A RANK BASED EDGE ENHANCEMENT FILTER

Donald G Bailey
Image Analysis Unit, Massey University, Palmerston North, New Zealand

ABSTRACT

Edge enhancement is an operation commonly employed in image processing. Usually the edge slope is steepened by the applying a linear filter that amplifies the high spatial frequency content of an image. This often results in the amplification of noise and in ringing or over-enhancement of edges that are already sharp.

The rank filter is a nonlinear filter. The output value for a pixel is obtained by ranking the pixel values within a neighbourhood in the input image and using the value from a specified rank position. Near an edge, low and high rank values represent of the dark and light sides of the edge respectively. Edges may be enhanced by specifying two rank positions and assigning the output value to that of the rank position which is closer in intensity to the centre pixel.

Strongest edge enhancement results from using extreme rank positions while noise sensitivity may be reduced by using rank positions close to the median. Excessive ringing is avoided since no new intensity values are generated. Successive applications of a rank-based edge enhancement filter converge to a stable root image.

INTRODUCTION

In image processing a local filter is any operator whose output for a pixel is a function of the input values within the neighbourhood of that pixel [1]. This neighbourhood can be thought of as a window that is scanned across the input image; each position contributes one pixel to the output image. Although the window may be any shape, it is usually square.

One useful nonlinear filter is the rank filter [2]. This filter works by ordering the N pixels within the window according to intensity, that is $(f_1, f_2, f_3, ..., f_N)$ where $f_1 \le f_2 \le f_3 \le ... \le f_N$ and selecting the intensity value from a specific position in this ranked list: $\text{rank}(i) = f_i$

An interesting property of rank filters from the point of view of edge enhancement is that near an edge, low rank values correspond to pixels from the dark side of the edge, while high rank values correspond to pixels from the light side of the edge. In fact, this property may be used to detect edges by taking the subtracting the intensity values for two selected positions [3].

DERIVATION OF EDGE ENHANCEMENT FILTER

When the edges within an image are blurred, whether by the optical system forming the image, or because of the sampling process, or as a result of processing to remove noise, it is often desirable to sharpen or enhance the edges in the image. One method which has been proposed [2] for this is to compare the original, centre pixel value with each of two rank values and selecting the nearer rank value for the output pixel:

$$\text{if } |f_j - f_{centre}| < |f_i - f_{centre}| \text{ then output} = f_j$$
$$\text{if } |f_j - f_{centre}| > |f_i - f_{centre}| \text{ then output} = f_i$$

When the window is close to an edge, the two rank values are considered to represent the intensities of the regions on either side of the edge. This operation assigns the centre pixel to the value of the region that it is most similar to in intensity. A problem is encountered if the original pixel value is exactly half way between the two selected rank values. In this case, a number of options are available:

- leave the value unchanged.
- use the median of the window as the output.
- select either f_i or f_j arbitrarily.
- select f_i if the original is odd or f_j if the original is even.

In this paper, the third option has been used. The selection is therefore:

if $|f_j - f_{centre}| < |f_i - f_{centre}|$ then output = f_j

if $|f_j - f_{centre}| \geq |f_i - f_{centre}|$ then output = f_i

This slightly favours the dark side of the edge, but is the easiest to implement. From this point of view, this filter may be thought of as a form of classification. It classifies pixels in the blurred area between two regions of different intensity into one of the two regions. This filter may be represented as $EN_N(j,i)$.

TEST IMAGES

The only other filter commonly used for edge enhancement is the linear filter, although gated linear filters may also be used [4]. Linear edge enhancement filters amplify the high spatial frequencies. This is usually achieved by subtracting a blurred version of the image from the original; a technique known as ‘unsharp masking’. The rank-based filter will be compared with the following two linear edge enhancement filters:

A) $\frac{1}{4}\begin{bmatrix} -1 & -1 & -1 \\ -1 & 12 & -1 \\ -1 & -1 & -1 \end{bmatrix}$ B) $\frac{1}{8}\begin{bmatrix} -1 & -1 & -1 \\ -1 & 16 & -1 \\ -1 & -1 & -1 \end{bmatrix}$

To determine and quantify the relative merits of the linear and rank-based filters, they were tested on images of blurred edges that had noise added. The test images were 10000 rows by 16 columns in size and contained a single step edge running vertically down the centre of the image. This edge image was blurred using a linear low pass filter and then Gaussian noise was added. Images of this form allow the effects of the edge enhancement to be determined easily by obtaining intensity statistics of each column in the image.

When the edges are enhanced, there are four parameters of the filtered images that are of particular interest. These may all be determined from the column statistics of the image:

- the amount of blur remaining in the edge (how well the edge was enhanced)
- the extent to which the edge is overenhanced (overshoot or ringing)
- the signal to noise ratio of the image away from the edge (noise amplification)
- the signal to noise ratio in the vicinity of an edge (effect of the edge on noise properties).

Edge blur is defined in this paper as the area between a blurred edge and an ideal edge, as illustrated in figure 1. This area is normalised by dividing by the height of the edge giving an indication of the width of blur (Blur = A/h). Note that overshoot, if present, does not reduce the amount of blur using this definition. The reduction of this parameter as a result of filtering gives an indication of the effectiveness of the filter at enhancing the edges. To reduce the effects of noise on this measurement, the column averages are used to indicate the average shape of the edge.

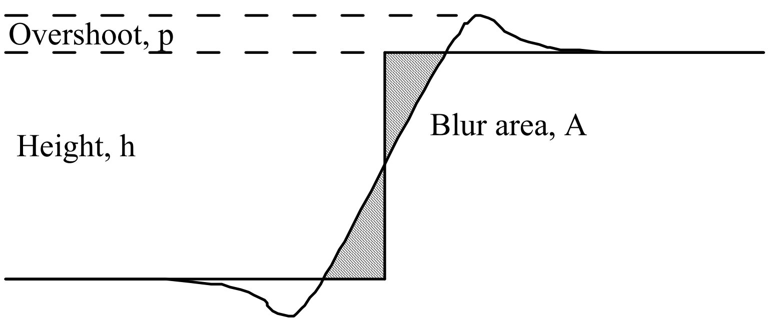

Figure 1: A blurred edge with overshoot, defining the overshoot and edge blur.

Overshoot is defined here as the height of the peak above the ideal edge normalised with respect to the edge height (Overshoot = p/h), as shown in figure 1. If overshoot is present, then this indicates that the filter has over-compensated for any blur in the original image. Again, since noise may result in isolated pixels being significantly higher than the edge height, the column averages are used.

The signal to noise ratio, in decibels, may be expressed as $SNR = 20 \log_{10}(h/\sigma)$ where σ is the standard deviation of the noise. The signal to noise ratio is measured away from the edge, since the selection process of the rank filters used to enhance the edge modifies the standard deviation of the noise near edges. The signal to noise ratio near the edge is also determined.

RESULTS

Four edge enhancement filters: $EN_9(9,1)$, $EN_9(8,2)$, $EN_9(7,3)$ and $EN_{25}(25,1)$ were compared with the two linear filters described above: LIN(A) and LIN(B).

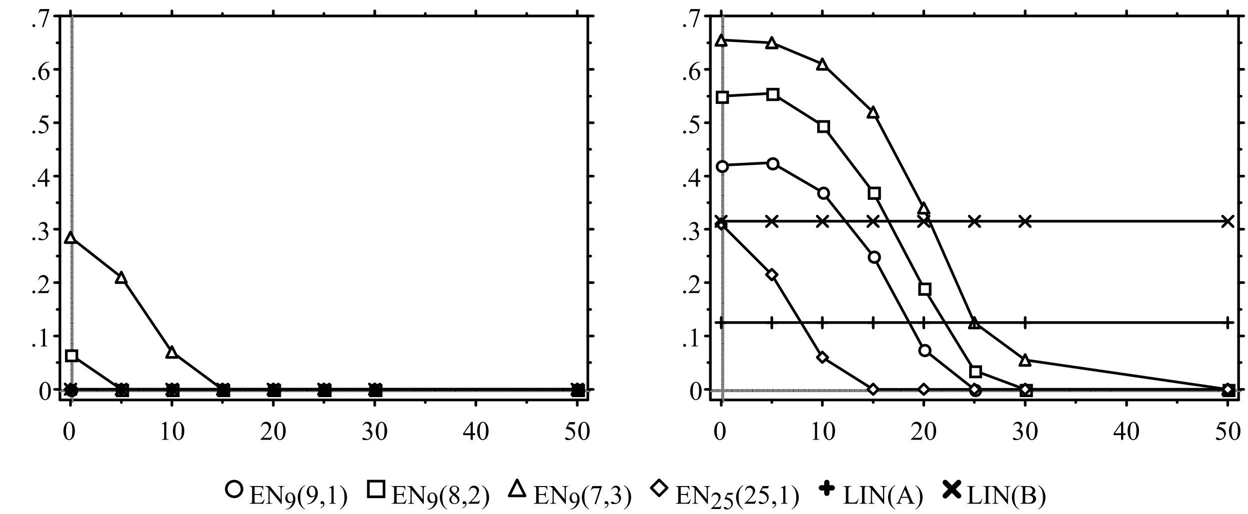

Figure 2: Edge blur as a function of signal to noise ratio.
Blur in original images: left graph 0 pixel; right graph 0.5 pixel.

Figure 2 shows the effects of the filters tested on edge blur. For high signal to noise ratio, the rank-based filters remove the blur, although as the noise level increases, the response deteriorates. This is especially so for filters that use rank values further from the extreme ranks. Such filters tend toward median filters which are known to blur edges slightly in the presence of noise [2]. A closer investigation of the intensity histograms of the columns adjacent to the edge revealed a distinct bimodal structure, especially for the filters using rank values close to the extremes. This is a result of the classification process. What this means is that the edge is not actually being blurred to the extent shown in figure 2, rather some of the pixels are being incorrectly classified. In other words, the intensity noise is being changed to edge position noise causing the average column intensity to indicate blur.

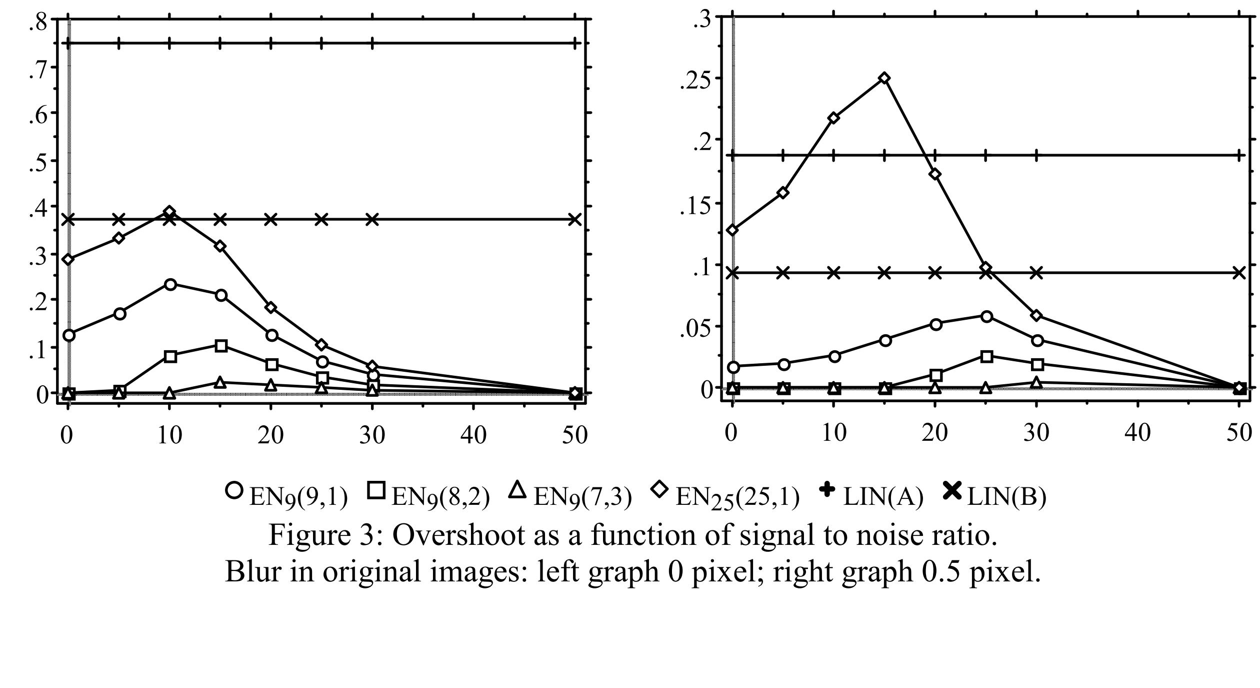

Figure 3: Overshoot as a function of signal to noise ratio.
Blur in original images: left graph 0 pixel; right graph 0.5 pixel.

With noise, the rank-based filters do exhibit a slight overshoot (see figure 3), especially filters that use the extreme pixel values. This is caused by selecting pixel values that are not truly representative of the intensities on either side of the edge. When the window is near an edge, the rank value representative of the opposite side of the edge is less likely to be selected. This effect is more pronounced with filters that use rank values near the extreme rank positions since such values are taken from the tails of the noise distribution and are therefore strongly affected by the noise characteristics. The overshoot is therefore worse with larger windows. As the noise level increases further, the overshoot reduces again since there is significant overlap in the distribution of pixel intensities from each side of the edge. The factor that ultimately limits the overshoot is that no new intensity values are generated by the rank-based filters.

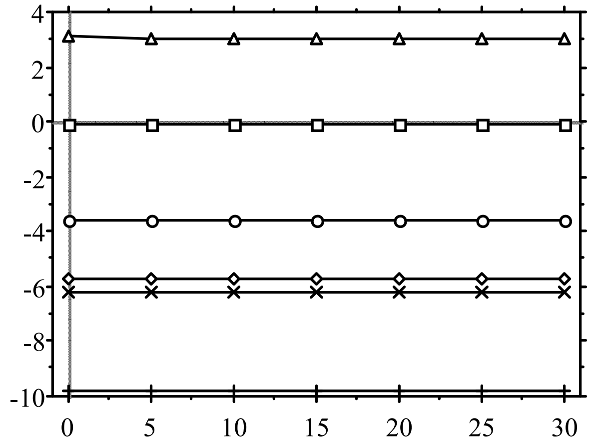

○ $EN_9(9,1)$ □ $EN_9(8,2)$ △ $EN_9(7,3)$ ◇ $EN_{25}(25,1)$ ✚ LIN(A) ✖ LIN(B)

Figure 4: Change in signal to noise ratio after filtering.

The change in signal to noise ratio away from an edge is independent of the original signal to noise ratio (figure 4). This is to be expected since there is no edge within the window, and the output characteristics depend only on the noise characteristics within the image. Figure 4 shows that filters using extreme rank values amplify the noise. Rank values closer to the median smooth the noise. It is interesting to note that there is virtually no change in the noise standard deviation as a result of filtering with the $EN_9(8,2)$ filter.

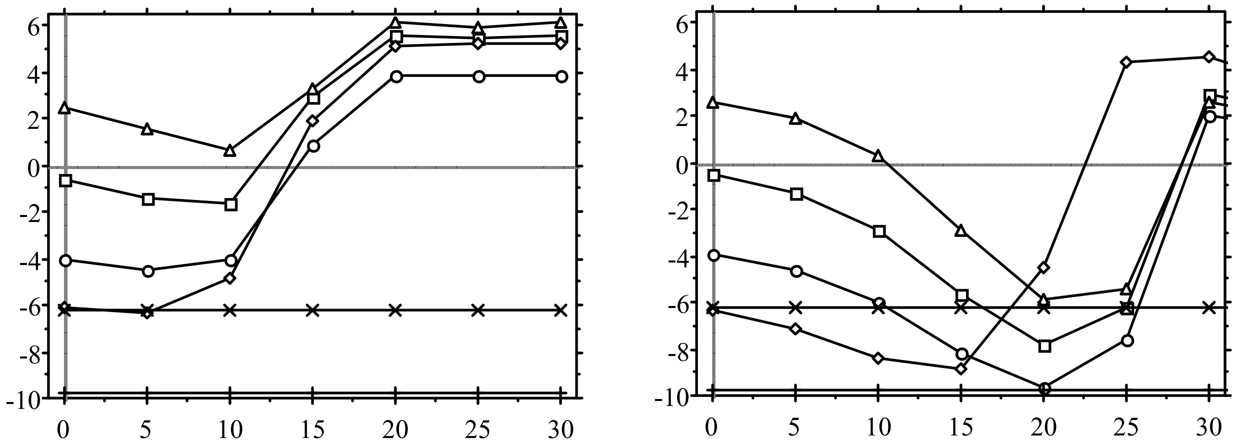

○ $EN_9(9,1)$ □ $EN_9(8,2)$ △ $EN_9(7,3)$ ◇ $EN_{25}(25,1)$ ✚ LIN(A) ✖ LIN(B)

Figure 5: Change in signal to noise ratio adjacent to an edge as a result of filtering.
Blur in original images: left graph 0 pixel; right graph 0.5 pixel.

Near an edge, the change in signal to noise ratio as a result of filtering with rank-based filters depends strongly on the noise level and edge blur (figure 5). Away from the edge, the histogram of intensities within a column of the test image is strongly bimodal, with the peaks having approximately the same area. Near an edge, the one peak has a significantly grater proportion of pixels than the other, and we are effectively measuring the standard deviation of a single peak. As the noise level increases, there is a greater overlap in the intensity distributions and the two peaks become more equally weighted. When the edge is blurred, this effect occurs at a lower noise level.

To illustrate the blurring and change in standard deviation near an edge, an intensity histogram was obtained of the pixels in the column adjacent to the edge after filtering an $EN_9(8,2)$ filter. As shown in figure 6, the histogram is stongly bimodal, with 73% of the pixels in the lower peak. The pixels within the upper peak are those that have been misclassified as a result of noise, and represent an edge position error of 0.27 pixel. This position error is causing the edge blur of 0.36 pixel shown in figure 2 and the 5.6 db drop in signal to noise ratio shown in figure 5. The pixels that were classified correctly (those in the lower peak) correspond to an edge blur of 0.03, and a 5.0 db improvement in SNR.

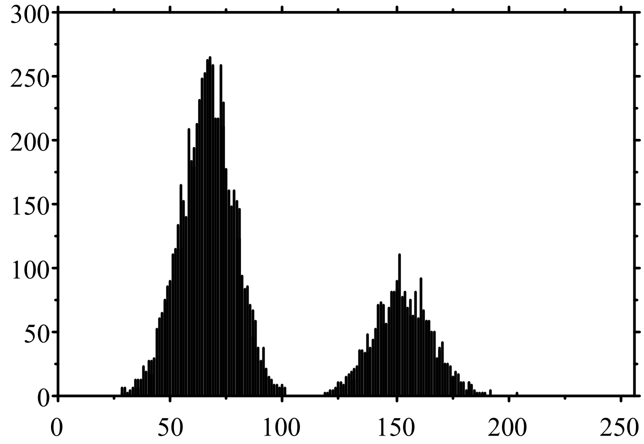

Figure 6: Pixel intensity histogram of the column adjacent to an edge after filtering with $EN_9(8,2)$. Input image has 0.5 pixel blur and 15db signal to noise ratio.

APPLICATIONS

An important application of edge enhancement filters is to enhance the local contrast in an image. When an image is captured, there is a certain amount of blur that results from the image capture process. The light sensing element for each pixel responds to the average incident light within a rectangular aperture. If an edge crosses a pixel, the pixel value will be between the intensities of the regions on either side of the edge. This results in an apparent blurring of the edge as shown in the left image of figure 7. Rank-based edge enhancement filters effectively assign such blurred pixels to the region on either side of the edge that it is most similar to. The result is a significant improvement in the local contrast of the image, as is seen in the centre image of figure 7.

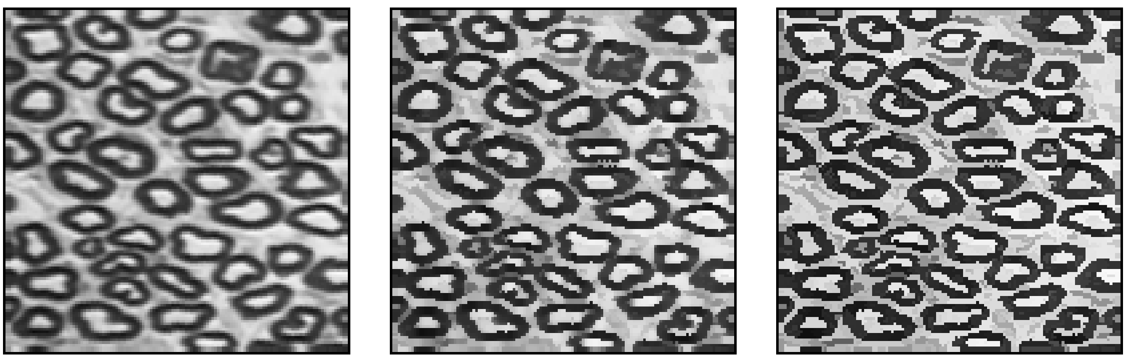

Figure 7: Left, an input image; centre, after a single application of $EN_9(8,2)$; right, after converging.

Repeated applications of a rank-based edge enhancement filters result in fewer changes with each successive application. This enables the filtering to be repeated if the first pass did not enhance the edges sufficiently. The key property which allows this is that edges are not overenhanced. If an edge is already sharp, it remains unaffected by the filter. The right image of figure 7 shows the result of filtering the left image with an $EN_9(8,2)$ filter until no further changes are made (31 iterations).

The change in the image resulting from successive applications of different filters was measured and is plotted in figure 8. The image size was 128x128 pixels and the change value plotted is the absolute difference between the images before and after each pass of the filter integrated over the complete image. The first two or three passes result in the most change. Filters that use a larger window or rank values closer to the median converge faster. A plateau effect appeared for the $EN_9(7,3)$ filter. In this case, the output of successive passes of the filter alternated between two states; two passes of the filter had no effect although each individual pass made minor changes to the image.

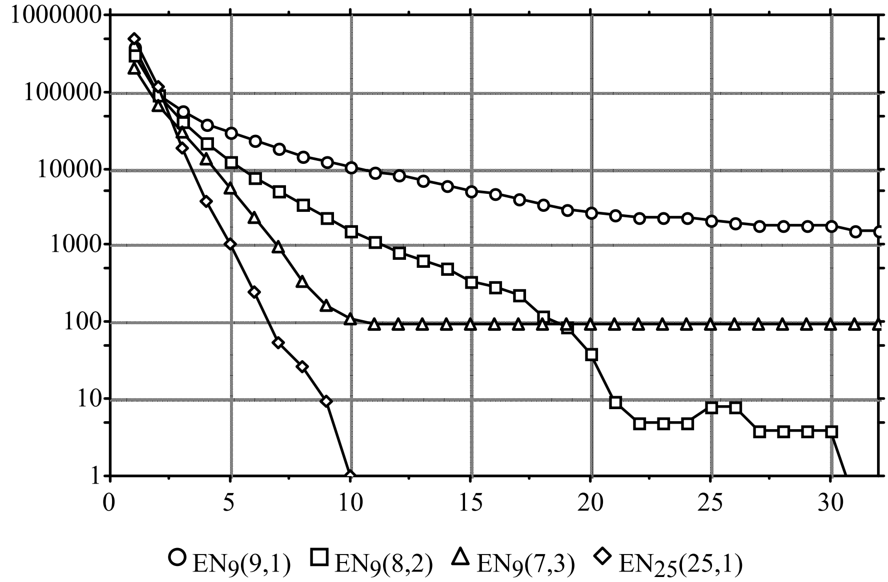

Figure 8: Change in a typical image after successive applications of the filters.

CONCLUSIONS

The rank-based filters investigated in this paper is very effective at enhancing edges within an image. Such filters remove edge blur by assigning pixels into one of the two regions on either side of the edge. Since no new pixel values are generated the edges are not overenhanced. Filters using rank values close to the extremes provide the best edge enhancement. The noise response of rank-based filters is superior to that of linear filters. For low signal to noise ratios, the intensity noise is transformed into noise in the position of the edge in the image after filtering. By carefully selecting the rank values the signal to noise ratio may be improved. These factors illustrate the compromise between the edge enhancement and noise performance. The properties of rank-based filters allow repeated application of the filters, which converge to a stable image.

Further investigation of the edge position noise introduced by the filters is warranted. The conditions under which the filters converge should also be studied in more detail.

REFERENCES

[1] Castleman K.R., “Digital image processing”, Prentice-Hall, Englewood Cliffs, NJ, USA (1979).

[2] Hodgson R.M., Bailey D.G., Naylor M.J., Ng A.L.M., and McNeill S.J., “Properties, implementations and applications of rank filters”, Image and Vision Computing, Vol 3, pp 3-14 (1985).

[3] Bailey D.G. and Hodgson R.M., “Range filters: local intensity subrange filters and their properties”, Image and Vision Computing, Vol 3, pp 99-110 (1985).

[4] Bailey D.G., Hodgson R.M., and McNeill S.J., “Local filters in digital image processing”, Proceedings of the 21st New Zealand National Electronics Conference, pp 95-100 (1984).