FROM CONTINUUM MODEL TO A DETAILED DISCRETE THEORY OF MEDIAN SHIFTS

E. R. Davies
Machine Vision Group, Department of Physics
Royal Holloway, University of London
Egham Hill, Egham, Surrey, TW20 0EX, UK
Tel: +44 1784 443497; fax: +44 1784 472794
e-mail: e.r.davies@rhbnc.ac.uk

ABSTRACT

This paper presents a new discrete theory of median shifts. It relates the theory to the older continuum model, predicts angular variations, and also gives an accurate figure for the active area of any discrete neighbourhood, thereby making the continuum model more accurate. The work shows that at low curvature values a quadratic law applies, this being followed by the previously known linear variation. It also explains why the observed linear behaviour is significantly larger than indicated by the continuum model. Overall, there is now very good agreement between the new discrete theory, the continuum model and the observed results.

1 INTRODUCTION

Median filters are well known for their ability to suppress noise in digital images without producing the blurring characteristic of mean and Gaussian filters. It has also been known for some time that median filters produce shifts of curved object boundaries and intensity contours [1, 2]. The actual shifts produced by such filters have been modelled using a continuum approximation and give rise to a simple formula for the shifts, namely:

$$D \approx \frac{1}{6}\kappa a^2 \qquad (1)$$

where κ is the curvature of any relevant intensity contour, and a is the radius of the neighbourhood, assumed to be nearly circular. While the agreement between experiment and theory provided by this model was not exact, it led to approximately the right variations, and the numerical values depended on a separate assessment of the effective size of the neighbourhood [1]. Semi-empirical estimates were used to obtain agreement between experiment and theory. However, the situation was not completely satisfactory, as absolute numerical agreement was not obtained by this approach, and with an operator as apparently simple as a median filter it should in principle be possible to obtain exact absolute agreement.

This is quite important as it has recently been shown that mean and mode filters also suffer from this phenomenon, so it is not a peculiarity of the median filter [3, 4]. Hence it is definitely of interest to see whether exact numerical agreement between experiment and theory can be attained in the median case, and also to see whether the relation between an exact theory and that offered by a continuum model can provide a rigorous measure of the effective size of the neighbourhood. It is also of interest to see the extent to which the non-circular neighbourhood shape arising from a discrete lattice of pixels affects the theory. These two factors are also important because relatively little work has been done to relate discrete to continuum concepts in the whole of image analysis, an important exception being the work of Kulpa and others on analogue boundary length estimation on a discrete lattice [5–7].

In the following section we determine accurate median shifts for 3×3 neighbourhoods. In Section 3 we extend this work by means of simulations, and obtain results for 5×5 neighbourhoods. In Section 4 we study the theory in more detail, and show how various intensity patterns affect the median shifts. Section 5 summarises the work and gives further insights into the operation of noise suppression filters.

2 BASIC THEORY OF MEDIAN SHIFTS IN 3×3 NEIGHBOURHOODS

In this section we assume that the tessellation is square and that each neighbourhood consists of an $n \times n$ array of pixels. We shall take each pixel intensity value to be a sample positioned at the centre of the pixel and obtained by averaging the intensities over the whole pixel, though we shall temporarily assume that the value is equal to the intensity at the pixel centre. We consider the case of a 3×3 neighbourhood, and assume that the underlying analogue intensity variation has contours of curvature κ and radius of curvature b, where $b = 1/\kappa$, as shown in Figure 1. First, it is clear that zero shift occurs for $\kappa = 0$. However, if a circular median intensity contour passes close to the centre of the neighbourhood with an orientation θ to the positive x-axis, we can identify several possibilities: (1) the median contour will or will not pass through the centre of the central pixel in the neighbourhood; (2) the median contour will pass through one or more than one pixel centre. Note that for most orientations θ the median contour will pass through only one pixel centre, which we shall call the median pixel centre. In any case the returned value of the median filter will be the intensity of the median pixel, and in the discrete case this may either be the central pixel or some other pixel.

The precise geometrical configuration is shown in Figure 2. In particular, O is the centre (0, 0) of the neighbourhood, and C is the centre (x_c, y_c) of the circle, which passes through the median pixel position (p, q). We now need to calculate the distance of the circle from O: this is the median shift D_θ. As the circle passes through (p, q):

$$(p - x_c)^2 + (q - y_c)^2 = b^2 \tag{2}$$

Substituting for y_c and solving for x_c gives:

$$x_c = \cos^2\theta\left\{(p + q \tan\theta) \pm \left[(p + q \tan\theta)^2 - (p^2 + q^2 - b^2)\sec^2\theta\right]^{1/2}\right\} \tag{3}$$

where the '+' sign gives the appropriate solution. We can now find the nearest point S (x_s, y_s) on the circle to O, and this leads to a median shift:

$$D_\theta = - x_s \sec\theta = b - x_c \sec\theta \tag{4}$$

Equations 3 and 4 now lead to the correct value for the median shift D_θ. Starting at $\kappa = 0$, the median will always be the centre pixel (0, 0), so D_θ will be zero for all θ. For a minute increase in κ, and assuming a very low value of θ, we see that the centre pixel can no longer be the median value, which will hence divert to (0, 1) as in Figure 3. This leads eventually to:

$$D_\theta = b - \sin\theta - (b^2 - \cos^2\theta)^{1/2} \tag{5}$$

For large b we find:

$$D_\theta \approx - \sin\theta + \frac{1}{2b}\cos^2\theta = - \sin\theta + \frac{1}{2}\kappa\cos^2\theta \tag{6}$$

Similar arguments for much larger angles lead to the angular variations shown in Figure 4. Clearly, these do not match what might have been expected from equation 1. A way forward is to find the mean value of D_θ as θ varies over the whole range $0 \le \theta \le \pi/4$. This leads to the following approximate result for low values of κ:

$$D \approx \left(\frac{1 + 2\sqrt{2}}{2\pi}\right)\kappa^2 \approx 0.61\ \kappa^2 \tag{7}$$

It will be seen from Figure 5 that this approximation is good up to $\kappa \approx 0.3$, and thereafter somewhat underestimates the variation in D, relative to values obtained from a full numerical computation. What is more interesting is the reason for the square-law variation. Examining Figure 4 more closely shows that the θ-variations are approximately piecewise linear functions of κ: in addition, they move laterally in a close to linear manner with variation in κ. As a result, integrating with respect to θ and averaging over this variable must yield an approximate square-law variation in κ. The limits to this variation occur when the 3-element piecewise linear approximation breaks down, and the θ-map as a whole starts increasing almost linearly with κ (this happens for $\kappa \approx 0.63$). However, the resulting small linear variation is insufficient in the case of a 3 × 3 neighbourhood to win over the still strong square-law variation – though, as we shall see below, the same is not true in the case of a 5 × 5 neighbourhood.

Next, it turns out that the situation is more complex for high values of κ. For example, when the orientation is sufficiently high, (0, –1) takes over from (1, –1) as the pixel giving the median value, while (0, 0) simultaneously drops out (Figure 3). Although other pixels might have been expected to be relevant, they correspond to curvatures so high that solutions do not exist for all θ. This means that the shift will be indefinable, since the small size of the circle will not encompass sufficient pixel centres in some orientations, and it will be ignored: i.e. the circular object will be eliminated by the median filter. Here we ignore such situations, and concentrate instead on determining mean values of D where D_θ is definable for all θ.

Overall, the important point it that it is possible to understand in detail the situation for a 3 × 3 neighbourhood. Unfortunately, the situation gets increasingly complex as more and more pixels come into play for larger neighbourhoods. Hence simulations will be seen to be important, as shown in the following section.

3 SIMULATIONS TO ESTIMATE MEDIAN SHIFTS IN 5 × 5 NEIGHBOURHOODS

An attempt was made to calculate median shifts in 5 × 5 neighbourhoods by the methods of the previous section. As a start, it is clear that for low κ and low θ the median pixel is (0, –1). Interestingly, this is below the centre pixel rather than above it (as happens with 3 × 3 neighbourhoods), and as a result D_θ increases at first rather than decreasing with increase in θ. For low κ, it is also clear that the next median pixel is (0, 2). However, the situation rapidly becomes complicated as θ increases further: accordingly, a simulation was carried out using the same concepts, to find the median shift for all κ and θ and its average values over all θ (Figures 6 and 7).

It will be apparent there is a gradual changeover from a square law to a linear law as κ increases (Figure 7). Again, it is clear that the square-law variation is due to the approximately linear forms of the D_θ variations, which, when averaged over θ, give a κ^2 characteristic. This changeover from a square law to a linear variation would be expected to become increasingly evident for larger neighbourhoods. Indeed, for larger neighbourhoods the range of the square-law variation will be reduced, and the most apparent variation will be the linear one. None of these factors was obvious from the older continuum model which merely predicted a simple linear variation.

4 SIMULATIONS WITH SUB-PIXEL INTENSITY VARIATIONS

While the previous simulations and theory related directly to binary and saturated grey-scale images, slight alteration of the methodology is required before normal grey-scale images can be coped with. In fact, the simplest approach is a direct simulation in which all sub-pixels are assigned appropriate grey-level values; next, these are averaged over each pixel; then the median intensity is determined, and the median shift is computed from the median contour position. To achieve this, the intensity patterns have to be decided in advance. Here we adopt the model of [1], and assume that there is a steady intensity gradient along a specific axis, with variations away from this axis being determined by contours of constant curvature, as indicated in Figure 1.

We carried out such a simulation, with 11×11 sub-pixels per pixel, this being found adequate for the task. The result of this work showed that for 3×3 neighbourhoods, the estimate given by the upper solid curve in Figure 5 is slightly low, by an amount which varies with the precise intensity pattern assumed, and in particular with the possible variation of κ on progressing from one intensity contour to another. In fact, this effect seemed to account for about half the discrepancy between the experimental and theoretical curves shown in Figure 5, the remaining differences probably being due to peculiarities of the data in [1]. Nevertheless, it appears that the experimental shifts reported in [1] could now reasonably be explained if the *precise* intensity patterns in the original image data were known.

It is of interest to compare the results obtained so far with those of the continuum model leading to equation 1. The theory presented above makes it clear that for a circular neighbourhood a should be interpreted as the mean distance from the centre to the *centres* of the outermost pixels. For a small square neighbourhood it seems reasonable to approximate by taking a as the mean distance to the square through the centres of the outermost pixels, which gives $a \approx 1.12a_0$, a_0 being the closest distance from the centre to the square. For a 5×5 neighbourhood, this leads to a shift $^1/_6\kappa \times 2.24^2$ which gives values ~35% lower than the linear variation predicted on the discrete model. However, for a 3×3 neighbourhood, we get a shift $^1/_6\kappa \times 1.12^2$ which gives values a factor ~3 lower than any possible linear variation that might be deduced on the discrete model. Clearly some explanation is needed of why the continuum model predicts these low values – especially in the latter case. In fact, it seems that the pixels in the 3×3 neighbourhood might well act as a thin ring of pixels, with the result that the linear variation will be closer to $^1/_2\kappa a^2$ than to $^1/_6\kappa a^2$, thereby fully accounting for the discrepancy. The effect would be expected to be smaller for a 5×5 neighbourhood, and would revert to normal for a continuum situation. This seems about as far as we can go to explain the small neighbourhood results on the continuum model.

Finally, we consider the results of the earlier experiments [1]. Their main relevance here is the total incapability of the earlier continuum model to explain in detail the very low values of shift observed in these data for low values of κ. However, this problem has now been entirely overcome, in that the linear variation is preceded by the hitherto unsuspected quadratic variation at low κ. Thus these experimental results are now largely explained, any remaining discrepancies being due to peculiarities of the data which have already been considered in [1].

5 CONCLUDING REMARKS

This paper has studied the shifts produced by median filters. The total incapability of the earlier continuum model to explain in detail the very low values of shift observed at low values of κ has been overcome by the use of accurate discrete models. In particular, it has been shown that the linear variation is preceded by a hitherto unsuspected quadratic variation at low values of κ. One odd circumstance is the fact that the predictions of the earlier continuum model are so small at high values of κ, but this has also been explained.

The derivation of the discrete models in this paper shows that the size of the neighbourhood to be assumed on the earlier continuum theory is not the apparent area of the neighbourhood, but the area bounded by the sampling points at the centres of the pixels. While the continuum theory would not be expected to give accurate results for small neighbourhoods, it should do so for large neighbourhoods if this factor is taken properly into account.

Overall, this paper has been able to relate the discrete and continuum theories of median shift and to use them to explain all observed results. There is no doubt that shifts are an important feature of median filters: interestingly they cannot be avoided merely by using simple alternative measures such as mean or mode filters.

REFERENCES

1. Davies, E.R. (1989) "Edge location shifts produced by median filters: theoretical bounds and experimental results", Signal Process., **16**, no. 2, pp. 83–96
2. Davies, E.R. (1997) *Machine Vision: Theory, Algorithms, Practicalities*, Academic Press (2nd edition), pp. xxxi + 750
3. Davies, E.R. (1991) "Median and mean filters produce similar shifts on curved boundaries", Electronics Lett., **27**, no. 10, pp. 826–828
4. Davies, E.R. (1997) "Shifts produced by mode filters on curved intensity contours", Electronics Lett., **33**, no. 5, pp. 381–382
5. Kulpa, Z. (1977) "Area and perimeter measurement of blobs in discrete binary pictures", Comput. Graph. Image Process., **6**, pp. 434–451
6. Proffitt, D. and Rosen, D. (1979) "Metrication errors and coding efficiency of chain-encoding schemes for the representation of lines and edges", Comput. Graph. Image Process., **10**, pp. 318–332
7. Davies, E.R. (1991) "Insight into operation of Kulpa boundary distance measure", Electronics Lett., **27**, no. 13, pp. 1178–1180

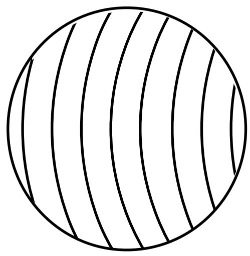

Figure 1 Idealised geometry for continuum model. Contours of constant intensity in an idealised circular neighbourhood on the continuum model.

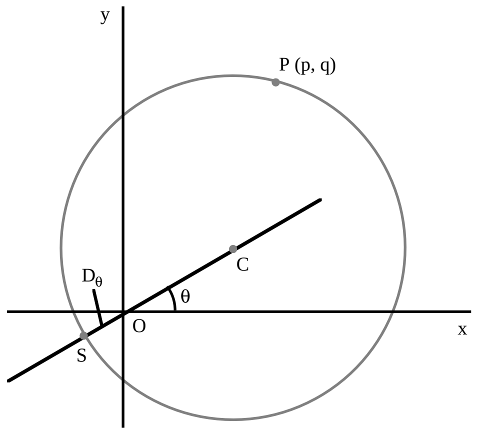

Figure 2 Geometry for calculation of median shifts in discrete model. The circular median contour passes though the pixel P at (p, q), and passes within a distance D_θ of the centre pixel at O, leading to a shift D_θ at angle θ.

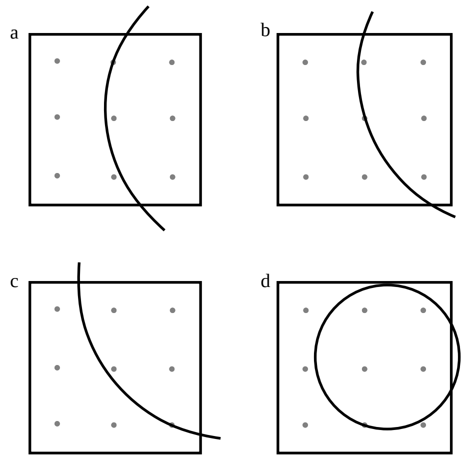

Figure 3 Phases in calculation of median shifts for 3 × 3 neighbourhood. The position of the median contour is shown in (a) for low κ and low θ, in (b) for low κ and moderate θ, in (c) for low κ and high θ, and in (d) for high κ and high θ.

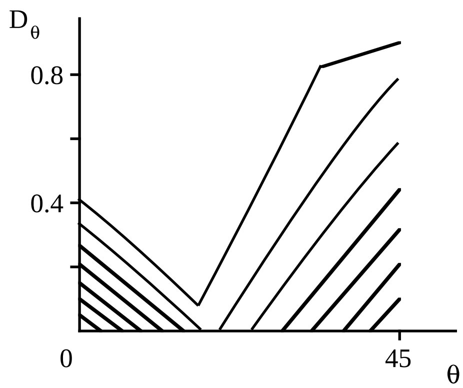

Figure 4 Angular variation of median shifts for 3 × 3 neighbourhood. The graphs show the variations in steps of 0.1 from κ = 0.1 (lowest) to κ = 0.7 (highest). Note that the lowest graphs are in three parts, the θ-axis constituting the middle part of the variation.

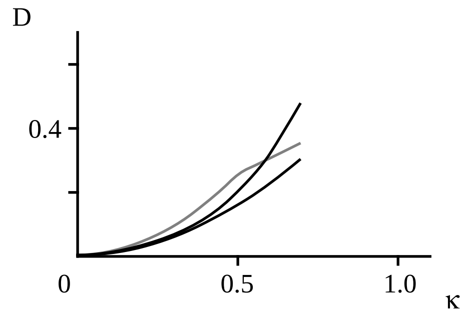

Figure 5 Comparisons of median shifts for 3 × 3 neighbourhood. The upper solid curve corresponds to the angular variations of Figure 4. The lower solid curve shows the approximate square-law model of equation 7. The dotted curve depicts the experimental data from [1].

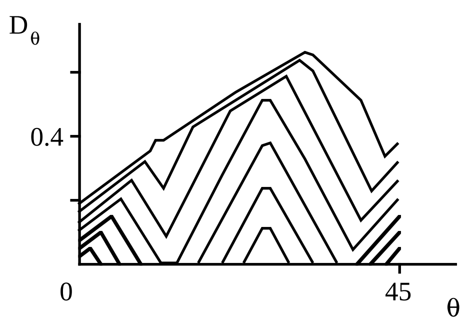

Figure 6 Angular variation of median shifts for 5 × 5 neighbourhood. The graphs show the variations in steps of 0.05 from κ = 0.05 (lowest) to κ = 0.35 (highest). Note that the lowest graphs are in five parts, the θ-axis constituting the other two parts of the variation.

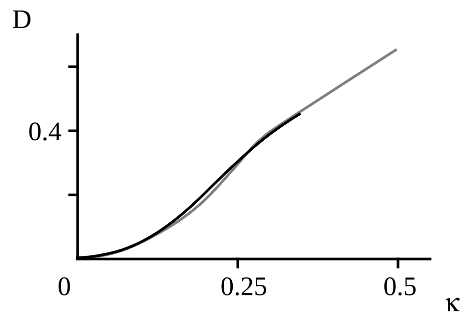

Figure 7 Comparisons of median shifts for 5 × 5 neighbourhood. The solid curve corresponds to the angular variations of Figure 6. The dotted curve depicts the experimental data from [1].