Vision Res. Vol. 32, No. 8, pp. 1565–1582, 1992
Printed in Great Britain.

0042-6989/92 $5.00 + 0.00

A Multi-Channel Approach to Brightness Coding

FRED KINGDOM,* BERNARD MOULDEN†

Received 21 March 1991; in revised form 20 November 1991; in final form 19 December 1991

A model of brightness coding is presented which is shown to predict the appearance of a number of classical brightness phenomena. The model is known as MIDAAS which stands for Multiple Independent Descriptions Averaged Across Scale. In common with many other approaches to brightness perception MIDAAS imputes to local feature detectors a central role in the computation of brightness. It also explicitly recognises the crucial importance to brightness perception of feature detectors operating at different spatial scales. The unique and definitive feature of the model however is the supposition that each scale of spatial filtering operates as if to generate its own description of the pattern of brightness relationships in the image. The final percept is then provided by the composite of those individual brightness descriptions. It is shown that MIDAAS provides a good account of a variety of Mach band phenomena, the conditions under which the Missing Fundamental illusion is observed, the effect of occluding bars on the apparent contrast of step edges, the Chevreul illusion, simultaneous brightness contrast and the non-linear appearance of high contrast sinusoidal gratings. The advantages of MIDAAS over other approaches to brightness perception is discussed, as well as its current limitations.

Brightness coding Multiple channels Mach bands Missing fundamental Simultaneous contrast

INTRODUCTION

The encoding of luminance changes, or contrasts, in the visual world is fundamental to vision. It enables objects to be defined, their relative reflectances estimated and their condition of illumination understood. Important clues to the mechanisms of contrast coding come from studies of the errors made by the visual system when computing the pattern of luminance variation across the image. Particularly dramatic instances of such errors are the class of phenomena known as brightness illusions, of which Mach bands (Mach, 1865), the Cornsweet illusion (Cornsweet, 1970) and Simultaneous Contrast (Brucke, 1865) are perhaps the best known examples. There is a long history of attempts to explain such phenomena, of which the most notable early example is Mach's classic explanation, in terms of lateral inhibition, of the illusory bands which bear his name. Although this explanation is now believed to be deficient on the grounds that it predicts Mach bands to be strongest at a step edge where in fact none are observed (Fiorentini, 1972; Ross, Holt & Johnstone, 1981; Ratliff, 1984), it nevertheless anticipated themes central to our current understanding of vision. More recently, brightness illusions have been studied in the context of linear systems analysis (Campbell, Howell & Robson, 1971; Campbell, Howell & Johnstone, 1978; Sullivan & Georgeson, 1977; Dooley & Greenfield, 1977), have been predicted in a number of computational models of both spatial vision (Watt & Morgan, 1985; Morrone & Burr, 1988) and reflectance recovery (Horn, 1974; Land, 1986; Hurlbert & Poggio, 1988; Blake, 1985), and have been the principal subject of models explicitly concerned with brightness coding (Arend, Buehler & Lockhead, 1971; Arend & Goldstein, 1987; Hamada, 1984; Cohen & Grossberg, 1984; Grossberg & Todorovic, 1988; Moulden & Kingdom, 1990).

This communication presents an approach to understanding brightness illusions in terms of the effects of spatial filtering at multiple spatial scales. The model we present draws upon a number of ideas developed in the MIRAGE model of Watt and Morgan (1985), but which differs from MIRAGE in at least one crucial way. MIRAGE is primarily concerned with how information at differential spatial scales may be combined by the visual system to provide meaningful information about the pattern of luminance discontinuities in the image. It makes the challenging suggestion that the convolution–responses of filters across all spatial scales are first summed non-linearly to produce a single composite convolution–response pattern. The shape of this composite response pattern is then interpreted to indicate the type of image feature (e.g. edge or bar) it is assumed to have originated from, as well as the features' associated attributes of position, blur and contrast. The approach has been particularly successful in predicting data on vernier acuity and edge blur discrimination (summarised

*McGill Vision Research Center, Department of Ophthalmology, Royal Victoria Hospital H4-14, 687 Pine Avenue West, Montreal, Quebec, Canada H3A 1A1.
†Department of Psychology, University of Western Australia, Nedlands, Perth, WA 6009, Australia.

by Watt, 1988). The model described here, known as MIDAAS, employs similar rules to that of MIRAGE to interpret the shape of filter convolution–responses. However MIDAAS differs crucially from MIRAGE in that those rules are applied *separately* to the filter responses at each spatial scale. The appearance of the image, according to MIDAAS, is then given by the linear average of those individual filter response pattern interpretations. It is this feature, that is the separate interpretation of the response pattern from each spatial scale of filtering, which we show to be a useful approach to predicting the appearance of brightness illusions.

We have previously shown that this approach can successfully predict quantitative data on the magnitude of the illusory brightness contrast in a variety of Craik–Cornsweet–O'Brien figures as a function of both the physical contrast and spatial scale of the stimuli (Moulden & Kingdom, 1990). In this communication we use demonstrations of brightness illusions as supportive evidence for our approach. Although MIDAAS is implemented as an exact model capable of quantitative prediction, we believe the demonstrations provided here are sufficient to illustrate the usefulness of its main themes. Nevertheless, this investigation should be regarded as a preliminary exercise anticipating a more detailed set of psychophysical investigations.

The full details of the model are provided in the next section. Following this the predictions of MIDAAS for the appearance of a number of demonstrations will be provided. The final Discussion section includes some comments on the parameters of the model, a comparison of MIDAAS with other approaches to brightness coding, and the model's current limitations.

THE MIDAAS MODEL

Figure 1 summarises the operation of the model, which consists of the following stages.

Stage 1. Implementation of the effects of light adaptation

In MIDAAS the effects of light adaptation are computed for each filter on a local basis by modifying the baseline gain of the filter continuously throughout the stimulus such that the gain of the filter at each point is inversely proportional to the average light level sampled by the filter's receptive field. If $G(\sigma, x)$ is the gain of a filter with space constant σ at x in the stimulus then

$$G(\sigma, x) = 1 \Big/ \left[K + \frac{1}{w(\sigma)} \int S(u) g(x - u)\, du \right]$$

where

$$g(x) = \Pi(x) \begin{cases} 1 & |x| < w(\sigma)/2 \\ 0 & |x| > w(\sigma)/2 \end{cases}.$$

K is a constant which sets an upper limit of $1/K$ on the gain as the mean sampled luminance drops to zero (e.g. see Shapley & Enroth-Cugell, 1984) with K being set to 10% of the mean luminance level of the stimuli employed here. $S(u)$ is the luminance of the stimulus at point u. $w(\sigma)$ is a measure of the overall width of the filter's receptive field, set equal to 8σ.

Stage 2. Spatial filtering at different spatial scales

We employ filters whose spatial weighting function $F(\sigma, u)$ is given by the second derivative of a Gaussian (2DG) multiplied by $1/\sigma$, where σ is the space constant of the filter. Thus

$$F(\sigma, u) = -1/\sigma \exp(-u^2/2\sigma^2) \cdot (u^2/\sigma^2 - 1).$$

The value of $1/\sigma$ can be thought of as the baseline gain setting of the filter and results in each filter producing the same peak amplitude–response to a sine-wave of its preferred spatial frequency. The space constants of the four filters are set at octave intervals. The filter with the largest space constant has a space constant one-sixteenth the width of a single period of the waveforms shown in Figs 3–9. If those stimuli were viewed such that a single cycle subtended 4 deg of visual angle (by holding the page about 18 in. away) then the space constants of the four filters would be approx. 16, 8, 4 and 2 min of arc, corresponding to receptive field centre sizes of approx. 36, 18, 9 and 4.5 min of arc respectively. Each

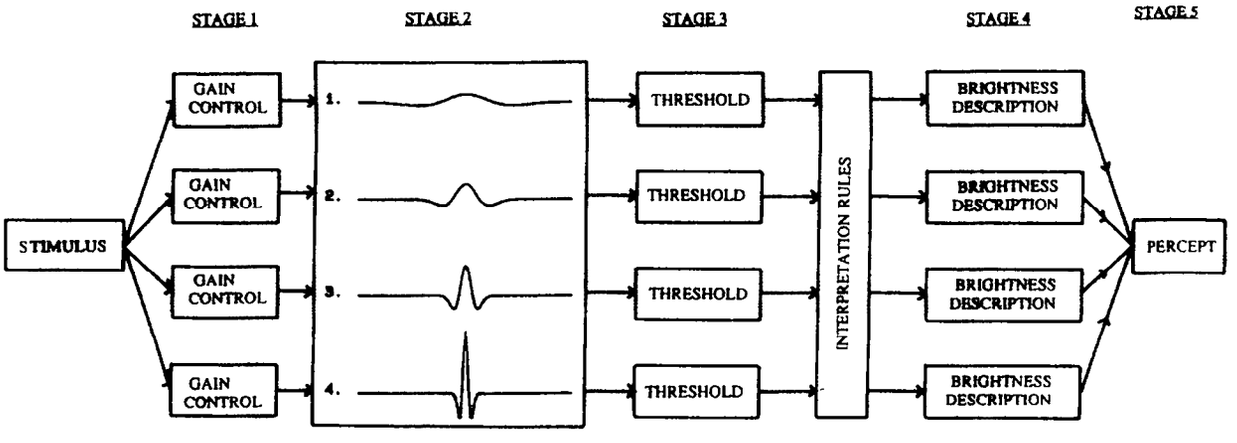

FIGURE 1. The MIDAAS model.

filter is separately convolved with the stimulus to produce a response $R(\sigma, x)$ thus:

$$R(\sigma, x) = G(\sigma, x) \int S(u)F(\sigma, x - u)\, du.$$

Stage 3. Thresholding and power transformation of the response of each filter

The thresholding stage is an essential requirement of the model, since it allows local regions of activity in each filter's response to be effectively isolated. Where the (absolute) response level of the filter falls below the threshold level t, the response is set to zero:

$$\text{If } R^+(\sigma, x) < t \; R'(\sigma, x) = 0 \text{ else } R'(\sigma, x) = R^+(\sigma, x)$$

$$\text{If } R^-(\sigma, x) < -t \; R'(\sigma, x) = 0 \text{ else } R'(\sigma, x) = R^-(\sigma, x).$$

The value of t is set to 5% of the response of the filters to a maximum contrast sinusoidal grating of preferred spatial frequency. The (absolute) filter output is also subjected to a power transformation. If $R(x)$ is the response of a given filter the power transformed response $R'(x)$ is given by

$$\text{sgn}[R(\sigma, x)] * |R(\sigma, x)|^n$$

with n being set to 0.7.

Stage 4. Generation of symbolic description of brightness changes at each spatial scale

Figure 2 illustrates the operation of rules to interpret the pattern of activity of each filter's thresholded response. The middle column of the figure shows three principal classes of response that are isolated, (a) biphasic, (b) triphasic and (c) monophasic. The class of response reflects the arrangement of adjacent lobes in the convolution output, where a lobe is defined as a region of non-zero activity bordered by zero crossings. The three classes of filter response are interpreted as indicating respectively the presence of (a) an edge, (b) a bar and (c) a bar, as in the MIRAGE model of Watt and Morgan (1985). In order to ascertain which of the three classes of response a given region of activity is a member, it is first necessary to have a rule which states what the minimum permitted distance between any two adjacent lobes must be before those lobes are interpreted as originating from different features in the image. The rule adopted here is that the distance between the near edges of two adjacent lobes must be greater than the width of the receptive field centre of the filter producing them. This width is 2σ.

The resulting symbolic brightness description of each feature, $B(\sigma, x)$ is shown in the right hand column of Fig. 2. For the "bar" responses the brightness description is shown as being isomorphic with the central lobe of the triphasic response. This is simply achieved by the appropriate half-wave rectification of the convolution–response (in the case of a monophasic response no half-wave rectification is necessary). In the case in which the central lobe of a triphasic response is "cusped", as would occur in the response of a filter to a relatively wide bar, the brightness profile is interpreted as a "filled-in" version of the lobe, through a linear interpolation between the two peaks of the cusped response. This is illustrated in Fig. 2(d).

The "edge" response is interpreted as a step function in brightness which is integrated along with any other computed brightness changes throughout the stimulus. In the implementation described here, the brightness profile of the actual step is set to be isomorphic with the filter response between the peak and trough of the biphasic response, but weighted such that the amplitude of the step is given by the amplitude of the larger of the two lobes in the biphasic response. This implementation rule thus incorporates both the "blur" and amplitude of the edge response into the computed brightness profile for a given scale of filtering.

Figure 2(e, f) illustrate the implementation of Stage 4 of the model for two stimuli in which it is not immediately obvious how such interpretation rules might operate. In Fig. 2(e) the stimulus comprises a bar superimposed upon an edge. For the particular filter employed, the resulting triphasic response is interpreted as *just* a bar, with no step function in brightness across it. At this scale of filtering therefore, the step function in luminance across the bar is simply not "seen". Were the filter employed to be one with a much larger space constant a biphasic response would have occurred, and the interpretation would be *just* a step function in brightness. This may appear at first sight to be simply a crude method of bypassing the problem of implementing a fully fledged deconvolution process on the filter response in order to correctly recover the luminance profile of the stimulus. However, it turns out to provide a remarkably simple account of some important brightness phenomena as demonstrated below. Finally, in Fig. 2(f) the stimulus is a Gabor patch and produces a filter response consisting of a series of abutting lobes alternating in polarity. This is interpreted by MIDAAS as constituting a series of interlaced triphasic responses. The feature and brightness interpretation of this pattern is thus a series of bars alternating in polarity.

Stage 5. Averaging brightness descriptions across spatial scales

The final stage of the model results in the actual predicted appearance of the stimulus, $A(x)$. It is obtained by taking the linear average of the brightness descriptions produced by each spatial scale of filtering. Thus:

$$A(x) = \frac{1}{n} \sum_i B_i(\sigma, x)$$

where n = the number of filters, equal to 4 in this implementation.

DEMONSTRATIONS

The acronym MIDAAS attempts to capture the definitive feature of the model, made explicit in Stages 4 and 5 of Fig. 1. As Stage 4 shows, the convolution

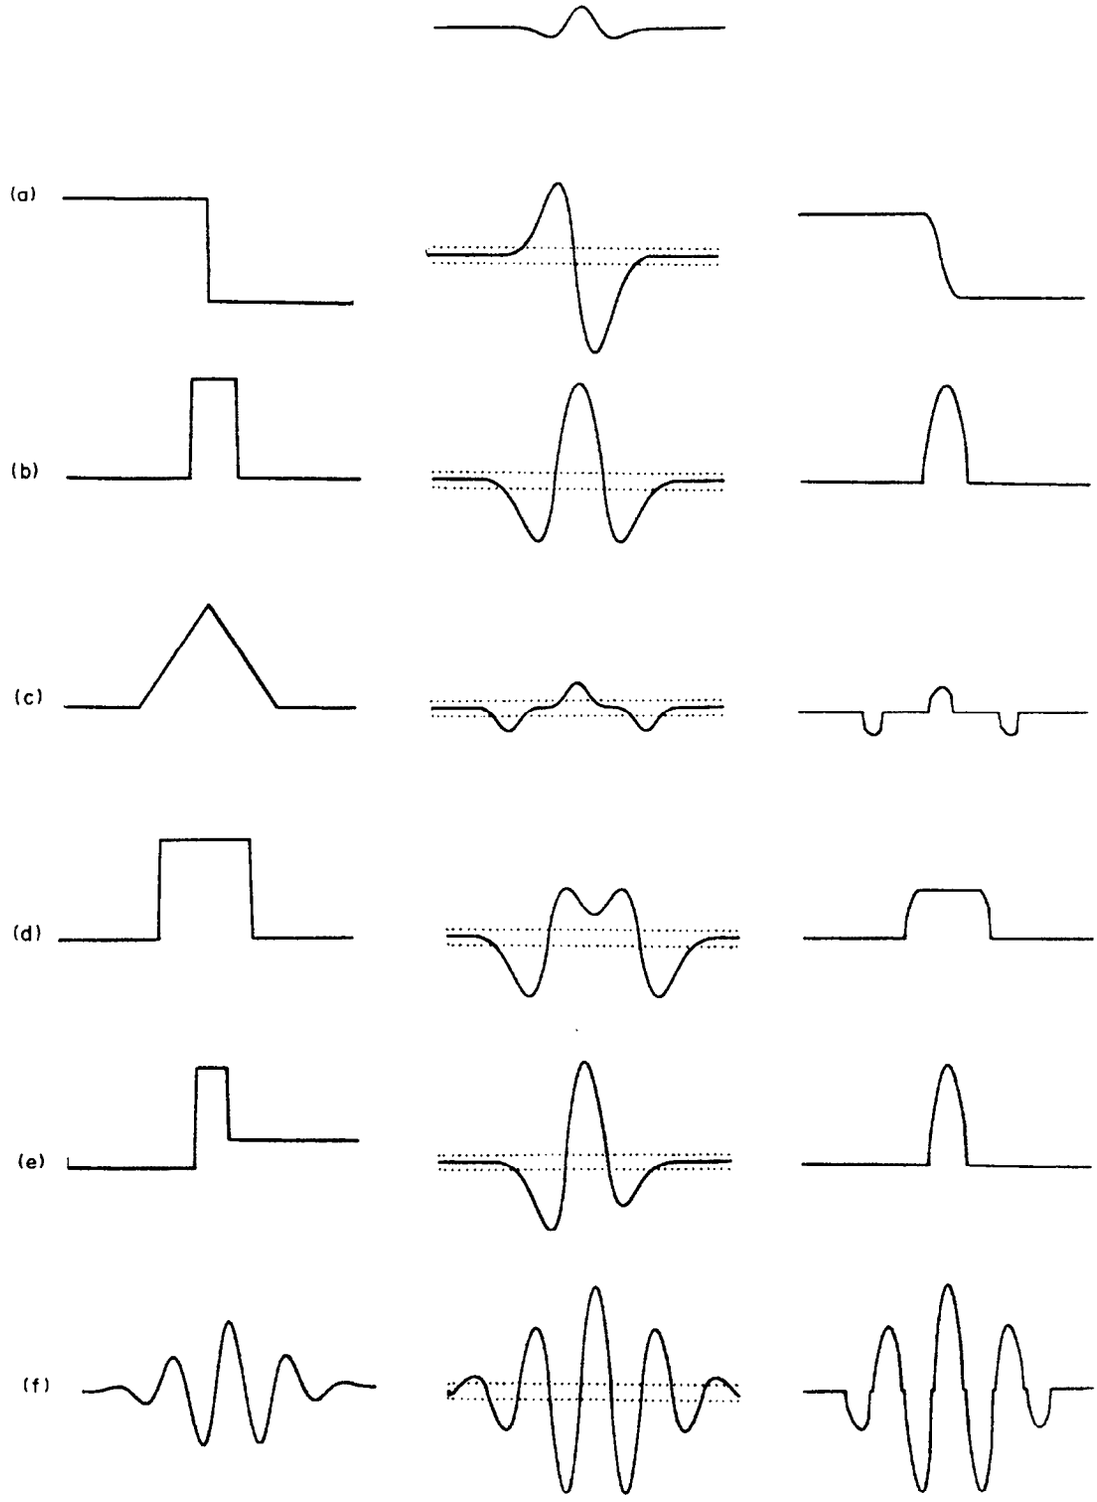

FIGURE 2. Responses (middle column) and brightness interpretations (right column) of a 2DG (second derivative of a Gaussian) filter (top) for various stimuli (left column).

response of each of a Multiple set of filters is Independently interpreted to provide an individual Description of the pattern of brightness relationships in the stimulus. The observed appearance of the stimulus is a result of Averaging those individual descriptions Across the different Scales of spatial filtering.

We now present a number of demonstrations of stimulus patterns and the MIDAAS predictions of their appearance. The figures were all generated using the VSG2 digital signal generator (Cambridge Research Systems) interfaced to a DEL 386 computer and displayed on a BARCO CD1551 video monitor. The VSG2 was programmed to produce a linear 12 bit (4096 grey levels) look-up-table by suitable selection of intensities from a 14 bit DAC (digital-to-analogue converter) calibrated using a UDT photometer. Unless otherwise specified stimulus contrast was set to 75%. The model is implemented in a fully automated algorithm which

operates on each stimulus pattern without prior knowledge.

Triangular wave (Fig. 3)

MIDAAS correctly predicts the attenuation of the ramps and the accentuation of the peaks and troughs (sometimes referred to as Mach bands) in the waveform. The operation of MIDAAS for this stimulus is shown below. In this and the remaining demonstrations the convolution outputs labelled 1–4 refer respectively to those produced by the four filters with space constants 16, 8, 4 and 2 min of arc. The brightness interpretation of each scale of filtering is shown alongside its convolution–response, with the predicted appearance of the stimulus pattern illustrated alongside its luminance profile. It must be emphasised that the brightness descriptions in Fig. 3 are symbolic in that they do not imply the existence of neural representations that are isomorphic with those descriptions.

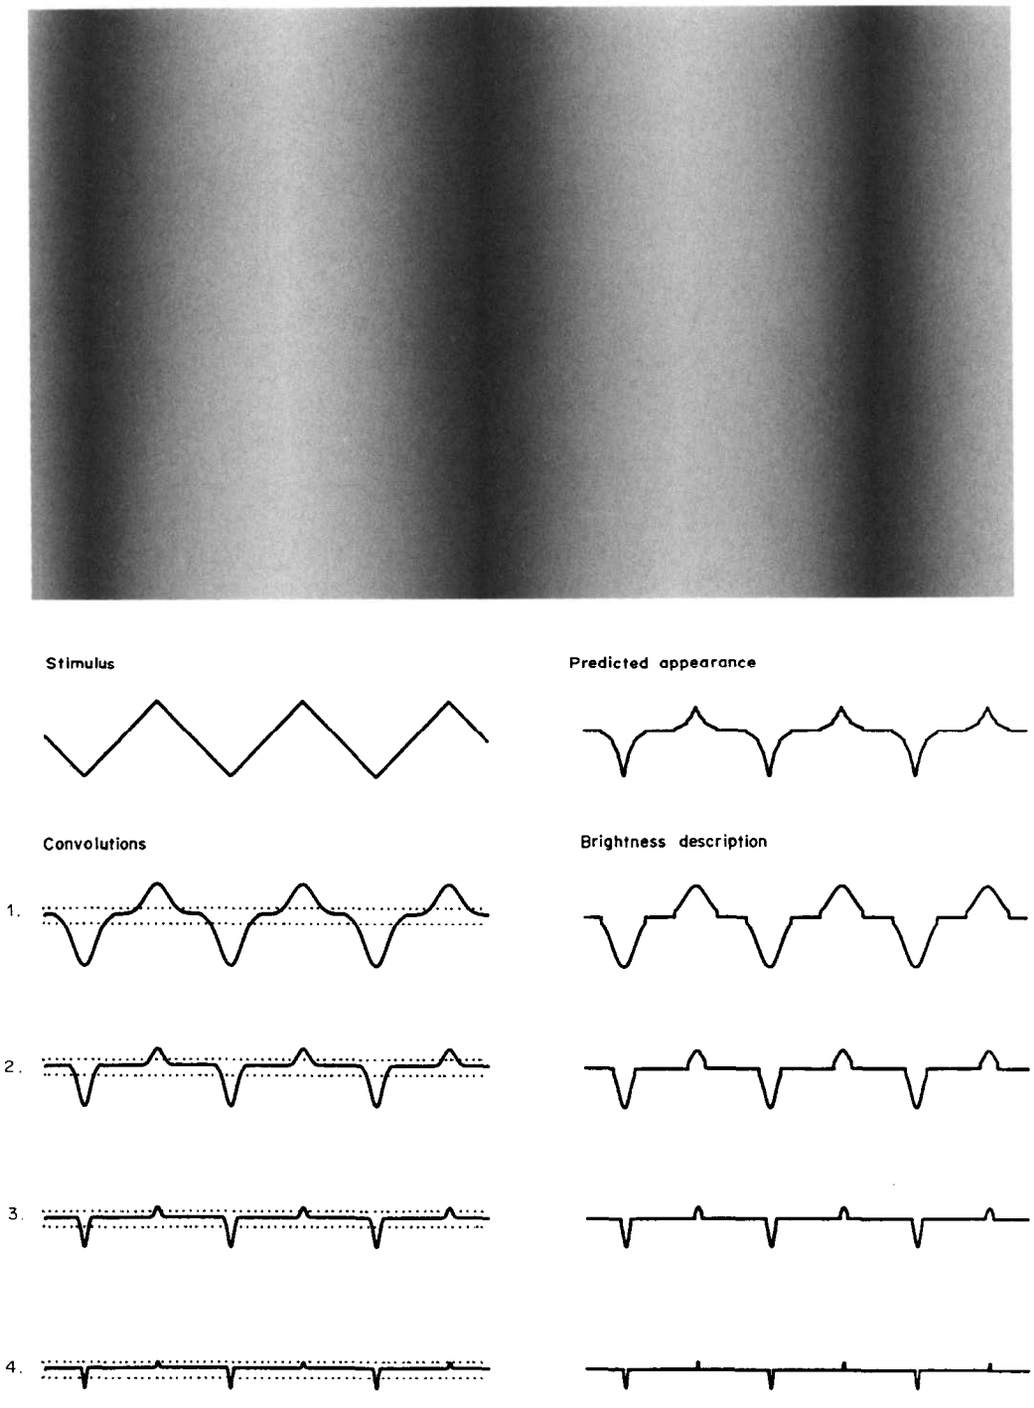

FIGURE 3. Triangular wave and MIDAAS prediction. The luminance profile of the stimulus is illustrated in top left of the figure. 1–4 represent the outputs of the four filters whose space constants decrease respectively in octave intervals. The symbolic interpretations of the pattern of brightness changes are shown alongside each filter output, with the final predicted appearance shown in the top right of the figure alongside the stimulus luminance profile.

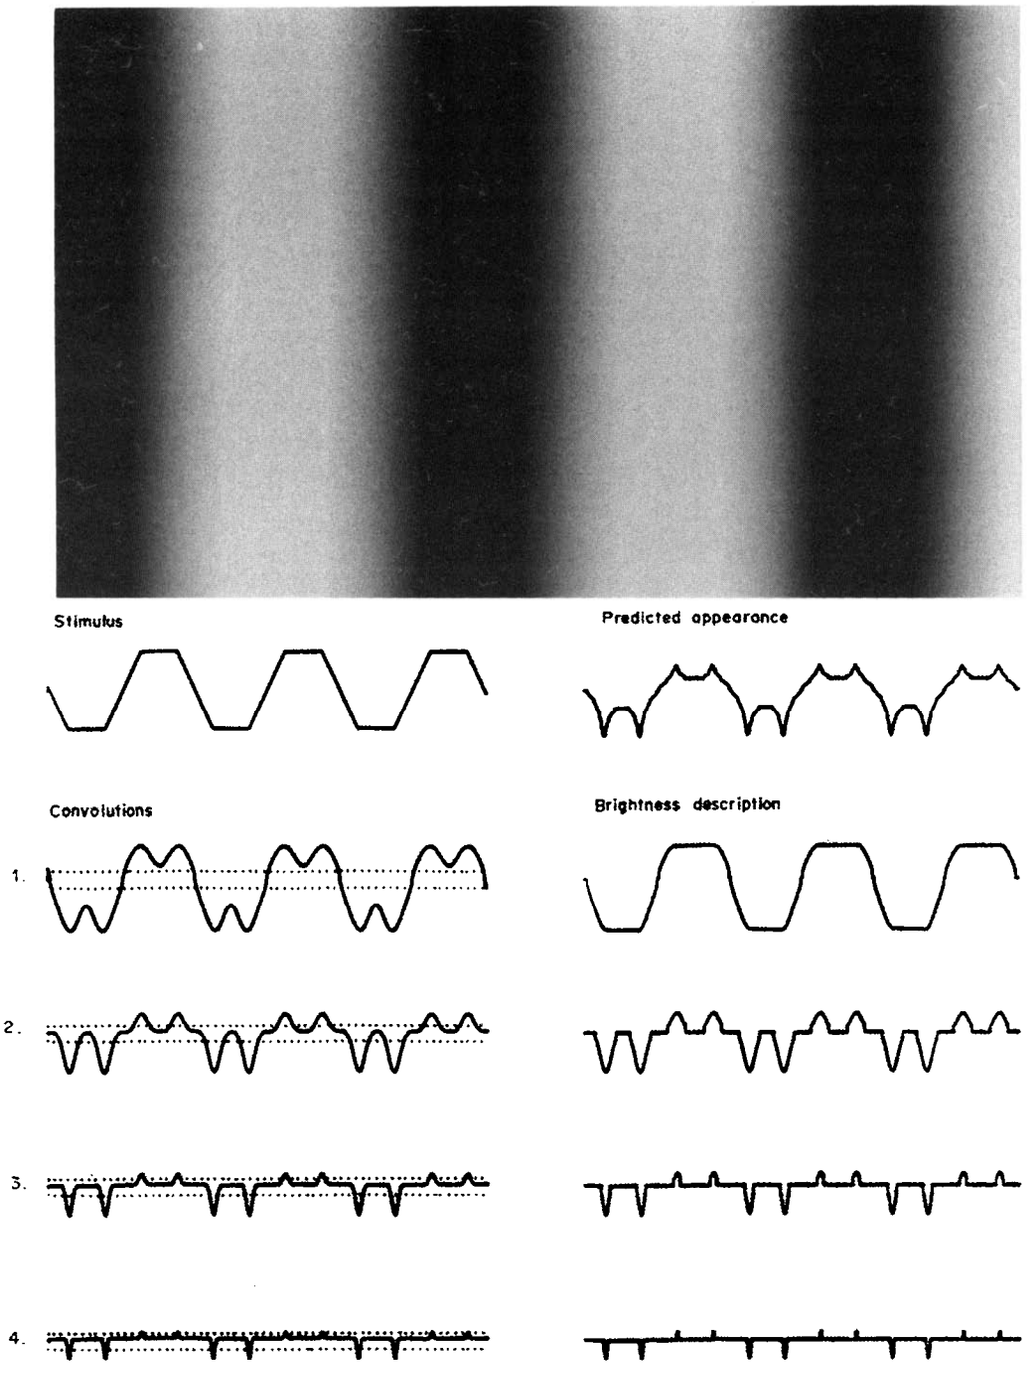

FIGURE 4. Trapezoidal wave showing Mach bands.

Trapezoidal wave (Fig. 4)

MIDAAS correctly predicts the presence of Mach bands at the feet and knees of the ramps in the trapezoidal wave. These result from the presence of monophasic responses in the convolution profiles of filters 2–4, which are interpreted as indicating the presence of bars. The filter with the largest space constant (1) is the only filter which "sees" the overall luminance modulation of the trapezoid and which renders that modulation visible in the final percept. As the spatial frequency of the trapezoid increases, MIDAAS correctly predicts the measured attenuation of the Mach bands (Ross, Morrone & Burr, 1989), because as one increases spatial frequency the monophasic filter responses which alone produce the Mach bands give way one by one (in scale order from coarse to fine) to triphasic responses which are interpreted as modulations at the same spatial scale as that of the trapezoid. Finally, MIDAAS correctly predicts that no Mach bands will appear in a square-wave (Fiorentini, 1972; Ross *et al.*, 1981; Ratliff, 1984). In the square-wave, biphasic responses (signalling edges) at the edges of the waveform will be produced by filters small in comparison to the scale of the square-wave, while triphasic responses (signalling bars in cosine phase with respect to the waveform), will be produced by filters large in comparison to the scale of the waveform.

Effect of occluding bars on the appearance of a low-contrast saw-tooth (Fig. 5)

Figure 5(a) and (b) show respectively a low contrast saw-tooth wave and the same stimulus with the addition of narrow dark contours lying along the sharp edges in the stimulus. Notice how the staircase like appearance of the saw-tooth is completely eliminated by the contours. This effect was first demonstrated by Tolhurst (1972) using a staircase function. MIDAAS predicts the effect because the triphasic responses produced by the contours in all the filters are interpreted only as bars without any step functions in brightness. As the contrast of the saw-tooth is increased relative to that of the added contours, more and more filters (starting with the largest) will produce biphasic rather than triphasic responses, adding "edge signals" to the final interpretation. Thus the appearance in these instances would be of an edge plus bar, with the apparent contrast of the edge however, still reduced compared with the situation where no bar were present. This accords with our observations. Other possible interpretations of this phenomena are discussed in the Discussion section.

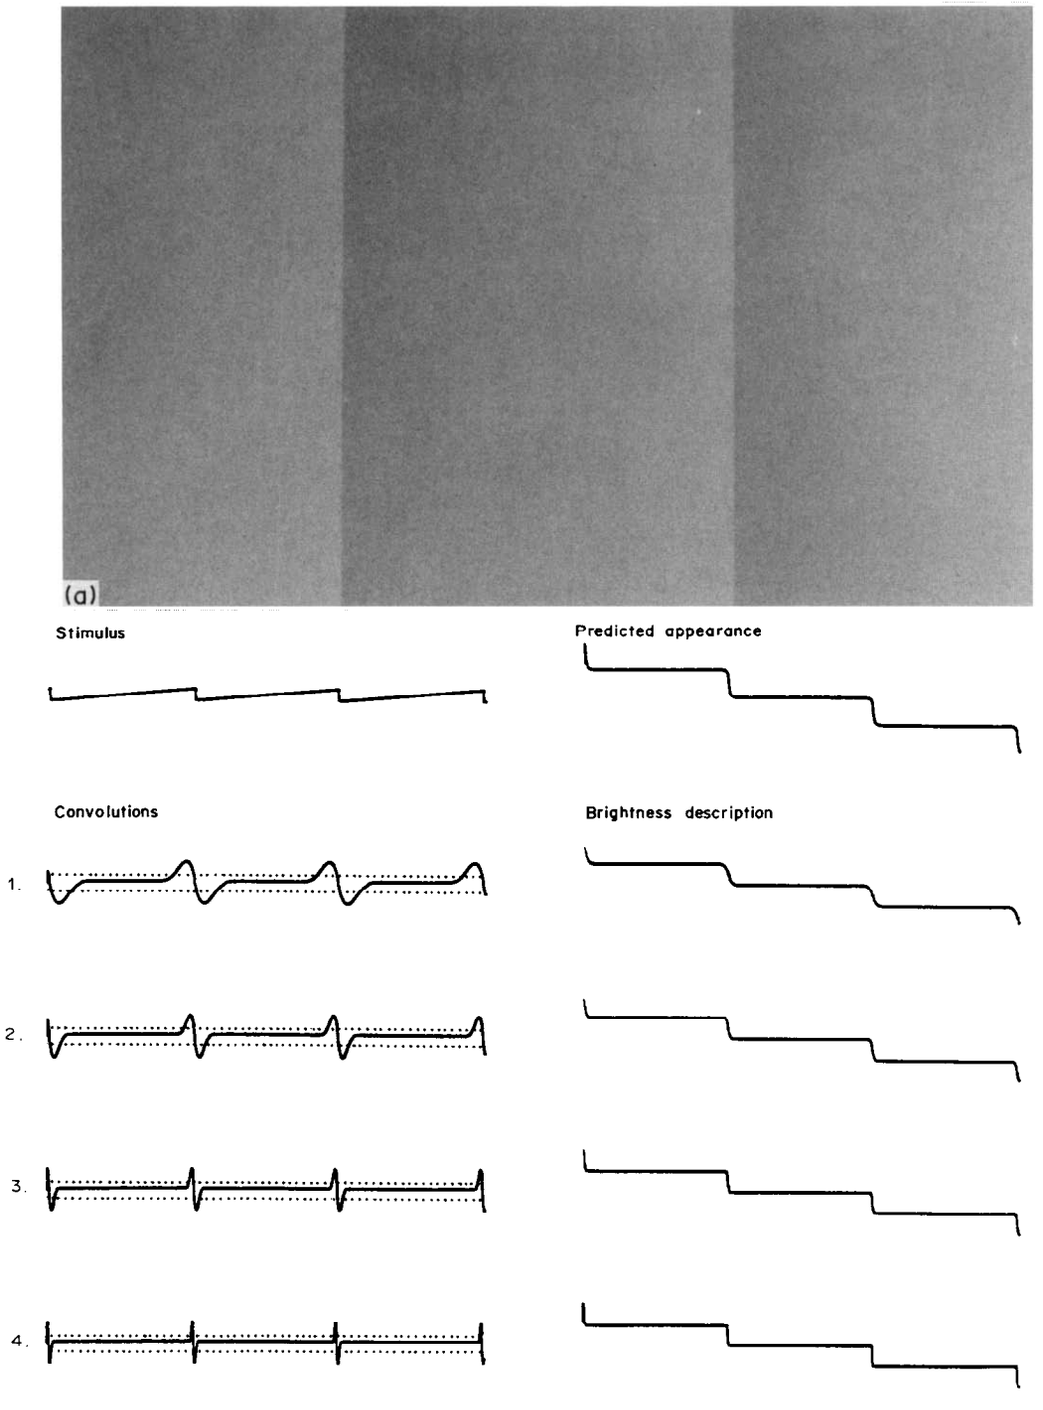

FIGURE 5(a). *Caption overleaf.*

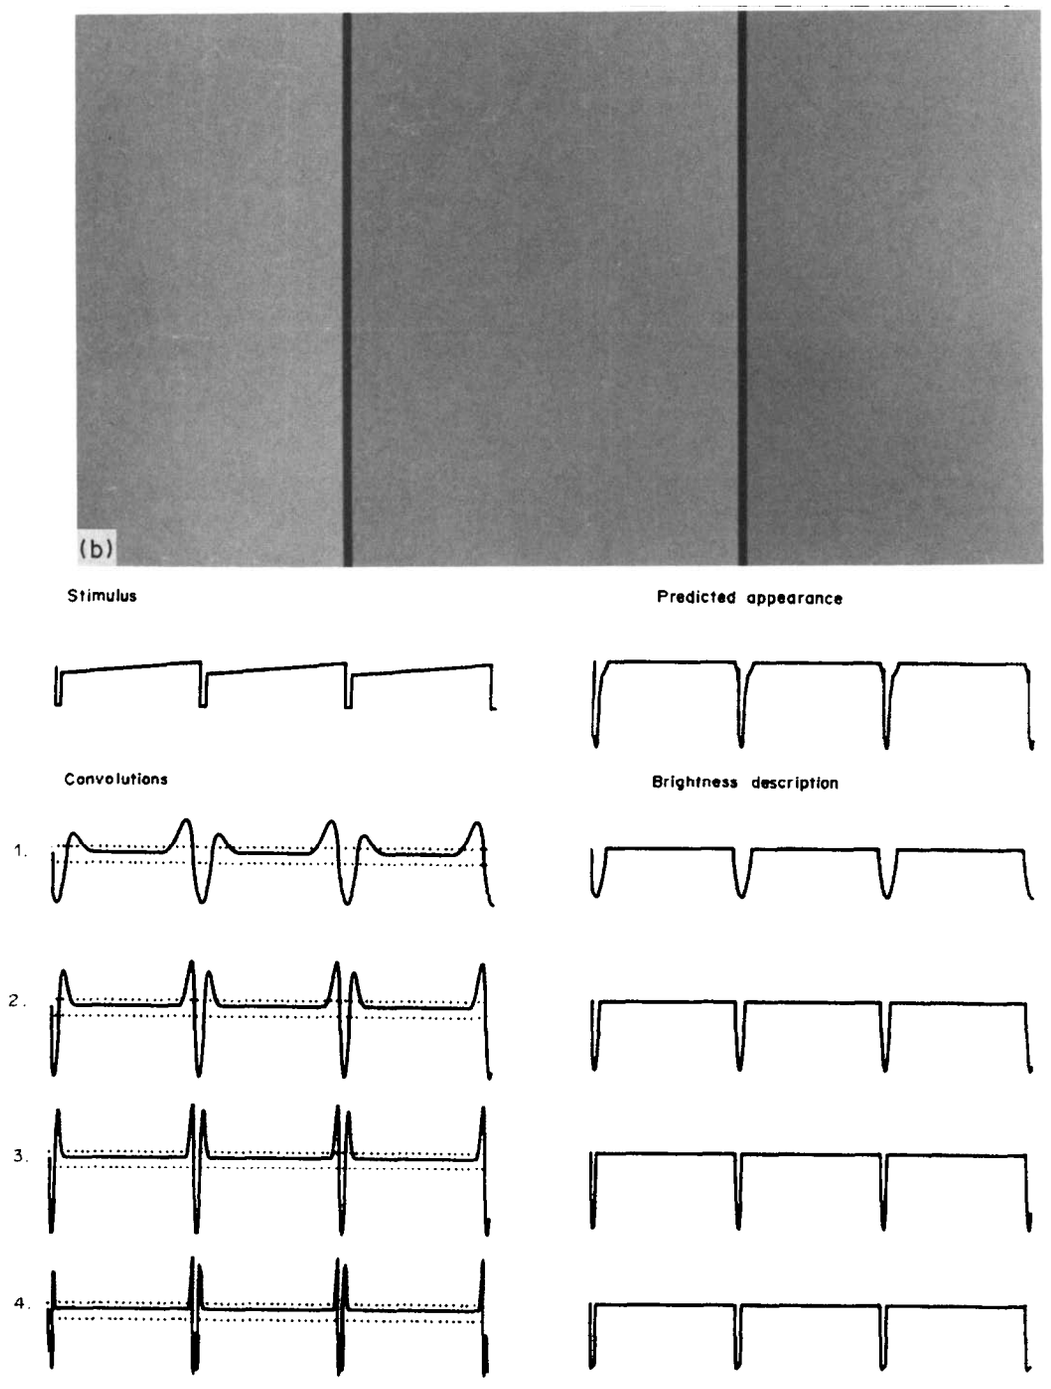

FIGURE 5(b)

FIGURE 5. (a) Low frequency, low contrast sawtooth, (b) with bars occluding the edges. The bars eliminate the staircase-like appearance, an effect predicted by MIDAAS.

Low and high contrast Missing Fundamental (MF) figures (Fig. 6)

The MF (missing fundamental) wave, first demonstrated by Campbell *et al.* (1971), consists of a square-wave less its fundamental harmonic component. At low contrasts [Fig. 6(a)] it is indistinguishable from a square-wave while its scalloped appearance is very visible at high contrasts [Fig. 6(b)]. MIDAAS predicts the square-wave appearance of the MF stimulus at low contrasts because only biphasic ("edge") responses are produced at the sharp discontinuities in the stimulus once a threshold has been imposed. At high contrasts triphasic responses are produced which are interpreted as bars, resulting in the more veridical appearance of the stimulus. The contrast at which the transition from a square-wave to a "cusped" appearance occurs becomes less as spatial frequency increases (Campbell *et al.*, 1978). This is also predicted by MIDAAS since biphasic responses will in general give way to triphasic responses as spatial frequency is increased. More detailed quantitative predictions of similar Craik–Cornsweet–O'Brien figures can be found in Moulden and Kingdom (1990).

Repeating staircase (Fig. 7)

Figure 7 shows a staircase in luminance which repeats itself. The appearance is that of a repeating saw-tooth profile, with the individual steps decreasing in apparent contrast as their luminance increases. The saw-tooth appearance of relatively high spatial frequency step functions is sometimes referred to as the Chevreul illusion (Cornsweet, 1970). The saw-tooth profile is predicted by MIDAAS because the symbolic interpretation of the filter II response is that of a series of bars alternating in polarity. Only the interpretations from the two smallest filters (3 and 4) interpret the figure more or less veridically. There is a small mismatch between the predicted and actual appearance in that the last step of each cycle is predicted to appear as a ramp, whereas it actually appears as a step. This occurs because the algorithm extrapolates between the peaks of a "cusped" lobe to give the brightness profile when the lobe is indicative of a bar [see Fig. 2(d)]. In the case of the last step in each cycle of the staircase in Fig. 7, the two peaks in the cusp are at different heights resulting in the incorrect prediction of a ramp appearance for that step. A similar mismatch occurs in Fig. 8 (see below), where some of the steps are incorrectly encoded as shallow ramps. This problem would be overcome if a symmetrical brightness profile was always computed from a cusped convolution response.

Simultaneous brightness contrast (Fig. 8)

The mid-grey bars in Fig. 8 have identical luminance yet appear different in brightness depending on their surrounding luminance. MIDAAS predicts simultaneous contrast effects principally because it gives more weighting to bars than equal amplitude edges in brightness computation. As with Fig. 7, there is a small mismatch between predicted and actual appearance in

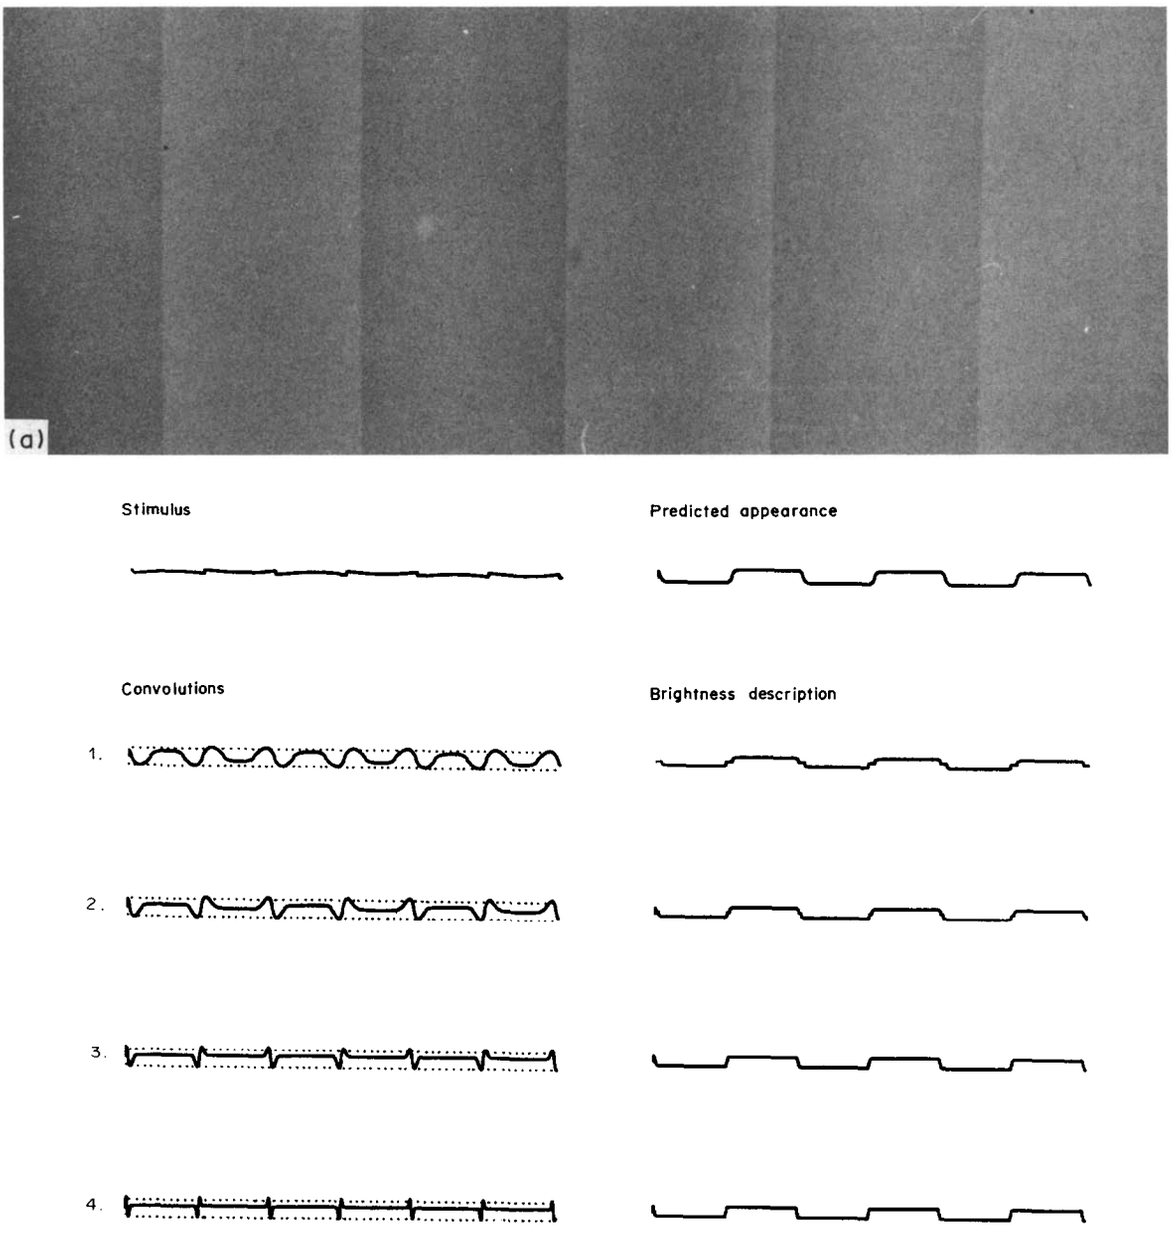

FIGURE 6(a). *Caption overleaf.*

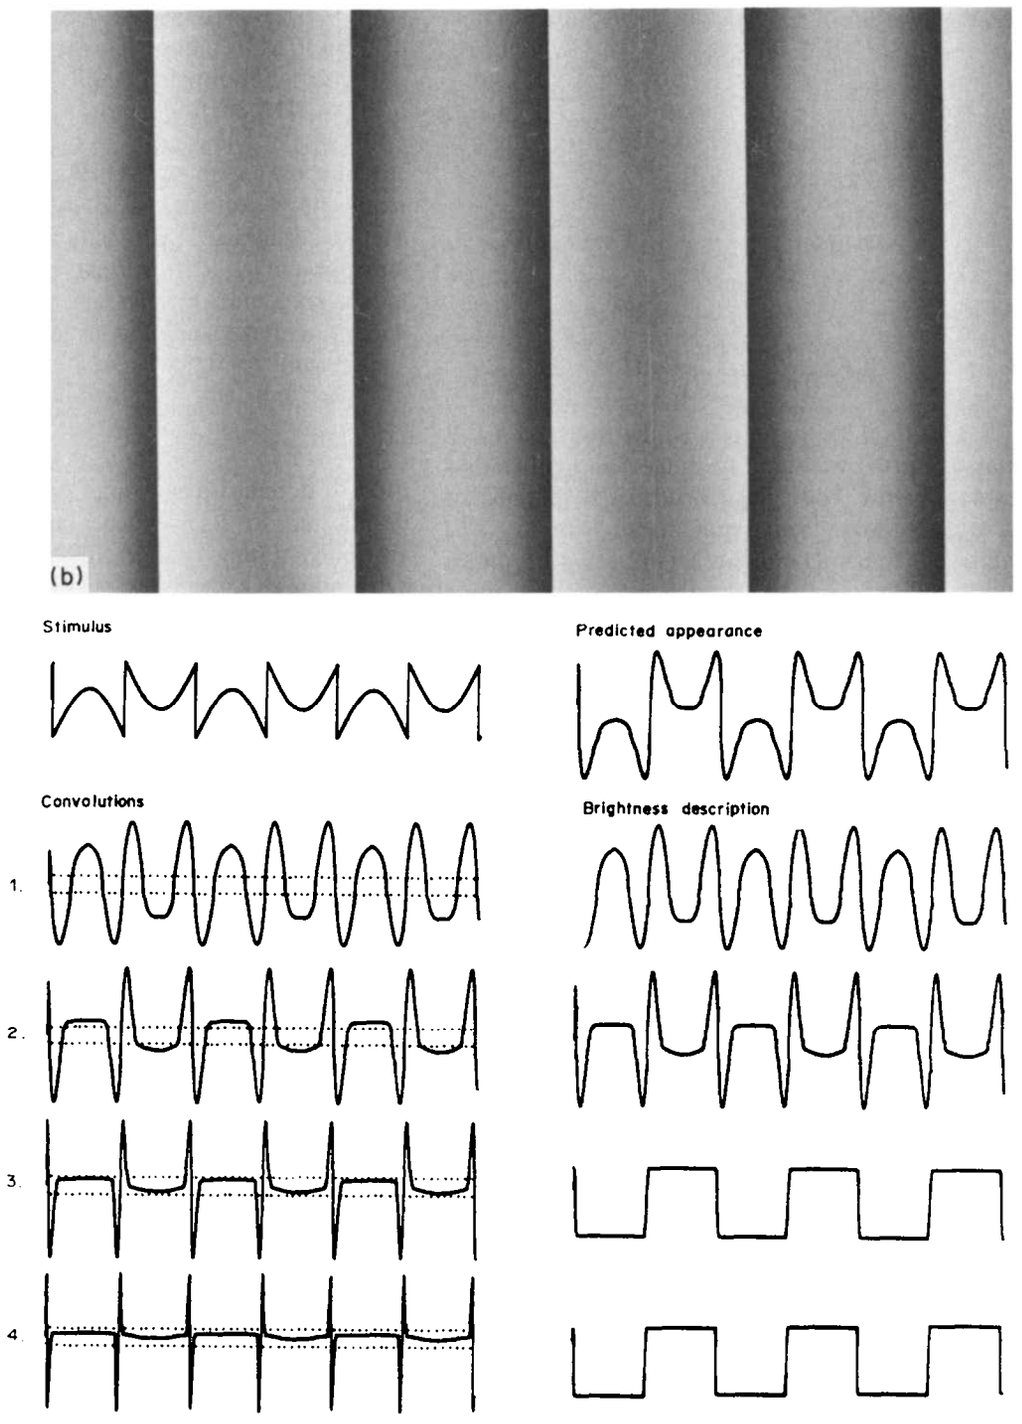

FIGURE 6(b)

FIGURE 6. Missing Fundamental wave at (a) low (about 5%) and (b) high (about 75%) contrasts. The change from a square-wave to a more veridical appearance is predicted by MIDAAS because biphasic, edge, responses in the filter outputs give way to triphasic, bar responses as contrast is increased.

the predicted ramp-like appearance of the individual bars within the stimulus.

This completes the demonstrations which illustrate the operation of the MIDAAS model. We now discuss some of the parameters of the model, compare MIDAAS with a number of other approaches to brightness coding and finally point out some of its current limitations.

DISCUSSION

Some comments on the parameters of MIDAAS

Stage 1. Both neurophysiological and psychophysical evidence supports the idea that retinal light adaptation is highly localised, very likely to within the summation areas of retinal ganglion cell receptive fields (Shapley & Enroth-Cugell, 1984; Burr, Ross & Morrone, 1985;

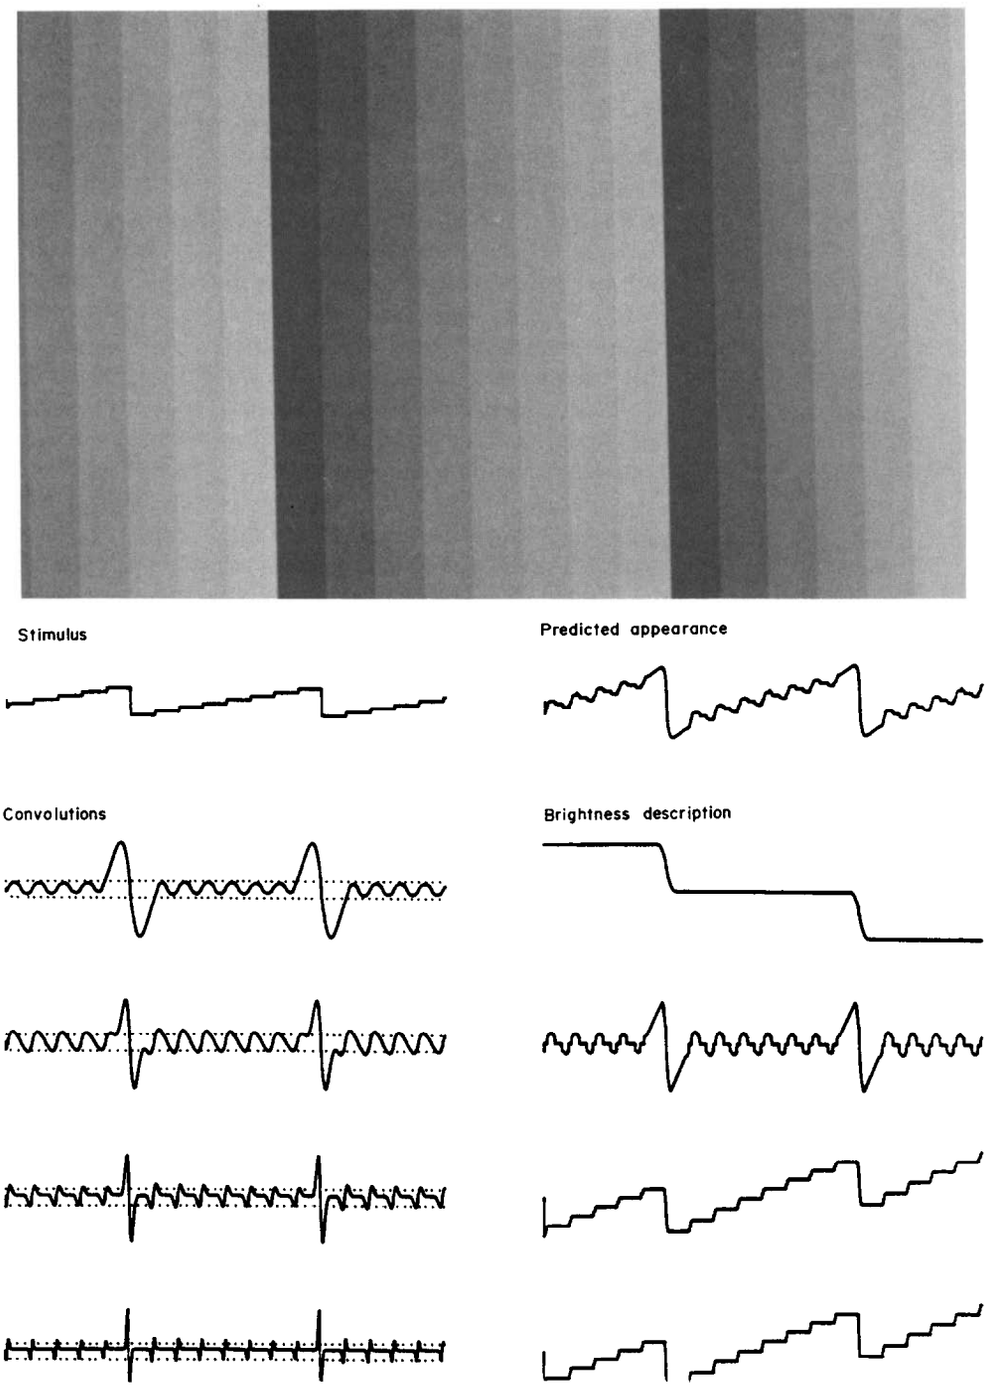

FIGURE 7. This pattern shows a high spatial frequency staircase superimposed on a lower spatial frequency sawtooth. The stimulus appears as a high spatial frequency sawtooth (the Chevreul illusion) riding on a lower spatial frequency sawtooth. MIDAAS predicts the sawtooth like appearance of the staircase because filter 2 produces triphasic, bar responses. Filters 1, 3 and 4 however all contribute to the presence of the lower frequency sawtooth component in the final appearance.

Cleland & Freeman, 1987). This idea has been incorporated into the MIDAAS model and accounts for its ability to predict the non-linear appearance of waveforms with high contrast. This non-linear appearance is most readily seen in a high contrast sine-wave (not reproduced here because of the confounding effects of non-linearities of photographic reproduction). At high contrasts, sine-waves appear markedly non-sinusoidal with broad "white" bars and narrow "dark" bars. The prediction of MIDAAS for a sine-wave is shown in Fig. 9. Because the gain of the filters is modified continuously throughout the stimulus according to the sampled mean luminance at each point, the effect is to attenuate local differences in luminance at the peaks of the waveform and accentuate differences at the trough. There are of course available compressive non-linear functions which can be applied to the stimulus prior to the stage of spatial filtering, for example the Naka-Rushton equation (Naka & Rushton, 1966; and see discussion in Shapley & Enroth-Cugell, 1984) or a

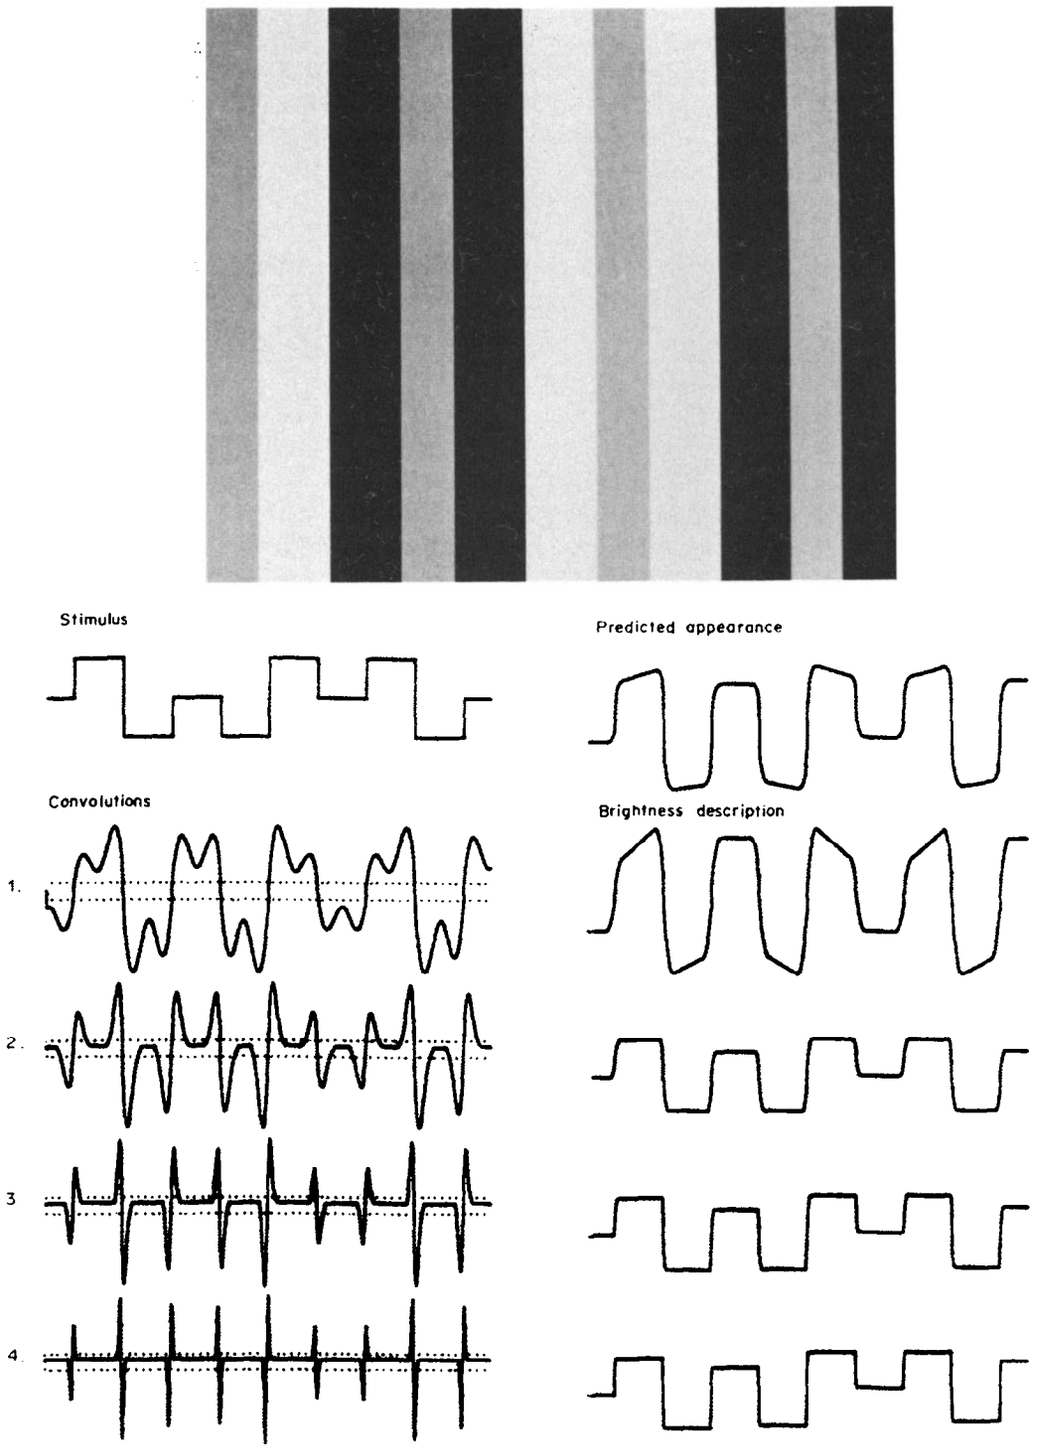

FIGURE 8. Simultaneous contrast and the MIDAAS prediction. They grey bars appear lighter when set against the black than when set against the white phases of the stimulus. MIDAAS predicts this effect principally because it gives more weighting to bar than to equal amplitude edge responses.

logarithmic transform (Kingdom & Moulden, 1991) which may also produce the predicted non-linear appearance of high contrast stimuli.

Stage 2. A 2DG (second derivative of a Gaussian) filter is an approximation to that of the receptive field profile of a retinal ganglion cell. Retinal ganglion cells are widely believed to carry out the initial stage of spatial filtering in vision and to be centrally involved in the mechanisms of retinal adaptation (Shapley & Enroth-Cugell, 1984). The information they provide to higher centres is thus crucial for brightness computation. Nevertheless the approach described in MIDAAS is not critically dependent on the employment of 2DG filters. For example, with different implementation rules (Stage 4), MIDAAS could be implemented with odd-or-even symmetric Gabor filters, which are commonly used for modelling the performance of cortical simple cells (Marcelja, 1980).

One could argue that the decision as to which type of filter to employ should be guided by what is known

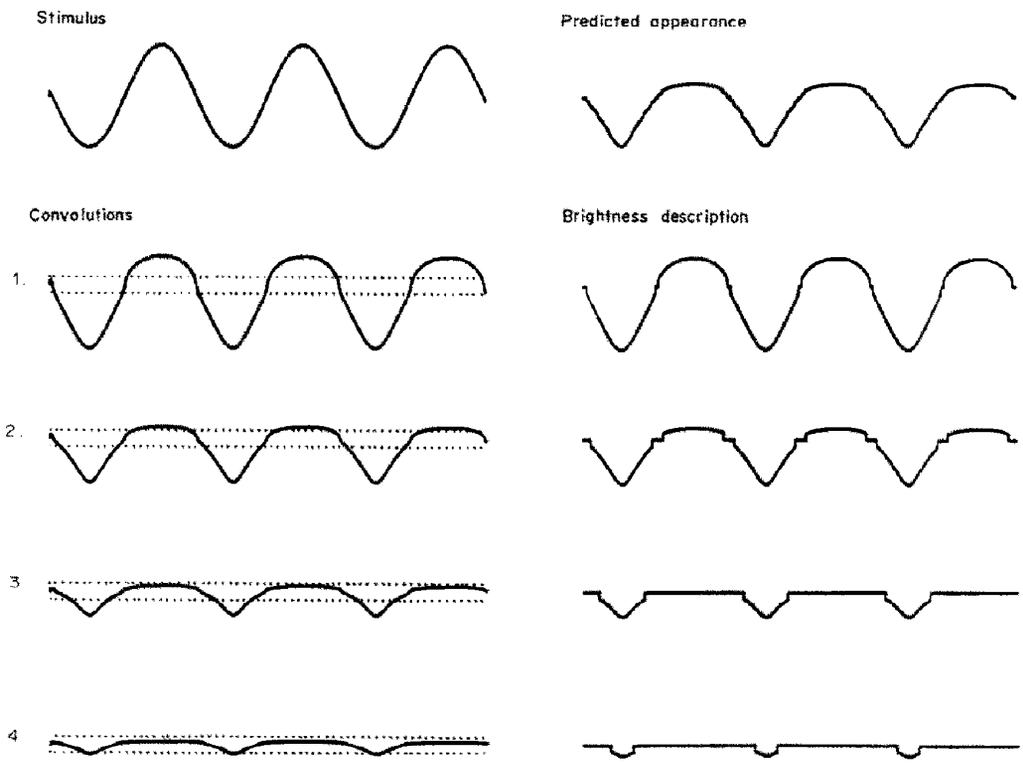

FIGURE 9. The prediction of MIDAAS for a high contrast (90%) sine-wave.

about the site of brightness computation in the visual pathway. It is unquestionably the case that the cortex is the site at which decisions are made about whether changes in brightness represent changes in reflectance as opposed to changes in illumination (Gilchrist, 1979; Land, Hubel, Livingstone, Perry & Burns, 1983). Most, but not all, evidence also points to the cortex as the site of brightness computation *per se* (see review by Kingdom & Moulden, 1989). However, as we have argued before, the initial retinal stage of spinal filtering must necessarily have consequences for brightness perception, even though those consequences may not be made explcit until later cortical stages (Moulden & Kingdom, 1989). With this in mind, and the fact that the decision as the type of filter to employ may in the end be somewhat arbitrary, we have adopted the 2DG filter because it represents the simplest of those spatial filters which have a proven physiological correlate.

What precisely is the importance of spatial filtering for MIDAAS if the exact type of filter is not critical? Firstly, band-pass spatial filtering, together with a procedure which is essentially the equivalent to a non-linear (i.e. logarithmic) transform of stimulus intensity, is arguably the means by which the visual system achieves lightness constancy (Land & McCann, 1971; Shapley & Enroth-Cugell, 1984). Secondly, and critically for the model described here, spatial filtering is the means of locating critical features in the image (such as edges and bars) at various spatial scales.

Finally, although we have only employed "on-centre" filters in the implementation used here, this does not mean that MIDAAS is inconsistent with the universally held view that the visual system contains both on- and off-centre units at various stages in the visual pathway. Indeed, in our study of brightness contrast in Craik–Cornsweet-O'Brien figures we implemented a version of MIDAAS in which the positive and negative components of the convolution image were carried separately by on- and off-centre 2DG filters whose outputs were half-wave rectified (Moulden & Kingdom, 1990). Moreover, we showed that the greater magnitude in induced brightness in negative compared with positive Cornsweet "cusp" stimuli, an effect first measured by Hamada (1985), could be accounted for if one supposed that the off-centre units had space constants on average 1.5 times those of the corresponding on-centre units. For the purpose of the demonstrations shown here we have not included this refinement for the sake of simplicity. Further research providing quantitative measurements of brightness contrast in a variety of stimuli will be needed to test whether such a putative difference between on- and off-centre filter sizes has any general significance.

Stage 3. In an earlier implementation of the model Moulden and Kingdom (1990) added noise to each filter output and located regions of activity whose mass or integral exceeded a given amount, as in the MIRAGE model of Watt and Morgan (1985). However, for the purposes of modelling the demonstrations here, a simpler "clipping" procedure is used to isolate local regions of response activity, in which any filter response less in absolute magnitude to that of the threshold value is set to zero.

Alternative approaches to MIDAAS

We turn now to a comparison of MIDAAS with other current approaches to brightness coding. We have pre-

viously categorized brightness models into essentially three classes: Contrast Sensitivity Function (CSF) models, Lightness Integration models and Edge Detector models (Kingdom & Moulden, 1989). MIDAAS falls into the third category, which we shall refer to below by the more general term "Feature Detector" models. We will discuss the other candidate models within this framework.

The CSF (Contrast Sensitivity Function) approach

It is axiomatic that low-frequency attenuation in contrast sensitivity is implicated in many brightness phenomena. For example, it easily accounts for the fact that low spatial frequency MF (Missing Fundamental) waveforms are indistinguishable from square-waves at low contrasts [Fig. 6(a)]. The fundamental harmonic in the square-wave, the presence of which alone physically distinguishes it from the MF waveform, is sufficiently attenuated by the visual system to render the harmonic undetectable (Campbell *et al.*, 1971, 1978; Sullivan & Georgeson, 1977). At high contrasts on the other hand [Fig. 6(b)], the fundamental component is independently detectable, and thus its absence in the MF stimulus becomes noticeable. While this explanation of the MF illusion is not incorrect, it is only partial since it does not offer an explicit mechanism by which both MF and square-wave stimuli have a square-wave appearance. To say that the low contrast MF stimulus "defaults" to a square-wave appearance (Campbell *et al.*, 1971) begs the question as to what mechanism implements this rule. This is necessary since the effects of low frequency attenuation on both MF and square-wave stimuli result in both their profiles looking more like a MF than a square-wave (Grossberg, 1983; Todorovic, 1987; Kingdom & Moulden, 1989). This highlights a more general problem with trying to "explain" brightness phenomena on the basis of the CSF. The CSF only provides information about the effects of low spatial frequency attenuation on the stimulus itself, and not on how the visual system interprets the neural image which reflects that attenuation. MIDAAS contains explicit rules for making such an interpretation.

Lightness integration models

This important class of brightness models originated from the need to account for lightness constancy and the earliest and best known model in this category is that of Land and McCann (1971) (for a review of this and other Lightness Integration models see Hurlbert, 1986). Lightness constancy is the apparent constancy of lightness and hue in the context of wide variations in the intensive and spectral content of illumination. All lightness integration models employ four stages. They are (i) a non-linear (usually log) transform of the luminance profile, (ii) differentiation of the (transformed) luminance profile, (iii) thresholding and (iv) reintegration. The imposition of a threshold removes any gradual luminance gradients in the image which most likely arise from illumination gradients.

Lightness integration models can predict brightness phenomena such as the Cornsweet illusion (Horn, 1974; Blake, 1985; Arend & Goldstein, 1987), although they appear not to be able to do so with any quantitative precision (Moulden & Kingdom, 1990). However, they do not predict important brightness phenomena such as simultaneous brightness contrast and Mach bands (Todorovic, 1987).

There are models of reflectance recovery which do not follow the multi-stage approach of lightness integration models, but which encode reflectance directly from the response of operators designed specifically to discount the effects of gradual illumination gradients (Land, 1986; Hurlbert & Poggio, 1988). These operators are characterised by having small receptive field centres with extensive receptive field surrounds, and both Land (1986) and Hurlbert and Poggio (1988) have shown that Mondrian patterns filtered with such operators show the predicted effects of simultaneous brightness contrast, and in the case of Land's (1986) operator, Mach bands as well. However, we have previously demonstrated that filtering a Cornsweet figure with such an operator does not predict the associated illusion, and thus models of brightness coding based on the use of operators with large surround-to-centre size ratios may be of limited applicability (Moulden & Kingdom, 1990).

Feature detector models

This approach imputes to local feature detectors a central role in the generation of the brightness description of an image, and as such MIDAAS may be classed as a feature detector model. We will examine four such models.

1. Grossberg and co-workers' spatial averaging model. The central feature of this approach (Cohen & Grossberg, 1984; Grossberg & Todorovic, 1988), and of a similar one by Hamada (1984), is that the brightness of a contour-bounded region in space is encoded by the average response of a spatial filter within that region. Since the correlate of brightness is thus an average filter response over a given area we refer to Grossberg's model as a "spatial averaging model". Grossberg and his co-workers place much emphasis on the presumed neurophysiological implementation of the spatial averaging process. In their model it occurs as a result of a diffusive spread, or "filling-in", of neural activity within the region of interest. This results in a pattern of homogenous neural activity which isomorphically defines the brightness profile on a point by point basis. We (Kingdom & Moulden, 1989) and others before (Foster, 1983; Laming, 1983) have argued against the logical requirement of such a filling-in process, whose proponents implicity assume that it is necessary to produce an isomorphic spatial mapping of internal response to percept (though for some evidence for filling-in processes see Paradiso & Nakayama, 1991). However, the suggestion that the spatial average in the filter response represents the code for brightness should not be seen as contingent on the existence of a filling-in mechanism, and is thus worth considering in its own right.

A necessary requirement of spatial averaging models is that the spatial filter is not completely band-pass: it must have a d.c. response. In Grossberg's model this is achieved by using a filter whose spatial weighting function is an approximation to the receptive field of a retinal ganglion cell, but whose excitatory centre is more heavily weighted than the surround, i.e. it is "unbalanced". This requirement is necessary because if the filter were balanced the brightness response for a homogeneous disc computed by spatial averaging would rapidly approach zero as the disc increased in size.

Grossberg's spatial averaging model qualitatively predicts a variety of brightness phenomena, and has the additional benefit of being able to operate on 2-D (two-dimensional) stimuli (in its present form MIDAAS only operated with 1-D stimuli). A particular advantage of Grossberg's model over lightness integration approaches is that, like MIDAAS, it can predict simultaneous brightness contrast.

There are, however, some difficulties with Grossberg's approach. Firstly, the model cannot account for what we have previously described as "transivity" effects in induced brightness (Kingdom & Moulden, 1989). Arend, Buehler and Lockhead (1971) first showed that the induced brightness on either side of a radial Cornsweet edge "carried across" suitably placed rings in those regions, as well as into the rings themselves. Since the spatial averaging process in Grossberg's model is contour-bounded, such an effect would not be predicted. MIDAAS, along with Lightness Integration models, predicts this effect because step changes in brightness are integrated throughout the image. Secondly, it is not clear how Grossberg's model, at least in its single filter version, could predict the presence of the Mach bands in for example, the trapezoid wave.

2. Watt and Morgan's MIRAGE model. This model (Watt & Morgan, 1985) presents a challenging view of how the outputs of filters tuned to different spatial scales might be combined to provide information about critical features in the image. It puts forward the novel suggestion that the outputs of different spatial scale filters are (non-linearly) combined *prior* to any interpretation of those filter outputs. Although MIRAGE was formulated principally to account for a wide range of data obtained from positional acuity and edge blur discrimination tasks [and is particularly noteworthy for its radically different interpretation of the results of threshold summation experiments from that conventionally offered (Graham, 1989)]. Watt and Morgan (1985) and Watt (1988) has shown how it can account for brightness phenomena such as Mach bands and the Chevreul illusion. As we have noted, MIDAAS draws on the implementation rules that MIRAGE employs to interpret the post-convolution image in terms of a pattern of critical features (edges and bars). The crucial difference between MIDAAS and MIRAGE lies in the fact that in MIRAGE the responses of the four filters are (non-linearly) combined *before* they are interpreted whereas in MIDAAS they are combined *after* they are interpreted. Thus, in MIRAGE a *single* interpretation of the pattern of brightness changes is produced from the combined filter response, whereas in MIDAAS *multiple* interpretations are produced, one for each filter, with the final percept being given by the average of those multiple interpretations.

The difference between these two approaches has important consequences for their ability to model brightness perception. For example as Fiorentini, Baumgartner, Magnusson, Schiller and Thomas (1990) have noted, MIRAGE can signal the presence of Mach bands at the knee and foot of a ramp, but not at the same time the presence of the ramp itself. Yet except when the gradient of the ramp is very shallow, the ramp is clearly visible along with the Mach bands. Fiorentini *et al.* (1990) also note that the composite percept of Mach bands and the ramp could be predicted if the MIRAGE interpretation rules were applied separately to the response of each filter rather than to the combined filter response. This is precisely what MIDAAS does.

A similar problem for MIRAGE would arise with the staircase figure in Fig. 7. Depending on the spatial scale of the stimulus, MIRAGE would either predict a series of bars (the Chevreul illusion) or a staircase, but unlike MIDAAS, not both simultaneously.

3. Morrone and Burr local energy model. This model (Morrone & Burr, 1988) offers an elegant and plausible explanation of how the visual system might detect both the presence and the location of salient features in the image, particularly edges and bars. Features are detected as peaks in "local energy", where local energy is defined as the square-root of the sums of squares of the responses of both odd- and even-symmetric filters. The nature of the feature is obtained by evaluating the relative responses of the even and odd detectors at that point. If the peak in local energy coincides with that of the even-symmetric filter, the stimulus is a bar; if it coincides with the peak of the odd-symmetric filter, the stimulus is an edge. The model is in keeping with psychophysical evidence for edge and bar detectors in the visual cortex (Burr, Morrone & Spinelli, 1989), and predicts the perceived magnitude of Mach bands under the conditions in which they occur (Morrone, Ross, Burr & Owens, 1986; Ross *et al.*, 1989).

A fundamental problem however with the model is that it predicts that a sinusoidal grating is featureless, since $(\sin^2 + \cos^2) = 1$ everywhere in the stimulus (first pointed out to us in personal communication by Mark Georgeson). Moreover, it cannot predict that the feature content of a stimulus could change with its physical amplitude. This is so because the position of the peaks in the odd-, even-symmetric and combined filter responses, upon which the computed feature content depends, is invariant with amplitude. Thus the model cannot predict the change in appearance of the MF (Missing Fundamental) waveform from that of a square-wave at low contrasts to that which is more or less isomorphic with its luminance profile at high contrasts (Fig. 6), a change which has a well-defined psychophysical threshold (Campbell *et al.*, 1971, 1978; Sullivan & Georgeson, 1977). MIDAAS predicts the change

because biphasic responses at each discontinuity at low contrasts signalling edges give way to triphasic responses at higher contrasts signalling bars.

4. Approaches involving interactions between "edge" and "bar" detectors. The effect of narrow contours on the perceived contrast of edges upon which they are superimposed [Fig. 5(b)], the effects of narrow contours on the width of nearby Mach bands as described below (Ratliff, Milkman & Kaufman, 1979; Ratliff, Milkman & Rennert, 1983; Ratliff, 1984), and the absence of Mach bands at a step edge, have been variously attributed to the idea of mutual inhibition between "edge" and "bar" detectors in the visual system (Tolhurst, 1972; Ratliff, 1984).

Tolhurst (1972) first suggested that the effectiveness of a superimposed contour in masking an edge (see Fig. 5) arises because strongly stimulated bar detector neurones inhibit the edge signal provided by less strongly stimulated edge detector neurones. On the other-hand, for a step edge on its own the reverse is true: the more strongly stimulated edge detectors inhibit the weakly stimulated bar detectors, preventing any illusory bars (Mach bands) being seen at the edge. Finally, with a luminance ramp, the optimally excited edge detectors would be in the centre of the ramp while the optimally excited bar detectors would be at the boundaries of the ramp. Because the two are spatially separated, there is less inhibition of the bar detectors by the edge detectors, and the illusory Mach bands appear at the foot and knee of the ramp.

In what has been argued to be a related phenomenon, Ratliff *et al.* (1979) first observed that the Mach bands at either end of a ramp could be attenuated by the presence of a bar positioned *half-way* up the ramp. This effect is also readily observed in a trapezoid, but when reproduced photographically appears markedly reduced and for this reason not reproduced here. Ratliff (1984) provided psychophysical measurements of this phenomenon and found the degree of Mach band attenuation to be inversely proportional to the contrast of a "biphasic" (a bright and a dark) bar that he positioned mid-way up the ramp. He went on to argue that the effect occurred because the bar located centrally in the ramp provided additional stimulation to "edge" detectors at this position, enabling them to inhibit the neighbouring "bar" detectors which produced the Mach bands more strongly. This resulted in the Mach bands being attenuated.

We have already shown how MIDAAS predicts the reduced apparent contrast of an edge resulting from the superimposition of a contour (see Fig. 5 and the accompanying explanation). MIDAAS also predicts the effect of nearby contours on the perceived magnitude of Mach bands described by Ratliff, and this is shown in Fig. 10. According to MIDAAS, the "culprit" is filter 2 which fails to produce monophasic responses at the foot and

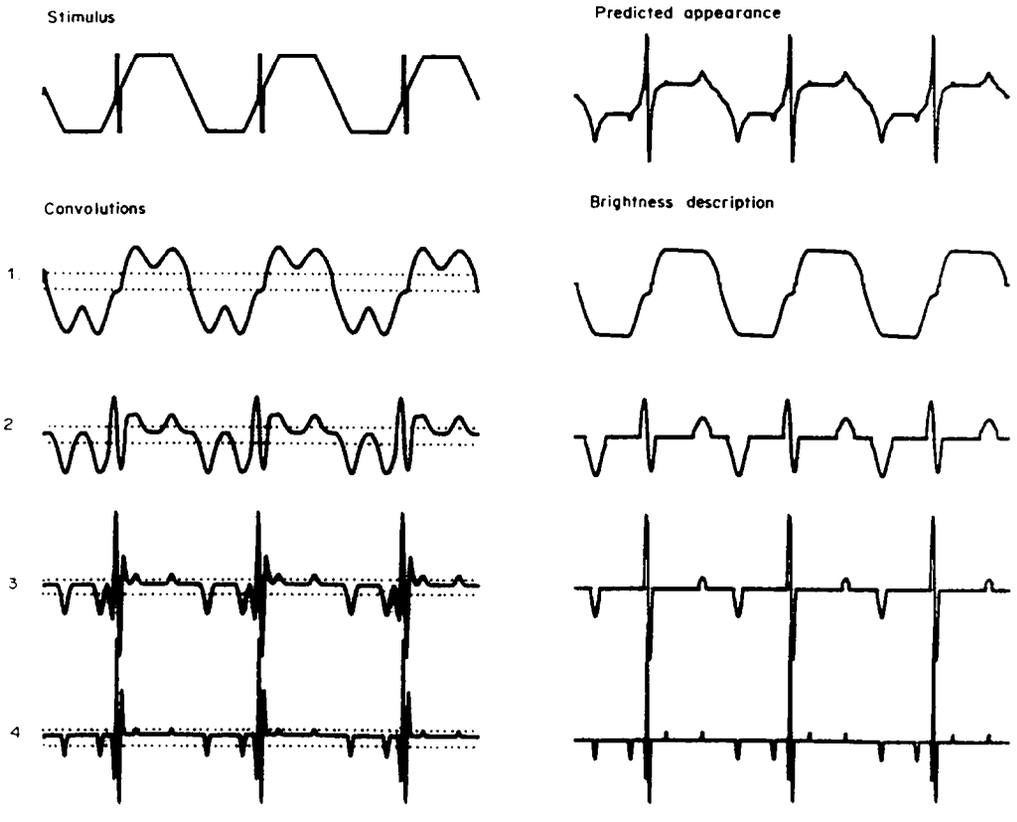

FIGURE 10. MIDAAS predicts attenuation of the Mach bands in a trapezoid wave when biphasic bars are positioned mid-way up the ramps. Note the difference in predicted amplitude of the Mach bands at the ends of the ramps between those with and those without biphasic bars. Each biphasic bar has the effect of removing the contribution of filter II to its adjacent Mach band signal.

knee of the ramps when the added biphasic bar is present. This results in only filters 3 and 4 contributing their monophasic responses to the final percept, with the result that the Mach band is predicted to be narrower in accordance with observation and psychophysical measurement.

Thus the attenuating effects of suitably positioned contours on apparent edge contrast and Mach bands, which according to some researchers impute inhibitory interactions between neurones specialised for detecting edges and those specialized for detecting bars, result according to MIDAAS from the combined effects of individual interpretation of filter outputs at different spatial scales.

Finally, it is worth noting another possible explanation for the effect of occluding bars on the apparent contrast of step edges. This effect may simply result from a non-linear transducer function for contrast (Legge & Foley, 1980; Wilson, 1980; Legge & Kersten, 1983; Whittle, 1986). It is possible that while in the case of the unoccluded edge the task is simply contrast detection, with the occluded edge the task is to discriminate between two contrasts, one at each edge of the occluding contour. For a given step edge contrast, the greater the amplitude of the occluding bar the smaller will be the perceived difference in contrast on each side of the bar. In the limit, if the amplitude of the occluding bar was great enough, the two contrasts at the edge of the contour would fall below their discrimination threshold.

Clearly, further research will be needed to test between the MIDAAS explanation and those alternatives just described.

CONCLUSION

We have shown that a model of brightness perception which includes (a) multiple spatial scale filtering and (b) separate symbolic descriptions of brightness from each spatial scale provides a good account of the pattern of brightness relationships in a range of one-dimensional brightness patterns. The model described here is not intended as a computational model of spatial vision providing an account of how the visual system computes critical stimulus features such as spatial position, local orientation, blur and contrast. It is intended as a preliminary exercise investigating the usefulness to understanding brightness perception of the challenging notion of individual filter response interpretation combined across spatial scale. The extent to which the approach of MIDAAS may be useful for understanding other spatial tasks can only be ascertained through further research.

Finally, are there any brightness phenomena which MIDAAS cannot predict? There are many brightness phenomena which are inherently two-dimensional and for which MIDAAS is thus currently inapplicable. These include White's effect (White, 1979; Moulden & Kingdom, 1989; Kingdom & Moulden, 1990), the grating induction effect [McCourt, 1982; and see Moulden and Kingdom (1991) for a quasi two-dimensional implementations of MIDAAS for grating induction], the Benussi ring and a number of two-dimensional Craik–Cornsweet–O'Brien figures (Todorovic, 1987). It is intended that a two-dimensional version of MIDAAS will be implemented in order to test its applicability to these phenomena.

REFERENCES

Arend, L. E. & Goldstein, R. (1987). Lightness models, gradient illusions and curl. *Perception and Psychophysics, 42*, 65–80.

Arend, L. E., Buehler, J. N. & Lockhead, G. R. (1971). Difference information in brightness perception. *Perception and Psychophysics, 9*, 367–370.

Blake, A. (1985). Boundary conditions for lightness computation in a Mondrian world. *Computer Vision, Graphics and Image Processing, 32*, 314–327.

Brucke, E. (1865), Uber Erganzungs und Contrastfarben. *Wiener Sitzungsber, 51*, 461–501.

Burr, D. C., Morrone, C. M. & Spinelli, D. (1989). Evidence for edge and bar detectors in human vision. *Vision Research, 29*, 419.

Burrr, D. C., Ross, J. & Morrone, M. (1985). Local regulation of luminance gain. *Vision Research, 25*, 717–727.

Campbell, F. W., Howell, E. R. & Johnstone, J. R. (1978). A comparison of threshold and suprathreshold appearance of gratings with components in the low and high spatial frequency range. *Journal of Physiology, 284*, 193–201.

Campbell, F. W., Howell, E. R. & Robson, J. G. (1971). The appearance of gratings with and without the fundamental Fourier component. *Journal of Physiology, London, 217*, 17–19.

Cleland, B. G. & Freeman, A. W. (1987). Visual adaptation is highly localised in the cat's retina. *Journal of Physiology, 404*, 591–611.

Cohen, M. A. & Grossberg, S. (1984). Neural dynamics of brightness perception: Features, boundaries, diffusion and resonance. *Perception and Psychophysics, 36*, 428–456.

Cornsweet, T. (1970). *Visual perception*. New York: Academic Press.

Dooley, R. P. & Greenfield, M. I. (1977). Measurements of edge-induced visual contrast and a spatial frequency interaction of the Cornsweet illusion. *Journal of the Optical Society of America, 67*, 761–765.

Fiorentini, A. (1972). Mach band phenomena. In Jameson, D. & Hurvich, L. M. (Eds), *Handbook of sensory physiology* (Vol. VII/4). Springer.

Fiorentini, A., Baumgartner, G., Magnusson, S., Schiller, P. H. & Thomas, J. P. (1990). The perception of brightness and darkness. In Spillman, L. & Werner, J. S. (Eds), *Visual perception: The neurophysiological foundations*. San Diego, Calif.: Academic Press.

Foster, D. H. (1983). Experimental test of a network theory of vision. *Behavioural and Brain Sciences, 6*, 664–665.

Gilchrist, A. (1979). The perception of surface white and blacks. *Scientific American, 24*, 38–97.

Graham, N. V. S. (1989). *Visual pattern analyzers*. Oxford: Clarendon and New York: Oxford University Press.

Grossberg, S. (1983). The quantized geometry of visual space: The coherent computation of depth, form and lightness. *Behavioural and Brain Sciences, 6*, 625–692.

Grossberg, S. & Todorovic, D. (1988). Neural dynamics of 1-D and 2-D brightness perception: A unified model of classical and recent phenomena. *Perception and Psychophysics, 43*, 241–277.

Hamada, J. (1984). A multi-stage model for border contrast. *Biological Cybernetics, 51*, 65–70.

Hamada, J. (1985). Asymmetric lightness cancellation in Craik–O'Brien patterns of negative and positive contrast. *Biological Cybernetics, 52*, 117–122.

Horn, B. K. P. (1974). Determining lightness from an image. *Computers and Visual Graphics Image Processing, 3*, 277–299.

Hubel, D. H. & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology, London, 195*, 215–243.

Hurlbert, A. (1986). Formal connections between lightness algorithms. *Journal of the Optical Society of America, A, 3*, 1684–1693.

Hurlbert, A. & Poggio, T. A. (1988). Synthesizing a color Algorithm from examples. *Science, 239,* 484–485.

Kingdom, F. & Moulden, B. (1989). Border effects on brightness: A review of findings, models and issues. *Spatial Vision, 3,* 225–262.

Kingdom, F. & Moulden, B. (1990). White's effect and assimilation. *Vision Research, 31,* 151–159.

Kingdom, F. & Moulden, B. (1991). A model for contrast discrimination with incremental and decremental test patches. *Vision Research, 31,* 851–858.

Laming, D. (1983). On the need for discipline in the construction of psychological theories. *Behavioural and Brain Sciences, 6,* 669–670.

Land, E. H. (1986). An alternative technique for the computation of the designator in the retinex theory of color vision. *Proceedings of the National Academy of Science, U.S.A., 83,* 3078–3080.

Land, E. H. & McCann, J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America, 61,* 1–11.

Land, E. H., Hubel, D., Livingstone, M., Perry, S. & Burns, M. (1983). Color-generating interactions across the corpus callosum. *Nature, 303,* 616–618.

Legge, G. E. & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America, 70,* 1458–1471.

Legge, G. E. & Kersten, D. (1983). Light and dark bars: Contrast discrimination. *Vision Research, 23,* 473–483.

Mach, E. (1865). Uber die Wirkung der raumlichen Vertheilung des Lichtreizes auf di Neztzhaut. Translated by Ratliff, F. (1965). *I.S.-B. Akad. Wiss. Wien, math.-nat. Kl, 54,* 303–322.

Marcelja, S. (1980). Mathematical description of the responses of simple cortical cells. *Journal of the Optical Society of America, 70,* 1297.

McCourt, M. E. (1982). A spatial frequency dependent grating induction effect. *Vision Research, 22,* 119–134.

Morrone, M. C. & Burr, D. C. (1988). Feature detection in human vision: A phase dependent energy model. *Proceedings of the Royal Society of London B, 235,* 221–245.

Morrone, M. C., Ross, J., Burr, D. C. & Owens, R. (1986). Mach bands depend on spatial phase. *Nature, 324,* 250–253.

Moulden, B. & Kingdom, F. (1989). White's effect: A dual mechanism. *Vision Research, 29,* 1245–1259.

Moulden, B. & Kingdom, F. (1990). Light–dark asymmetries in the Craik–Cornsweet–O'Brien illusion and a new model of brightness coding. *Spatial Vision, 5,* 101–121.

Moulden, B. & Kingdom, F. (1991). The local border mechanism in grating induction. *Vision Research, 31,* 1999–2008.

Naka, K. I. & Rushton, W. A. H. (1966). S-potentials from luminosity units in the retina of fish (Cyprinidae). *Journal of Physiology, 185,* 587–599.

Paradiso, M. A. & Nakayama, K. (1991). Brightness perception and filling-in. *Vision Research, 31,* 1221–1236.

Ratliff, F. (1965). *Mach bands.* San Francisco, Calif.: Holden–Day.

Ratliff, F. (1984). Why Mach bands are not seen at the edges of a step? *Vision Research, 24,* 163–165.

Ratliff, F., Milkman, N. & Kaufman, T. (1979). Mach bands are attenuated by adjacent bars or lines. *Journal of the Optical Society of America, 69,* 1444.

Ratliff, F., Milkman, N. & Rennert, N. (1983). Attenuation of Mach bands by adjacent stimuli. *Proceedings of the National Academy of Science, 80,* 4554–4558.

Ross, J., Holt, J. J. & Johnstone, J. R. (1981). High frequency limitation on Mach bands. *Vision Research, 21,* 1165–1167.

Ross, J., Morrone, M. C. & Burr, D. C. (1989). The conditions under which Mach bands are visible. *Vision Research, 29,* 699–715.

Shapley, R. & Enroth-Cugell, C. (1984). Visual adaptation and retinal gain controls. In Osbourne, N. N. & Chader, G. J. (Eds), *Progress in retinal research* (pp. 263–346). Oxford: Pergamon Press.

Sullivan, G. D. & Georgeson, M. A. (1977). The missing fundamental illusion: Variation of spatio-temporal characteristics with dark adaptation. *Vision Research, 17,* 977–981.

Todorovic, D. (1987). The Craik–O'Brien–Cornsweet effect: New varieties and their theoretical implications. *Perception and Psychophysics, 42,* 545–560.

Tolhurst, D. J. (1972). On the possible existence of edge detector neurons in the human visual system. *Vision Research, 12,* 797–804.

Watt, R. J. (1988). *Visual processes: Computational, psychophysical and cognitive research.* London: Erlbaum.

Watt, R. J. & Morgan, M. J. (1985). A theory of the primitive spatial code in human vision. *Vision Research, 11,* 1661–1674.

White, M. (1979). A new effect on perceived lightness. *Perception, 8,* 413–416.

Whittle, P. (1986). Increments and decrements: Luminance discrimination. *Vision Research, 26,* 1677–1691.

Wilson, H. R. (1980). A transducer function for threshold and suprathreshold human vision. *Biological Cybernetics, 38,* 171–178.

Acknowledgements—This research was supported by the Wellcome Trust (U.K.) and by the Royal Victoria Hospital Research Institute, Montreal, Canada.