction discriminations, but slight biases in the subjects'

performance revealed that information about coarse and fine discriminations was derived from different populations of neurons (Jazayeri and Movshon, 2007).

Furthermore, we cannot exclude the possibility that, under particular circumstances, decoding schemes other than the maximum likelihood mechanism may be utilized. For example, when integrating visual and auditory information in a target localization task, subjects probably use a hybrid decoding strategy. Specifically, they combine the maximum likelihood decoding algorithm with a tendency to rely on visual information rather than on auditory information, which is a form of the winner-take-all model (Battaglia et al., 2003).

In addition, implementations of the maximum likelihood algorithm in models of neural networks require multiple iterations to converge (Deneve et al., 1999). Thus, the implementation of this algorithm in the brain may be more time consuming than implementation of the vector averaging or winner-take-all decoding algorithms and may require more information about the tuning properties of the encoding cells than these alternative algorithms (Oram et al., 1998). Thus, the brain may employ different decoding schemes, depending on the task being performed and on the information available.

3.3. The effect of stimulus directional variance on the oblique effect in direction discrimination

In addition to accounting for our psychophysical results, our model also produces novel predictions. Specifically, as the directional variance of the stimulus increases, absolute thresholds in the motion direction discrimination task should also increase, but the directional anisotropy of these thresholds (the difference in thresholds for cardinal and oblique stimuli) should decrease. This prediction was not tested in our behavioral experiments, as subjects performed the direction discrimination task only in the zero directional variance condition.

However, this prediction has been tested in a previous study (Dakin et al., 2005). Human observers performed a motion direction discrimination task in both cardinal and oblique directions, and the directional variance in the stimulus was manipulated. Two results were obtained in Dakin et al.'s study that are pertinent to this discussion: 1) as the directional variance of the stimulus increased, the threshold of direction discrimination increased. That is, the task became more difficult. 2) As the directional variance of the stimulus increased, the oblique effect decreased. That is, the difference between direction discrimination thresholds in the oblique and the cardinal directions became smaller. At sufficiently large directional variances, the oblique effect was completely abolished. Both of these results are captured in the results of the simulations we conducted.

Despite the match between our model and the empirical results obtained by Dakin et al., there are differences in interpretation between our study and that of Dakin et al. They interpreted their results within the framework of an equivalent noise model, which assumes that direction discrimination thresholds reflect the sum of the noise that exists in the stimulus (the variance in the motion directions of the elements) and the internal noise (in the representation of the stimulus by the visual system). In contrast, our computa-

tional modeling results suggest that the relationship between variance of motion direction in the stimulus and direction anisotropy in motion direction discrimination thresholds can be accounted for by a combination of directional anisotropies in stimulus encoding and a maximum likelihood decoding strategy. Both the equivalent noise model of Dakin et al. and our computational model of V1 and MT provide an excellent description of the behavioral results (Fig. 4).

However, the two models make different predictions regarding the existence of physiological directional anisotropies in area MT. Our model posits that an oblique effect should be present in the tuning of MT neurons (inherited from the V1 neurons), while Dakin et al. conclude that no such oblique effect should exist in MT. Their interpretation relies on the assumption that when the standard deviation of the directions of motion in the stimulus is small, there is no need to integrate over many different elements, and the reliability of the representation will therefore be limited by the fidelity of the responses of cells in V1. When the standard deviation of the stimulus is larger, cells in MT must integrate the directions of motion of all the elements within their receptive fields. Thus, Dakin et al. reason that when integration over many elements is required in order to determine the direction of motion of the pattern, the threshold results from the activity of MT cells. However, more recent physiological recordings have shown that MT neurons only integrate elements within their receptive fields when these elements spatially overlap (Majaj et al., 2007).

3.4. Task dependence of the oblique effect for motion

Our model posits that the behavioral oblique effect is a consequence of an anisotropy in the primary representation of motion stimuli and of the decoding mechanism applied to this representation. If the oblique effect is indeed a consequence of the primary representation of the stimulus, it should also be present for other tasks involving motion perception. However, directional anisotropies have not been found for detection of coherent motion in a field of incoherently moving dots (Gros et al., 1998) or for speed discrimination (Westheimer, 2003). One interpretation of the lack of oblique effect in these tasks is that encoding of the motion stimulus by the visual system is isotropic, and the anisotropy only results from decoding in higher-level areas (Westheimer, 2003). We have not simulated either motion detection or speed discrimination tasks in the characterization of our model. However, the model is constructed such that the total firing rate of the population of MT cells is invariant to motion direction. This invariance is consistent with the lack of directional anisotropy in tasks requiring detection of coherent motion. Simulations of these tasks will be the topic of further studies of the model.

3.5. Summary and conclusions

We have described a novel oblique effect in motion perception in which the tuning width of adaptation is different for oblique and cardinal directions. In addition, we constructed a computational model of encoding and decoding of motion direction information in the visual system. Our model accounts for four

distinct psychophysical findings: 1) an oblique effect in the strength of motion adaptation (Figs. 3 and 5), 2) an oblique effect in the width of direction tuning of motion adaptation (Figs. 3 and 5), 3) an oblique effect in motion direction discrimination thresholds (Fig. 2 and the points in Fig. 4A corresponding to zero directional variance in the stimulus), and 4) the directional tuning of the oblique effect in motion direction discrimination, originally described by Dakin et al. (2005) (Fig. 4).

The model accounts for these findings only when a specific combination of encoding anisotropies and decoding mechanism is implemented. On the encoding side, more V1 units represent cardinal directions than oblique directions, and the units coding for cardinal directions are more narrowly tuned than those coding oblique directions. These directional anisotropies are inherited by MT units through the pattern of connectivity between V1 and MT. On the decoding side, a statistically optimal maximum likelihood decoding algorithm is used to read out the information from the population of MT units. These modeling results emphasize the significance of addressing both encoding and decoding of stimulus information when describing complex perceptual phenomena.

4. Experimental procedures

4.1. Subjects

Subjects were 16 young adults (6 female, mean age 24.1 ± 3.3 years) with normal or corrected-to-normal vision. All subjects were naïve to the purpose of the experiment and had no prior experience in performing psychophysical tasks. All subjects provided written informed consent, and the experimental protocols were approved by the Committee for the Protection of Human Subjects at the University of California, Berkeley.

4.2. Stimuli and experimental procedures

Stimuli were produced using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) for Matlab (Mathworks, Natick, MA) on Macintosh OS 10 (Apple, Cupertino, CA). The stimuli were presented on a Multisync FE992 CRT monitor (NEC, Tokyo, Japan) at a screen resolution of 600×800 and a refresh rate of 85 Hz. The edges of the screen were obscured with a circular cardboard aperture that eliminated any cues that could have been provided by the monitor frame. Similarly, a circular fixation point was used to eliminate the possibility of orientation cues. Subjects were seated comfortably and used a chin rest to insure consistent presentation of the stimuli.

4.2.1. Experiment 1 — motion direction discrimination

Stimuli were random dot kinetograms (RDK). The RDKs were presented within a circular annulus covering 1.0 – 3.1° radius from the fixation point. The RDKs always contained 100% coherent motion. However, dots moved to another position in the annulus after a lifetime of two monitor refresh frames in order to prevent the possibility of extracting the direction of the stimulus by tracking a single dot. Dots were approximately square and were 4.8 arcmin in size. The dot density

was approximately 2 dots/degree², and the dot velocity was 13° /s.

In each trial (see Fig. 1A for the trial structure), subjects viewed a 500 ms RDK stimulus followed by a 200 ms interstimulus interval and then another 500 ms RDK stimulus. The stimuli moved in the same direction in half of the trials and moved in different directions, separated by a small angle, α , in the other half. In the trials for which two different directions were shown, the first or the second stimulus (randomly selected for each trial) had a standard motion direction that was maintained throughout a testing block. For the remaining trials in which the stimuli had the same direction, the two stimuli could be the standard for that block, $\text{standard} + \alpha$, or $\text{standard} - \alpha$. Subjects were asked to respond whether the stimuli were moving in the same or in different directions within a 625 ms response period. They received auditory feedback following each trial.

A brief training session was administered prior to the first testing session to verify that the subjects understood the instructions and to acclimate the subjects to the task. Then, four testing sessions were administered. Each testing session was divided into 8 blocks of 50 trials each. In each block, the standard stimulus was kept constant and was one of the cardinal directions or one of the off-cardinal diagonals (oblique directions), randomly assigned to each block. The difference between the standard and comparison stimulus was adjusted in each trial according to the QUEST algorithm, and the threshold in each block (for $\sim 80\%$ correct performance) was estimated according to this algorithm (Watson and Pelli, 1983). Thresholds were defined as the average of the four estimates for each of the eight directions. Subjects were given the opportunity to rest between blocks.

4.2.2. Experiment 2 — MAE tuning width measurement

RDKs were presented in a circular region with a radius of 4° around the fixation point. Each block of 50 trials began with presentation of an adapting stimulus for 40 s (Fig. 1C). Then, at the beginning of each trial, the adapting stimulus was presented for an additional 4 s (top-up adaptation). Following a 50 ms interstimulus interval, a test stimulus was presented for 500 ms. Subjects were required to indicate whether the stimulus was moving in the same direction as the adapting stimulus or in the opposite direction (Blake and Hiris, 1993). Responses were collected during a 625 ms response interval. The motion coherence of the test stimulus was adjusted in each trial according to the responses in previous trials, based on a 2-up/1-down staircase (converging on $\sim 70\%$ correct performance). The threshold was then determined by fitting a cumulative Weibull distribution to all the trials in the assessment. The goodness of fit was determined for each psychometric function, and thresholds that did not conform to a Weibull function (Evans et al., 1989) were excluded from additional analysis. Specifically, we excluded all thresholds that did not have at least a 99% probability of coming from a Weibull distribution (Watson, 1979).

A brief training session was administered for each subject to verify that he or she understood the instructions and was acclimated to the task. Then, each subject participated in four testing sessions. In two of the testing sessions, the adapting direction was the oblique direction for which the subject's

threshold in Experiment 1 was highest. In the other two sessions, the adapting direction was the cardinal direction for which the subject's threshold in Experiment 1 was lowest. The order of administration of these different directions was counterbalanced across subjects. In each testing session, 8 blocks were administered. Blocks differed in the variance of directions of motion that were present in the adapting stimulus. Stimulus variance was manipulated by assigning a motion direction to each dot in the RDK from a Gaussian distribution (Fig. 1B). The distribution mean was the adapting stimulus direction, and the standard deviation of the Gaussian determined the directional variance for that block of trials. The standard deviations used were 0 (no variance, all dots moved in the same direction), 2.8125, 5.625, 11.25, 16.875, 22.5, 45, and 90°.

4.3. Computational model of motion processing in visual cortex

4.3.1. General model structure

The model consisted of one layer representing motion direction-selective V1 cells and one representing MT cells. V1 units projected in a feedforward manner to the MT units, whose firing rate was determined from the activity of their inputs from V1 cells and the strength of the synaptic connections between each V1 cell and each MT cell. Finally, the direction of the stimulus was decoded from the activity across the population of MT cells using a maximum likelihood procedure.

Following the physiological evidence from primary visual cortical neurons (Li et al., 2003), directional anisotropies were implemented in the V1 layer. Thus, there were more cells representing cardinal directions in the V1 layer, and the mean tuning width of these cells was narrower than the tuning width of cells representing oblique directions. The MT cells inherited these anisotropies through a homogenous set of connections between V1 and MT.

The different stimulus variance conditions were simulated by providing each V1 unit with an instance of one direction of motion for each iteration of the model. This simulated the RDK used in the experiments, under the assumption that each dot in the RDK excited one V1 unit. The directions of motion of the inputs to the V1 cells were drawn from a distribution of directions, and the variance of this distribution corresponded to the directional variance of the stimulus. Before being passed to the MT units via the synaptic connections between the layers, the output of every V1 unit was normalized by the sum of the activity of the entire population of V1 cells. Then, activity in each MT unit was computed, based on the activity in the connected V1 cells. Finally, the stimulus direction was decoded from the activity of the population of MT units using a maximum likelihood procedure. We adopted the convergence level used by Rust et al. (2006) of 12 V1 cells for each MT cell. Our complete model contained 32 MT units and 384 V1 units.

4.3.2. Generating model activity

In the first layer, representing V1 cells, the firing rate of each unit as a function of stimulus direction was described by a circular Gaussian distribution, also known as a von Mises

function (Patel and Read, 1996). This is a bell-shaped tuning curve of the form:

$$f(\theta) = a_{V1} e^{\frac{\cos(\theta-\theta_0)}{Z}} + b_{V1} \quad (1)$$

θ_0 is the unit's preferred direction, b_{V1} is the spontaneous rate of the unit (set to 10 Hz for all V1 units), a_{V1} is the maximal stimulus-evoked response to a stimulus moving in the preferred direction (set to 100 Hz for all V1 units), and $Z = 1/(\sigma \cdot 360 \cdot \text{Bessel}(1/\sigma))$. Z determines the width of tuning. $\text{Bessel}(x)$ is a Bessel function of the first kind of x and σ , is the tuning width of each V1 cell, which was set according to the cell's preferred direction:

$$\sigma(\theta_0) = \gamma(1 - \cos(4\theta_0)) + \epsilon \quad (2)$$

γ is a parameter that determines the ratio between the maximal tuning width (occurring in the cells tuned to oblique directions) and the minimal tuning width (occurring in cells tuned to cardinal directions). This minimal tuning width is represented by ϵ and was set to 45°. In addition, V1 cells were distributed unevenly along the different directions, according to the following equation:

$$\rho(\theta_0) = \frac{\beta(1 - \cos(4\theta_0)) + \delta}{360} \quad (3)$$

β is a parameter that determines the ratio between the densest representation (the difference in degrees between cells with preferred directions around the cardinal directions) and the sparsest representation (the difference between cells with preferred directions around oblique directions), and δ corresponds to the smallest difference between the preferred directions in the representation of cardinal directions, set to 1. In the results presented here, β and γ were set to 1.2 and 2, respectively. However, as verified by an extensive study of the parameter space, similar results were obtained over a range of values of β and γ .

In each trial, the firing rate of each cell was determined by randomly choosing θ from a distribution with the mean set to be either a cardinal or an oblique direction (in different runs of the model simulation) and with a standard deviation corresponding to the directional variance of the stimulus tested in that trial. This θ determined the mean response of the cell, based on its directional tuning curve. The activity of the cell in each trial was then determined according to a Poisson distribution:

$$P(r|\theta) = \frac{f(\theta)^r}{r!} e^{-f(\theta)} \quad (4)$$

In practice, r was determined for each V1 unit and each trial by drawing a number from a Poisson random number generator with mean equal to the firing rate of the cell in response to that trial's stimulus, or $f(\theta)$.

Before passing the V1 outputs to the MT cells, a static non-linearity (a squaring) was applied to the output of the V1 cells, and divisive normalization was applied to this squared output:

$$\hat{r}_i(\theta) = \frac{r_i(\theta)^2}{\sum_{j \in V1} r_j(\theta) + \zeta} \quad (5)$$

where ζ is a parameter which controls the relative contribution of the other V1 units to reducing activity of a given V1 unit. An exploration of different values of ζ verified that the results are qualitatively the same as long as the output of this stage was between the noise level (b_{V1} in Eq. 1) and the gain of the firing rate (a_{V1} in Eq. 1).

The connectivity for each pair of V1 and MT cells was defined according to a von Mises function:

$$w_{ij} = a_{MT} e^{\frac{\cos(\theta_i - \theta_j)}{Z}} + b_{MT} \quad (6)$$

i is an index of the MT cell and j is an index of the V1 cell, $a_{MT} = a_{V1} = 100$ Hz and $b_{MT} = b_{V1} = 10$ Hz. As in Eq. 1, Z determines the direction tuning width of the connectivity between V1 and MT. Z was set such that the direction tuning width of the connectivity was always equal to 45° , independent of the preferred direction of the units. The inputs to each cell were set such that the sum of the synaptic weights to the population of MT cells was the same for all possible stimulus directions.

The firing rate of each MT cell, $f_i(\theta)$, was then determined by summing over all of its inputs from V1:

$$f_i(\theta) = \sum_{j \in V1} w_{ij}(\theta) r_j(\theta) \quad (7)$$

The activity for each MT cell and each trial was then determined by drawing a number from a Poisson number generator, as in Eq. 4. The resulting profile of activity in MT is the population code on which decoding then proceeds.

4.3.3. Decoding

Decoding of the direction of motion from the activity of the population of MT cells was done according to a statistically optimal scheme. Under this scheme, we are interested in finding the motion direction θ which is maximally likely, given a certain distribution of activity in the population of MT units, r_{MT} . That is:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \{P(\theta|r_{MT})\} \quad (8)$$

Here, $\hat{\theta}$ is the θ that maximizes the likelihood of θ given a particular r_{MT} . However, the functional form of this likelihood is unknown. Bayes's theorem states that the likelihood of θ given r_{MT} and the inverse likelihood, of r_{MT} given θ , or $P(r_{MT}|\theta)$, are closely related:

$$P(r_{MT}|\theta) = \frac{P(\theta|r_{MT})P(r_{MT})}{P(\theta)} \quad (9)$$

Therefore, assuming that the prior probability distributions for both θ and r_{MT} are flat (the probability of activity is the same for all units and no direction is more likely to appear than any other direction):

$$\underset{\theta}{\operatorname{argmax}} \{P(\theta|r_{MT})\} = \underset{\theta}{\operatorname{argmax}} \{P(r_{MT}|\theta)\} \quad (10)$$

As described above (Section 4.3.2 Generating model activity), the actual likelihood functions of activity in the MT cells, given θ , are independent Poisson processes.

Therefore, the likelihood of the population activity of all the MT cells, given θ , a sum of these probabilities, is also a Poisson distribution (Pitman, 1993) which resembles the Poisson distribution that characterizes the firing rate of individual units (Eq. 4):

$$P(r_{MT}|\theta) = \frac{f_{MT}(\theta)^{r_{MT}}}{r_{MT}!} e^{-f_{MT}(\theta)} \quad (11)$$

In practice, we approximate and maximize the following log likelihood function (Jazayeri and Movshon, 2006; Seung and Sompolinsky, 1993), which is derived by taking the log of Eq. 11:

$$\begin{aligned} \operatorname{LogL}(\theta) &= \sum_{i \in MT} \log P(r_i|\theta) \\ &= \sum_{i \in MT} r_i \log f_i(\theta) - \sum_{i \in MT} f_i(\theta) - \sum_{i \in MT} \log(r_i!) \end{aligned} \quad (12)$$

The last term can be dropped, as it does not depend on θ . The second term can also be dropped, as the model was constructed so that the total firing rate does not depend on the direction of the presented stimulus. Specifically, the sum of the inputs to the population of MT cells was set so that it was independent of stimulus direction. Therefore, Eq. 12 reduces to:

$$\operatorname{LogL}(\theta) = \sum_{i \in MT} r_i \log f_i(\theta) \quad (13)$$

The tuning curves of cells in the MT layer, $f_i(\theta)$, result from both the tuning curves of the V1 cells as well as the connectivity between the V1 and MT layers. Therefore, we do not have an analytical form of these tuning curves. However, there is an empirical form of the tuning curve for each MT cell which can be derived by summing the V1 tuning curves and assigning a weight to each V1 tuning curve corresponding to the strength of the V1/MT synapse:

$$f_i^{emp}(\theta) = \sum_{j \in V1} w_{ij} f_j(\theta). \quad (14)$$

These empirically derived tuning curves resemble the form of the circular Gaussian V1 tuning curves described by the von Mises function (Eq. 1). Therefore, we also approximated the tuning curve for each MT cell by a von Mises function:

$$f_i(\theta) \approx a_i e^{\frac{\cos(\theta - \theta_0)}{Z}} + b_i \quad (15)$$

The parameter a_i , describing the tuning width of the unit, was estimated from the width of the empirically derived tuning curve (Eq. 14) at half maximum. Following Jazayeri and Movshon (2006), the approximate log likelihood function is then the log of each unit's tuning curve, weighted by the inverse of its relative tuning width and by the activity of the unit in a given trial. This quantity was summed over the population of units in MT:

$$\operatorname{LogL}(\theta) \approx \sum_{i \in MT} r_i \frac{\kappa_i}{\sum_{j \in MT} \kappa_j} \cos(\theta - \theta_i) \quad (16)$$

where r_i is the activity of the cells for a given trial and κ_i is the inverse of the tuning widths as estimated from the empirical tuning curves (a_i in Eq. 15). In this sum, the cells

with smaller tuning widths are weighted more heavily than the cells with larger tuning widths. For each iteration of the model, we used an unconstrained nonlinear optimization algorithm (implemented as the Matlab function `fminsearch`) to find a value of θ that maximizes this log likelihood function, given the population response in MT.

4.3.4. Estimating the MAE strength

In order to simulate the MAE, we presented the model with a stimulus of 0% coherence following an adapting stimulus. The response amplitude of most MT cells decreases following adaptation to motion stimuli (Petersen et al., 1985). This reduction in amplitude may actually be greater when the adapting direction is slightly different from the preferred direction (Kohn and Movshon, 2004). However, for simplicity, we modeled adaptation in MT such that the adaptation in each unit was proportional to the response to the adapting stimulus. Thus, the strongest adaptation occurred in cells tuned to the adapting stimulus. The activity in response to the probe stimulus (with 0% coherence) was then calculated for the V1 cells and propagated to the cells in the MT layer (for simplicity, we assumed that V1 cells do not adapt). The firing rate in each MT cell in response to the probe stimulus was calculated based on a combination of its V1 inputs and the adaptation state of the MT cell. Thus, each cell's response to the adapting stimulus was multiplied by a factor that determines the strength of adaptation. The value of this factor can vary rather substantially without significantly affecting the results, as long as the firing rates do not become negative. For each cell, the firing rate during adaptation was subtracted from the firing rate that would have been obtained in the post-adaptation probe stimulation, had there been no adaptation. The actual firing is then derived from the Poisson distribution, as described above (Eqs. 4 and 11). The MAE strength was estimated from the ratio of the likelihood of the adapting stimulus and the stimulus moving in the opposite direction. This ratio can be calculated from the difference between the log likelihood functions of the two directions (Jazayeri and Movshon, 2006). We compared the likelihood ratios to the value of the likelihood ratio when the adapting stimulus had a standard deviation of 0° .

4.3.5. Estimating the direction discrimination threshold

In each presentation of a stimulus to the model, the θ that maximized the log likelihood function (Eq. 13) was considered to be the direction perceived by the observer for that stimulus presentation. However, due to variability of the neuronal responses across presentations of the same stimulus, this maximum likelihood direction varied between presentations of the same stimulus. The level of variability limited the fidelity of the representation of motion direction. This limit corresponds to the threshold obtained from the two alternative forced choice procedure employed in our psychophysical experiments (Fig. 1A).

We quantified this variability by estimating a distribution of the differences between the perceived stimuli in two consecutive presentations of the same stimulus. The model was iterated 100 times. For each iteration, the same stimulus was presented twice in succession, and the observed stimulus was decoded for each presentation (see Section 4.3.3 Decoding). The difference between the two decoded stimuli was denoted $\Delta\theta$. Over the 100 iterations of

the model, we constructed a distribution of $\Delta\theta$. We defined the median of the distribution of $\Delta\theta$ for a given condition (direction and standard deviation) to be the threshold for that condition (Han et al., 2007; Kim and Bao, 2008), in order to satisfy the 50% guess rate in the two alternative forced choice paradigm used in the psychophysical measurements of direction discrimination thresholds.

4.3.6. Alternative decoding mechanisms

Two alternative decoding mechanisms were tested. Vector averaging involves computation of a weighted average over all of the cells in area MT. Each cell contributed a vector pointing in the direction best represented by its inputs and proportional in size to its firing rate. Then, the vectors were summed. The direction of this summed vector was considered to be the predicted stimulus direction. This mechanism has been proposed for other neural populations (Georgopoulos et al., 1986) and for MT neurons, under some conditions (Nichols and Newsome, 2002; Zohary et al., 1996). The strength of the MAE was estimated in a manner similar to the one described above (Section 4.3.4 Estimating the MAE strength). A stimulus with 0% coherence was presented to a population of cells, following adaptation. Then, the MAE strength was considered to be the relative length of the projection of the population vector in the direction opposite to the adapting direction.

The other alternative decoding mechanism we considered was a winner-take-all mechanism. Here, the decoding of MT activity occurs by identifying the cell with the most activity and assigning the predicted direction to the optimal direction for that cell, computed from the cell's inputs. This has also been suggested to be a decoding mechanism of activity in MT under certain circumstances (Salzman and Newsome, 1994; Zohary et al., 1996; Nichols and Newsome, 2002). Here, the MAE strength corresponded to the ratio of the activity in the cell preferring the adapting direction and the activity in the cell preferring the direction opposite to the adapting direction.

Acknowledgments

We acknowledge Shradha Sanghvi's assistance in data collection and analysis. Her participation in this project was funded by an NEI T35 grant through the School of Optometry, University of California, Berkeley.

We thank Dennis Levi as well as Shaowen Bao and members of his lab for fruitful discussions, Baowang Li and Ralph Freeman for sharing their knowledge about direction-selective cells, and Ayelet Landau, Dennis Levi and Thomas Lauritzen for comments on a previous version of the manuscript. Finally, we thank Steven C. Dakin, Isabelle Mareschal, and Peter J. Bex for permission to reproduce a figure from their paper (Dakin et al., 2005).

REFERENCES

- Albright, T.D., 1984. Direction and orientation selectivity of neurons in visual area MT of the macaque. *J. Neurophysiol.* 52, 1106–1130.

- Anstis, S., Verstraten, F.A.J., Mather, G., 1998. The motion aftereffect. *Trends Cogn. Sci.* 2, 111–117.
- Appelle, S., 1972. Perception and discrimination as a function of stimulus orientation: the “oblique effect” in man and animals. *Psychol. Bull.* 78, 266–278.
- Ball, K., Sekuler, R., 1982. A specific and enduring improvement in visual motion discrimination. *Science* 218, 697–698.
- Barlow, H.B., Hill, R.M., 1963. Evidence for a physiological explanation of the waterfall phenomenon and figural after-effects. *Nature* 200, 1345–1347.
- Battaglia, P.W., Jacobs, R.A., Aslin, R.N., 2003. Bayesian integration of visual and auditory signals for spatial localization. *J. Opt. Soc. Am. A* 20, 1391–1397.
- Blake, R., Hiris, E., 1993. Another means for measuring the motion aftereffect. *Vis. Res.* 33, 1589–1592.
- Brainard, D.H., 1997. The Psychophysics Toolbox. *Spat. Vis.* 10, 433–436.
- Carandini, M., Heeger, D.J., Movshon, J.A., 1997. Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* 17, 8621–8644.
- Churchland, A.K., Gardner, J.L., Chou, I., Priebe, N.J., Lisberger, S.G., 2003. Directional anisotropies reveal a functional segregation of visual motion processing for perception and action. *Neuron* 37, 1001–1011.
- Cohen, E.H., Zaidi, Q., 2007. Fundamental failures of shape constancy resulting from cortical anisotropy. *J. Neurosci.* 27, 12540–12545.
- Dakin, S.C., Mareschal, I., Bex, P.J., 2005. An oblique effect for local motion: psychophysics and natural movie statistics. *J. Vis.* 5, 878–887.
- De Valois, R.L., Yund, E.W., Hepler, N., 1982. The orientation and direction selectivity of cells in macaque visual cortex. *Vis. Res.* 22, 531–544.
- De Valois, R.L., Cottaris, N.P., Mahon, L.E., Elfar, S.D., Wilson, J.A., 2000. Spatial and temporal receptive fields of geniculate and cortical cells and directional selectivity. *Vis. Res.* 40, 3685–3702.
- Deneve, S., Latham, P.E., Pouget, A., 1999. Reading population codes: a neural implementation of ideal observers. *Nat. Neurosci.* 2, 740–745.
- Ernst, M.O., Banks, M.S., 2002. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 429–433.
- Evans, J.W., Johnson, R.A., Green, D.W., 1989. Two- and three-parameter Weibull goodness-of-fit tests. *USDA Forest Products Laboratory Research Paper, FPL-RP-493*, pp. 2–3.
- Furmanski, C.S., Engel, S.A., 2000. An oblique effect in human primary visual cortex. *Nat. Neurosci.* 3, 535–536.
- Furmanski, C.S., Schluppeck, D., Engel, S.A., 2004. Learning strengthens the response of primary visual cortex to simple patterns. *Curr. Biol.* 14, 573–578.
- Georgopoulos, A.P., Schwartz, A.B., Kettner, R.E., 1986. Neuronal population coding of movement direction. *Science* 233, 1416–1419.
- Gold, J.I., Shadlen, M.N., 2001. Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5, 10–16.
- Gros, B.L., Blake, R., Hiris, E., 1998. Anisotropies in visual motion perception: a fresh look. *J. Opt. Soc. Am. A* 15, 2003–2011.
- Han, Y.K., Köver, H., Insanally, M.N., Semerdjian, J.H., Bao, S., 2007. Early experience impairs perceptual discrimination. *Nat. Neurosci.* 10, 1191–1197.
- Hiris, E., Blake, R., 1992. Another perspective on the visual motion aftereffect. *Proc Natl. Acad. Sci. U. S. A.* 89, 9025–9028.
- Horn, B.K.P., 1986. *Robot Vision*. MIT Press, Cambridge, MA.
- Hubel, D.H., Wiesel, T.N., 1959. Receptive fields of single neurones in the cat’s striate cortex. *J. Physiol.* 148, 574–591.
- Huk, A.C., Ress, D., Heeger, D.J., 2001. Neuronal basis of the motion aftereffect reconsidered. *Neuron* 32, 161–172.
- Jazayeri, M., Movshon, J.A., 2006. Optimal representation of sensory information by neural populations. *Nat. Neurosci.* 9, 690–696.
- Jazayeri, M., Movshon, J.A., 2007. Integration of sensory evidence in motion discrimination. *J. Vis.* 7, 1–7.
- Kim, H., Bao, S., 2008. Distributed representation of perceptual categories in the auditory cortex. *J. Comput. Neurosci.* 24, 277–290.
- Kohn, A., Movshon, J.A., 2004. Adaptation changes the direction tuning of macaque MT neurons. *Nat. Neurosci.* 7, 764–772.
- Li, B., Peterson, M.R., Freeman, R.D., 2003. Oblique effect: a neural basis in the visual cortex. *J. Neurophysiol.* 90, 204–217.
- Majaj, N.J., Carandini, M., Movshon, J.A., 2007. Motion integration by neurons in macaque MT is local, not global. *J. Neurosci.* 27, 366–370.
- Nichols, M.J., Newsome, W.T., 2002. Middle temporal visual area microstimulation influences veridical judgments of motion direction. *J. Neurosci.* 22, 9530–9540.
- Oram, M.W., Földiák, P., Perrett, D.I., Sengpiel, F., 1998. The ‘Ideal Homunculus’: decoding neural population signals. *Trends Neurosci.* 21, 259–265.
- Pack, C.C., Livingstone, M.S., Duffy, K.R., Born, R.T., 2003. End-stopping and the aperture problem: two-dimensional motion signals in macaque V1. *Neuron* 39, 671–680.
- Patel, J.K., Read, C.B., 1996. *Handbook of the Normal Distribution*. CRC Press, New York.
- Pelli, D.G., 1997. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* 10, 437–442.
- Petersen, S.E., Baker, J.F., Allman, J.M., 1985. Direction-specific adaptation in area MT of the owl monkey. *Brain Res.* 346, 146–150.
- Peterson, M.R., Li, B., Freeman, R.D., 2004. The derivation of direction selectivity in the striate cortex. *J. Neurosci.* 24, 3583–3591.
- Pitman, J., 1993. *Probability*. Springer, New York.
- Pouget, A., Dayan, P., Zemel, R., 2000. Information processing with population codes. *Nat. Rev., Neurosci.* 1, 125–132.
- Ross, H.E., Woodhouse, J.M., 1979. Genetic and environmental factors in orientation anisotropy: a field study in the British Isles. *Perception* 8, 507–521.
- Rust, N.C., Mante, V., Simoncelli, E.P., Movshon, J.A., 2006. How MT cells analyze the motion of visual patterns. *Nat. Neurosci.* 9, 1421–1431.
- Salzman, C.D., Newsome, W.T., 1994. Neural mechanisms for forming a perceptual decision. *Science* 264, 231–237.
- Seung, H.S., Sompolinsky, H., 1993. Simple models for reading neuronal population codes. *Proc. Natl. Acad. Sci. U. S. A.* 90, 10749–10753.
- Simoncelli, E.P., Heeger, D.J., 1998. A model of neuronal responses in visual area MT. *Vis. Res.* 38, 743–761.
- Timney, B.N., Muir, D.W., 1976. Orientation anisotropy: incidence and magnitude in Caucasian and Chinese subjects. *Science* 193, 699–701.
- Watson, A.B., 1979. Probability summation over time. *Vis. Res.* 19, 515–522.
- Watson, A.B., Pelli, D.G., 1983. QUEST: a Bayesian adaptive psychometric method. *Percept. Psychophys.* 33, 113–120.
- Westheimer, G., 2003. Meridional anisotropy in visual processing: implications for the neural site of the oblique effect. *Vis. Res.* 43, 2281–2289.

- Xu, X., Collins, C.E., Khaytin, I., Kaas, J.H., Casagrande, V.A., 2006. Unequal representation of cardinal vs. oblique orientations in the middle temporal visual area. *Proc. Natl. Acad. Sci. U. S. A.* 103, 17490–17495.
- Xu, X., Anderson, T.J., Casagrande, V.A., 2007. How do functional maps in primary visual cortex vary with eccentricity? *J. Comp. Neurol.* 501, 741–755.
- Zohary, E., Case, M.O., Braddick, O.J., 1996. Integration across directions in dynamic random dot displays: vector summation or winner take all? *Vis. Res.* 36, 2321–2331.