Vision Research 48 (2008) 1374-1382

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Centre-surround effects on perceived orientation in complex images

Erin Goddard^{a,b,*}, Colin W.G. Clifford^a, Samuel G. Solomon^b

^a School of Psychology, The University of Sydney, Sydney, Australia

^b Disciplines of Physiology, Anatomy and Histology, School of Medical Sciences and Bosch Institute, The University of Sydney, Sydney, Australia

A R T I C L E I N F O

ABSTRACT

Article history: Received 21 October 2007 Received in revised form 22 February 2008

Keywords: Tilt illusion Contextual modulation Natural scenes Visual cortex Spatial vision Using the simultaneous tilt illusion [Gibson, J., & Radner, M. (1937). Adaptation, after-effect and contrast in the perception of tilted lines. *Journal of Experimental Psychology*, *12*, 453–467], we investigate the perception of orientation in natural images and textures with similar statistical properties. We show that the illusion increases if observers judge the average orientation of textures rather than sinusoidal gratings. Furthermore, the illusion can be induced by surrounding textures with a broad range of orientations, even those without a clearly perceivable orientation. A robust illusion is induced by natural images, and is increased by randomising the phase spectra of those images. We present a simple model of orientation processing that can accommodate most of our observations.

© 2008 Elsevier Ltd. All rights reserved.

VISION

1. Introduction

Our current understanding of the visual system is based to a large extent on the measurement of physiological and perceptual responses to reduced, simplified stimuli, such as sinusoidal gratings. In a recent review of the current understanding of V1, Olshausen and Field (2005) suggested that for the highly non-linear visual system, sinusoidal gratings have no particular significance, and that the body of experimental work on V1 is biased by the number of studies which use sinusoidal gratings and other reduced stimuli to infer the response properties of V1 cells. While the processing of bars, spots and gratings by V1 cells may be relatively well described, we cannot necessarily use these neural responses to predict responses to more complex visual stimuli, including natural scenes (Burr, Morrone, & Maffei, 1981; David, Vinje, & Gallant, 2004).

There is psychophysical evidence to suggest that the visual system is tuned to scene statistics that are characteristic of natural scenes. One such statistical commonality among natural scenes is their amplitude spectra, which typically vary with spatial frequency (f) such that

amplitude
$$(f) \propto f^{-\alpha}$$
 (1)

where, across natural images, α tends to fall within a range of about 0.7–1.7, with an average value of approximately 1. For example, average α values of 0.94 (van der Schaaf & van Hateren, 1996),

* Corresponding author. E-mail address: ering@psych.usyd.edu.au (E. Goddard). 1.16 (Tadmor & Tolhurst, 1994) and 0.91 (Ruderman & Bialek, 1994) have been reported.

Stimuli with this spatial frequency structure are more easily discriminated on the basis of spatial frequency content (Tadmor & Tolhurst, 1994). Adaptation to a series of natural images selectively reduces sensitivity to low and medium spatial frequencies (Webster & Miyahara, 1997), and perceived contrast is maximally suppressed by a surrounding stimulus when the inducing stimulus has an α value of 1, regardless of the α value of the central test patch (McDonald & Tadmor, 2006). Thus not only are simple, reduced stimuli unlike those encountered in the natural world, they are also unlikely to optimally engage mechanisms of the visual system which are under scrutiny.

These findings highlight the need for models of the visual system that can account for physiological and perceptual responses to natural stimuli. To relate the processing of simple visual stimuli to that of complex stimuli, we should test whether models that can account for responses to simpler stimuli can predict perceptual and physiological responses to them. Here we ask whether an underconstrained model of a visual illusion, the simultaneous tilt illusion, which is well documented with bars and gratings, can also be used to account for illusions induced by more complex textures.

The simultaneous tilt illusion refers to the impact of nearby lines or gratings on the perceived orientation of contours. The impact of the nearby stimulus depends on the orientation difference between it and the test: small differences repulse the perceived orientation of the test away from that of the surrounding stimulus, while larger differences can attract (Wenderoth & Johnstone, 1988; Westheimer, 1990). The neural basis for this illusion has been modelled as contextual modulation of those cells whose respon-



^{0042-6989/\$ -} see front matter \odot 2008 Elsevier Ltd. All rights reserved. doi:10.1016/j.visres.2008.02.023

siveness contributes to the perceived orientation of the test, where the presence of the surrounding stimulus either reduces their responsiveness, shifts their orientation preference, or broadens their bandwidth. In each case this contextual modulation is dependent upon the orientation of the surrounding lines or gratings, and each manipulation is capable of shifting the population response and predicting the tilt illusion with simple stimuli (Clifford, Wenderoth, & Spehar, 2000; Coltheart, 1971; Gilbert & Wiesel, 1990; Jin, Dragoi, Sur, & Seung, 2005). Here we ask whether a simple model of orientation processing, similar to these, can account for the tilt illusion with textures, which share some properties that are typical of natural images.

We produced textures whose two-dimensional Fourier spectra were defined by a 1/f distribution of spatial frequencies at each orientation, similar to that of natural images, and a broad range of orientations. We measured the tilt illusion using these broadband textures as the inducing and test stimuli, and asked whether our model, which predicts the illusion for gratings, could account for our results. We generated predictions of this model for both gratings and textures. We evaluated the extent to which the model of the illusion with gratings can predict the illusion with textures, by tracking the performance on textures of the parameter sets that provide the best 10% of model fits for gratings. The model, which considers the amplitude spectra of these textures, predicts most of our observations. We then measured the tilt illusion induced when using segments of natural images as surrounding stimuli. These image segments, unlike broadband textures, contain clear contours: the tilt illusion induced by them was mostly, but not entirely, predicted by their amplitude spectra.

2. Methods

2.1. Visual stimuli

Images were generated and displayed with Matlab (v7.0, MathWorks, Natick, MA) using Open GL and routines from PsychToolbox (Brainard, 1997; Pelli, 1997) on a G5 Macintosh computer driving a nVidia GeForce 6600LE graphics card. The images were displayed on a gamma-corrected ViewSonic G810-6 cathode ray tube monitor, refreshed at a rate of 75 Hz and viewed from a distance of 0.57 m. To remove visual cues to vertical the subject viewed the screen in a darkened room, and matt black cardboard with a circular window covered all but a circular central portion (radius 10.75 deg) of the screen. The perceived tilt of a central circular surface (outer radius 1.5 deg) was measured in the presence of an abutting annular surface (outer radius 7.25 deg). The remainder of the screen was held at the mean luminance of ~50 cd/m².

In the first experiment the centre and annular surfaces were sinusoidal gratings (spatial frequency 3 cycles/deg). The spatial phase of both the central and annular gratings was randomly selected on each stimulus presentation. In the subsequent experiments with broadband textures, central and annular surfaces were the real component of the inverse Fourier transform of a square matrix of complex numbers; the amplitude of these complex numbers defined the amplitude spectrum of the surface; the angle of these numbers in the complex plane defining the phase spectrum. Each point in the phase spectrum was drawn from a uniform random distribution ranging from 0 to 2π , and was regenerated for every stimulus presentation. The two-dimensional amplitude spectrum was specified as the product of a 1/*f* distribution over spatial frequency for each orientation, and a Gaussian distribution over orientation; frequency being the polar distance, and orientation the polar angle relative to the central element of the matrix.

The power at each orientation was defined by a wrapped normal distribution with a standard deviation ranging from 3.125 to 50 deg and a peak at vertical. At large standard deviations, these distributions are non-zero at the orientation orthogonal to the peak, so they must be wrapped; this is not possible analytically, and was approximated as described by Dakin, Mareschal, and Bex (2005). We also used annular surfaces with 'notched' orientation structure, which were specified by a flat distribution, from which a wrapped normal distribution was subtracted. The amplitude spectrum was then restructured (using *fftshift2* in Matlab) before combination with the phase spectrum. We used OpenGL commands to appropriately rotate the central and annular surfaces on the video card.

Natural image segments used in the final experiment were selected from 100,000 randomly chosen 256×256 pixel segments taken from the first 1000 images in van Hateren's database (van Hateren & van der Schaaf, 1998). Each image segment was decomposed using Fourier analysis into its amplitude and phase spectra. We compared the amplitude spectrum of each image segment with each of the amplitude spectra used to generate the broadband textures. For each pair of amplitude spectra, we assessed their similarity by calculating the correlation

coefficient between the pixel values in the two spectra. Four of the sixteen image segments with the highest correlation coefficients were selected for each orientation bandwidth, such that each image segment was of a separate scene (*r* values ranged from 0.82 to 0.94). These images can be downloaded from van Hateren's database (van Hateren & van der Schaaf, 1998) and located using the image numbers and segment locations in Appendix 1. Annular surfaces always had the amplitude spectrum of the original natural image segment, and either the original phase spectrum (natural image condition), or a randomly generated phase spectrum (phase randomised condition). The central (test) surface was a broadband texture with a wrapped normal distribution of orientations of 12.5 deg standard deviation, as described above.

Since we randomly generated the phase matrix of each broadband texture, there was some variation in the distribution of pixel intensities of each broadband texture, even between textures with the same amplitude spectrum. To remove this variation, the surfaces were normalised such that the root mean square (RMS) contrast of the pixel intensities was 0.25. RMS contrast has been shown to correlate well with the detectability of visual stimuli, and is used here to approximately normalise the surfaces for perceived contrast (Bex & Makous, 2002).

2.2. Subjects and procedure

A total of 11 subjects (18–39 years old, 3 male) participated, including two of the authors (EG, CC). At least 4 subjects participated in each experiment, and each experiment included at least 3 subjects who were naïve to the purpose of the investigation. All subjects had normal or corrected to normal vision. In accordance with the guidelines of the Human Research Ethics Committee of The University of Sydney, human subjects gave informed written consent before participating. All subjects were free to withdraw from the study at any time.

On each trial, a central fixation spot, displayed for 0.3 s, preceded the stimulus, which was also presented for 0.3 s. Subjects were required, by responding with a key-press, to report the average orientation of the central patch as tilted clockwise or counterclockwise of vertical. In each experiment, the orientation of the surround denotes how far the texture in the surrounding annulus was rotated from vertical. Each session contained equal numbers of trials for annuli that were tilted clockwise or counterclockwise of vertical, to avoid a build up of adaptation to one average orientation. The orientation at which subjects were equally likely to report the central test patch as tilted clockwise or counterclockwise of vertical (the point of subjective vertical) was estimated in sessions of 120 trials. Each session included four randomly interleaved Bayesian adaptive staircases (Kontsevich & Tyler, 1999), two with clockwise- and two with counterclockwise-annuli, each consisting of 30 trials. The staircases provided two estimates of the tilt illusion in each session (each estimate is half the difference between perceived vertical obtained from one staircase with clockwise- and another with counterclockwise-annuli). For each observer, each data point is the average of at least two sessions (ie. at least eight staircases).

2.3. Model

Perceived average orientation is computed from a population of orientation selective detectors whose receptive fields lie over the central surface, and whose tuning is sensitive to the pattern in the annular surface. We call the mechanism that provides this sensitivity to the annular surface the 'surround'. The detectors' response to the central stimulus (C) is divided by its surrounds response to the annular at stimulus (S), so the net response (R) is

$$R = \frac{C}{1 + k \cdot S} \tag{2}$$

where *k* is the overall strength of the surround. *C* is the correlation of the detectors' tuning curve and the orientation distribution in the central stimulus (θ_c); *S* is the correlation of the tuning curve of the surround and the distribution of orientations in the annular surface (θ_s). The tuning curve of each detector, $vm(\theta_{max}, \alpha)$, is a von Mises distribution (Mardia & Jupp, 2000) with a peak at the preferred orientation of that detector (θ_{max}), and a full width at half height of α . The model does not allow for variation in bandwidth across detectors (Clifford, Wenderoth, & Spehar, 2000), so

$$C_{(\theta_{\max},\alpha,\theta_c)} = vm_{(\theta_{\max},\alpha)} * \theta_c \tag{3}$$

where * denotes correlation. The tuning curve of the surround is the difference of two von Mises distributions with different bandwidths (β and γ), which generally produces a surround whose sensitivity across orientations looks like a 'Mexican hat'. The tuning curve of the surround is symmetric around the preferred orientation of the detector. Each detector is therefore maximally suppressed by surrounds of its preferred orientation; other orientations suppress the response to a lesser extent, and can instead facilitate the response.

$$S_{(\theta_{\max},\beta,\gamma,q,\theta_s)} = (vm_{(\theta_{\max},\beta)} - q \cdot vm_{(\theta_{\max},\gamma)}) * \theta_s$$
(4)

The model is analogous to the surrounds defined in physiological investigations of V1 of cat and primate, but we do not mean it to be seen as a direct implementation, and there are some potential discrepancies. In V1 the orientation of the surround that is most effective at suppressing responses can depend on the orientation of the stimulus used to evoke responses from the classical receptive field, at least for complex cells; this may not be the case for simple cells (Cavanaugh, Bair, & Movshon,

2002; Muller, Metha, Krauskopf, & Lennie, 1999). Facilitation (or disinhibition) by annular surrounds is consistently noted in the literature (Levitt & Lund, 1997; Walker, Ohzawa, & Freeman, 2002) but tends to be weaker than our model requires, and might depend on exact stimulus configurations. The orientation bandwidth, and overall shape, of surround suppression differs from study to study. While individual cells in visual cortex show the kinds of behaviour our model requires, the available data do not quantitatively constrain a choice of model.

In total, there are five free parameters in the model: the bandwidths of the distributions, α , β and γ ; the relative strength of the excitation and inhibition from the surround, q; and k, the overall strength of the surround.

To extract the predicted orientation, we compared the responses of the detectors in the model with a series of response templates, each with the same shape but centred on a different stimulus orientation (in steps of 0.1 deg). We assumed that, regardless of the stimulus, the distribution of responses across the population of detectors should have the same template, the tuning curve of the detectors (Dayan & Abbott, 2001). For each template, we calculated the probability of the population response, assuming a Poisson likelihood function (where response variability is proportional to response amplitude). The predicted orientation is the peak of the template which fits the population response with maximum likelihood. The exact physiological mechanism by which such an estimate may be calculated is unclear, but relevant neural models have been described (Deneve, Latham, & Pouget, 1999; Jazaveri & Movshon, 2006). In the simulations here we set the maximum amplitude of the distribution *R* to 90 and added a spontaneous activity of 10, as in Dakin et al. (2005); the same was done to the template. The peak and spontaneous activities are somewhat arbitrary, but resemble the firing rates (in impulses/s) of neurons in V1 and extrastriate cortical areas (Dean, 1981; Tolhurst, Movshon, & Dean, 1983). We first generated predictions across a coarse and even sampling of the parameter space, and then concentrated our search around that which gave the best prediction (as recorded in Appendix 2); not all the parameters are orthogonal, so the reported values are not guaranteed to be tightly constrained.

3. Results

For four observers we measured the simultaneous tilt illusion with sinusoidal gratings. Fig. 1C shows for two observers the subjective vertical of a central test grating, measured in the presence of an annular grating at each of several orientations. The shape of the curve replicates the findings of previous studies using gratings (Wenderoth & Johnstone, 1988) or bars (Westheimer, 1990). That is, for small angular separations between inducer and test we found a repulsion illusion, which for all observers peaked in magnitude at an inducer orientation between 10 and 20 deg from vertical (see average data in Appendix 3). For larger angular separations between inducer and test, we also found a smaller attraction illusion for most observers, which peaked in magnitude at an inducer orientation of 80 degrees from vertical.

3.1. Tilt illusion with broadband stimuli

We next measured the tilt illusion induced by broadband textures, using annular surfaces with a narrow (3.1 deg SD) distribution of orientations, and central (test) surfaces with either a narrow (3.1 deg SD) or broad (12.5 deg SD) distribution of orientations. The tilt illusion induced by the broadband textures is qualitatively similar to that for the gratings (Fig. 1D). In these observers and in all others tested (see average data in Appendix 3) there are some dominant trends. The orientation of the annulus that elicited the greatest illusion depends on the bandwidth of the central stimulus: the peak shifted to an annulus orientation further from vertical as the central stimulus bandwidth increased from 3.125 to 12.5 deg, but this was a small effect (a maximum of 10 deg in our coarsely sampled range of orientations). Additionally, as the bandwidth of the distribution of orientations in the central surface increased, the entire set of data points tended to shift upward, towards a repulsion illusion.

For two orientations of the annulus we parametrically varied the orientation distributions over the centre and surround (Fig. 2B and C, average data in Appendix 3). With the annulus at 15 deg from vertical, increasing the bandwidth of the annular surface decreases the magnitude of the repulsion illusion. At 75 deg from vertical, increasing the bandwidth of the annulus decreases the magnitude of the attraction illusion, and further increasing it induces a repulsion illusion. This is not a pure scaling of the magnitude of the illusion, so it is not simply caused by a decrease in the impact of the annulus when its average orientation is obscured by power at other orientations.

3.2. Tilt illusion with notched broadband stimuli

For the broadband stimuli described above, the average orientation of the inducing stimulus became less discernable as the distribution of orientations in the texture widened. However, this change was always coupled with the orientation distribution varying less steeply. We dissociated these two effects by measuring the tilt illusion for annuli with 'notched' orientation structure, in which the average orientation was sometimes not discernible, due to the power present at almost all orientations (see Fig. 3A). vet which often induced a robust illusion. Each notch texture has power at all orientations, except within a 'notch' where the power is described by an inverted Gaussian. For Gaussian broadband textures the average orientation of the texture is less apparent for large bandwidths-that is, when the slope of the amplitude orientation is shallow. This is not the case for notched textures: for narrow notch bandwidths the amplitude-orientation slope is steep, but observers find it difficult to report the average orientation.

We measured the tilt illusion induced by annular notch textures for two notch orientations: 15 and 75 deg from vertical. Across the four observers tested, annular textures with a notch centred on 75 deg did not induce a robust illusion, except at the widest notch bandwidth, which induces a small illusion, in which the perceived orientation of the central stimulus was attracted towards the orientation of the notch (see Fig. 3C, average data in Appendix 3). Surround textures with a notch centred on 15 deg induce a robust attraction illusion for all bandwidths of the notch, as seen in Fig. 3B. In this case the relationship between notch bandwidth and the attraction illusion was not monotonic: increasing the orientation notch bandwidth caused an increasing then decreasing attraction illusion.

For sinusoidal or broadband stimuli, surfaces with an average orientation of 15 deg induce a repulsion illusion of greater magnitude than the attraction illusion induced by surfaces with an average orientation of 75 deg. Here we found that textures with a notch at 15 deg (and hence an average orientation of -75 deg from vertical), induce an attraction illusion of greater magnitude than textures with a notch at 75 deg (and hence an average orientation of -15 deg from vertical). This implies that impact of the surround on the perceived orientation of the test is greatest not when the average orientation of the test, but when the relative power at nearby orientation changes steeply around the orientation of the test. Additionally, this illusion shifts further towards an attraction illusion when observers are asked to judge the orientation of a central test texture with a larger bandwidth.

3.3. Modelling the tilt illusion with broadband stimuli

Overall, these results demonstrate that the tilt illusion can be induced by surfaces that do not necessarily contain clear contours, and for some without a clearly perceptible average orientation. The tilt illusion induced by these broadband stimuli cannot be explained simply in terms of the extent to which the average orientation of either the test or inducer texture is obscured by power at other orientations. We therefore asked whether a simple model of orientation processing, capable of predicting the tilt illusion for sinusoidal gratings, was also capable of predicting the tilt illusion for these broadband textures: this model does not depend on the detection of contours or objects, but rather the Fourier amplitude spectra.



Fig. 1. Simultaneous tilt illusion as a function of surround orientation. (A and B) A sinusoidal grating (A) and an example of a broadband texture (B). To the right of each texture the relative amplitude of each texture in the Fourier plane are plotted across orientation and spatial frequency. (C and D) Average illusory tilt of the central surface for observers AW and EG, as a function of surround orientation. The magnitude of the illusion gives the amount to which the central patch was rotated away from vertical before the observer reported that the orientation of the central surface was vertical. Positive values of the illusion refer to illusory tilt of the central surface away from the orientation of the surround (repulsion illusion). Negative values of the illusion refer to illusory tilt of the central surface away from the orientation of the surround (repulsion illusion). Negative values of the illusion refer to illusory tilt of the central surface away from the orientation of the surround (repulsion illusion). Negative values of the illusion refer to illusory tilt of the central surface away from the orientation. It is possible to a surface towards the orientation of the surround (attraction illusion). Error bars indicate 95% confidence intervals of the mean illusion for each observer. Lines join the predictions of the best fitting model, obtained for observer AW with parameter values of a = 30, b = 50, g = 140, k = 0.4, q = 0.7, and for observer EG with parameter values of a = 30, b = 70, g = 120, k = 0.33, q = 0.90. In (C) both central and surrounding suffaces were sinusoidal gratings, in (D), both central and surrounding suffaces were broadband textures, with an orientation bandwidth of 3.125 deg in the surround, and either 3.125 or 12.5 deg in the centre.

The perception of orientation is usually thought to be dependent on a population of orientation selective detectors: the response properties of these detectors, whose receptive fields are sensitive to the central (test) stimulus, are altered by the presence of annular stimuli (giving rise to the tilt illusion), or are altered by adaptation (leading to the tilt aftereffect). In both cases the perceptual phenomena can be understood as the impact of the surround or adapting stimulus on the detector: either reducing the responsiveness of detectors, changing their orientation preference, or broadening their bandwidth (Clifford et al., 2000; Coltheart, 1971; Gilbert & Wiesel, 1990; Jin et al., 2005; Regan & Beverley, 1985); each of these manipulations is capable of shifting the activity of the population and therefore changing perceived orientation. Here we have modelled the influence of the surround as suppressing or facilitating the responsiveness of the detector. For each detector the influence of the surround depends on the orientation of the annular stimulus, and is most suppressive for annuli at the preferred orientation of the detector.



Fig. 2. Simultaneous tilt illusion as a function of surround orientation bandwidth. (A) Examples of broadband textures of increasing orientation bandwidth (3.125, 12.5, 25 and 50 deg). To the right of each texture the relative amplitude of each texture in the Fourier plane is plotted across orientation. (B and C) Average illusory tilt of the central surface for observers AW and EG, as a function of surround orientation bandwidth, conventions as in Fig. 1. In (B) and (C), the average surround orientations were 15 and 75 deg from vertical, respectively.

The model of the population response has five free parameters that define the shape of the tuning curve for the detector and surround, and the strength of the surround's input to the detector. The predicted tilt illusion was determined from the population response through template matching (see Section 2), but similar results were obtained when we used the population vector and winner-take-all methods. The lines on plots of results in Figs. 1–3 show the predictions of the best-fitting set of parameters, for each of two observers who completed all experimental conditions (AW and EG). For both observers the model was fit to the entire set of data, across the six conditions, as we wanted to know if a single set of model parameters could predict all the observations. In both cases, the model predictions from a single set of parameters generally fit well the observations for both narrow and broadband stimuli. The centre and surround tuning curves specified by these best fitting sets of model parameters are plotted in Appendix 4, for observers AW and EG.

As expected, the model fit is improved if only the data from a single condition is considered. The capacity of the model to fit data from a single condition is illustrated in Appendix 3, where the model was fit to each condition separately. For example, the best fitting overall models capture the shape of the curve, but not the magnitude of the illusion in Fig. 3 for AW, but the model is nevertheless capable of predicting both the shape and magnitude of this effect (see Appendix 3).



Fig. 3. Simultaneous tilt illusion as a function of surround orientation notch bandwidth. (A) Examples of broadband 'notched' textures of increasing orientation notch bandwidth (3.125, 12.5, 25 and 50 deg). To the right of each texture the relative amplitude of each texture in the Fourier plane is plotted across orientation. Note that the average orientation of each notched texture is orthogonal to the notch. (B and C) Average illusory tilt of the central surface for observers AW and EG, as a function of surround notch bandwidth, conventions as in Fig. 1, with the exception that the illusion is plotted with regard to the notch, rather than the average, orientation of the surround. In (B) and (C), the surround notch orientations were15 and 75 deg from vertical, respectively.

We calculated the mean square error (MSE) between the measured and predicted data points in each experiment, for each of 20,940 combinations (see Appendix 2 for details) of the five free model parameters, as a measure of the goodness of fit for each set of model parameters. The model provides reasonable predictions for both sinusoidal gratings and broadband textures, but this does not mean that parameters estimated from gratings are capable of predicting the illusions obtained with broadband textures. (Within each experiment many combinations of model parameters were capable of predicted the illusions observed). We would like to know whether the same sets of parameters are likely to have predictive power for both gratings and textures. We therefore determined the performance, for textures, of the parameter sets that best predicted the grating induced illusions (we identified from fits to the grating-induced illusions the best 10% of the parameter sets). For each experiment we then compared the predictive power of two groups of parameter sets: all the parameter sets tested (20,940 sets), and the 2094 sets whose predictions were best for the grating-induced illusion. We chose to compare the distributions of MSE values obtained in each group, as illustrated in Fig. 4, where these distributions are plotted for an example experimental condition for observer AW. The distribution of MSE values for model fits to sinusoidal gratings (A) and the varying orientation bandwidth condition with broadband textures, as plotted in Fig. 2B, (B) are shown in blue. The distributions of MSE values for the 10% of model parameters best fitting the tilt illusion with sinu-



Fig. 4. Mean Squared Error (MSE) values for model fits to the tilt illusion with sinusoidal gratings and broadband textures, for observer AW. The distribution of MSE values for model fits to sinusoidal gratings (A) and the varying orientation bandwidth condition with broadband textures, as plotted in Fig. 2B, (B) are shown by the upward going histograms (blue). The distributions of MSE values for the 10% of model parameters best fitting the tilt illusion with sinusoidal gratings are shown by the downwardly going histograms (red). Note the magnified scale. Some of the model parameter sets (1.5% of all model parameter sets for sinusoidal gratings, and 2.5% for the broadband textures) had a MSE value of greater than 10, and are excluded from the plot. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

soidal gratings are shown in red. In (A), the red histogram is a scaled version of the left end of the blue histogram. If the same 10% of model fits were the best fitting model parameters in (B), then the red histogram in (B) would be a magnified version of the left end of the blue histogram, as in figure (A). If the best fitting models in (A) performed no better than average in (B), we would expect the size and shape of the blue and red histograms in (B) to be the same.

The histograms in (B) are somewhere between these two extremes: the fact that the centre of mass of the red histogram is shifted further to the left than that of the blue one implies that the best fitting model parameters for sinusoidal gratings also performed better than average in predicting the illusion with broadband textures. We found a similar pattern of results for the data from other experimental conditions, and for observer EG. In each case we tested the significance of this shift: all distributions had a positive skew, so we compare the median, rather than the mean value of the two distributions in each case. The median MSE for the best grating sets was always lower than that of the total population of parameter sets (p < .01, Wilcoxin rank sum), so we can reject the hypothesis that they perform no differently to the entire set. We interpret this as evidence that the model parameters which best predict the illusion induced by gratings are also generally those that best predict the illusion induced with broadband stimuli.

3.4. Tilt illusion with natural image segments

Since the model considers only the amplitude spectra of textures, it predicts that the tilt illusion will be the same for textures with the same amplitude spectrum but different phase spectrum. We therefore measured the tilt illusion with twelve different natural image segments in the annulus, and with phase-randomised versions of these. Image segments were selected from the van Hateren database (see Section 2) such that their amplitude spectrum correlated strongly with one of those used to generate the broadband textures; four images were chosen for each of the three amplitude spectra with different bandwidths. Examples of the images selected, and their phase-randomised counterparts, are shown in Fig. 5A and B, respectively. Since we wanted to maximise the magnitude of the illusion so as to detect small differences, the centre test was always a broadband texture, as above, with a bandwidth of 12.5 deg. The surrounding image segment was rotated so that its average orientation was at 15 deg from vertical. Fig. 5C shows the tilt illusion induced by the surrounding natural image segments for observer AW; Fig. 5D shows the average of three observers. For all of the images, in all of the observers, the illusion was repulsive if present; as with the synthetic textures, images with larger orientation bandwidths induced tilt illusions of smaller magnitude than those with narrow bandwidths.

Randomizing the phase of the natural image segments tended to increase the magnitude of the tilt illusion. For observer AW (Fig. 5C), the illusion was significantly greater for the phase randomised images than for the natural images ($F_{(1,3)}$ = 38.50, 42.00 and 14.35, p < .01, for orientation bandwidths of 3.1, 12.5 and 50.0 deg, respectively; two-way ANOVA). This was also the case for the other three observers (all but three points were significantly different). On average, across observers, for inducing images with narrow orientation bandwidth (3.1 deg) the phase randomised images certainly induced greater illusions ($F_{(1,3)}$ = 8.98, p < .01). For images with a broader orientation bandwidth (12.5



Fig. 5. Simultaneous tilt illusion as a function of surround orientation bandwidth, with surrounds comprised of natural image segments or their phase randomised counterparts. (A) Examples of three natural image segments used, whose Fourier amplitude spectra correlate strongly with those of three of the broadband textures of increasing orientation bandwidth (3.125, 12.5 and 50 deg, from left to right). (B) Phase randomised versions of the natural image segments in (A). (C and D) Average illusory tilt of the central surface (a broadband texture of 12.5 deg orientation bandwidth) for observer AW (C) and averaged across four observers (D), as a function of surround orientation bandwidth. Conventions as in Fig. 1, with the exception that here lines join data points, they do not show model predictions, and in (D) the error bars indicate 95% confidence intervals of the mean illusion across observers.

or 50 deg), randomizing the image phases increased the magnitude of the illusion, but this effect was not significant when pooled across observers ($F_{(1,3)} = 2.66, 1.05, p = .12, .32$, respectively).

4. Discussion

Neural and perceptual responses to simple visual stimuli cannot necessarily be generalised to predict responses to complex visual stimuli (David et al., 2004; Olshausen & Field, 2005). By testing the simultaneous tilt illusion with both sinusoidal gratings and broadband textures, we are asking how well a simple model of orientation processing can be generalised to account for the illusion in both cases.

We demonstrate here that the tilt illusion is not a perceptual phenomenon unique to bars or gratings, but can be induced by stimuli with broad distributions of orientations and spatial frequencies. The magnitude and direction of the tilt illusion induced by these broadband stimuli is not predicted by the presence of contours or the clarity of the average orientation in the centre and surround surfaces. Instead, we find that the influence of the surround on the centre is best understood in terms of the overall distributions of orientations in the surfaces.

When the central stimulus is narrowband all detectors that respond to the stimulus have similar preferred orientations. When the bandwidth of the central stimulus is larger, the stimulus recruits detectors whose preferred orientation lies further from the mean orientation of the central stimulus: for some this is close to the mean orientation of the surrounding stimulus. Our model suggests that the influence of the surround on these detectors is substantial, and so the surround has a greater influence on the population response to the central stimulus, and therefore induces a greater illusory tilt of the central stimulus.

In testing the tilt illusion induced by natural image segments, we find that for stimuli with non-random phase the tilt illusion shows similar trends to the illusion induced by the broadband textures. There is, however, a small but significant discrepancy between the tilt illusion induced by the natural images and their phase-randomised counterparts, and we must reject the hypothesis that phase randomisation has no effect on the tilt illusion. Nevertheless the mechanism remains unclear: our results may reflect the impact of contours on the illusion, but they might also reflect the spatial distribution of orientation structure. Phase-randomised natural images lack recognisable objects (Piotrowski & Campbell, 1982), and do not contain as many contours, but additionally these surfaces have a more homogenous local orientation structure than their natural image counterparts. We think it likely that the spatial extent of the detector's surround does not extend to the entire size of the annulus in our stimulus. If this is the case, then the orientation structure falling over the surround may vary between the natural image and the phase-randomised version, and this might contribute to the changes in the illusion induced by phase randomisation. Our experiments do not distinguish this: further investigation is required to determine whether the impact of phase-randomisation on the tilt illusion is independent of its impact on the local amplitude spectrum. If it were not, this may imply a phase-independent surround mechanism.

Modelling the tilt illusion provides a means of testing whether our understanding of this illusion with gratings can be generalised to account for the illusion with broadband textures. If we are unable to generalise the model from gratings to broadband textures, then the model can be rejected as incorrect. Alternatively, the model may describe only the 'special case' of sinusoidal gratings, and hence be uninformative about those mechanisms underlying orientation processing that are engaged under natural viewing conditions. If the model can be generalised to predict the illusion for broadband textures, we can use these results to further constrain the plausible parameters of the model.

Here the tilt illusion was modelled by considering the population response of a group of orientation selective detectors whose response properties are not inconsistent with known physiology of orientation selective cells in the primate visual system (Hubel & Wiesel, 1968; Shapley, Hawken, & Ringach, 2003; Webb, Dhruv, Solomon, Tailby, & Lennie, 2005). We found that this model, which considers only the distribution of power across orientation in test and inducing stimuli, can account for the majority of the observed tilt illusion for both sinusoidal gratings and broadband textures. Furthermore, the best fitting model parameters for the illusion with gratings were also better at accounting for the illusion with broadband textures.

5. Conclusions

We show that a simple model of orientation processing predicts perceptual responses, namely, the tilt illusion, for both gratings and broadband textures. Although the model considers only the Fourier amplitude spectra of images, it accounts for the majority of the illusion induced by narrow and broadband textures, and natural and phase-randomised images.

Acknowledgments

We thank P. Bex for help in stimulus generation. Support from the National Health and Medical Research Council of Australia, Grants 457337 and 211247 to S.G.S. and an Australian Research Fellowship from the Australian Research Council to C.W.G.C.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.visres.2008.02.023.

References

- Bex, P. J., & Makous, W. (2002). Spatial frequency, phase, and the contrast of natural images. Journal of the Optical Society of America A Optics Image Science and Vision, 19, 1096–1106.
- Brainard, D. H. (1997). The psychophysics toolbox. Spatial Vision, 10, 433-436.
- Burr, D., Morrone, C., & Maffei, L. (1981). Intra-cortical inhibition prevents simple cells from responding to textured visual patterns. *Experimental Brain Research*, 43, 455–458.
- Cavanaugh, J. R., Bair, W., & Movshon, J. A. (2002). Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *Journal of Neurophysiology*, 88, 2547–2556.
- Clifford, C. W., Wenderoth, P., & Spehar, B. (2000). A functional angle on some aftereffects in cortical vision. Proceedings of Biological Sciences, 267, 1705–1710.
- Coltheart, M. (1971). Visual feature-analyzers and after-effects of tilt and curvature. *Psychological Review*, 78, 114–121.
- Dakin, S. C., Mareschal, I., & Bex, P. J. (2005). Local and global limitations on direction integration assessed using equivalent noise analysis. *Vision Research*, 45, 3027–3049.
- David, S. V., Vinje, W. E., & Gallant, J. L. (2004). Natural stimulus statistics alter the receptive field structure of v1 neurons. *Journal of Neuroscience*, 24, 6991–7006.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience*. Cambridge, MA: MIT Press. Dean, A. F. (1981). The variability of discharge of simple cells in the cat striate cortex. *Experimental Brain Research*, 44, 437–440.
- Deneve, S., Latham, P. E., & Pouget, A. (1999). Reading population codes: a neural implementation of ideal observers. *Nature Neuroscience*, 2, 740–745.
- Gilbert, C. D., & Wiesel, T. N. (1990). The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Research*, 30, 1689–1701.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195, 215–243.
- Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience*, 9, 690–696.
- Jin, D. Z., Dragoi, V., Sur, M., & Seung, H. S. (2005). Tilt aftereffect and adaptationinduced changes in orientation tuning in visual cortex. *Journal of Neurophysiology*, 94, 4038–4050.
- Kontsevich, L. L., & Tyler, C. W. (1999). Bayesian adaptive estimation of psychometric slope and threshold. Vision Research, 39, 2729–2737.
- Levitt, J. B., & Lund, J. S. (1997). Contrast dependence of contextual effects in primate visual cortex. Nature, 387, 73–76.
- Mardia, K. V., & Jupp, P. E. (2000). Directional statistics. Chichester, UK: Wiley.
- McDonald, J. S., & Tadmor, Y. (2006). The perceived contrast of texture patches embedded in natural images. *Vision Research*, 46, 3098–3104.
- Muller, J. R., Metha, A. B., Krauskopf, J., & Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science*, 285, 1405–1408.
- Olshausen, B. A., & Field, D. J. (2005). How close are we to understanding v1? Neural Computation, 17, 1665–1699.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. Spatial Vision, 10, 437–442.
- Piotrowski, L. N., & Campbell, F. W. (1982). A demonstration of the visual importance and flexibility of spatial-frequency amplitude and phase. *Perception*, 11, 337–346.
- Regan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. Journal of the Optical Society of America A, 2, 147–155.
- Ruderman, D. L., & Bialek, W. (1994). Statistics of natural images: Scaling in the woods. *Physical Review Letters*, 73, 814–817.
- Shapley, R., Hawken, M., & Ringach, D. L. (2003). Dynamics of orientation selectivity in the primary visual cortex and the importance of cortical inhibition. *Neuron*, 38, 689–699.
- Tadmor, Y., & Tolhurst, D. J. (1994). Discrimination of changes in the second-order statistics of natural and synthetic images. Vision Research, 34, 541–554.
- Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research*, 23, 775–785. van der Schaaf, A., & van Hateren, I. H. (1996). Modelling the power spectra of
- natural images: Statistics and information. Vision Research, 36, 2759–2770. van Hateren, J. H., & van der Schaaf, A. (1998). Independent component filters of
- van Hateren, J. H., & van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings* of Biological Sciences, 265, 359–366.
- Walker, G. A., Ohzawa, I., & Freeman, R. D. (2002). Disinhibition outside receptive fields in the visual cortex. *Journal of Neuroscience*, 22, 5659–5668.
- Webb, B. S., Dhruv, N. T., Solomon, S. G., Tailby, C., & Lennie, P. (2005). Early and late mechanisms of surround suppression in striate cortex of macaque. *Journal of Neuroscience*, 25, 11666–11675.
- Webster, M. A., & Miyahara, E. (1997). Contrast adaptation and the spatial structure of natural images. Journal of the Optical Society of America A Optics Image Science and Vision, 14, 2355–2366.
- Wenderoth, P., & Johnstone, S. (1988). The different mechanisms of the direct and indirect tilt illusions. *Vision Research*, 28, 301–312.
- Westheimer, G. (1990). Simultaneous orientation contrast for lines in the human fovea. Vision Research, 30, 1913–1921.