

Estimates of Edge Detection Filters in Human Vision.

William McIlhagga,

Bradford school of Optometry and Vision Science,
University of Bradford,
Richmond Road,
Bradford Bd7 1DP,
England.

Email: w.h.mcilhagga@bradford.ac.uk

Abstract.

Edge detection is widely believed to be an important early stage in human visual processing. However, there have been relatively few attempts to map human edge detection filters. In this study, observers had to locate a randomly placed step edge in brown noise (the integral of white noise) with a $1/f^2$ power spectrum. Their responses were modelled by assuming the probability the observer chose an edge location depended on the response of their own edge detection filter to that location. The observer's edge detection filter was then estimated by maximum likelihood methods. The filters obtained were odd-symmetric and similar to a derivative of Gaussian, with a peak-to-trough width of 0.1 to 0.15 degrees. These filters are compared with previous estimates of edge detectors in humans, and with neurophysiological receptive fields and theoretical edge detectors.

1. Introduction.

Edges are an important feature of the retinal image because they indicate the position of object boundaries and shadows. For that reason, edge detection has long been considered a vital first step in visual processing. Neurons sensitive to edges are common in the visual cortex (Hubel & Wiesel, 1962, 1968) and their receptive fields have been mapped in detail. However, less has been done to map the “receptive fields” or templates that underlie the detection of edges in humans. Previous psychophysical investigations of edge detectors have used indirect methods, such as subthreshold summation (Kulikowski & King-Smith, 1973; Shapley & Tolhurst, 1973); or have concentrated on demonstrating the existence of odd-symmetric detectors without characterizing their spatial properties (Burr, Morrone, & Spinelli, 1989; Stromeyer & Klein, 1974). Here I use a method based on classification images (Murray, 2011) to map the templates used in edge detection and localization.

Classification images were introduced by Beard and Ahumada (1998). The idea is that when noise is added to a stimulus, that noise sometimes takes on the aspect of what the observer is looking for when they perform a psychophysical task. By correlating observer responses with the noise, it is possible to determine what observers are really looking for when they perform a visual task.

Typically, the averaged noise over one response type (yes, or correct) and the averaged noise of the other response type (no, or incorrect) are subtracted to form an image of the points in the stimulus the observer uses to perform the task (Murray, Bennett, & Sekuler, 2002). This is equivalent to the Fisher discriminant, hence the name “classification” image (because the Fisher discriminant is a tool for statistical classification). Here, however, we use a more general maximum likelihood technique. Nonetheless, we will still refer to the estimated observer templates as classification images.

In the experiment, observers had to detect and locate a horizontal step edge by clicking a mouse at its perceived location. The step edge was embedded in brown noise with frequency spectrum proportional to $1/f^2$. Brown noise was used because it is ecologically relevant (natural images have a $1/f^2$ power spectrum (Burton & Moorhead, 1987; Field, 1987)), and because, unlike white noise, it is the kind of noise that yields *localized* optimal edge detectors (McIlhagga, 2011). The probability that the observer clicked at a particular location was assumed to be a function of the edge detector output at that location. Using maximum likelihood estimation, the filter that best fitted the observer responses was estimated. The filters that were found are like derivative of Gaussian filters, with a peak-to-trough width of 0.1 to 0.15 degrees. These filters are similar to those found by Shapley and Tolhurst (1973).

2. Methods

2.1 Experimental Procedure

On each trial, observers were shown a 10 degree tall and 4.5 degree wide stimulus consisting of a horizontal step edge embedded in horizontal brown noise. The step edge was always dark above and light below, and could appear anywhere in the central vertical 5 degrees of the stimulus. The brown noise was generated by a cumulative sum of white noise samples with a standard deviation of 0.002 in contrast units, where the contrast of a point with luminance L is given by $L/L_{mean} - 1$. That is, the brown noise at scan line y is given by $\sum_{j < y} w_j$, where w_j is a sample of white Gaussian noise for scanline j . The brown noise was then shifted so that the mean was zero. An example stimulus is shown in Figure 1.

The stimulus stayed on screen until the observer moved a mouse pointer to click where they believed the edge to be. If the observer clicked within 0.25 degrees of the true edge location, they were deemed correct. The edge contrast was controlled by a staircase. If the observer was deemed correct twice in a row, the edge contrast was reduced by 20%; if deemed incorrect once it was increased by 25%. This staircase was used to control the contrast of the edge to a point where the task was moderately difficult, and not for the purpose of threshold calculation. Following the observer's response, there was a 1 second delay before the next stimulus was presented.

Five observers A, C, H, T, and W participated in the experiment (W is the author). All were aware of the purpose of the experiment. Observers C, H, and W also collected data using white noise instead of brown noise (observers A and T were unavailable for the white noise experiment). A full set of data was collected over a few days, in 12 experimental blocks consisting of 150 trials. At the beginning of each block, observers were shown a high contrast step edge without noise, so they knew what they were looking for. The first 10 trials in each block were discarded prior to analysis. The experiment complied with University of Bradford Ethics Procedures and was conducted in accordance with the Code of Ethics of the World Medical Association (Declaration of Helsinki).

2.2 Calibration & Apparatus.

Stimuli were displayed on a Sony Multiscan E450 CRT monitor driven by a Bits++ device in Colour++ mode (Cambridge Research Systems Ltd. Kent, UK). In Colour++ mode, adjacent 8-bit pixels in the frame buffer are paired to yield 16 bits per pixel for each electron gun, and the 14 most significant bits are passed to a D/A converter. Stimuli were calculated and displayed by Matlab (MATLAB Release 2007b, The MathWorks, Inc., Natick, Massachusetts, United States), using the Psychophysics Toolbox (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997). The gamma of the monitor was

measured using a ColorCal meter (Cambridge Research Systems, Kent, UK) and linearized with a lookup table. Display resolution was 1024 by 768 pixels, and the monitor was viewed at a distance of 1 metre. Angular resolution was 50.86 pixels per degree.

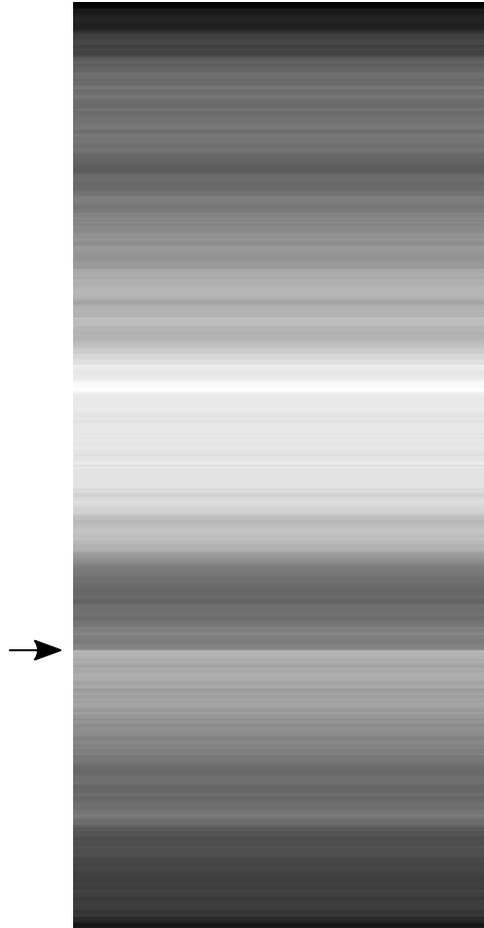


Figure 1: Example stimulus. The stimulus was 10 degrees tall and 4.5 degrees wide. A step edge (marked here by the arrow) could appear anywhere in the central vertical 5 degrees.

2.3 Data Analysis

Observer responses were analysed by assuming that they first convolved the stimulus with an edge detection filter f_x to yield a response

$$r_x^{(i)} = s_x^{(i)} * f_x$$

Here $r_x^{(i)}$ is the filter response at position x on trial i , $*$ indicates convolution, and $s_x^{(i)}$ is the stimulus contrast at position x in that trial. This is diagrammed in Figure 2 (a) and (b).

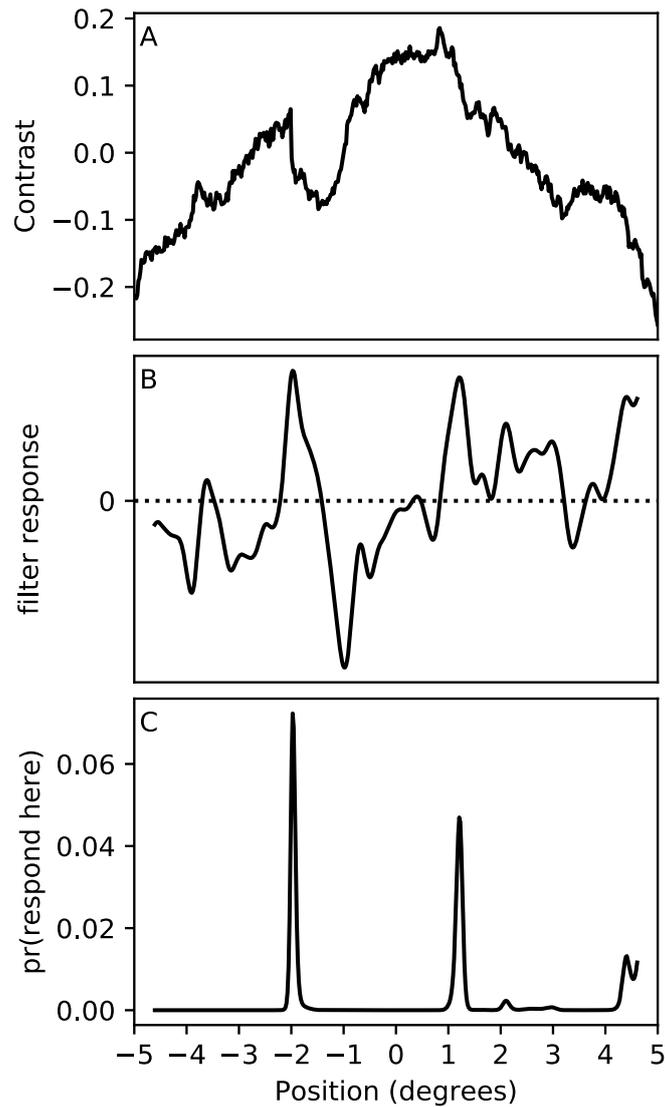


Figure 2: The edge detection model. Panel A (top) shows the contrast of a step edge embedded in brown noise as a function of position. This is the contrast profile of the stimulus in Figure 1. This contrast profile is convolved with an edge detection filter to yield a filter response shown in Panel B (middle). The filter response is transformed into a probability that the observer locates the edge by applying a softmax function (panel C, bottom). The true edge location is at -2 degrees, and this is the most likely response for this filter, but a response at about 1.2 degrees is also possible.

The most likely location for the edge is the point where the response $r_x^{(i)}$ is maximized. However, the filter response in Figure 2(b) has several local maxima, and it is possible that the observer might instead choose a local maximum instead of the global one. Thus, rather than being entirely deterministic, it is assumed that chosen location is selected randomly, based on the filter response. If the observer chose location $e^{(i)}$ on trial i , we assumed that the probability of this is given by the “softmax” transformation of filter outputs

$$p_i = pr(\text{observer chose location } e^{(i)} \text{ on trial } i) = \frac{\exp(r_e^{(i)})}{\sum_x \exp(r_x^{(i)})}$$

The softmax function, on the right hand side of this equation, is a generalization of the logistic function typically used when there is a multinomial response (McCullagh & Nelder, 1989), as is the case here, because the observer response is one out of a possible 510 pixel locations. The probabilities are shown in Figure 2(c) based on the responses in Figure 2(b). The true edge location is the most likely choice, but there is an alternative location with nearly the same probability. Note that the probabilities for each location are small, because there are so many possible locations, but the overall probability for a click in the region of -2 degrees in Figure 2(c) is 0.42.

The log-likelihood of the observer's responses over the entire set of trials is given by

$$\mathcal{L} = \sum_i \log p_i$$

Since this log-likelihood is a function of the filter values f_x , we can estimate the filter by maximizing this log-likelihood (or equivalently minimizing the negative log-likelihood). The maximum likelihood estimate of the filter \hat{f}_x will sometimes be referred to as the classification image.

2.4 Computational Considerations.

Efficient optimization of a function like \mathcal{L} requires the gradient $\partial\mathcal{L}/\partial f_x$ of \mathcal{L} with respect to the elements of the filter f_x . In the absence of the gradient, the optimization will be extremely slow, perhaps prohibitively so. Unfortunately, the gradient of \mathcal{L} in this case is very difficult to derive, either algebraically or symbolically.

However, automatic differentiation libraries allow us to compute the gradient of \mathcal{L} precisely and efficiently. I used the autograd library (Maclaurin, Duvenaud, & Adams, 2015) for Python. With this library, you simply write the likelihood function in Python, and the library creates a new function which computes the gradient of the likelihood. The likelihood function and the automatically computed gradient can be used with the SciPy optimization library (Jones, Oliphant, Peterson, & others, 2001) to efficiently optimize the log likelihood \mathcal{L} with respect to the filter values f_x . Each optimization took only a few minutes on a Microsoft Azure Notebook.

The full optimization code and the Azure "library", including all data, can be found at <https://notebooks.azure.com/william-mc/libraries/edge-detection>. This can be "cloned" if you have a free Azure Notebooks account. The code is in the file named "edge_templates.ipynb" in the library.

Optimization of \mathcal{L} yielded reasonably smooth filter estimates for observers H, T, and W, but the filters for observers A and C were more jagged and difficult to interpret. Accordingly, a smoothness penalty (Hastie & Tibshirani, 1986), given by

$$\pi = \frac{\sum_x (f_{x-1} - 2f_x + f_{x+1})^2}{\sum_x f_x^2}$$

was imposed on observer A and C. The numerator of this penalty is the sum of squares of the second derivative of the filter. The smoothness is modified by dividing by the power of the filter to ensure that the penalty is unchanged when multiplying the filter values by a constant.

Filter values f_x were selected to maximize the penalized likelihood $\mathcal{L} - 25\pi$. Different penalty weights were tried and produced essentially the same results. For consistency, the same smoothness penalty was applied to observers H, T and W, although it made little difference to their filters.

3. Results.

Figure 3 plots the accuracy of observer clicks as a function of step edge contrast. Accuracy (measured by the median deviation of the mouse click from the true edge location) improves with contrast. This is most likely because the deviations come from a mixture of two distributions: a single atom at a deviation of zero, when the observer sees the edge and clicks on it, and a triangular distribution of deviations when the observer misses the true edge and picks a ‘false’ edge (a triangular distribution is the difference between a uniformly distributed edge location and a uniformly distributed incorrect mouse click). The atom and triangular distributions are mixed by the probability of seeing the true edge. At very low contrast, where the probability of seeing the edge is low, the distribution of deviations is mostly triangular. At high contrast, where the probability of seeing the edge is near 100%, the distribution of deviations is concentrated at zero. The distribution of deviations is further blurred by motor error in placing the mouse on screen. At very high contrasts, the observer almost always saw the true edge location, so the median deviation is mostly motor error.

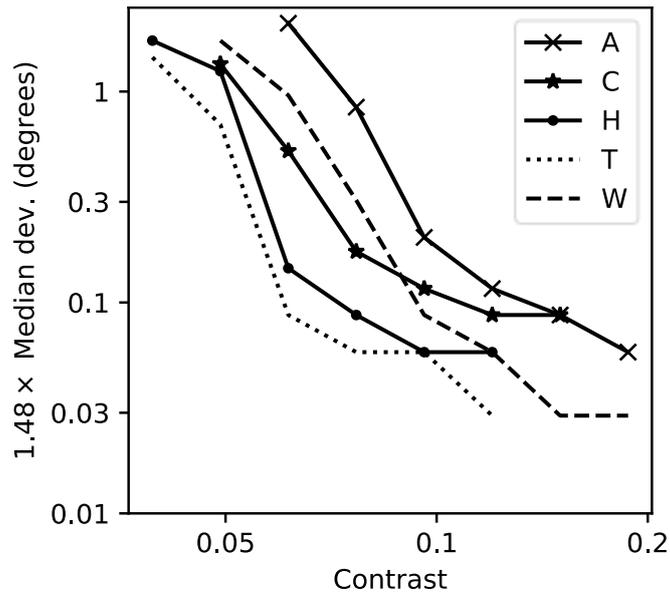


Figure 3: The median deviation between the observer's click location and the true edge location as a function of step edge contrast. The factor of 1.48 converts the median deviation into a standard deviation when the clicks are distributed normally.

Figure 4 plots the estimated edge detection filters \hat{f}_x for the five observers used in the brown noise condition. The estimation algorithm requires one to select a filter size a priori and will estimate the best-fitting filter once that is given. Different filter sizes were tried but filters larger than 48 pixels wide simply contained larger flanking areas of noisy near-zero values. Thus, only the results for 48 pixel filters (about 0.94 degrees wide) are shown in Figure 4. All the filters in Figure 4 have a similar shape, roughly a derivative of Gaussian, with peak-to-trough widths of about 0.1 degrees (H, T and W) and 0.17 degrees (A and C). The heights of the filters, which are different for the different observers, depend on how well the filter accounts for observer responses, since higher filters yield a larger response and hence a more sharply peaked probability distribution of mouse clicks following the softmax transformation.

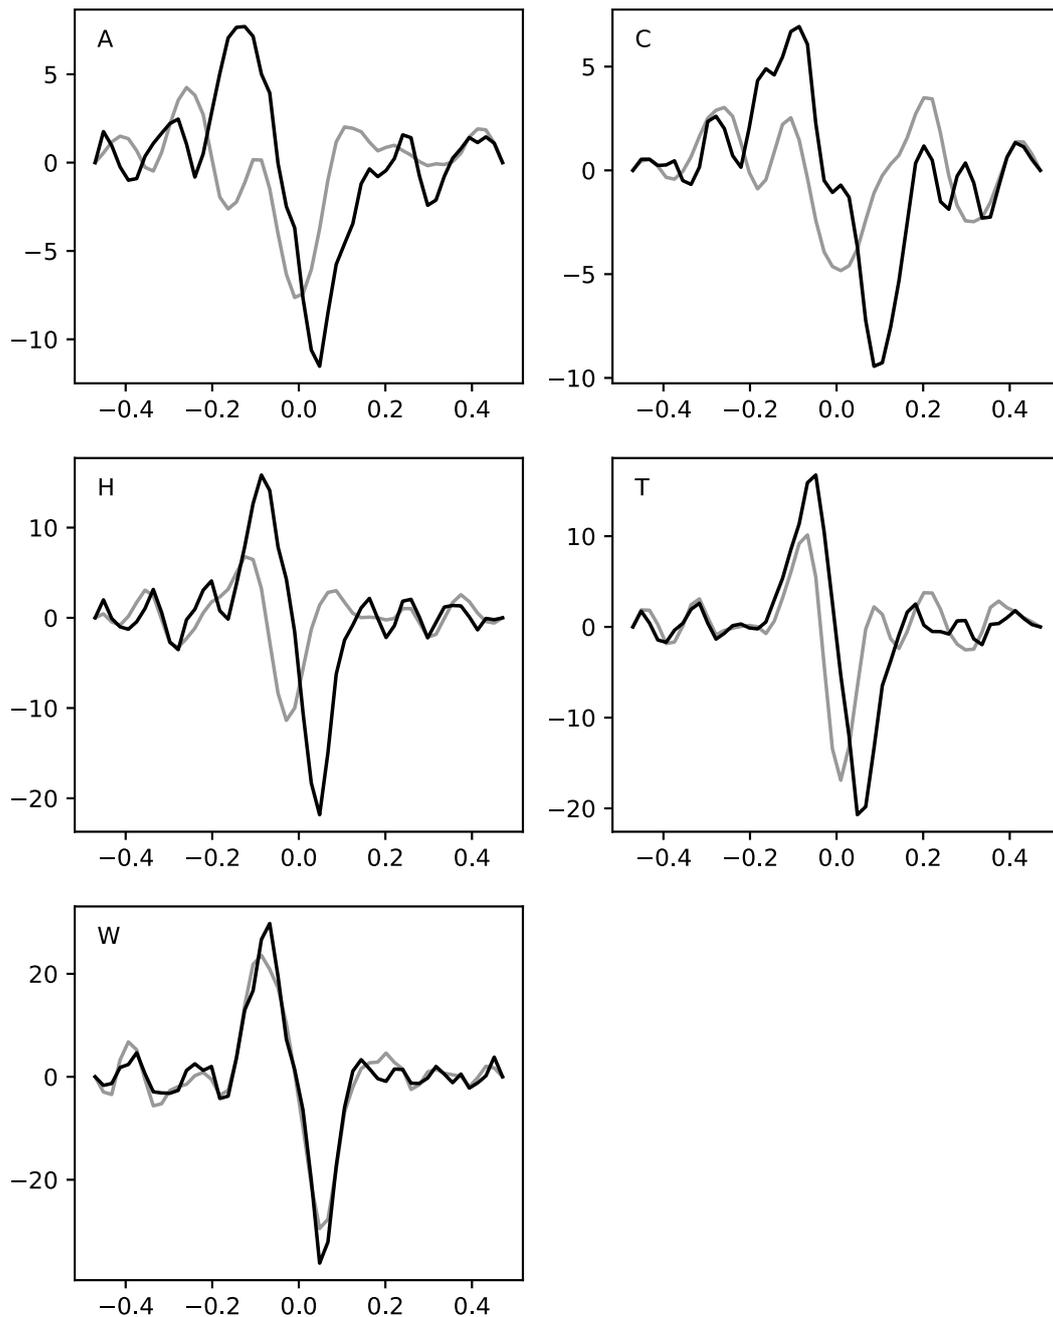


Figure 4: Estimated edge detection filters for all observers. Black lines give the filters estimated from all trials, and grey lines the filters estimated from just the wrong trials. Note different y axis scales in the three panels.

The edge detection filters are estimated from both correct and incorrect responses. Since the observers were often quite accurate, it is conceivable that the filters simply reflect what any reasonably performant system would do to detect edges, rather than specifically the human visual system. To check this, the analysis was redone using only the incorrect responses (that is, where the observer chose some location more than 0.25 degrees from the true edge location). The filters estimated from incorrect responses only (using the same smoothness penalty) are also shown in

Figure 4, as grey lines. For observers H, T and W, the incorrect-response filters are very similar to the all-response filters. For observers A and C, the incorrect filters are somewhat different. This may be due to more inconsistency in these observers, because some of their incorrect responses are difficult to understand on reviewing the stimulus.

An edge detector is explicitly a spatial filter whose purpose is to detect a spatially localized feature in an image. However, human spatial channels are often characterized in terms of their spatial frequency. To aid comparison with previous work, the amplitude spectra of the filters from Figure 4 are shown in Figure 5. In the Fourier domain, these filters are bandpass with a full width at half height around 3 cycles/degree, and a peak at 2-3 cycles/degree.

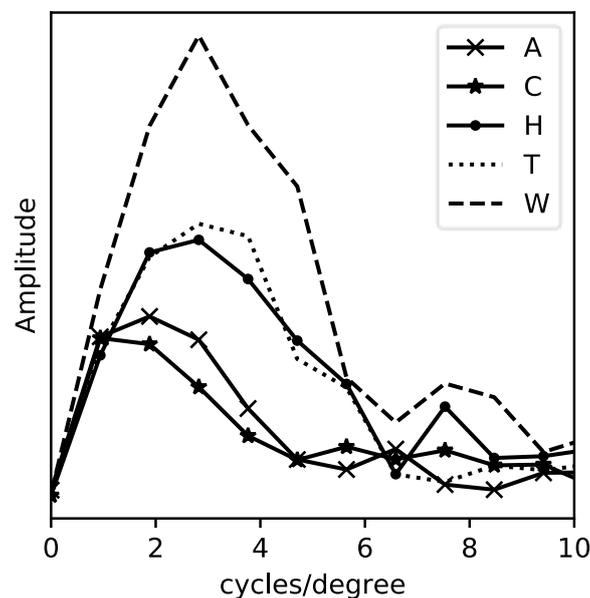


Figure 5: Fourier amplitude of the edge detection filters from Figure 4.

3.1 White Noise.

Although brown noise is the appropriate noise to use when trying to map localized edge detectors, for reasons given earlier, it is interesting to compare the template for brown noise with that found using white noise, and previously it has been shown that changes in noise correlation structure yields changes in classification images (Abbey & Eckstein, 2000). Only observers C, H, and W collected data using white noise. The accuracy of their responses is shown in Figure 6 (corresponding to Figure 3). Their filters estimated in white noise are shown in Figure 7, with brown noise filters included for comparison.

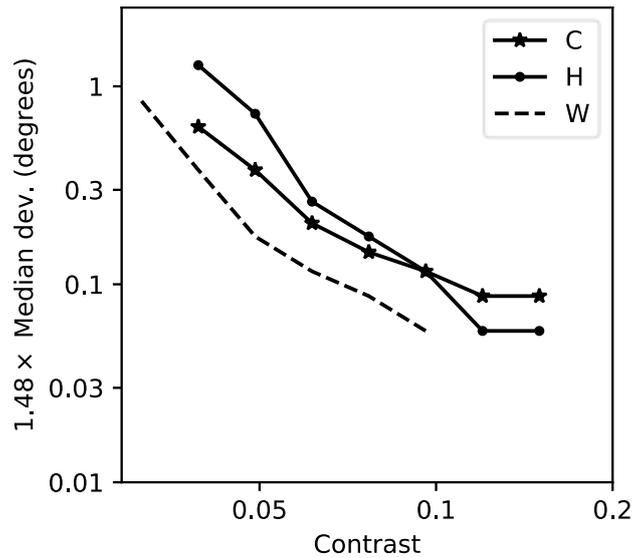


Figure 6: Observer accuracy in white noise (c.f. Figure 3).

The filter shapes shown in Figure 7 are blurred DISEF filters (McIlhagga, 2011; Shen & Castan, 1992). A DISEF filter (Derivative of the Infinite Symmetric Exponential Function) is given by

$$DISEF(x) \propto -\text{sign}(x) \exp\left(-\frac{|x|}{s}\right)$$

which is sharply discontinuous at $x = 0$. The filters in Figure 7 are not, so to fit them, the DISEF filter was blurred with a Gaussian function. The best fit blurred DISEF filters are shown by the dotted lines in Figure 7.

The DISEF filter is the optimal edge detection filter in a mixture of brown and white noise (McIlhagga, 2011). The scale factor s depends on the ratio of the white and brown noise power in the image, or more generally, the extent to which the image has a flat power spectrum or a $1/f^2$ power spectrum. A small value for s indicates more brown than white noise, while a large value indicates the opposite. In the limit of zero brown noise, $s = \infty$ and the DISEF filter becomes a step edge, which is the unconstrained optimal edge detector found by Canny (1986).

The DISEF filters in Figure 7 suggest that the observers are acting as if there is some component of image power with a $1/f^2$ profile. They are not completely wrong in this – the step edge they are trying to detect has a $1/f^2$ power spectrum. However, observer W has a smaller value for the DISEF scale s than the other two observers, so shape of the DISEF filter may be more influenced by an individual *a priori* bias towards assuming the presence of $1/f^2$ noise or power in images.

The blurring of the DISEF filter must be, in part at least, due to motor error. The best fitting values for the Gaussian blur standard deviation in the filters was 0.061° , 0.036° , and 0.039° for observers C,

T, and W respectively. This is consistent with, but about 0.7 times smaller than, the median errors of 0.087° , 0.058° , and 0.058° given in Figure 6 at the highest contrasts.

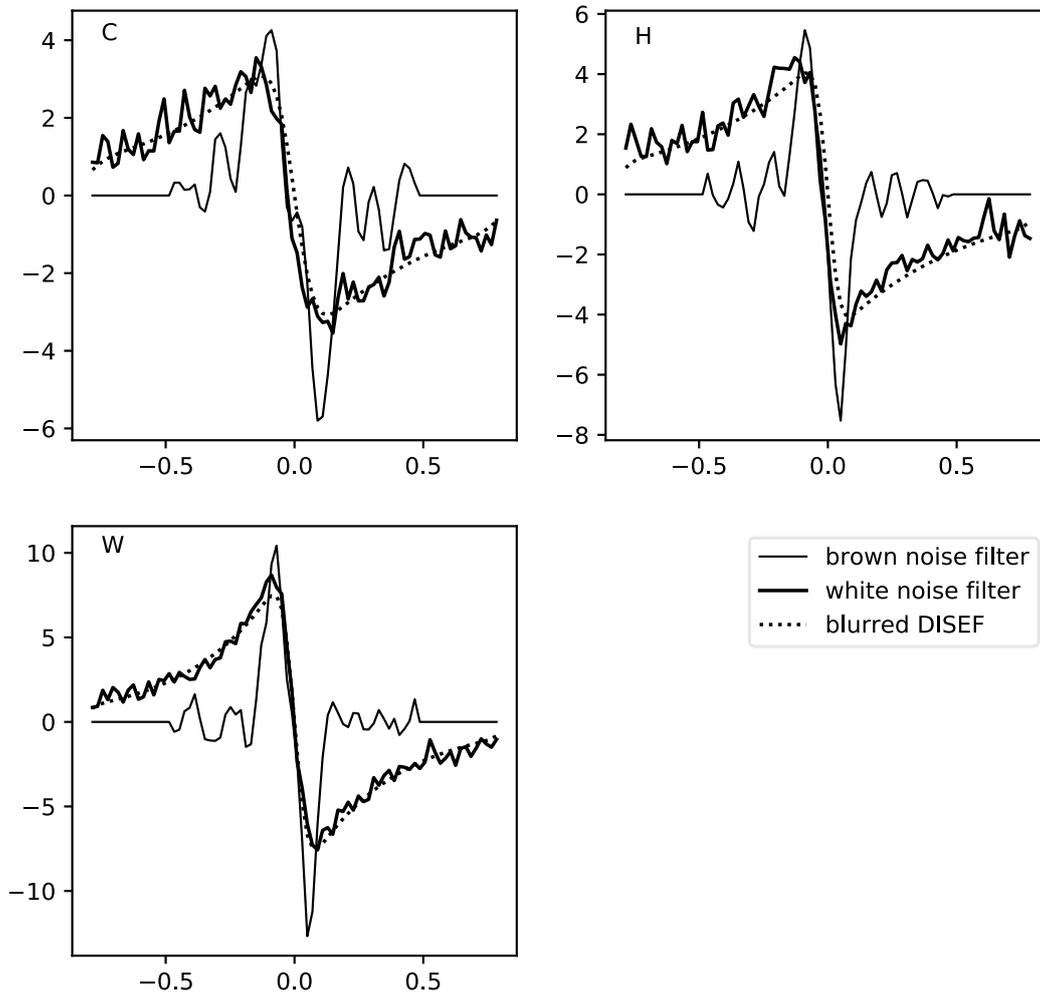


Figure 7: The edge detection filter in white noise (solid line) compared to the edge detection filter in brown noise (thin line) for observers C, H, and W replotted from Figure 4. The replotted brown noise filters are scaled for easy comparison. The dotted line shows the best-fit blurred DISEF filter, which is described in the text.

4. Discussion

We have shown that detection of step edges in human vision can be modelled by a linear filter, roughly the same shape as a derivative of Gaussian, with a peak-to-trough width of 0.1 to 0.17 degrees. How does this compare with previous studies?

4.1 Previous Psychophysical Estimates of Edge Detectors.

Previous psychophysical estimates of edge detection filters in humans have relied on indirect approaches. Shapley and Tolhurst (1973), in one of a series of experiments, measured edge

detection thresholds in the region of a subthreshold line. They found an asymmetrical map of edge sensitivity which can be interpreted as an edge detection filter with a peak to trough width in the region of 0.04 degrees. Figure 7 shows the line sensitivity profile (digitized from Shapley and Tolhurst (1973) Figure 3) with edge detectors of observers H, T and W scaled for comparison. The profile in Shapley and Tolhurst (1973) is narrower than the filters found here.

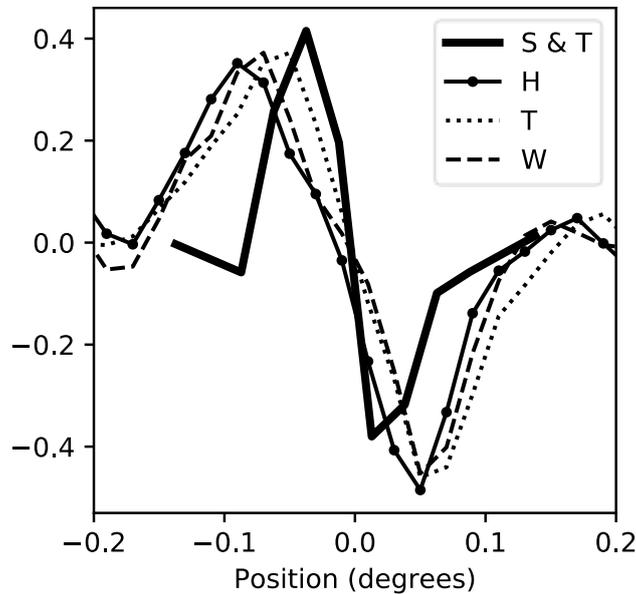


Figure 8: Comparison of the edge detection filters found in this study for observers H, T, and W (who had the narrowest filters) with the profile from Shapley & Tolhurst (1973). The Shapley and Tolhurst profile is given by the solid line labelled “S & T”; the observer filters by dashed and dotted lines. The edge detection filters were scaled to the height of the Shapley & Tolhurst profile for ease of comparison.

Kulikowski and King-Smith (1973) also used a subthreshold summation approach. In one experiment they measured edge detection thresholds in the presence of a subthreshold grating and obtained a grating sensitivity profile for the edge detector (their Figure 6). This peaked at 3 cycles/degree, the same as found here (my Figure 5). They also replicated Shapley and Tolhurst (1973), measuring the threshold of an edge in the presence of a subthreshold line. This gave them an edge detector profile with a peak to trough width of approx. 0.1 degrees (their Figure 7).

For a derivative of Gaussian filter with scale σ_f , the peak-to-trough width is $2\sigma_f$. Thus the filter found by Shapley and Tolhurst (1973) has a standard deviation of about 0.02 degrees, and that found by Kulikowski and King-Smith (1973) has a standard deviation of about 0.05 degrees. The peak-to-trough widths of the filters in Figure 4 are 0.18, 0.18, 0.14, 0.10, and 0.12 degrees, for observers A, C, H, T, and W respectively. However, these are contaminated by the motor error shown in Figure 3. If we assume that the motor error is Gaussian, with standard deviation σ_m , then a

Gaussian derivative filter with scale σ_f will appear to have a scale of $(\sigma_f^2 + \sigma_m^2)^{1/2}$, and a peak-to-trough width twice that. The motor errors from the highest contrast in Figure 3 are $\sigma_m = 0.058, 0.087, 0.058, 0.029,$ and 0.029 degrees respectively. From this, we can infer that the edge detection filter scales σ_f are $0.066, 0.013, 0.036, 0.039, 0.051$ respectively. The average of these is 0.041 degrees. This accords well with Kulikowski and King-Smith but is about twice as large as Shapley and Tolhurst's result.

Stromeyer and Klein (1974) used subthreshold summation of gratings at different phases to show that odd-symmetric filters were involved in detecting the gratings, but no sizes for these filters were given. Tolhurst and Dealy (1975) suggested edge (and bar) detectors were needed to detect their stimuli, but also did not give any filter sizes. Burr, Morrone, and Spinelli (1989) found evidence for even and odd symmetric filters in phase discrimination tasks, but again gave no filter sizes.

Elder and Sachs (2004) measured detection efficiency for edges of different width and aspect ratio in white noise. They found that a specific kind of multiscale edge detector model, using derivative of Gaussian filters with hyperbolic scaling of length and width, fitted their data best. The range of filter scales needed was from $\sigma_x = 7$ to 15 arc min. For derivative of Gaussian filters, the peak-to-trough widths are double the scale, or 14 to 30 arc min (0.23 to 0.5 degrees). These filter widths are larger than found here. This may be because they used white noise, as white noise yields wider filters, as shown in Figure 7, whereas brown noise will favour filters matched to the edge scale (McIlhagga, 2011).

McIlhagga and May (2012) derived classification images for blur discrimination of edges in white noise. The classification images looked like derivative of Gaussians, but were much narrower than found here, with filter scales of $\sigma_f = 0.018, 0.047,$ and 0.026 degrees, for observers KAM, TS, and WHM respectively. These are smaller than the filter widths inferred earlier, but the task is different. In McIlhagga and May (2012), observers had to distinguish the blur of two super-threshold edges. The resultant classification images in McIlhagga and May are, to a first approximation, a difference of two derivative of Gaussian functions at the same location but with different scales (both KAM and WHM show sidelobes which could occur from that differencing). The peak-to-trough width of such a difference would be smaller than the peak-to-trough width of the narrower Gaussian. However, this would also not be enough to explain the difference between their results and the current experiment; thus, it is likely there are even finer scales of edge detector than $\sigma_f = 0.05$ degrees, which simply do not appear in (or cannot be reliably estimated from) the present data.

Neri (2011) looked at global context effects on the shape of edge detectors using a classification image approach. In this study, the edges were embedded in a white noise patch which was inserted

into a manipulated natural image. The edge detectors that were inferred had a peak-to-trough width of about 0.3 degrees, two or three times larger than the filters plotted in Figure 4, but perhaps a similar size to the white noise filters shown in Figure 7. The task used by Neri is somewhat different to a straight detection task, however; observers had to decide whether the orientation of the edge probe was the same or different to the orientation of the surrounding image context. In addition, the edge patches in Neri's study were presented somewhat peripherally. In the current study, observers could move their eyes over the entire stimulus, and accurately clicking on the edge location requires foveation of the edge. Regardless, Neri (2011) found powerful effects on edge detector shape that cannot be measured using the methods in the present study.

4.2 Comparison with Cortical Receptive Fields.

There is of course no necessary relationship between the filters found in this study and individual V1 or V2 receptive fields (Neri & Levi, 2006). Some electrophysiology studies have shown that the receptive fields of many simple cells are more complicated than just a linear filter, so a straightforward comparison with these sorts of simple cells is impossible. Nonetheless, it is interesting to compare the filters found here with receptive field in primate V1.

Hubel & Weisel (1962, 1968) of course, started the idea that neurons in V1 could be thought of as edge detectors. However, they did not provide any quantitative measurements of receptive field size. Foster, Gaska, Nagler, & Pollen (1985) found the most common peak spatial frequency sensitivity of V1 parafoveal neurons in monkey was around 3 cycles/degree, and the receptive fields were usually 1.5 - 2 cycles wide. The peak accords with the amplitude spectrum in Figure 5, but the filters found here don't have any sidebands, and are only about 1 cycle wide.

Hawken & Parker (1987) modelled V1 receptive fields of simple cells by separated difference of Gaussians. The separation between the Gaussians is roughly equivalent to the peak-to-trough width of the filter. They found separations in the range 2 – 12 arc min (0.03 to 0.2 degrees) which encompasses the filter widths found here.

4.3 Computational Theories.

A number of models of edge detection have been proposed for human vision (Morgan, 2011; Watt & Morgan, 1985) most notably the N1 and N3+ models of Georgeson and colleagues (Georgeson, 1994; Georgeson & Freeman, 1997; Georgeson, May, Freeman, & Hesse, 2007; Hesse & Georgeson, 2005). These models suggest that humans analyse the image with a set of multiscale filters. The sizes of the individual filters aren't very important, so long as they cover the scales of edges in natural images. This study only finds a single filter, but that is because observers only had to look for a step

edge which has a single scale. Possibly a different experimental design might enable us to find filters at different scales.

5. Conclusion.

This study has provided direct evidence for edge detection filters in human vision using a classification-image or reverse-correlation methodology. The filters found are consistent with edge detection filters found previously using indirect methods, and with receptive field sizes in macaque V1. The width of the filters is slightly larger than the filters found in some previous studies (McIlhagga & May, 2012; Shapley & Tolhurst, 1973) but consistent with others (Kulikowski & King-Smith, 1973). This could be because our measure of motor error is simply not accurate enough to reveal the smallest scales, because the task recruited larger scale filters as well as small scale ones (perhaps because observers were not required to maintain fixation), or because the smallest scale filters have a poorer sensitivity than slightly larger ones, and so are not picked up by the estimation process.

This study has also shown that with the appropriate software, complex models of human vision are very easy to fit to data.

6. References

- Abbey, C. K., & Eckstein, M. P. (2000). Estimates of human-observer templates for a simple detection task in correlated noise. In *Medical Imaging 2000: Image Perception and Performance* (Vol. 3981, pp. 70–78). International Society for Optics and Photonics.
<https://doi.org/10.1117/12.383092>
- Beard, B. L., & Ahumada, A. J. (1998). Technique to extract relevant image features for visual tasks. In *Proceedings of SPIE* (pp. 79–85). San Jose, CA, USA. <https://doi.org/10.1117/12.320099>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436.
<https://doi.org/10.1163/156856897X00357>
- Burr, D. C., Morrone, M. C., & Spinelli, D. (1989). Evidence for edge and bar detectors in human vision. *Vision Research*, 29(4), 419–431. [https://doi.org/10.1016/0042-6989\(89\)90006-0](https://doi.org/10.1016/0042-6989(89)90006-0)
- Burton, G. J., & Moorhead, I. R. (1987). Color and spatial structure in natural scenes. *Applied Optics*, 26(1), 157–170. <https://doi.org/10.1364/AO.26.000157>

- Canny, J. (1986). A Computational Approach to Edge Detection. *IEEE Trans. Pattern Analysis Machine Intelligence*, 8(6), 679–698. <https://doi.org/10.1109/TPAMI.1986.4767851>
- Elder, J. H., & Sachs, A. J. (2004). Psychophysical receptive fields of edge detection mechanisms. *Vision Research*, 44(8), 795–813. <https://doi.org/10.1016/j.visres.2003.11.021>
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America. A, Optics and Image Science*, 4(12), 2379–2394. <https://doi.org/10.1364/JOSAA.4.002379>
- Foster, K. H., Gaska, J. P., Nagler, M., & Pollen, D. A. (1985). Spatial and temporal frequency selectivity of neurones in visual cortical areas V1 and V2 of the macaque monkey. *The Journal of Physiology*, 365, 331–363. <https://doi.org/10.1113/jphysiol.1985.sp015776>
- Georgeson, M. A. (1994). From filters to features: location, orientation, contrast and blur. *Ciba Foundation Symposium*, 184, 147–165; discussion 165-169, 269–271. <https://doi.org/10.1167/7.13.7>
- Georgeson, M. A., & Freeman, T. C. (1997). Perceived location of bars and edges in one-dimensional images: computational models and human vision. *Vision Research*, 37(1), 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(13), 1–21. <https://doi.org/10.1167/7.13.7>
- Hastie, T., & Tibshirani, R. (1986). Generalized Additive Models. *Statistical Science*, 1(3), 297–318.
- Hawken, M. J., & Parker, A. J. (1987). Spatial Properties of Neurons in the Monkey Striate Cortex. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 231(1263), 251–288. <https://doi.org/10.1098/rspb.1987.0044>
- Hesse, G. S., & Georgeson, M. A. (2005). Edges and bars: where do people see features in 1-D images? *Vision Research*, 45(4), 507–525. <https://doi.org/10.1016/j.visres.2004.09.013>
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, 160(1), 106-154.2.

- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, *195*(1), 215–243.
<https://doi.org/10.1113/jphysiol.1968.sp008455>
- Jones, E., Oliphant, E., Peterson, P., & others. (2001). SciPy: Open Source Scientific Tools for Python. Retrieved March 15, 2018, from
- Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? In *Perception ECVF Abstract Supplement* (Vol. 36).
- Kulikowski, J. J., & King-Smith, P. E. (1973). Spatial arrangement of line, edge and grating detectors revealed by subthreshold summation. *Vision Research*, *13*(8), 1455–1478.
[https://doi.org/10.1016/0042-6989\(73\)90006-0](https://doi.org/10.1016/0042-6989(73)90006-0)
- Maclaurin, D., Duvenaud, D., & Adams, R. (2015). Autograd: Reverse-mode differentiation of native python. In *ICML workshop on Automatic Machine Learning*. Retrieved from
<https://github.com/HIPS/autograd>
- McCullagh, P., & Nelder, J. A. (1989). *Generalized Linear Models, Second Edition*. Taylor and Francis.
- McIlhagga, W. (2011). The Canny Edge Detector Revisited. *International Journal of Computer Vision*, *91*, 251–261. <http://dx.doi.org/10.1007/s11263-010-0392-0>
- McIlhagga, W., & May, K. A. (2012). Optimal edge filters explain human blur detection. *Journal of Vision*, *12*(10). <https://doi.org/10.1167/12.10.9>
- Morgan, M. J. (2011). Features and the 'primal sketch'. *Vision Research*, *51*(7), 738–753.
<https://doi.org/10.1016/j.visres.2010.08.002>
- Murray, R. F. (2011). Classification images: A review. *Journal of Vision*, *11*(5).
<https://doi.org/10.1167/11.5.2>
- Murray, R. F., Bennett, P. J., & Sekuler, A. B. (2002). Optimal Methods for Calculating Classification Images: Weighted Sums. *Journal of Vision*, *2*(1). <https://doi.org/10.1167/2.1.6>
- Neri, P. (2011). Global Properties of Natural Scenes Shape Local Properties of Human Edge Detectors. *Frontiers in Psychology*, *2*. <https://doi.org/10.3389/fpsyg.2011.00172>

- Neri, P., & Levi, D. M. (2006). Receptive versus perceptive fields from the reverse-correlation viewpoint. *Vision Research*, 46(16), 2465–2474. <https://doi.org/10.1016/j.visres.2006.02.002>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10(4), 437–442. <https://doi.org/10.1163/156856897X00366>
- Shapley, R. M., & Tolhurst, D. J. (1973). Edge detectors in human vision. *The Journal of Physiology*, 229(1), 165–183.
- Shen, J., & Castan, S. (1992). An optimal linear operator for step edge detection. *CVGIP: Graphical Models and Image Processing*, 54(2), 112–133. [https://doi.org/10.1016/1049-9652\(92\)90060-B](https://doi.org/10.1016/1049-9652(92)90060-B)
- Stromeyer, C. F., & Klein, S. (1974). Spatial frequency channels in human vision as asymmetric (edge) mechanisms. *Vision Research*, 14(12), 1409–1420. [https://doi.org/10.1016/0042-6989\(74\)90016-9](https://doi.org/10.1016/0042-6989(74)90016-9)
- Tolhurst, D. J., & Dealy, R. S. (1975). The detection and identification of lines and edges. *Vision Research*, 15(12), 1367–1372. [https://doi.org/10.1016/0042-6989\(75\)90192-3](https://doi.org/10.1016/0042-6989(75)90192-3)
- Watt, R. J., & Morgan, M. J. (1985). A theory of the primitive spatial code in human vision. *Vision Research*, 25(11), 1661–1674. [https://doi.org/10.1016/0042-6989\(85\)90138-5](https://doi.org/10.1016/0042-6989(85)90138-5)