PROPERTIES OF THE MODE FILTER WHEN APPLIED TO COLOUR IMAGES

D Charles and E R Davies

Royal Holloway, University of London, UK

ABSTRACT

This paper explains the problems underlying the design of mode filters and describes a new mode filter algorithm that can be applied to colour images. It demonstrates the effectiveness of the new filter for enhancing colour images. It also compares the edge shifting properties of mode filters with those of median filters and shows that the results are completely in line with theoretical predictions. It also finds that mode filters are far superior to median filters at handling images containing high levels of impulse noise.

1. INTRODUCTION

Image noise suppression is a highly important process and much work has been carried out to design effective algorithms for this purpose (Davies, 1). The median filter reigned supreme in this area for a good many years, as, unlike the mean and Gaussian smoothing filters, it does not blur images significantly. While mean and median filters are well known, it is natural to ask about filters employing that other well known centre seeking statistic – the mode. In fact mode filters appear relatively rarely in the literature (e.g. Coleman and Andrews, 2; Davies, 3; Evans and Nixon, 4), and this is probably the result of the difficulty of estimating the mode in the sparse distributions that arise with relatively small window operators. However, the work that has been done in this area indicates that mode filters should not be thought of as noise suppression filters but rather as image enhancement filters.

The implementation problem that arises with mode filters was solved some time ago by employing a substantially modified version of the median filter, and as such was applicable only to grey-scale images. As the technique seems to have proved its value in this capacity, it was felt to be worthwhile to try to devise a version for enhancing colour images. Unfortunately, the median filter involves sorting operations that are only defined in one dimension. This means that the strategy for designing a colour mode filter has to be considered carefully. However, we have found a suitable strategy, and the method is described in Section 3 of this paper.

Section 2 of the paper describes the original approach to mode filter design and Section 3 shows how this has been extended to cope with colour images. Sections 4–6 describe the enhancement, geometric (edge-shifting) and noise suppression properties of the mode filter, and show that the mode filter has the unexpected property of being able to eliminate high levels of impulse noise considerably more effectively than the median filter.

2. PROBLEMS OF MODE FILTERING

As remarked in Section 1, it is difficult implementing an effective mode filter because the small window sizes normally used in image processing lead to sparsely populated distributions. As a result, the highest location in the distribution may be the mode in a rigorous technical sense, yet it may not be the *underlying* mode, i.e. it may not represent a statistically robust estimate of the mode. This means that some smoothing of the local intensity distribution will be needed before the mode can realistically be extracted. However, the amount of smoothing is unclear, and must vary radically from one region to another and between images.

A way of overcoming the problem was described by Davies (3). The starting point is the following procedure applied to a bimodal distribution D: (a) assume the main mode M is known; (b) find the end N of D that is nearer to M; (c) move the distance d from N to M, and move a further distance d to a point T; (d) truncate the part of D beyond T: this will eliminate the minor mode, and the main mode can be found by averaging the remaining part of D; this is more accurate than just locating the maximum of D because any noise on the peak tends to be ignored. Next, note that the parameter d is normally larger for the median than the mode, so less of D is truncated, making the procedure 'safer' but still highly effective in eliminating the minor mode: further iterations can eliminate it entirely. In practice, the resulting 'truncated median filter' has been found to be highly effective for enhancing grey-scale images.

The means by which the technique enhances images is by ignoring the onset of edges and thus maintaining the same intensity right up to the edge: a similar phenomenon occurs on the other side of the edge, and so the edge region is narrowed and 'crispened'. This is the sense in which the image is enhanced.

3. THE NEW COLOUR MODE FILTER

The first problem with any attempt to extend the truncated median filter concept to a 3-D colour space is defining the median filter in that space. In fact it is known how to do this – by aiming to find a data point (the '3-D median') that minimises the sum of the absolute distances $|d_{ij}|$ to the other data points. This effectively eliminates the first part of the problem of devising a mode filter algorithm. The next part of the problem is how to eliminate any outlier region representing a subsidiary mode. To achieve this the following strategy, which is described and illustrated (Figure 1) just for the 2-D case, was devised:

1. Determine the median M_1 of the input distribution.
2. Find the point O which maximises the sum-of-distances cost function.
3. Move from O to M_1 through a vector distance **v**, move a further distance **v** to a location P: this will be outside the input distribution.
4. Find the data point T closest to P and measure the distance r from M_1 to T.
5. Truncate the distribution at distance r from M_1.
6. Find the median M_2 of the new distribution.
7. If $M_1 = M_2$, terminate the procedure: otherwise go to step 2.

As for the original truncated median filter, step 7 is only needed if the process is to be iterated to completion.

Like the 1-D version, the 2-D and 3-D mode filters are cautious and safe. They converge to the truncated median filter in the 1-D case, and provide a useful generalisation which should constitute an effective colour image enhancement filter.

4. ENHANCEMENT PROPERTIES

The colour mode filter has been tested on a number of off-camera digital images, and in all cases a useful degree of enhancement was achieved. Here space only permits its effect on a single colour image to be demonstrated: a 5 × 5 window was employed in this instance (Figure 2). For comparison, the effect of applying a 5 × 5 vector median filter (Astola et al, 5) is also shown. Both filters show similar loss of detail because of the 'softening' effect sometimes noticed with the median filter, but in neither case does this correspond to the blurring obtained with a mean filter.

When colour images are processed by median filters one colour dimension at a time, the phenomenon of 'colour bleeding' commonly occurs: this arises as edges are differentially shifted in the different colour dimensions by any noise spikes that may be present. The vector median filter largely eliminates this problem, and likewise the 3-D mode filter described in Section 3 has been found to give rise to negligible colour bleeding. This resistance to colour bleeding is intrinsic to the design of the algorithm, as it automatically orientates itself in colour space to find the optimal direction for suppressing outliers.

5. EDGE SHIFTING PROPERTIES

Amongst the most important properties of median filters is their ability to eliminate small spots and noise, though accompanying this capability is the possibility of noise-induced edge shifts (Bovik et al, 6) and the property of slightly reducing the sizes of convex objects. This latter property can be traced to the shift of edges towards the local centre of curvature, and it has been shown (Davies, 7) that the shift approximates to ${}^1/_6\kappa a^2$, where κ (= $1/b$) is the local curvature and a is the radius of the median window. These geometric properties are key to fully understanding the median filter: here it is relevant to determine to what extent the mode filter has identical or similar properties.

As in earlier work to measure the edge shifts produced by median filters, we take a number of grey-scale circles of varying curvature and apply the relevant filter. We measure the differences in area before and after filtering. Note that the method of measuring the sizes of the processed circles is to measure their integrated intensities, and thus deduce the new effective radius values. The same precautions are taken as in Davies (7) to attain sub-pixel accuracy.

Figure 3 shows the general features of the shift variations for a given size of window. Notice that the median shift (lowest dark curve) follows a simple variation right up to the situation (uppermost curve) where the circles are entirely eliminated by the filter, whereas the mode curve suddenly deviates upwards (grey section): the foregoing models seem incapable of explaining this sudden upward shift, and a new concept is needed for the purpose. The simplest explanation of why the mode shift starts increasing sharply involves considering the slope of the edge of the circular object. The basic situation is that the mode filter must examine the two regions within the window that lie on either side of the edge slope region (regions B and T in Figure 4), and the shift corresponds to the distance between the centre of the edge and the centre of the window for which these two regions have equal area.

In the case of finite edge slope, the condition under which the median will completely eliminate circles is when the outer radius of the circle, $b + t/2$, corresponds to half the area of the window:

$$\pi(b + t/2)^2 = \pi a^2/2 \tag{1}$$

which leads to the median breakpoint (Figure 3):

$$b = a/\sqrt{2} - t/2 \tag{2}$$

whereas the condition for the mode to completely eliminate the circle is when the area outside the edge slope region is equal to that inside it:

$$\pi a^2 - \pi(b + t/2)^2 = \pi(b - t/2)^2 \tag{3}$$

so we get the upper mode breakpoint (this is the case depicted in detail in Figure 4):

$$b = (a^2/2 - t^2/4)^{1/2} \tag{4}$$

Clearly, the median and upper mode breakpoints are the same when $t = 0$.

So far the slope model has not led to an explanation of the lower mode breakpoint, where the mode shift starts deviating from the median shift. However, this is easily solved, as it is the point where the inner edge for the first time falls completely within the window, thus causing a discontinuity in the gradient of the mode characteristic. It is not possible to obtain a closed formula for this situation, but it can be calculated numerically in any given case. However, before such a calculation can be made, it is necessary to estimate the edge width t. Conveniently, this can be achieved by using the median filter breakpoint values, obtained using equation (2) to determine t.

Tests with 5×5 to 13×13 mode filters led to a mean t of 1.45 ● 0.04. Substituting for t in equation (4), and also using it for the calculation referred to above, led to the mode breakpoint values listed in Table 1. Considering the simplicity of the model, the degree of agreement between estimated and observed values is good. The closeness of the agreement confirms that the mode shifts are now well understood: they also relate well to the situation for the median filter.

TABLE 1 - Curvature breakpoints for the mode filter. The observed values are listed with estimated values appearing in brackets.

Window size	Lower breakpoint	Upper breakpoint
5×5	0.50 (0.49)	0.59 (0.60)
7×7	0.40 (0.38)	0.45 (0.43)
9×9	0.30 (0.29)	0.32 (0.31)
11×11	0.25 (0.24)	0.27 (0.26)
13×13	0.21 (0.21)	0.22 (0.22)

TABLE 2 - Normalised errors for various window sizes

Window size (square window)	Median NMSE	Mode NMSE
5×5	0.049	0.042
7×7	0.042	0.024
9×9	0.039	0.013
11×11	0.038	0.007

6. NOISE PROPERTIES

Tests showed that the median is marginally better than the mode at eliminating small amounts of impulse noise. This is natural as the median is buried at the centre of the local intensity distribution and is thus more resistant to impulse noise. However, when the window contains two or more noise pixels at one end of the distribution, the median will move, whereas the mode is able to eliminate them. Thus the mode tends to be better at eliminating higher levels of impulse noise.

This is illustrated by the following simulation example. We generated a plain image of constant colour on which 70% impulse noise was applied, the impulse noise being randomly (uniformly) distributed between 0 and 255 in each colour dimension (we shall refer to this as type A noise). The effectiveness of both types of filter for this image is shown in Table 2 (a 5×5 window was used for this test). Quite clearly, the mode is always more effective in this case, and the relative effectiveness increases markedly with window size.

The same test was carried out for a variety of off-camera images (this time using 9×9 windows), with the results shown in Figure 5. The superiority of the mode filter at high impulse noise levels seems quite definite, and considering its relative simplicity, it may even be more competitive for time-critical applications working in heavy noise than some of the latest 'switched' noise suppression filters.

7. CONCLUDING REMARKS

This paper has explained the problems of designing mode filters and has described a new mode filter that can be applied to colour images. It has demonstrated the effectiveness of the new filter for enhancing colour images. It has also investigated the geometric properties of mode filters and in particular has shown that they shift low curvature edges by similar amounts to median filters. However, the edge shifts suddenly rise sharply for certain curvatures, and become much more drastic than for median filters. This property is in line with the capability of mode filters for ignoring quite large proportions of any window, and indeed proportions significantly larger than the 50% which median filters can ignore. This same property permits mode filters to ignore very large amounts of impulse noise well in excess of 70% of any image, and at these levels the mode filter is far superior to the median filter in its noise suppression properties. Nevertheless, the median filter is better at coping with low levels of impulse noise: this is reasonable as the mode filter was originally conceived as an image enhancement filter rather than a noise suppression filter.

In general, median and mode filters should not be compared solely from the point of view of noise or edge shift, as they have different aims, different specifications and different combinations of properties. These and all other filters should be included in the library of available algorithms and brought out only when their individual capabilities match the needs of the particular application.

REFERENCES

1. Davies ER, 1997, *Machine Vision: Theory, Algorithms, Practicalities*, Academic Press, London, 2nd edition

2. Coleman GB and Andrews HC, 1979, "Image segmentation by clustering", Proc IEEE, 67, 773–785

3. Davies ER, 1988, "On the noise suppression and image enhancement characteristics of the median, truncated median and mode filters", Pattern Recogn Lett, 7, 87–97

4. Evans AN and Nixon MS, 1995, "Mode filtering to reduce ultrasound speckle for feature extraction", IEE Proc Vision Image Signal Process, 142, 87–94

5. Astola J, Haavisto P and Neuvo Y, 1990, "Vector Median Filters", Proc IEEE, 78, 678–689

6. Bovik AC, Huang TS and Munson DC, 1987, "The effect of median filtering on edge estimation and detection", IEEE Trans PAMI, 9, 181–194

7. Davies ER, 1999, "Image distortions produced by mean, median and mode filters", IEE Proc Vision Image Signal Process, 146, 279–285

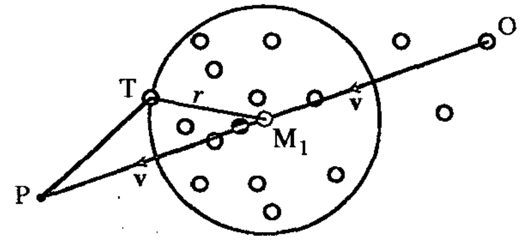

Figure 1. Method of truncating the local distribution in 2-D, extendible to 3-D colour spaces. Empty circles are samples, shaded circle is M_1, black circle is M_2.

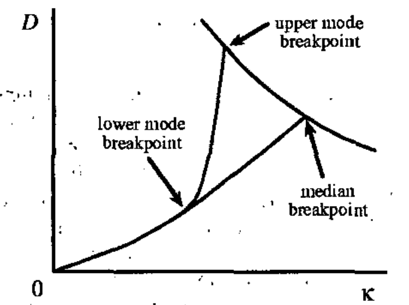

Figure 3. General features of median and mode shift variations. The median shift is indicated by the lowest dark curve. The mode shift is indicated by the rapidly rising grey section: further down it coincides with the median shift. The uppermost curve shows the maximum shift, given by $D = b = 1/\kappa$.

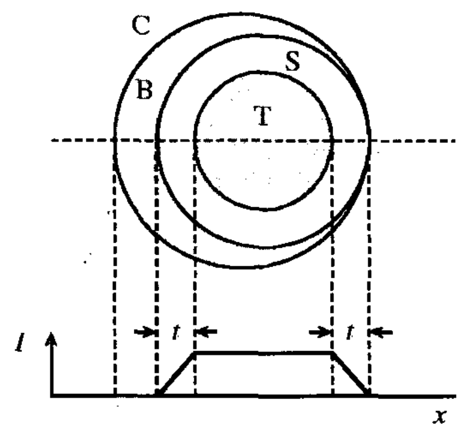

Figure 4. Circular object with a broad slope edge S of width t, lying just within a circular window C. The two shaded regions T (top of object) and B (background) have equal area. Shown below is the side view along the middle of the window which represents the intensity I on passing right through the circular object.

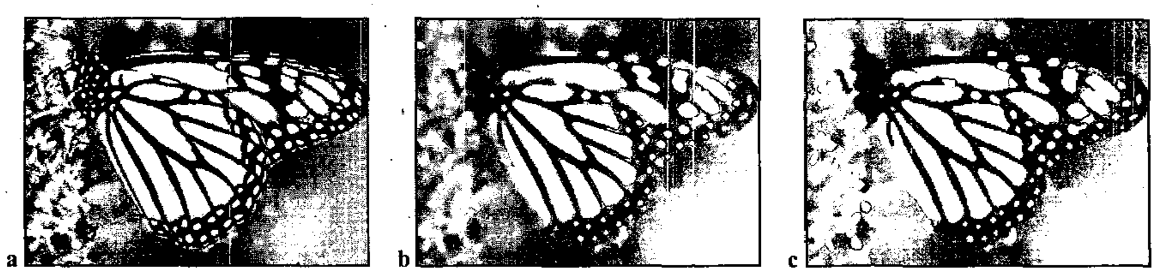

Figure 2. Enhancing effect of the mode filter. (a) Original image (from the Waterloo repertoire). (b) Result of applying 5 × 5 vector median filter. (c) Result for 5 × 5 mode filter.

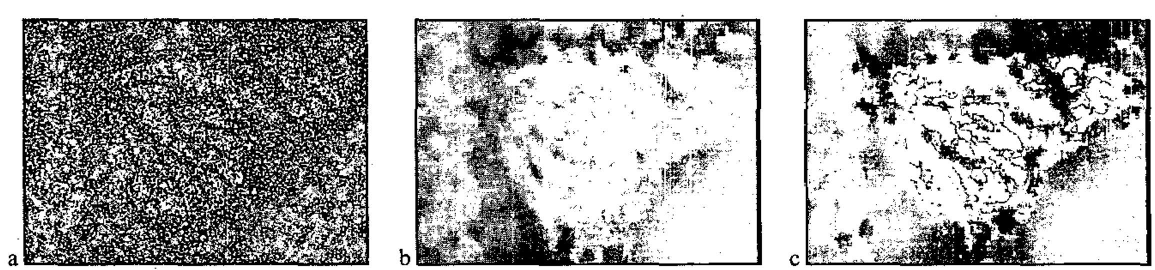

Figure 5. Noise removal capability of the mode filter. (a) Figure 2(a) contaminated with 70% type A impulse noise. (b) Effect of applying 9 × 9 vector median filter. (c) Result for 9 × 9 mode filter.