

Contents lists available at ScienceDirect

Vision Research

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

Mach edges: Local features predicted by 3rd derivative spatial filtering

Stuart A. Wallis, Mark A. Georgeson

School of Life and Health Sciences, Aston University, Birmingham B4 7ET, UK. Email: s.a.wallis2@aston.ac.uk

ARTICLE INFO

Article history:

Received 17 October 2008

Received in revised form 8 April 2009

Keywords:

Human vision
Psychophysics
Edge localisation
Spatial derivatives
Mach bands

ABSTRACT

Edges are key points of information in visual scenes. One important class of models supposes that edges correspond to the steepest parts of the luminance profile, implying that they can be found as peaks and troughs in the response of a gradient (first derivative) filter, or as zero-crossings in the second derivative (ZCs). We tested those ideas using a stimulus that has no local peaks of gradient and no ZCs, at any scale. The stimulus profile is analogous to the Mach ramp, but it is the luminance gradient (not the absolute luminance) that increases as a linear ramp between two plateaux; the luminance profile is a blurred triangle-wave. For all image-blurs tested, observers marked edges at or close to the corner points in the gradient profile, even though these were not gradient maxima. These Mach edges correspond to peaks and troughs in the third-derivative. Thus Mach edges are inconsistent with many standard edge detection schemes, but are nicely predicted by a recent model that finds edge points with a 2-stage sequence of 1st- then 2nd-derivative operators, each followed by a half-wave rectifier.

1. Introduction

Edges are key points of information in visual scenes. But despite much research it remains uncertain how edges are extracted from the eyes neural output. It is widely accepted that the retinal image is filtered by even- and odd-symmetric spatial filters at various scales, early in the visual pathway (Burr, Morrone & Spinelli, 1989; Field & Nachmias, 1984; Hubel & Wiesel, 1968; Pollen & Ronner, 1981; Ringach, 2002) but how the filters are used in feature detection remains an open question. Early psychophysical work proposed an edge-detector role for odd-symmetric filters and a bar-detector role for even-symmetric filters (Kulikowski & King-Smith, 1973; Shapley & Tolhurst, 1973). This could be true, but as many have pointed out, such a simple interpretation is incomplete because the 'edge-detectors' also respond to bars, and the 'bar-detectors' respond to edges (Fig. 1), with peak responses offset to left or right of the feature in question. The interpretive parsing rules of MIRAGE (Watt & Morgan, 1985) and the quadratic summation of even and odd responses in the local energy model (Morrone & Burr, 1988) were both motivated by the need to resolve that ambiguity.

One attractive general view of receptive fields in early vision is that they act as spatial derivative operators (Lindeberg, 1994; Marr & Hildreth, 1980; ter Haar Romeny, 2003; Watt & Morgan, 1985; Young, 1985; Young & Lesperance, 2001). For example, summation of the output of the two regions of the odd-symmetric receptive field shown in Fig. 1 (left) is equivalent to obtaining the difference in luminance between these two regions. If this filter is convolved with a 1-D image, then the output at each point is proportional to the spatial luminance gradient (the 1st derivative), after a degree of smoothing that is determined by the scale (size)

of the receptive field (Fig. 1, middle row). By a similar argument (with an increasing number of receptive field regions) filters can be obtained that compute the 2nd, 3rd or any higher derivative.

Edge detection models based on derivative computation often suppose that edges correspond to the steepest parts of the luminance profile, implying that they can be found as peaks and troughs in the response of a 1st derivative (gradient) filter (Bergholm, 1987; Canny, 1986; Korn, 1988; Sarkar & Boyer, 1991; Zhang & Bergholm, 1997), or as zero-crossings in the second derivative (Elder & Zucker, 1998; Georgeson, 1992; Marr & Hildreth, 1980; Watt & Morgan, 1985).

The third spatial derivative has until recently played an ancillary role in edge-detection schemes. Zero-crossings (ZCs) in the 2nd derivative occur both at maxima and minima of gradient magnitude (Clark, 1989). Clark regarded the minima (points




Figure 1: Receptive fields of odd-symmetric and even-symmetric filters (left, in plan view and cross-section), and their responses (right) to the blurred bar and edge shown in the top row. Note that each filter responds to both features.

of locally shallowest slope) as spurious edges, and showed from standard calculus how ZCs could be classed as ‘real’ or spurious from the sign of the product of the 1st and 3rd derivatives at the ZC. A negative sign identifies a real edge while a positive sign identifies a spurious one. Thus the 3rd derivative was used to categorize edges found at ZCs in the 2nd derivative.

Important theoretical developments, in the framework of Gaussian scale-space theory, were made by Lindeberg (1998) who used peaks in the multi-scale 1st derivative to find edge locations, and then used the multi-scale 3rd derivative to determine the strength and blur of each edge. This was followed by a step that provided a measure of the saliency of each edge by integrating edge strength along the contour. By selecting a restricted number (such as 100) of the most salient curves, the algorithm produced an effective line drawing of a variety of test images.

In this paper we ask whether the 3rd derivative may be used directly for edge-finding in human vision. At first sight, it appears unpromising because a 3rd derivative filter generates not only a peak at the edge location, but also a pair of flanking troughs, thus apparently signalling two spurious edges of opposite polarity adjacent to the ‘real’ edge (Fig. 2A). Georgeson, May, Freeman & Hesse (2007) however, showed that decomposing the 3rd derivative into two stages overcomes the problem of spurious edges (Fig. 2B). This scheme - a 1st derivative operator followed by an inverted 2nd derivative, with half-wave rectification on the output of each stage - creates a nonlinear channel sensitive to edges of a given polarity, but the peak response to a preferred edge is unaffected by the nonlinearities (Fig. 2B).

The multiscale model (called N3+) based on this approach predicted very well the perceived blur of a wide variety of edge-like waveforms, including sinusoids. May & Georgeson (2007) further showed that the addition of a smooth, threshold-like suppression of small values at the first rectifier accounted well for the finding that reducing contrast made blurred edges appear sharper. The

success of N3+, however, did not rule out a simpler multiscale gradient model (N1), based solely on the first derivative, which performed fairly well in edge-finding and blur coding in many circumstances. Our aim therefore was to devise a more definitive experimental test across the family of derivative-based models of edge-finding outlined above.

The idea is inspired directly by the phenomenon of Mach Bands, whose well-known ramp waveform has no peaks or troughs in luminance, but does have a peak and a trough in the 2nd derivative at the perceived location of the bright and dark bands. This logic can be shifted up by one derivative order, to test the role of 3rd derivative extrema in edge detection. The first derivative of our stimulus is defined by a Mach ramp between two plateaux, thus ensuring that the 3rd derivative (rather than the 2nd) has a peak and a trough at the ends of the ramp. If edges are reliably seen at these points, in the absence of gradient maxima, then we propose that they be called ‘Mach Edges’, by direct analogy with Mach Bands.

Our experiments are therefore a search for Mach Edges. We designed luminance waveforms that contained peaks in the 3rd derivative but had no corresponding peaks in the 1st derivative nor zero-crossings in the 2nd derivative, at any scale. We then used the feature-marking method (Hesse & Georgeson, 2005) to determine whether edges were perceived in these stimuli, and if so where. We consider a simplified version of the N3+ model that uses filters at a single fine scale, which we shall refer to as the sN3+ model (‘s’ meaning ‘single-scale’). Since the N3+ and sN3+ models produce very similar predictions for our stimuli, we can simplify the multi-scale aspect of N3+ while retaining the ability to test its use of derivative filters and half-wave rectification.

2. Experiment 1

The purpose of this experiment is to test for the existence of Mach Edges: that is, to determine whether edges are reliably seen at or near 3rd derivative extrema in luminance waveforms designed to have no corresponding gradient maxima.

2.1 Stimulus design

Since the absence of peaks in the 1st derivative was of prime importance, the starting point was to create a peak-free waveform representing the gradient profile, and then integrate it to form the luminance profile. The gradient profile (Fig. 3) was a single period of a trapezoidal wave whose ramps were 1, 2, 4, 8, 16, 32 or 64 pixels wide. This was integrated to form the luminance profile of the vertical 1-D test image. The luminance profile can equivalently be described as a triangle-wave blurred by a box function whose width ranged from 1 to 64 pixels. We shall refer to this width as blurwidth.

The stimulus design is analogous to the Mach Band stimulus, but it is the luminance gradient (not the absolute luminance) that increases as a linear ramp between two plateaux. Two example images, their luminance waveforms and first three derivatives are shown in Fig. 3. A graph of the luminance profile reveals no obvious edge locations: a uniform positive gradient shades smoothly into a uniform negative one. It has no peak in the 1st derivative and no ZC in the 2nd derivative, so no edges are predicted by models based on these derivative features. That remains true at all filter scales, because Gaussian smoothing does not introduce any new peaks (Babaud, Witkin, Baudin & Duda, 1986; Koenderink, 1984; Lindeberg, 1990; Yuille & Poggio, 1986). However, Fig. 3 shows that there is a sharply localized peak and trough in the 3rd derivative corresponding to the corner points in the gradient profile. If such peaks and troughs are taken as the signature of edges, then human observers should see two edges, of opposite polarity, at these locations.

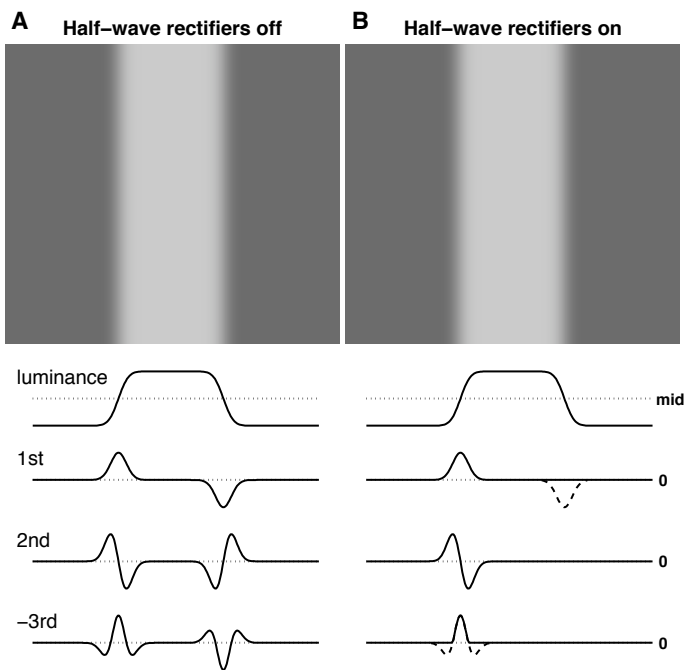


Figure 2: Two blurred edges and their first 3 derivatives. The 3rd derivative is shown inverted. A) Sequence of derivatives computed without half-wave rectification. B) Sequence of derivatives computed with half-wave rectification after the 1st and -3rd derivatives were obtained. Responses suppressed by rectification are shown by dashed lines. Key property of this nonlinear model (sN3+) is that a single response peak occurs at the positive-going (dark-to-light) edge location. A second, complementary channel is needed for edges of the opposite polarity (cf. Fig. 9). (Note: Each derivative was computed in Matlab by convolution with the small-scale, 3-point gradient operator whose weights were [-0.5, 0, 0.5]).




Figure 3: Two images from experiment 1 (blurwidths 2 and 32 pixels, peak central), their luminance waveforms and first three (linear) derivatives. The 1st derivative profiles have the form of a Mach ramp.

2.2 Method

Image arrays were generated in Matlab on a Macintosh G4 computer and displayed using PsychToolbox software on an Eizo 6600-M greyscale monitor, calibrated and gamma-corrected using a Minolta LS110 digital photometer. A Cambridge Research Systems Bits++ box was used in Mono++ mode to render 14-bit greyscale resolution.

Images had one of two polarities: ‘peak central’ (see Fig. 3), or a contrast inversion of this named ‘trough central’. Image size was 256 by 256 pixels and subtended 4° at the viewing distance of 131.6 cm. Test images had Michaelson contrasts of 0.2 or 0.4, and were surrounded by a full-screen ($16^\circ \times 12^\circ$) mid-grey of luminance 40.7 cd/m^2 . They were displayed flashing (on 0.3s, off 0.6s) in order to reduce the build-up of afterimages that would cause instability and possible shifts in edge location (Georgeson & Turner, 1985). The inter-stimulus display was a full screen of mid-grey.

The task was to indicate the position and polarity of all edges seen in each image. Their position was identified by moving a marker across the image and pressing a button when the marker was over an edge. A second button-press indicated the polarity of the edge as either light-to-dark (LD) or dark-to-light (DL). Once all the perceived edges had been marked, the observer initiated the next trial. The marker consisted of two black dots, each 1 pixel wide by 3 pixels high. One dot was centred 32 pixels (0.5 deg) above, and the other 32 pixels below, the horizontal midline of the image. The observer was instructed to fixate midway between the two dots. The starting position of the marker alternated between left and right on successive trials, and was 64 pixels (1 deg) from the image border. Its movement was constrained to the central 2.5 degrees of the image. The 28 conditions (7 blurwidths, 2 phases, 2 contrast levels) were presented in randomised order, blocked by contrast. This procedure was repeated a further two times in one experimental session, which took about 30 minutes to complete.

The three observers (SAW, DHB and TAY) were all experienced psychophysical observers and had normal uncorrected vision. They viewed the display binocularly with natural pupils in a darkened room, with the head supported by a chin-and-forehead rest. They each completed three sessions, giving a total of 9 repetitions of each condition per observer.

2.3 Results: Mach Edges

Observers reliably saw pairs of edges at positions to the left and right of the luminance peaks and troughs in each image. We shall refer to these edges perceived without gradient peaks as Mach edges. Plots of perceived edge position against blurwidth are shown in Fig. 4 (symbols). Data were similar for all 3 observers and group means are shown. The perceived separation between those edges increased markedly with blurwidth.

Solid curves in Fig. 4 trace the positions of peaks in the sN3+ output derived from the luminance profile (i.e. with half-wave rectification after the 1st and 3rd derivative operation, cf. Fig. 2B). For this class of waveforms (but not in general), the outcome is almost identical to that for the linear 3rd derivative without rectifiers. The overall correspondence between observed edges and 3rd derivative extrema is strikingly close. In contrast, standard models based on the lower derivatives predict no edges here. A 1st derivative filter at a broad scale does have peaks at the mid-points of the luminance ramps (at position $x=64$ pixels), but these points were not marked as edges. At small blurwidths, the observed edges were systematically further apart than the 3rd derivative extrema by about 3.3 min arc, considered further below.

Because it is linear, the 3rd derivative operator predicts the same edge positions for peak-central and trough-central conditions, but with reversed polarity, and this light-dark symmetry is not affected by the rectification stages of the sN3+ model. Figure 4 shows, however, that in the experiment the two edges were seen as slightly further apart in the peak-central than the trough-central conditions, by an average of 1.9 min arc. This appears to be an example of the Helmholtz irradiation effect which may arise from compressive nonlinearity in the retinal response to luminance (Georgeson & Freeman, 1997; Mather & Morgan, 1986).

2.3.1. No effect of edge length or marker location

We wondered whether the greater-than-predicted separation between edges at small blurwidths might arise from an influence of larger receptive fields in peripheral vision, given that the image height was quite large (4 deg). If this were so, we might expect the Mach edges to appear closer together when the image was truncated to exclude the peripheral contribution.

Methods were as above, except (i) image height was reduced to 8, 16 or 32 pixels (still 256 pixels wide), (ii) only one contrast level (0.4) was used, (iii) only four blur-widths were used (1, 4, 16 and 64 pixels), (iv) marker-spots were 8 pixels above and below the image borders, on the mid-grey background. Results (Fig. 5) were very similar to those of the main experiment (Fig. 4) and any effect of image truncation was small. Rather than being closer together, the perceived edges were, if anything, a little further apart when most truncated (triangles in Fig. 5). Thus the contribution of peripheral retina does not appear to be crucial. A further experiment that used a single marker-spot on the image mid-line gave almost identical results (not shown), suggesting that the placement of the marker-spots on the image, or on the grey background, was also not an important factor.

2.3.2. Optical and neural blur

The sN3+ predictions (Figs. 4 & 5) were computed at a single fine scale, neglecting any impact of optical blur (Campbell & Gubisch, 1966; Williams, Brainard, McMahon & Navarro, 1994) or neural ‘intrinsic’ blur (Levi & Klein, 1990a; Levi & Klein, 1990b). To gauge the likely influence of blur, we applied Gaussian blur to the luminance profile before computing the sN3+ peaks. Gaussian blur had little effect on the predicted position of the widely separated Mach edges (high blurwidth images), but for low blurwidth images response peaks were shifted away from centre.

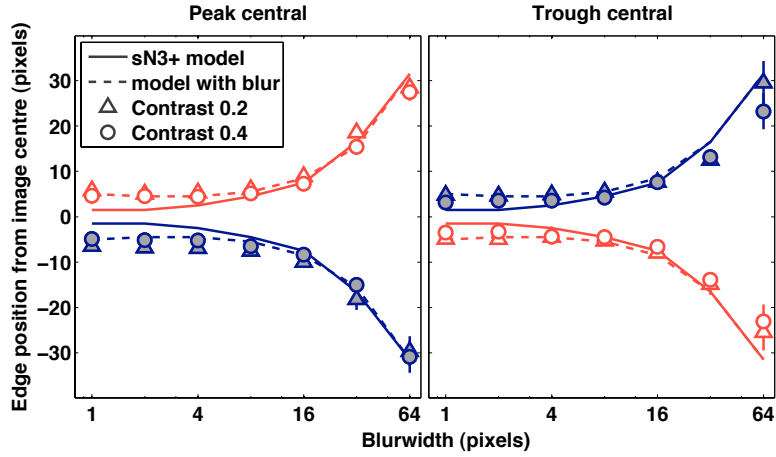


Figure 4: Experiment 1. Perceived edge locations as a function of blurwidth. Symbols show group mean data at two contrast levels, for dark-to-light edges (DL, filled symbols) and light-to-dark edges (LD, open symbols). Positions of peaks in the two nonlinear 3rd derivative channels are shown by solid curves. Dashed curves show the effect of including simulated Gaussian blur ($\sigma=4.6$ arc mins) before the derivative operators. 64 pixels = 60 arc mins. Left: peak-central waveform; right: trough-central waveform. Error-bars (± 1 standard error) are plotted behind symbols and show between-observer variation ($n=3$).

The dashed lines in Fig. 4 show that the simulated blur gave an improved and very close match between predictions and data (without blur: rms error = 3.1 pixels, Pearsons $X^2 = 1540$; with blur: rms error = 2.3 pixels, Pearsons $X^2 = 823$; this improvement in the X^2 goodness-of-fit, distributed as chi-square with 1 d.f., is hugely significant, $p < 0.00001$). The best fitting blur (with lowest rms error) was $\sigma=4.6$ min arc. This is far too large to represent dioptric blur alone, but could reflect the scale of the filter used by the observer in this task.

2.3.3. Pre-cortical filtering?

Odd-symmetric filters (including 1st and 3rd derivatives) necessarily have oriented receptive fields, and so presumably would be implemented by cortical neurons. We should not ignore, however, the possible role of earlier filtering in the retina or LGN. We found that applying a broad, mildly bandpass Difference-of-Gaussian (DoG) filter (Fig. 6A), similar in shape to the contrast sensitivity function, followed by gradient peak detection, did enable Mach edges to be detected and did predict the observed edge positions well (Fig. 6B). This is not too surprising, because to the extent that the DoG filter emulates a second-derivative operator (Marr & Hildreth, 1980), this revised gradient (DoG+1st derivative) model is analogous to a 3rd derivative filter. The

gradient filter alone does not predict Mach edges. We therefore devised a further experiment to distinguish between the revised-gradient and sN3+ models.

3. Experiment 2

The aim of this experiment is to distinguish between the two competing models (sN3+ vs DoG+1st derivative) that can account for the results of experiment 1. The success of the DoG+1st derivative model rests specifically on the low frequency attenuation provided by the DoG's inhibitory surround. Without this attenuation the 1st derivative alone, as we have seen, does not predict Mach edges. Experimentally, we can attempt to by-pass this low-frequency attenuation in two ways: (i) by using a brief presentation (e.g. 50 ms) that reduces the relative attenuation of sensitivity to low spatial frequencies in the CSF (Legge, 1978), and (ii) by shrinking the spatial period of the triangle-wave, and thus increasing the fundamental frequency to (say) 4-6 c/deg, near the peak of the MTF (Fig. 6A). The test image then contains only high frequencies ($\geq 4c/deg$), which are not subject to low frequency attenuation by the early pre-filter. If Mach edges remain visible in high frequency, blurred triangle-wave gratings then they are unlikely to arise from DoG+1st derivative filtering.

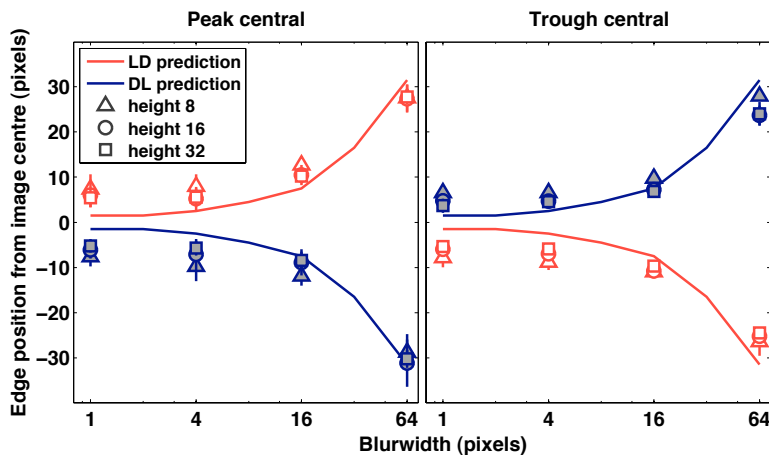


Figure 5: Mean observed positions of dark-to-light (DL) and light-to-dark (LD) edges in blurred triangle-wave images of 3 different heights (8, 16, 32 pixels). Other conventions as Fig. 4.

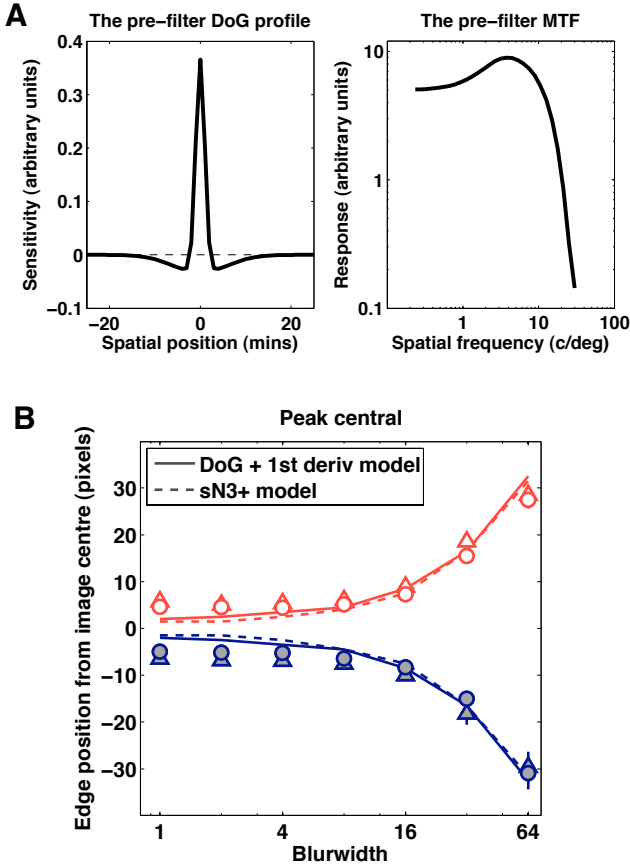


Figure 6: (A) The difference of Gaussian (DoG) receptive field profile and its MTF. The equation for this 1-D DoG function was $\text{DoG}(x) = G(x, \sigma_c) - K \cdot G(x, \sigma_s)$, where $G(x, \sigma) = \exp(-x^2/(2\sigma^2))/\sqrt{2\pi\sigma^2}$, with $\sigma_c = 1'$, $\sigma_s = 6'$, $K = 0.5$. These parameters correspond closely to the median values given by Croner & Kaplan (1995) for P-cells in the monkey central retina, allowing for the translation from 2-D to 1-D. (B) Results of the peak-central condition (Fig. 4) with predictions from the sN3+ (dashed curve, rms error = 3.1 pixels) and the DoG+1st derivative model (solid curve, rms error = 2.9 pixels).

3.1 Method

Test images (Fig. 7A) were a high spatial frequency version of those used in experiment 1. One set contained 16 cycles of a blurred triangle-wave (512 pixels wide; period 32 pixels), generated as in experiment 1. Blurwidth was 1, 2, 4 or 8 pixels with the usual two polarities (peak or trough central). Since each period occupied only 1/8 the number of pixels used in experiment 1, each level of blurwidth shown here is equivalent, in its effect on waveform shape, to 8 times the previous amount of blurring.

These waveforms were also used to create a second set of images, by shifting their Fourier phases through 90 degrees, while leaving the amplitude spectrum unchanged. This produced blurred waveforms that were akin to a square-wave, but whose amplitude spectrum (before the blurring) declined as $1/f^2$, rather than $1/f$, where f is spatial frequency of the (odd) harmonics. For brevity, we refer to these as ‘square-waves’ though strictly they are not. The triangle-wave image (Fig. 7A, left) appears to contain a thin light (or dark) bar of high contrast at the centre of each wide light (or dark) bar. These thin bars are not apparent in the square-wave image (Fig. 7A, right). Predictions of the two models (sN3+ and DoG+1st derivative) will be considered later.

The test images subtended $2.67^\circ \times 2.67^\circ$, with a fundamental frequency of 6 c/deg, at a viewing distance of 383cm (observer SAW). For the second observer (SEW) viewing distance was reduced to 255cm (fundamental frequency 4 c/deg) because she was unable to discriminate between the triangle-wave and square-wave images at 6 c/deg.

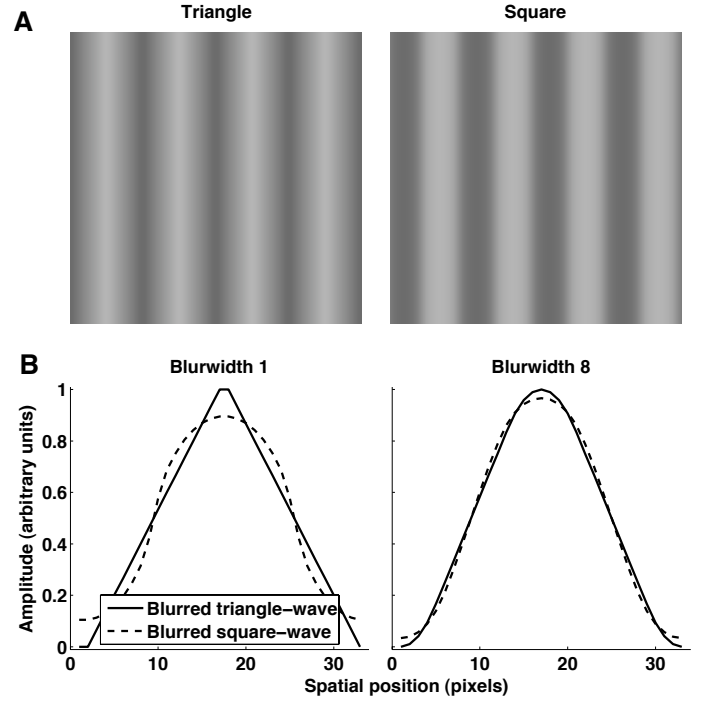


Figure 7: Experiment 2. (A) Four periods of the two types of images used. Here blurwidth=4, period=32 pixels. (B) One period of the waveforms at blurwidth 1 and blurwidth 8. These were the smallest and largest blurs used in experiment 2.

The test images were surrounded by a full screen of mid-grey luminance and viewed in a darkened room. RMS contrast (the ratio of the standard deviation of the luminance profile to its mean) was the same for both triangle-wave and square-wave images at a given blurwidth. A consequence of matching the RMS contrasts was that Michelson contrast for the triangle-wave images was 0.4, while that of the square-wave images was 0.32, 0.32, 0.33 and 0.37 for blurwidths 1, 2, 4 and 8 respectively. A subsidiary experiment controlled this factor by matching Michelson contrast instead of RMS contrast.

Before data collection began, each observer was shown a pair of 3 c/deg stimuli (blurwidth 8 pixels), and it was confirmed that a pair of closely spaced edges was easily visible in the centre of each half-period of the triangle-wave image but not in the square-wave image.

The task was a single-interval procedure requiring a yes-no decision about the presence or absence of the Mach edges. Each trial consisted of a single stimulus presentation of 50ms, preceded and followed by a full-screen of mid-grey showing a central small fixation dot (4×4 pixels = 1.25×1.25 mins). The observer had unlimited time to indicate the presence or absence of a central pair of edges in one or more half-periods of the grating. The inter-trial interval was at least 1s. Each of the 4 blurwidths, 2 image-types and 2 polarities were shown 15 times in a randomised block. The two observers each completed 9 blocks. The first block from each subject was discarded as practice. No feedback was given about the correctness of response because this experiment was concerned with the perception of Mach edges, rather than the ability to distinguish between the triangle-wave and square-wave images per se, which could be based on other image properties (such as the edge blur of the wide bars in the image). Nevertheless, we can make use of signal detection measures (d') to quantify the reliability with which Mach edges were reported.

3.2 Results

For each observer, the proportion of ‘yes’ responses was similar for the peak- and trough-central images, so data were pooled across both polarities. Z-scores corresponding to the proportion

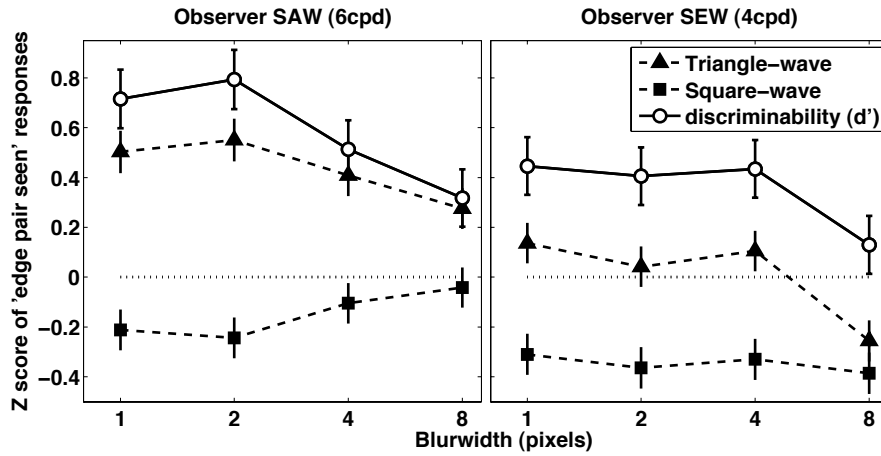


Figure 8: Experiment 2. Z-scores corresponding to the proportion of Mach edge responses (filled symbols), and discriminability (d' , open symbols), averaged across stimulus polarity, plotted as a function of blurwidth. Each filled symbol represents 240 trials and error bars (± 1 sd) are derived from the expected variance in binomial sampling. Each open symbol (d') represents the difference of the two z-scores, and error bars here are the square root of the sum of the expected variances of the triangle and square-wave z-scores at a given blurwidth.

of 'yes' (Mach edge) responses for each condition, averaged across stimulus polarity, are shown as filled symbols in Fig. 8. Both observers had a significantly higher proportion of 'yes' responses for the triangle-wave than the square-wave images, except perhaps at the greatest blur. Z-scores overall were higher for SAW than SEW, implying a lower criterion on the internal decision axis for SAW.

Treating 'yes' responses to the triangle stimulus as Hits, and those to the square stimulus as False Alarms, the discriminability index (d') was calculated in the standard way as $Z(\text{yes}|\text{triangle})$ minus $Z(\text{yes}|\text{square})$, plotted as open symbols in Fig. 8. The d' values were all significantly greater than 0 (except for SEW, blur 8), implying that both observers reliably associated Mach edges with the triangle-wave rather than the square-wave test images. Both observers showed a general trend of decreasing discriminability as blurwidth increased. At each blurwidth, d' values for the highly practised observer (SAW) were higher than for SEW.

To test whether discrimination might be cued by the lower Michaelson contrast of the square-wave images, observer SAW

repeated the experiment with images whose Michaelson contrast was always 0.4. Results were similar to those from the original image set, implying no artefactual effect of contrast difference.

3.3 Modelling

The results of experiment 2 show that Mach edges were reliably reported in blurred triangle-wave images that were high-spatial-frequency, short-duration versions of the images used in experiment 1. The use of high-frequency gratings is expected to by-pass the influence of the DoG pre-filter's inhibitory surround, as outlined above, and so offer a critical test between two models. Figure 9 shows the DoG+1st derivative output and the sN3+ output for this experiment, at blurwidth 4. For the square-wave (Fig. 9, right) both models produced peaks and troughs, and hence predicted edges, separated by half a period. These are not Mach edges. For the triangle-wave (Fig. 9, left), the DoG+1st model again predicted half a period separation. The sN3+ model however yields a closely spaced peak-and-trough pair separated by about 1/4 period. These are Mach edges. It is evident from Fig. 9

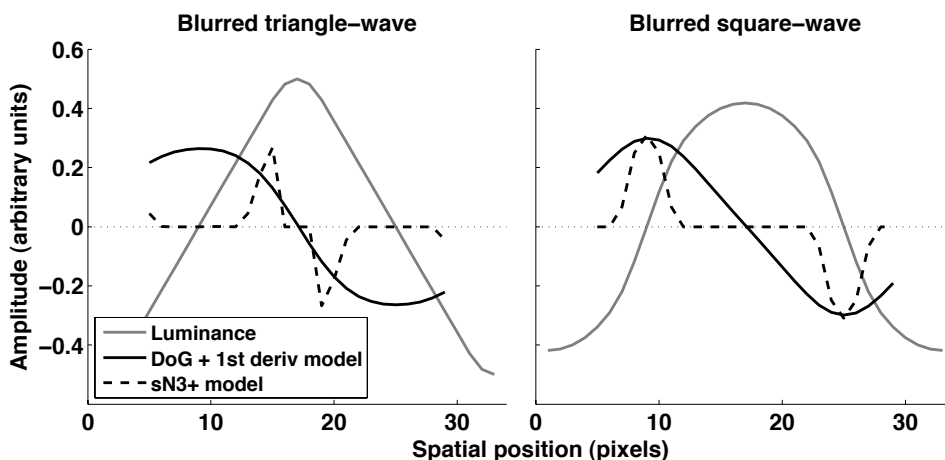


Figure 9: Experiment 2. One cycle of the luminance waveforms at blurwidth 4, with predictions of the DoG+1st derivative model and the sN3+ model. [Note: For the latter, positive values are the output of the nonlinear cascade of filters that detects dark-to-light edges (Fig. 2, right); negative values are the inverted output of the opposite channel whose 1st derivative stage is inverted, and which thereby detects light-to-dark edges.] The DoG+1st model did not predict Mach edges for this experiment.

that the spatial arrangement of edges seen in both the square- and triangle-wave cases is explained by extrema in the sN3+, while the DoG+1st derivative fails to account for perception in the triangle-wave case.

4. Discussion

Experiment 1 showed that observers reliably marked the position and polarity of Mach edges. These are visible edges seen at points on a luminance waveform where there was no peak in the 1st derivative and no zero-crossing in the 2nd derivative. Standard edge detectors that are based on these derivative features must therefore have great difficulty in accounting for Mach edges, but we have shown that using peaks and troughs in the 3rd derivative to locate edges readily predicts the occurrence, location and polarity of the Mach edges. The predicted and observed locations agreed especially well when a plausible amount of Gaussian smoothing was introduced (Fig. 4) that might represent the scale of the most sensitive filter for these stimuli.

We found that one way of rescuing the 1st derivative approach was to introduce a centre-surround (DoG) filter (Fig. 6), perhaps representing pre-cortical filtering, before the gradient operator. But, in a critical test, this model did not predict the Mach edges that were reliably observed at high spatial frequencies in experiment 2, whereas the 3rd derivative model did do so. The ‘rescue’ of the DoG+1st model might be extended to higher spatial frequencies by making the pre-filter even smaller, but the cost of this ad hoc modification would be to make the model increasingly similar to a 3rd derivative, thus underlining our point that peaks in the luminance gradient are insufficient to account for human edge-finding.

Our main aim in this paper was to contrast edge-finding models based on different orders of spatial derivative: 1st, 2nd and 3rd. Of these candidates, it seems clear that only the 3rd derivative offers a straightforward account of Mach edges. Of course, the Mach edges do leave some ‘fingerprint’ in the lower derivatives: they sit at or near the corner points in the gradient profile and at or near the abrupt steps in the 2nd derivative (Fig. 3). But we emphasize (a) that those features in the 1st and 2nd derivative are not the ones that have been widely proposed as edge markers (namely, gradient peaks or 2nd derivative ZCs), and (b) that those features have to be made explicit in some way, and finding peaks in an appropriate higher derivative (the 3rd) seems a straightforward and general way to do so. It is general because, as well as locating Mach edges, the peaks and troughs in the 3rd derivative also locate the more familiar (sharp or blurred) step edges that do give rise to gradient peaks and ZCs (Fig. 2). Thus the 3rd derivative seems the most parsimonious because the same rule accounts for Mach edges and step edges, but this is not true for the 1st and 2nd derivatives where different rules would be needed.

The linear 3rd derivative operator produces too many peaks and troughs for step edges (Fig. 2, left), but this need not trouble us, because the two-stage nonlinear 3rd derivative (Georgeson et al., 2007) solves that problem in a physiologically plausible way (Fig. 2, right), and gives very accurate predictions about perceived edge location and blur, without introducing any other difficulties that we are aware of. For the Mach edge (blurred triangle) waveforms, even the linear 3rd derivative gives a single peak (or trough) at the observed edge (Fig. 3), and so the presence or absence of these interesting nonlinearities is immaterial for those waveforms. In a broader context, the half-wave rectifiers are crucial in making the 3rd derivative a viable and general basis for edge detection.

Acknowledgement

SAW was supported by an EPSRC doctoral training grant to Aston University.

References

- Babaud, J., Witkin, A.P., Baudin, M., & Duda, R.O. (1986). Uniqueness of the Gaussian kernel for scale-space filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8, 26-33.
- Bergholm, F. (1987). Edge focusing. *IEEE Transactions On Pattern Analysis and Machine Intelligence*, 9, 726-741.
- Burr, D.C., Morrone, M.C., & Spinelli, D. (1989). Evidence for edge and bar detectors in human vision. *Vision Research*, 29, 419-431.
- Campbell, F.W., & Gubisch, R.W. (1966). Optical quality of the human eye. *Journal of Physiology*, 186, 558-578.
- Canny, J. (1986). A computational approach to edge-detection. *IEEE Transactions On Pattern Analysis and Machine Intelligence*, 8, 679-698.
- Clark, J.J. (1989). Authenticating edges produced by zero-crossing algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11, 43-57.
- Croner, L.J., & Kaplan, E. (1995). Receptive fields of P-ganglion and M-ganglion cells across the primate retina. *Vision Research*, 35, 7-24.
- Elder, J.H., & Zucker, S.W. (1998). Local scale control for edge detection and blur estimation. *IEEE Transactions On Pattern Analysis and Machine Intelligence*, 20, 699-716.
- Field, D.J., & Nachmias, J. (1984). Phase reversal discrimination. *Vision Research*, 24, 333-340.
- Georgeson, M.A. (1992). Human vision combines oriented filters to compute edges. *Proceedings of the Royal Society of London Series B-Biological Sciences*, 249, 235-245.
- Georgeson, M.A., & Freeman, T.C.A. (1997). Perceived location of bars and edges in one-dimensional images: Computational models and human vision. *Vision Research*, 37, 127-142.
- Georgeson, M.A., May, K.A., Freeman, T.C.A., & Hesse, G.S. (2007). From filters to features: Scale-space analysis of edge and blur coding in human vision. *Journal of Vision*, 7, 1-21.
- Georgeson, M.A., & Turner, R.S.E. (1985). Afterimages of sinusoidal, square-wave and compound gratings. *Vision Research*, 25, 1709-1720.
- Hesse, G.S., & Georgeson, M.A. (2005). Edges and bars: where do people see features in 1-D images? *Vision Research*, 45, 507-525.
- Hubel, D.H., & Wiesel, T.N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195, 215-243.
- Koenderink, J.J. (1984). The Structure of Images. *Biological Cybernetics*, 50, 363-370.
- Korn, A.F. (1988). Toward a symbolic representation of intensity changes in images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10, 610-625.
- Kulikowski, J.J., & King-Smith, P.E. (1973). Spatial arrangement of line, edge and grating detectors revealed by subthreshold summation. *Vision Research*, 13, 1455-1478.
- Legge, G.E. (1978). Sustained and transient mechanisms in human-vision - temporal and spatial properties. *Vision Research*, 18, 69-81.
- Levi, D.M., & Klein, S.A. (1990a). Equivalent intrinsic blur in amblyopia. *Vision Research*, 30, 1995-2022.
- Levi, D.M., & Klein, S.A. (1990b). Equivalent intrinsic blur in spatial vision. *Vision Research*, 30, 1971-1993.
- Lindeberg, T. (1990). Scale-space for discrete signals. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 234-254.
- Lindeberg, T. (1994). *Scale-space theory in computer vision*. Dordrecht: Kluwer.
- Lindeberg, T. (1998). Edge detection and ridge detection with automatic scale selection. *International Journal of Computer Vision*, 30, 117-154.
- Marr, D., & Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London Series B-Biological Sciences*, 207, 187-217.
- Mather, G., & Morgan, M. (1986). Irradiation - implications for theories of edge localization. *Vision Research*, 26, 1007-1015.
- May, K.A., & Georgeson, M.A. (2007). Blurred edges look faint, and faint edges look sharp: The effect of a gradient threshold in a multi-scale edge coding model. *Vision Research*, 47, 1705-1720.
- Morrone, M.C., & Burr, D.C. (1988). Feature detection in human vision - a phase-dependent energy model. *Proceedings of the Royal Society of London Series B - Biological Sciences*, 235, 221-245.
- Pollen, D.A., & Ronner, S.F. (1981). Phase-relationships between adjacent simple cells in the visual cortex. *Science*, 212, 1409-1411.

- Ringach, D.L. (2002). Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of Neurophysiology*, 88, 455-463.
- Sarkar, S., & Boyer, K.L. (1991). On optimal infinite impulse response edge detection filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13, 1154-1171.
- Shapley, R.M., & Tolhurst, D.J. (1973). Edge detectors in human vision. *Journal of Physiology-London*, 229, 165-183.
- ter Haar Romeny, B. (2003). *Front-end vision and multi-scale image analysis*. Dordrecht: Kluwer.
- Watt, R.J., & Morgan, M.J. (1985). A theory of the primitive spatial code in human vision. *Vision Research*, 25, 1661-1674.
- Williams, D.R., Brainard, D.H., McMahon, M.J., & Navarro, R. (1994). Double-pass and interferometric measures of the optical quality of the eye. *Journal of the Optical Society of America A*, 11, 3123-3135.
- Young, R.A. (1985). *The Gaussian derivative theory of spatial vision: analysis of cortical cell receptive field line-weighting profiles*. Warren, Mich.: General Motors Research Laboratories.
- Young, R.A., & Lesperance, R.M. (2001). The Gaussian Derivative model for spatial-temporal vision: II. Cortical data. *Spatial Vision*, 14, 321-389.
- Yuille, A.L., & Poggio, T.A. (1986). Scaling theorems for zero crossings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8, 15-25.
- Zhang, W., & Bergholm, F. (1997). Multi-scale blur estimation and edge type classification for scene analysis. *International Journal of Computer Vision*, 24, 219-250.