

# ESTIMATING ASYMMETRIC NOISE

Alison Noble & Michael Brady

Robotics Research Group,  
Oxford University, OX1 3PJ.

---

*Recent work is summarised that characterises deterministically and statistically the performance of morphological and rank order filters. We propose adaptively filtering noise that is asymmetric by one of the biased morphological filters  $co$  and  $oc$  whose average is known to be an unbiased estimator of a signal in symmetric noise. Significant differences in filtered images are determined using a nonparametric statistical test. Preliminary results illustrating the theory and applying it to texture segmentation are presented.*

---

This paper arises from a recent study of the noise suppression characteristics of morphological filters [1]. Noise suppression filters based on the morphological opening and closing transformations have been used for a number of years for pre-detection smoothing in industrial vision [2] and for a variety of medical applications. From a practical viewpoint, morphological filters are good at suppressing impulse noise, and are edge-preserving so do not suppress high frequencies. Since they preserve structure, they perform better than rank-order filters, such as the median, on two-dimensional image structures. Morphological filters can be efficiently implemented relative to mean and median filters: for example, all of the filters discussed in this paper have been implemented on our Laboratory's Datacube image processing system [4]. Several examples of the application of morphological filters are presented in [5,1]. Figure 1 shows a typical result of filtering a range image with a sequence of successively larger *comoc* operators. The *comoc* filter is the average of the open-closing filter *oc* and the clos-opening filter *co*, that is  $comoc = (co + oc)/2$ . The iterated *comoc* performs significantly better than: edge-preserving smoothness, weak membrane, and median filters.

The application of morphological filters to image processing has for too long preceded the development of theory (the same comment applies to early work in other areas of vision research such as edge detection). As morphological and linear filters are based on quite different mathematical theories it has proven difficult to relate precisely their performance. Recently, however, there has been some considerable progress in the development of the theory, specifically concerning two aspects of morphological filtering:

- deterministic analysis of performance; and
- the statistical properties of those filters.

First, Maragos has studied the clos-opening *co* and open-closing *oc* morphological filters. The *co* filter over-

estimates a one-dimensional signal, while the *oc* filter underestimates it, and in both cases they trim sharp peaks and valleys. He has shown that *co* and *oc* transform signals in a similar manner to the median and he derived relationships between rank-order, linear and some morphological filters [6,7]. The value of these studies and those in [8], is that they have furthered our understanding of the shape properties of morphological filters (signals invariant to filters of different geometry and size).

Recently, Arce and Stevenson studied the statistical properties of morphological filters [9]. More precisely, they derived the output cumulative pdf that results from applying the *co* filter to a signal whose assumed input cumulative pdf is  $F_x^n(t)$ , where  $t$  is a (threshold) signal level. [1] describes the relevant mathematics, and derives several new results for the cases where  $F_x^n(t)$  is uniform, Gaussian, and Rayleigh (for example, the assumed noise model for the magnitude of an image gradient [10].)

Taken altogether, the results of Maragos, Arce, and Noble quantify the noise suppression performance of filters on a *constant* signal corrupted with noise from the corresponding distribution. The take-home messages can be succinctly stated:

1. Morphological filters smooth a noise-corrupted constant signal less than the median which in turn is less good than linear averaging for Gaussian noise. This underlies the structure preserving properties of the morphological filters relative to median and mean.
2. Both median and morphological filters are good for heavy-tailed noise distributions and the suppression of outliers. They remove 'impulses' of extent less than  $n$  where the filter window size is  $2n + 1$ . This implies that they are good at eliminating point impulses (salt-and-pepper noise) and missing line noise. By contrast, a moving average filter performs badly on this task.
3. The application of a single morphological filter, such as *co* and *oc*, introduces a *bias* with respect to the median. For symmetric input noise, the *comoc* is an unbiased signal estimator.
4. Morphological filters differ from the median in that reapplication of a morphological filter produces no further change in the signal structure. Morphological and median filters preserve edges and, suitably chosen, can preserve line structure [9]. In particular, "constant sets" are preserved while "impulse sets" are transformed into constant sets.

- Rank-order filters, closely related to morphological filters, play an analogous role for the mean absolute error criterion that linear filters do for the mean square error criterion [11]. In particular, the median filter gives the minimum mean absolute error estimate of a constant signal corrupted by zero mean symmetric noise.

For our present purposes, the principal implication of these properties is that it suggests that morphological filters should be well suited for applications which can exploit the trade off between good noise suppression (properties 1, 2, 3) and structural preservation (property 4).

## An Illustrative Example

We noted above (see for example Figure 1) that the *comoc* filter performs efficiently and well on a variety of images. Property 3 was the starting point of the current investigation. If the input cumulative pdf is symmetric, then it is appropriate to use the *comoc* filter. This is often the case in image and signal processing, which explains the good performance shown in Figure 1. In particular, it is usually the case for a step edge that separates two un-textured regions.

Suppose, however, that the noise distribution is *asymmetric*, which might be the case for “edges” that separate textured regions. Consider, for example, the one-dimensional signal shown in Figure 2. The underlying signal is constant, with value 127. For  $x \leq 64$ , the noise is a zero-mean (symmetric) Gaussian process  $G(0, \sigma)$ ; but for  $x > 64$ , the noise is “rectified”  $|G(0, \sigma)|$ . As can be seen from Figure 2b, the *oc* filter approximates the signal more closely than the mean, median, *co*, or *comoc*, indeed it is close to the true value 127. On the other hand, the *co* filter smooths the signal, but still essentially follows the noise. The *comoc* performs more poorly than the *oc* in this case. The observations reported earlier, and shown in [1] are illustrated in Figure 2c by the performance of the *comoc* filter on the symmetric noise.

How could one detect, and possibly localise, the Gaussian step in Figure 2? Clearly, we seek to measure significant change in the output of the various symmetric and asymmetric filters after their application to an image. A key aspect of our approach is that we pay minimal attention to numerical values of pixels, as would be required for parametric statistics (eg the variance and mean), and statistical tests derived from them (eg the ‘t’-test, or ‘F’-test). Instead, morphology emphasises *relative* values and rank orders (eg *min*, *max*). This is, of course, the underlying reason why morphological filters are related to rank-ordered filters, but are better at preserving image structure. Consequently, we turn to *non*-parametric statistics [12].

We use the (Wilcoxon-)Mann-Whitney test, which we invoke as a NAG routine (there is no obvious advantage in using the more powerful Kolmogorov-Smirnov test for our purpose). Roughly, it is the nonparametric analogue

of the Student ‘t’-test. The Mann-Whitney test is used to test the null hypothesis  $H_0$  of identical distributions. To apply the test in a horizontal direction (for example) at a pixel  $(x, y)$  in a filtered image, we compare the  $7 \times 3$  window to the left of the pixel with that to the right of the pixel. Let the window to the left be called  $A$ , and that to the right be  $B$ . The values in both  $A$  and  $B$  are jointly sorted in increasing rank. The Mann-Whitney  $U$  statistic is defined as the number of times a value in  $B$  precedes one in  $A$  in this sequence. Efficient techniques exist to compute this sum [12]. According to the critical value of  $U$  corresponding to a preselected significance level  $\alpha$  of its probability distribution, we choose to reject  $H_0$  in favour of one, of two possible, one-sided alternative hypotheses; (1)  $F_A > F_B$ , or (2)  $F_B > F_A$ . The alternative 1, for example, implies that the values in  $A$  are stochastically larger than those in  $B$  which is a family of alternatives including as a special case  $\mu_A > \mu_B$ . Figure 3 shows the result of applying the Mann-Whitney test (with  $\alpha$  chosen to be 0.002) to an “absolute Gaussian step” image of which Figure 2 shows a single row. The *oc* filter picks out the step change (as is foreshadowed in Figure 2d) much more successfully than the *co* filter. Note that the *comoc* filter would perform similarly to *co*, and hence so would the median.

## Experiments

The remainder of this initial study addresses the following problems:

- Can textures give rise to asymmetric noise distributions?
- If so, do the biased *oc* and *co* filters give better noise suppression than the *comoc* filter without sacrificing localisation?
- How do the biased filters perform relative to the median filter?
- Can one determine, in practice, that the input cumulative noise distribution is best considered asymmetric, and adaptively filter the image structure by the *co*, *oc*, or *comoc*?
- Is it useful to treat texture boundaries as transitions between two, possibly asymmetric, noise distributions, and process them accordingly?

Positive answers are presented to each of these questions.

## Signal Structure Over Scale

For reasons of space, we confine attention to the texture boundary shown in Figure 4. As noted in the previous section, our goal is *not* to determine texture boundaries by parameter identification of the flanking textures followed by a significance test on parameter differences [13]. To give a feel for the complexity of the signal analysis, however, we show in Figures 5 representative cross sections of (what we know to be) the flanking textures.

We first illustrate the observed behaviour of morphological filters, on a real texture.

Experimental filter characteristics were determined as follows. Each image was smoothed with a one-dimensional filter applied in the horizontal direction (vertical for the rotated image example). Filter sizes of equivalent filters are for the mean and median  $2n+1$  and the morphological filters  $n$ . The average output after smoothing, computed over a  $32 \times 32$  neighbourhood was used to quantify the level of bias of each filter. As can be seen from the Figure 6a, changes in bias of morphological filters over scale characterise the underlying signal structure in a quite different manner from the mean, median or *comoc* filters. As one would expect, similar trends in behaviour are found using rank-order filters [14].

Reduction in variance was measured by the mean absolute deviation (MAD) between the original and filtered versions. In these experiments, the MAD was computed over the filter support  $2n+1$  and then averaged over the same neighbourhood as used to estimate the average value. The usual measure of dispersion for a distribution is the root mean square sum of deviations, which is least if the deviations are measured from the mean. The mean absolute deviation is least when the deviations are taken from the median. If the distribution is symmetric or moderately skewed, and the deviations taken from the median, the MAD gives a biased estimate of the standard deviation. For the Normal case, correction factors to convert the sum into an unbiased estimate can be looked up in Statistical Tables. In interpreting the MAD graphs in Figure 6b, it should be remembered that the *co* and *oc* filters are biased signal estimators which influences the value of the MAD. Clearly, the MAD for the symmetric filters (mean, median *comoc*) will also be influenced by the bias introduced by asymmetric distributions. Observe that in all the examples the MAD of the *comoc* filter is always less than at least one of the *co* and *oc* filters, as predicted by our analysis.

## Detecting Texture Differences

Figure 7 shows a collage of four textures. The nonparametric approach proposed in this paper works directly from the intensity surface. In contrast, most schemes proposed in the computer vision literature for texture segmentation, operate upon edge features [10,15]. To illustrate the difficulty of applying such a technique to stochastic textures, Figure 8 shows the edges found by the Canny detector. Separating the wool from the wood would be difficult.

Next, the image shown in Figure 7 is processed in turn using the biased filters *oc* and *co* of size  $n = 3$ . The Wilcoxon-Mann-Whitney test is applied as described in the previous section (with a significance level  $\alpha = 0.001$ ). The minimum of the probability levels from both filters is used. Pixels where  $F_A > F_B$  are shown in black, and where  $F_B > F_A$  are shown in white. This corresponds to

the contrast sign of “edges”, but here the measurements are of the statistical significance of the difference rather than of magnitude (size). The result is shown in Figure 9. Figure 9c indicates the spatial organisation of sites at which there is significant sample difference.

## Adaptive Smoothing

Figure 10a shows the “smoothest” filtered image resulting from applying the biased filters. More precisely, let the sample sets  $A$  and  $B$  consist of point differences between the median and the *oc* and *co* responses respectively. We apply the Mann-Whitney test to assert the “best” underlying point estimate of the underlying signal, that is, the *co* response if  $F_A > F_B$  or the *oc* response if  $F_B > F_A$ . Otherwise we accept the hypothesis that the original signal was symmetric and use the median response. Figure 10b is the absolute value of the difference between the original image (Figure 7) and the smoothed version shown in Figure 10a. It shows that the textural variation has been effectively separated out from the underlying smooth grey value surface, from which the texture boundary can be found.

## Concluding Remarks

The paper began with an analysis of asymmetric noise and the detection of steps using the nonparametric Mann-Whitney test. The previous section applied the technique to texture segmentation. We suggest that a texture boundary can be constrained to lie between the positions predicted by what are known a priori to be biased filters. The uncertainty in the position of a texture, its poor localisation relative to a step that separates un-textured edges, is consistent with intuition, and corresponds to a pair of one-sided statistical analyses of the signal.

In this paper, we have concentrated on processing techniques rather than the specific application to texture segmentation. Considerably more work needs to be done to analyse texture boundaries and to compare the performance of our technique to those previously published [16,17]. There are many limitations to this preliminary report on our approach and technique. For example:

- We have concentrated on one-dimensional analyses, seeking significant change in the horizontal direction. We need to extend the technique to two dimensions, seeking the most significant directional change.
- We treat each row of the image independently. Instead, we need to add the constraint of edge continuity.
- As usual, there is an uncomfortable number of thresholds, significance levels, and scales. Though the treatment is quite robust to changes, choices need to be rationalised, and possibly determined automatically.
- We suggest that the MAD is a useful criterion for choosing between the *co* and *oc* filtered images. This requires further investigation.







Figure 1: (a) A range image taken by a triangulation laser range finder developed jointly with NEL [Brady et. al. 1988]. Black points correspond to no data, either because of specular reflections or because of laser shadowing effects. (b) The result of smoothing the image in a. with a comoc filter. (image size  $220 \times 220$ )



2(d)




Figure 2: (a) A one-dimensional signal with symmetric and asymmetric noise. The underlying signal is constant with zero-mean additive Gaussian noise for  $x \leq 64$  and, for  $x > 64$ , everywhere positive noise whose amplitude is the modulus of a Gaussian process. (b) The co filter primarily responds to the noise whereas the oc approximates the signal more closely. Compare, for example the circled regions. (c) The observations reported earlier, and shown in [Noble 1989] are illustrated by (c) the performance of the comoc filter on the symmetric noise. (d) The filter output characteristics in the assymetric noise (left of this caption), show that the oc follows the underlying constant signal more closely.




Figure 3: (a) The result of applying the Mann-Whitney test after filtering an absolute Gaussian step with a 3x3 filter. White points are deemed significant difference between the populations A and B (see text). (b) The same process applied to the co filter. (image size 54 x 54, filter size is  $n = 3$ )




Figure 4: A texture boundary between wood and wool (image size 154 x 196).




Figure 5: Representative profiles through the texture at locations A and B shown in Figure 4.

## References

- [1] Noble, J.A. "Descriptions of Image Surfaces," *D.Phil thesis, Oxford University (in preparation)*.
- [2] Sternberg, S.R. "Grayscale Morphology," *Computer Vision, Graphics, and Image Processing*, Vol. 35, No. 3, (1986), pp. 333-355.
- [3] Brady, J.M., McIvor, A.M., Brint, A., Foulkes, P., Scott, G. & Dickson, W. "Vision and the Oxford AGV," *Proceedings of Image Processing '88*, London, November 1988, pp. 267-286.
- [4] Reid, I. "Interpreting Range Images," *Technical Report, Oxford University (in preparation)*.





Figure 6: (a) The average output and, (b) the MAD, for the filters discussed in this paper applied to the wood texture of Figure 4.




Figure 7: A texture collage consisting of (clockwise from top left) patches of wool, wood, bread, and fur (image size 176 x 176).




Figure 8: Canny edges for the texture shown in Figure 7 ( $\sigma = 1.5$ , hysteresis threshold ratio 3 : 1). Distinguishing wool from wood is quite difficult using an edge-based image representation.




Figure 9: (a) Sites of significant change in filtered samples of textures shown in Figure 7, where  $F_A > F_B$  (b) Similar, but where  $F_B > F_A$ . (c) Significant changes corresponding to cases (a) and (b). (filter size  $n = 3$ ,  $\alpha = 0.001$ )

[5] Noble, J.A. "Morphological Feature Detection," in *Proceedings of the IEEE 2<sup>nd</sup> International Conference on Computer Vision*, 1988, pp.112-116.

[6] Maragos, P. & Schafer, R.W. "Morphological Filters Part I - Their Set-Theoretic Analysis and Relations to Linear Shift-Invariant Filters," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. 35, No. 8 (1987),

[7] Maragos, P. & Schafer, R.W. "Morphological Filters Part II - Their Relations to Median, Order Statistics, and Stack Filters," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. 35, No. 8 (1987), pp.1170-1184.

[8] Serra, J. "Image Analysis and Mathematical Morphology: theoretical advances," *Academic Press*, Vol. 2, (1988).




Figure 10: (a) The "smoothest" filtered image produced from the image shown in Figure 7. (b) The absolute value of the difference between (a) and Figure 7. (filter size  $n = 5$ )

[9] Stevenson, R.L. & Arce, G.R. "Morphological Filters: Statistical and Further Syntactic Properties," *IEEE Transactions on Circuits and Systems*, Vol. 34, No. 11, (1987), pp.1292-1305.

[10] Voorhees, H. & Poggio, T. "Detecting Textons and Texture Boundaries in Natural Images," *Proceedings of the IEEE 1<sup>st</sup> International Conference on Computer Vision*, 1987, pp.250-258.

[11] Coyle, E.J. "Rank Order Operators and the Mean Absolute Error Criterion," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. 36, No. 1 (1988), pp.63-76.

[12] Siegel, S. "Nonparametric Statistics for the Behavioral Sciences", McGraw-Hill, (1956), Chapter 6.

[13] Kashyap, R.L. & Eom, K.-B. "Texture Boundary Detection Based on the Long Correlation Model," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 11 No. 1 (1989) pp.58-67.

[14] Heygster, G. "Rank Filters in Digital Image Processing," *Computer Vision, Graphics, and Image Processing*, Vol. 19, No. 2, (1981), pp.148-164.

[15] Vilenrotter, F.M., Nevatia, R. & Price, K.E. "Structural Analysis of Natural Texture," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 8 No. 1 (1986) pp.76-89.

[16] de Souza, P. "Edge Detection Using Sliding Statistical Tests," *Computer Vision, Graphics, and Image Processing*, Vol. 23, No. 1, (1983), pp.1-14.

[17] Malik, J. & Perona, P. "A Computational Model of Texture Segmentation," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1989, pp.326-332.