THE ROYAL SOCIETY

doi 10.1098/rspb.2000.1198

A functional angle on some after-effects in cortical vision

Colin W. G. Clifford[1*], **Peter Wenderoth**[1] **and Branka Spehar**[2]

[1]*Department of Psychology, Macquarie University, Sydney, New South Wales 2109, Australia*
[2]*Department of Psychology, University of New South Wales, Sydney, New South Wales 2052, Australia*

The question of how our brains and those of other animals code sensory information is of fundamental importance to neuroscience research. Visual illusions offer valuable insight into the mechanisms of perceptual coding. One such illusion, the tilt after-effect (TAE), has been studied extensively since the 1930s, yet a full explanation of the effect has remained elusive. Here, we put forward an explanation of the TAE in terms of a functional role for adaptation in the visual cortex. The proposed model accounts not only for the phenomenology of the TAE, but also for spatial interactions in perceived tilt and the effects of adaptation on the perception of direction of motion and colour. We discuss the implications of the model for understanding the effects of adaptation and surround stimulation on the response properties of cortical neurons.

Keywords: tilt after-effect; tilt illusion; motion; colour; computational model; decorrelation

1. INTRODUCTION

Through the study of illusions and after-effects, psychophysicists aim to gain insight into the mechanisms of visual processing. Evolution has tailored our sensory systems to provide us with the information we need to function in our environment. However, our vision can sometimes give rise to illusions or misperceptions. It is widely believed that when we experience perceptual illusions our sensory systems are actually functioning in the normal way (Barlow & Foldiak 1989). Illusions and after-effects can thus be valuable diagnostic tools, in the sense that they are potentially very informative about the mechanisms of visual processing and the status of competing theories.

Here, we are concerned primarily with the perception of spatial orientation (tilt). The extraction of information about spatial orientation within an image is an important step in the recovery of the three-dimensional structure of objects and the layout of our surroundings. The phenomenology of the tilt after-effect (TAE) can be summarized as follows. Prolonged exposure to an orientated pattern affects the perceived orientation of a subsequently observed pattern (Gibson & Radner 1937). For adapting orientations of 0–50°, a vertical test appears to be repelled away from the adapter in orientation, with the strongest effect occurring between 10° and 20°. For larger angles there is a smaller attraction effect, such that a vertical test appears rotated towards the adapter. The strongest attraction effect is observed between 75° and 80° (Wenderoth & Johnstone 1987) and occurs robustly when the display includes relatively large adapting gratings or long adapting and test lines (Gibson & Radner 1937; Kohler & Wallach 1944; Morant & Harris 1965). Gibson & Radner explained the TAE in terms of the visual system's tendency towards the 'norms' of spatial orientation: horizontal and vertical. However, TAEs to horizontally or vertically orientated adapting patterns with tilted test stimuli (Kohler & Wallach 1944) cannot be explained in this way. Indeed, the similarity of the angular dependence of the TAE for vertical and oblique test stimuli (Mitchell & Muir 1976) suggests that the effect is better understood in terms of relative rather than absolute or normative orientations.

Convergent evidence from physiology, computational theory and psychophysics indicates that the mechanisms underlying orientation perception are closely related to those in other visual modalities. It is known that the primate primary visual cortex contains columns of orientation-selective cells (Hubel & Wiesel 1962). This columnar architecture is also found elsewhere in the visual cortex and beyond (Mountcastle 1997). The middle temporal area, for example, contains columns of cells tuned to direction of motion (Albright *et al.* 1984). From a computational perspective, the problem of extracting information from the image signal can be cast as the recovery of orientation information in various domains (Adelson & Bergen 1991). Motion, for example, can be considered as orientation in space–time. Psychophysically, the domains of spatial orientation (Gibson & Radner 1937) and direction of motion (Schrater & Simoncelli 1998) share a common phenomenology of repulsive and attractive adaptation effects, with prolonged exposure to a moving pattern affecting the perceived direction of subsequent motion (Levinson & Sekuler 1976; Patterson & Becker 1996; Schrater & Simoncelli 1998). Adaptation has also been shown to affect colour appearance, such that prolonged exposure to an isoluminant field modulating in colour over time alters the perceived hue of subsequently presented chromatic stimuli (Webster & Mollon 1991, 1992).

We propose a model suggesting that misperceptions of orientation result from self-calibration and decorrelation (Barlow & Foldiak 1989) in its neural representation. These functional principles are motivated by two main considerations. First, the visual system must be self-calibrating in its mapping of environmental stimulation onto patterns of neural activity (Andrews 1964). Self-calibration is the property of a system to change itself in response to changes in the environment (recalibration)

---

[*] Author for correspondence (colinc@vision.psy.mq.edu.au).

Received 15 March 2000 *Accepted* 2 June 2000

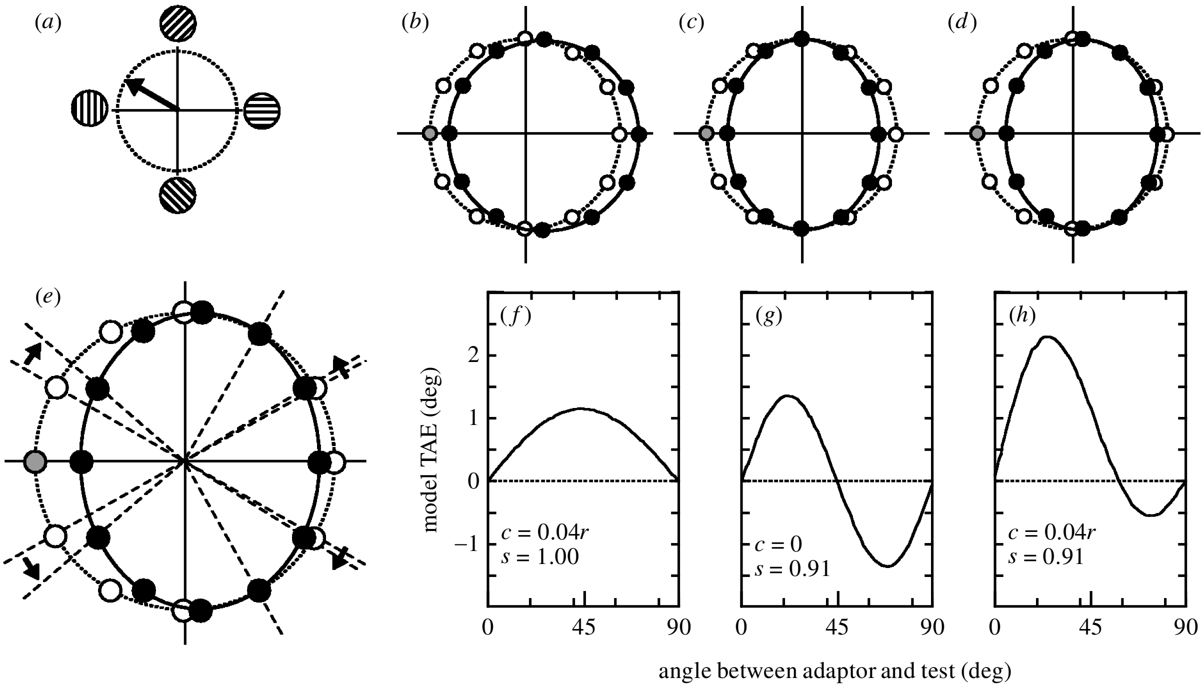

Figure 1. Description of the model. (*a*) Schematic diagram of the representation of spatial orientation used by the model. The model response is a vector in a two-dimensional space. The direction of this response vector represents the perceived tilt. In this representation, horizontal and vertical are opposites. Thus, 180° in model space corresponds to 90° in tilt. The effects of (*b*) centring, (*c*) scaling and (*d*) centring and scaling in combination: white circles represent the response vector for a range of orientations in the unadapted state; black circles show the response after prolonged exposure to the orientation denoted by the grey circle. Adaptation is modelled as a shift in the origin of the stimulus–response mapping and a rescaling of the adapting dimension. The degree of adaptation is controlled by two parameters: a centring parameter, c, and a scaling parameter, s. The origin of the stimulus–response mapping shifts by an amount cr. The adapting dimension is rescaled by a factor s. The geometry illustrated in (*e*) allows us to relate the perceived orientation of the test stimulus with (black circles) and without (white circles) adaptation. Predicted angular tuning functions of the TAE due to the effects of (*f*) centring, (*g*) scaling and (*h*) a combination of centring and scaling.

and to adjust to perturbations within the system in an unchanging environment (error correction). Second, adaptation tends to optimize the use of the limited dynamic range of the visual pathways for the coding of visual stimuli (Laughlin 1989) by reducing the transmission of redundant information (Attneave 1954). Dynamic range optimization reduces redundancy in the responses of individual sensory neurons, maximizing the effective bandwidth available for the transmission of novel information about the stimulus (Srinivasan *et al.* 1982). The principle of redundancy reduction can be extended from single neurons to populations of neurons by adaptively decorrelating (Barlow & Foldiak 1989; Atick *et al.* 1993) or orthogonalizing (Kohonen & Oja 1976; Zaidi & Shapiro 1993) their responses.

2. A MODEL OF CORTICAL ADAPTATION

The mechanisms of self-calibration in our population of model neurons are termed centring and scaling. The population response is modelled as a vector whose direction represents the perceived tilt (figure 1*a*). In this representation, horizontal and vertical are opposites, so 180° in model space corresponds to 90° in tilt. Centring operates to set the zero point of the population response according to the prevailing stimulus distribution, while scaling serves to decorrelate the model's responses to those stimuli. Centring can be thought of as adaptation to the centroid of the stimulus distribution and scaling as adaptation to variation along a particular stimulus dimension, such as horizontal–vertical orientation. The effects of centring and scaling, individually and in combination, are illustrated schematically in figure 1*b*–*d*. Without loss of generality, it can be assumed that the response vector to the adapting stimulus in Cartesian coordinates is $(r, 0)$ in the population(s) of neurons sampling that region of the image. The effect of the adapter is to shift (centre) the stimulus–response mapping of the test stimulus directly away from the orientation of the adapter by an amount cr (figure 1*b*) and to scale the (adapting) x-dimension by a factor s (figure 1*c*). Adaptation is modelled as a combination of these two operations (figure 1*d*). The geometry illustrated in figure 1*e* allows us to relate the perceived orientation of the test stimulus with (θ_A) and without (θ_0) adaptation according to the following equation:

$$\sin(\theta_A) = \frac{\sin(\theta_0)}{\sqrt{\left(s \times \cos(\theta_0) - \left(\frac{c}{r}\right)^2\right)^2 + \sin^2(\theta_0)}}. \qquad (1)$$

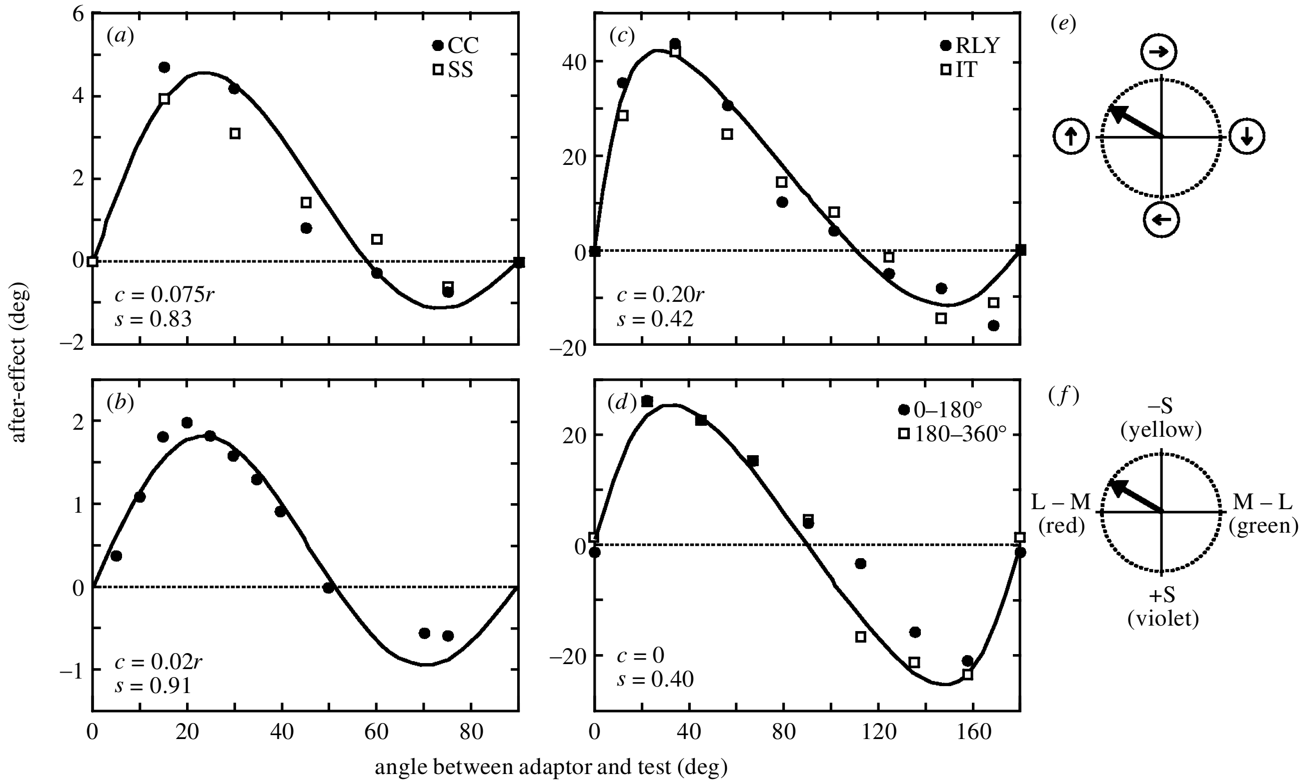

Figure 2. Predictions of the model applied to data on (*a*) the tilt after-effect, (*b*) the tilt illusion, and the effects of adaptation on the perception of (*c*) direction of motion and (*d*) hue. Data are redrawn from Clifford *et al.* (2000), Westheimer (1990, fig. 1, averaged across subjects), Schrater & Simoncelli (1998, fig.1) and Webster & Mollon (1992, fig. 7), respectively. Webster & Mollon used adapting stimuli modulating in colour over time between isoluminant red and green, for which no net centring effect would be expected, so in this case the model's centring parameter was set to zero. (*e*) The representation of direction of motion used by the model. The direction of the response vector corresponds to the perceived direction of motion. (*f*) The representation of the isoluminant plane used by the model. The direction of the response vector corresponds to the perceived hue of the stimulus.

Centring alone produces only repulsive interactions between stimuli of different orientations (figure 1*f*). Scaling can produce repulsion or attraction, depending on the angular difference between the orientated patterns in question (figure 1*g*). The effects of centring and scaling combine to generate angular tuning functions of the form shown in figure 1*h*.

3. SIMULATIONS OF PSYCHOPHYSICAL PERFORMANCE

Figure 2*a*,*b* shows fits of the model to data from the TAE and its spatial analogue, the tilt illusion (TI). In the TI, the presence of an orientated surround stimulus biases the perceived orientation of a simultaneously presented test (Gibson & Radner 1937; Wenderoth & Johnstone 1987; Westheimer 1990). The TAE and TI show a very similar dependence on the angle between the adapting (inducing) stimulus and the test (O'Toole & Wenderoth 1977) and are believed to engage similar mechanisms (Wenderoth & Johnstone 1988).

The tuning curve of the direction after-effect (DAE) reported by Schrater & Simoncelli (1998) is strikingly similar to that of the TAE (figure 2*a*,*c*). The DAE is repulsive for small (0–100°) differences in direction between adapter and test, with a smaller attraction effect for larger differences. The largest repulsion is evident at around 30° and the largest attraction at 150–160°. We note that the 'critical' values for the DAE are consistently around twice those for tilt. Representing tilt in vector form highlights the potential for similar strategies in the coding of tilt (figure 1*a*) and direction of motion (figure 2*e*) but, while orientations differing by 90° (e.g. horizontal–vertical) are opposites in tilt, opposites in direction of motion are 180° apart (e.g. up–down). Hence, the model predicts that interactions in the motion domain occur at twice the angular difference of their tilt analogues.

To establish the analogy between the TAE and the effect of adaptation on colour appearance, it is first necessary to define the appropriate colour space. A colour space can be defined by one luminance axis and two chromatic axes (MacLeod & Boynton 1979; Derrington *et al.* 1984). Measures of threshold sensitivity reveal two cardinal chromatic axes (Krauskopf *et al.* 1982): RG and YV (red–green and yellow–violet). Lights represented by points on the RG axis are indistinguishable to the S-cones in the retina and differ only in L- and M-cone excitation. Points on the YV axis are indistinguishable to L- and M-cones and differ only in S-cone excitation. All lights of a given luminance occupy a single plane in this space (figure 2*f*). Webster & Mollon (1991, 1992) demonstrated that adaptation to an isoluminant field modulating in colour over time affects the appearance of subsequently presented stimuli of the same contrast and luminance such that perceived hue is repelled away from the adapting axis. Temporal modulation of the adapting stimulus isolates cortical mechanisms by avoiding adaptation at the receptor level. The strongest repulsion occurs when the adapting and test stimuli differ by an angle of about 22.5° in colour space, with little or no effect when they are 90° apart (figure 2*d*). In the context of colour vision, centring corresponds to adaptation to the mean chromaticity of the prevailing stimulus distribution, while scaling is adaptation to the contrast along a particular (not necessarily cardinal) chromatic axis. We find that the angular dependence of the shift in perceived hue can be modelled by scaling in the absence of centring. This is a consequence of adapting to a modulation between opposites (red–green), since any centring effects would cancel out. The tilt analogue of this experiment is adaptation to alternately presented plaid components, which exhibits a similar angular dependence (S. Smith, C. W. G. Clifford and P. Wenderoth, unpublished data).

4. RELATIONSHIP TO EXISTING THEORIES

We have argued that after-effects in the perception of tilt, motion and colour are a consequence of adaptive coding strategies employed by our visual systems, and that the proposed mechanisms of centring and scaling also underlie spatial interactions in perceived tilt. Centring and scaling have a functional basis in maximizing the information content of the population response, and are analogous to the centring and scaling transformations applied to data prior to regression analysis (Draper & Smith 1998). Centring is essentially an error-correcting distribution shift (Levinson & Sekuler 1976; Attneave 1954; Mather 1980) while scaling is equivalent to the transformation proposed by Atick *et al.* (1993) to underlie the effect of cortical adaptation on colour appearance. In the language of control theory, centring is a form of additive (subtractive) gain control while scaling is divisive (multiplicative) in nature (Carandini & Heeger 1994). While previous models have attempted to apply the concepts of decorrelation and information maximization to cortical adaptation (Dong 1996; Wainwright 1999), none provides a functional account of the existence and relative magnitude of the repulsive and attractive effects.

The functional decomposition into centring and scaling operations proposed here is substantiated by psychophysical observations. Morant & Harris (1965) measured the TAE relative to the orientation of an objectively vertical line in the opposite hemifield, rather than relative to subjective vertical. They reasoned that the angular tuning function of the TAE measured in this way would reduce effects operating over a wide spatial extent, as these would affect both the test and comparison stimuli. They reported a purely repulsive TAE under these conditions, as would be predicted if centring were the dominant effect. By subtracting this angular tuning function from that measured in the normal way, they inferred the shape of the angular tuning curve of the TAE presumably due to the second, more global, process. This curve had a form consistent with the operation of a scaling process. For the TI, Wenderoth & Johnstone (1987) also found that minimizing low-level contributions reduced the magnitude of the repulsion effect to the level of the attraction effect, as would be predicted if only a scaling process were still operating.

5. IMPLICATIONS FOR UNDERSTANDING THE RESPONSE PROPERTIES OF CORTICAL NEURONS

The effects modelled here have all been demonstrated to have a cortical locus. First, the orientation tunings of the TAE and TI appear inconsistent with the properties of cells in the sub-cortical visual pathway, giving those effects a striate or extrastriate origin (Georgeson 1973; Ware & Mitchell 1974; Wenderoth & Johnstone 1987). Second, a direction after-effect has been reported for stereoscopic (Cyclopean) motion, placing its substrate at or after the point of binocular combination in the primary visual cortex (V1) (Patterson & Becker 1996). Third, the chromatic adaptation procedure used by Webster & Mollon (1991) was designed to isolate post-receptoral mechanisms, and lateral geniculate neurons would have been little affected by a uniform field modulating slowly in chromaticity (Derrington *et al.* 1984), implying a cortical locus for the effects on perceived hue.

Centring and scaling are proposed to operate amongst populations of neurons. It is nonetheless possible to speculate upon their implementation at the level of single neurons, where physiological data is available on the effects of adaptation and surround modulation in various areas of the visual cortex (Blakemore & Tobin 1972; Gilbert & Wiesel 1990; Allman *et al.* 1985; Zeki 1983). Following Gilbert & Wiesel (1990), we simulated on computer the response properties of a set of 18 orientation-selective filters peaking 10° apart. The response, $f(\theta)$ of each filter as a function of orientation, θ, is defined by an equation of the form

$$f(\theta) = \alpha \exp(\beta(\cos(\theta - \theta_0) - 1)), \tag{2}$$

where α is the peak response, β controls the width of the tuning curve and θ_0 is the peak tuning. This function is known as the circular normal distribution (Pouget *et al.* 1998), a periodic function with a profile very similar to a Gaussian distribution. In the unadapted state, all the filters have parameter values of $\alpha = 1.0$ and $\beta = 5.0$, giving a full width at half height of just over 30° (figure 3*a*), well within the range reported for simple cells in V1 (Watkins & Berkley 1974). Each filter acts as a 'labelled line', representing its preferred orientation. The response of each filter is a vector pointing in the direction representing the filter's preferred orientation with a length proportional to the filter's response. The response of the population is the vector sum of the individual filters' responses, which produces veridical estimates at all orientations (figure 3*b*). To simulate the effects of an orientated inducing (adapting or surrounding) stimulus we adjust the filter responses in two ways. First, we inhibit the responses of filters tuned to nearby orientations (figure 3*c*). This shifts the population response away from the inducing orientation. The peak responses after centring are given by

$$\alpha(\theta_0) = \alpha[1 - \lambda(1 + \cos\theta_0)], \tag{3}$$

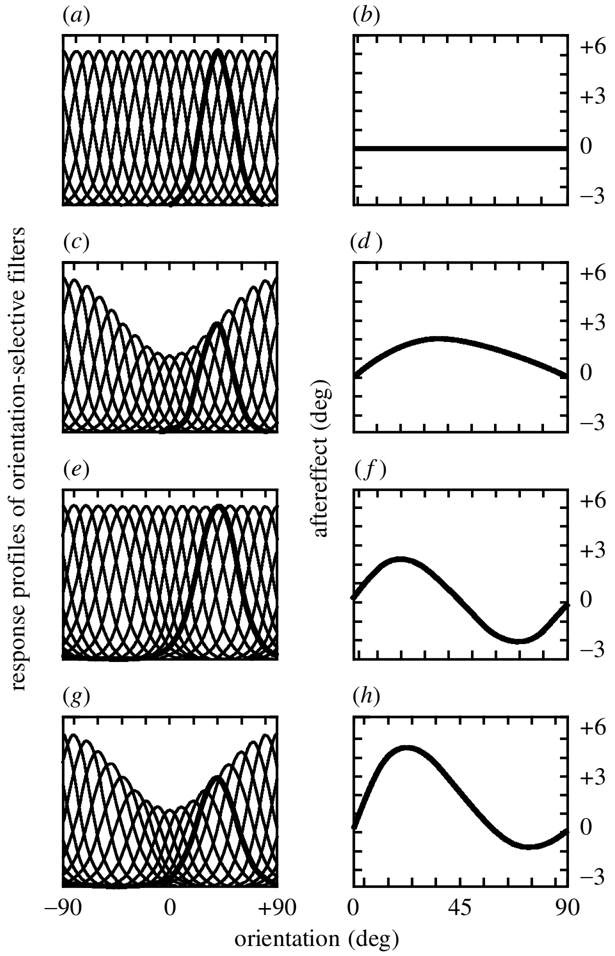

Figure 3. Possible neural implementation of the model. (*a*) The responses of 18 orientation-selective filters peaking 10° apart. The tuning curve of the filter peaking at +40° is shown in bold for illustrative purposes. In the unadapted state, the filters have equal peak responses and a full width at half height of 30°, which produces veridical estimates at all orientations (*b*). To simulate the effects of an orientated inducing stimulus we adjust the filter responses in two ways. (*c*) First, we inhibit the responses of filters tuned to nearby orientations. (*d*) This shifts the population response away from the inducing orientation, approximating centring. (*e*) Second, we increase the orientation bandwidth of filters with preferred tunings away from the inducing axis in the orientation domain. (*f*) The effect of this broadening resembles scaling. (*g*) In combination, inhibition around the inducing orientation and broadening away from the inducing axis produce (*h*) an angular tuning function of the form observed psychophysically.

where λ controls the degree of centring. For the simulations shown in figure 3*c*,*g*, $\lambda = 0.25$. When the width of inhibition is broad, the effect closely approximates centring (compare figures 1*f* and 3*d*). Second, we increase the orientation bandwidth of filters with preferred tunings away from the inducing axis in the orientation domain. For example, if the inducer is vertical, the inducing axis will be vertical–horizontal, and bandwidth will be broadened most for filters tuned to oblique orientations (figure 3*e*). The orientation bandwidths after scaling are given by

$$\beta(\theta_0) = \beta[1 - \mu(1 - 2\cos\theta_0)], \qquad (4)$$

where μ controls the degree of scaling. For the simulations shown in figure 3*e*,*g*, $\mu = 0.20$. It should be noted that for each filter the peak tuning, θ_0, is unaffected by adaptation. In isolation, the effect of this broadening closely resembles scaling (compare figures 1*g* and 3*f*). Reducing the breadth of filters responding to horizontal and vertical orientations produces a similar result. In combination, inhibition around the inducing orientation and broadening away from the inducing axis produce an angular tuning function of the form observed psychophysically (figure 3*h*).

Centring and scaling are implemented in our set of orientation-selective filters in ways consistent with the physiological response properties of V1 neurons. The broadly tuned profile of lateral inhibition amongst our filters resembles that observed experimentally in the presence of an orientated inducing stimulus surrounding the receptive field (Blakemore & Tobin 1972). Orientated surround stimuli have also been shown to induce changes in orientation bandwidth in the response properties of V1 cells (Gilbert & Wiesel 1990). Broadening of the orientation bandwidth of filters tuned to unstimulated orientations could result from a reduction in intracortical inhibition from cells suppressed by the inducer. Such intracortical feedback has been proposed to sharpen the tuning curves of V1 cells (Somers *et al.* 1995).

The functional bases of the proposed centring and scaling operations are self-calibration and decorrelation of the responses to the prevailing stimulus distribution. Recently, direct evidence of decorrelation amongst the responses of populations of V1 cells has been reported (Muller *et al.* 1999). Muller *et al.* found that the peak orientation tuning of complex cells, but not of simple cells, depended on the orientation of the adapting stimulus, although adaptation did affect the responsiveness of simple cells. They point out that their results are consistent with adaptation at the level of simple cells, which in turn drive complex cells. That account is entirely in accord with the model proposed here, in which adaptation alters the responsiveness and orientation bandwidth of simple cells but not their peak tuning. The changes in simple-cell responses affect the way in which stimulus orientation is mapped onto the model response vector, causing biases in perception.

6. SUMMARY

In summary, we have demonstrated that application of the principles of self-calibration and decorrelation in the cortical representation of visual information can account for the existence and form of temporal and spatial interactions observed psychophysically in a number of visual domains. While these principles operate at the level of neuronal populations, we have tentatively linked self-calibration and decorrelation at the single-neuron level with changes in responsiveness and orientation bandwidth, respectively. The ideas forwarded here will motivate further work to clarify the links between visual after-effects and their neural substrates, and to establish whether the proposed coding principles generalize beyond the visual cortex.

We are grateful to John Ross and Derek Arnold for comments on draft versions of this manuscript. This work is supported by Australian Research Council Large Grants to P.W. and C.W.G.C. and a Macquarie University Research Fellowship to C.W.G.C.

REFERENCES

Adelson, E. H. & Bergen, J. R. 1991 The plenoptic function and the elements of early vision. In *Computational models of visual processing* (ed. M. S. Landy & J. A. Movshon), pp. 3–20. Cambridge, MA: MIT Press.

Albright, T. D., Desimone, R. & Gross, C. G. 1984 Columnar organization of directionally selective cells in visual area MT of the macaque. *J. Neurophysiol.* **51**, 16–31.

Allman, J., Miezin, F. & McGuinness, E. 1985 Direction- and velocity-specific responses from beyond the classical receptive field in the middle temporal visual area (MT). *Perception* **14**, 105–126.

Andrews, D. P. 1964 Error-correcting perceptual mechanisms. *Q. J. Exp. Psychol.* **16**, 104–115.

Atick, J. J., Li, Z.-P. & Redlich, A. N. 1993 What does post-adaptation appearance reveal about cortical color representation? *Vision Res.* **33**, 123–129.

Attneave, F. 1954 Some informational aspects of visual perception. *Psychol. Rev.* **61**, 183–193.

Barlow, H. B. & Foldiak, P. 1989 Adaptation and decorrelation in the cortex. In *The computing neuron* (ed. R. Durbin, C. Miall & G. J. Mitchison), pp. 54–72. Wokingham, UK: Addison-Wesley.

Blakemore, C. & Tobin, E. A. 1972 Lateral inhibition between orientation detectors in the cat's visual cortex. *Exp. Brain Res.* **15**, 439–440.

Carandini, M. & Heeger, D. J. 1994 Summation and division by neurons in visual cortex. *Science* **264**, 1333–1336.

Clifford, C. W. G., Smith & Wenderoth, P. 2000 Inter- and intra-fold tilt after-effects. (In preparation.)

Derrington, A. M., Krauskopf, J. & Lennie, P. 1984 Chromatic mechanisms in lateral geniculate nucleus of macaque. *J. Physiol.* **357**, 241–265.

Dong, D. W. 1996 Associative decorrelation dynamics in visual cortex. In *Lateral interactions in the cortex: structure and function* (ed. J. Shirosh, R. Miikkulainen & Y. Choe). Austin, TX: The UTCS Neural Networks Research Group.

Draper, N. R. & Smith, H. 1998 *Applied regression analysis*. New York: Wiley.

Georgeson, M. A. 1973 Spatial frequency selectivity of a visual tilt illusion. *Nature* **245**, 43–45.

Gibson, J. J. & Radner, M. 1937 Adaptation, after-effect, and contrast in the perception of tilted lines. I. Quantitative studies. *J. Exp. Psychol.* **20**, 453–467.

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 Res.* **30**, 1689–1701.

Hubel, D. H. & Wiesel, T. N. 1962 Receptive fields, binocular architecture and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106–154.

Kohler, W. & Wallach, H. 1944 Figural aftereffects: an investigation of visual processes. *Proc. Am. Phil. Soc.* **88**, 269–357.

Kohonen, T. & Oja, E. 1976 Fast adaptive formation of orthogonalizing filters and associative memory in recurrent networks of neuron-like elements. *Biol. Cybern.* **21**, 85–95.

Krauskopf, J., Williams, D. R. & Heeley, D. W. 1982 Cardinal directions of color space. *Vision Res.* **22**, 1123–1131.

Laughlin, S. B. 1989 Coding efficiency and design in visual processing. In *Facets of vision* (ed. D. G. Stavenga & R. C. Hardie), pp. 213–234. Berlin: Springer.

Levinson, E. & Sekuler, R. 1976 Adaptation alters perceived direction of motion. *Vision Res.* **16**, 779–781.

MacLeod, D. I. A. & Boynton, R. M. 1979 Chromaticity diagram showing cone excitation by stimuli of equal luminance. *J. Opt. Soc. Am.* **69**, 1183–1186.

Mather, G. 1980 The movement aftereffect and a distribution-shift model for coding the direction of visual movement. *Perception* **9**, 379–392.

Mitchell, D. E. & Muir, D. W. 1976 Does the tilt aftereffect occur in the oblique meridian? *Vision Res.* **16**, 609–613.

Morant, R. B. & Harris, J. R. 1965 Two different aftereffects of exposure to visual tilts. *Am. J. Psychol.* **78**, 218–226.

Mountcastle, V. B. 1997 The columnar organization of the neocortex. *Brain* **120**, 701–722.

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.

O'Toole, B. & Wenderoth, P. 1977 The tilt illusion: repulsion and attraction effects in the oblique meridian. *Vision Res.* **17**, 367–374.

Patterson, R. & Becker, S. 1996 Direction-selective adaptation and simultaneous contrast induced by stereoscopic (Cyclopean) motion. *Vision Res.* **36**, 1773–1781.

Pouget, A., Zhang, K., Deneve, S. & Latham, P. E. 1998 Statistically efficient estimation using population coding. *Neural Computation* **10**, 373–401.

Schrater, P. R. & Simoncelli, E. P. 1998 Local velocity representation: evidence from motion adaptation. *Vision Res.* **38**, 3899–3912.

Somers, D. C., Nelson, S. B. & Sur, M. 1995 An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.* **15**, 5448–5465.

Srinivasan, M. V., Laughlin, S. B. & Dubs, A. 1982 Predictive coding: a fresh view of inhibition in the retina. *Proc. R. Soc. Lond.* B **216**, 427–459.

Wainwright, M. 1999 Visual adaptation as optimal information transmission. *Vision Res.* **39**, 3960–3974.

Ware, C. & Mitchell, D. E. 1974 The spatial frequency selectivity of the tilt aftereffect. *Vision Res.* **14**, 735–737.

Watkins, D. W. & Berkley, M. A. 1974 The orientation selectivity of single neurons in cat striate cortex. *Exp. Brain Res.* **19**, 433–446.

Webster, M. A. & Mollon, J. D. 1991 Change in colour appearance following post-receptoral adaptation. *Nature* **349**, 235–238.

Webster, M. A. & Mollon, J. D. 1992 The influence of contrast adaptation on color appearance. *Vision Res.* **34**, 1993–2020.

Wenderoth, P. & Johnstone, S. 1987 Possible neural substrates for orientation analysis and perception. *Perception* **16**, 693–709.

Wenderoth, P. & Johnstone, S. 1988 The different mechanisms of the direct and indirect tilt illusions. *Vision Res.* **28**, 301–312.

Westheimer, G. 1990 Simultaneous orientation contrast for lines in the human fovea. *Vision Res.* **30**, 1913–1921.

Zaidi, Q. & Shapiro, A. G. 1993 Adaptive orthogonalization of opponent-color signals. *Biol. Cybern.* **69**, 415–428.

Zeki, S. M. 1983 Colour coding in the cerebral cortex. *Neuroscience* **9**, 741–765.