SPECKLE FILTERING IN ULTRASOUND IMAGES FOR FEATURE EXTRACTION

A N Evans and M S Nixon

Department of Electronics and Computer Science, University of Southampton, UK

INTRODUCTION

In ultrasonic imaging, speckle degrades images by concealing fine structures and reducing the signal to noise ratio (SNR). The theoretical aspects of speckle, both laser and ultrasound, have been dealt with by several authors including Goodman [7], Burckhardt, [3] and Wagner [15], although it should be noted that in any practical ultrasound imaging system effects such as nonlinearities and envelope smoothing can result in deviation from theoretical models (Tuthill *et al* [13] and Waag *et al* [14]).

Removing speckle improves the SNR and provides a more suitable basis for computer based image interpretation using feature extraction techniques, for example the work of Nixon *et al* [11] and Helin and Ayache [8]. The effect of a speckle reduction technique on the position of edges is of critical importance if the resulting image is to be used as an input to a feature extraction stage.

Various techniques have been applied to reduce the speckle content of ultrasonic images including compounding in which multiple images taken from varying angles, or with different centre frequencies, are averaged and phase-based methods after Forsburg *et al* [6] where phase and amplitude characteristics are directly manipulated in individual returning A-lines.

There are many established image processing techniques that can be used to reduce the noise in images of which the median filter is arguably the most popular nonlinear method. However, although it successfully removes speckle noise it also modifies edges present within the image. For fully developed speckle, when the number of scatterers present within the imaging system resolution cell is large, the image statistics are constant but as the number of scatterers decrease and at boundaries between features the local statistics vary and can be used to control the degree of filtering.

The ability to filter heavily in uniform regions while preserving edges is found in the adaptive class of filters, resulting in the adaptive weighted median filter of Loupas *et al* [10] and the unsharp masking filter of Bamber and Daft [1]. These filters use local image statics to vary the output between the original value and the mean or median of a local region. In both cases the ratio of the local mean to the local variance controls the proportion of each, with the output tending towards the original value where the ratio is low, ie at an edge, thus improving the performance of the filter at discontinuities. Unfortunately as the complexity of the filter increases so does the computational cost, motivating the development of specialist hardware (Loupas [9] and Bamber and Phelps [2]) if a real-time implementation is to be achieved, which is desirable as ultrasound itself is a real-time imaging system. In addition, both filters require prior choice of a controlling parameter which in practice requires an initial pass of the filter over the image. Although for series of similar images a common constant can be used, these adaptive filters require some degree of user interaction and are not therefore completely automatic.

This paper introduces use of the truncated median filter of Davies [5] to filter ultrasound images, contrasting the results achieved with those of other current filters for speckle reduction. The filters used for comparison were a standard median filter, the adaptive median filter [10] and the adaptive unsharp masking filter [1].

The motivation behind this research is to determine whether a simple non-adaptive filter can achieve acceptable results with a lesser degree of computational complexity than more advanced adaptive filters with fully automatic implementation.

FILTERING TECHNIQUES

The weighted median filter extends the simple concept of the median filter by the introduction of weight coefficients. This operation is performed by multiplying each term of the ranked pixel intensity values, $\{f_1, f_2, \ldots f_N\}$, by its respective weight, denoted $w_1, w_2, \ldots w_N$, to form an extended sequence $\{w_1 f_1, w_2 f_2 \ldots w_N f_N\}$. For a square window with sides of length $2K+1$ the weight coefficient $w_{i,j}$ at position (i,j) is given by:

$$w_{i,j} = [W_{(K+1,K+1)} - cd\sigma^2/\bar{x}]$$

where $W_{(K+1,K+1)}$ is the central weight and the square brackets, [ ], denote the nearest integer to

the term inside the brackets or zero if the result of the addition within them is negative; c is a scaling constant and d is the distance of the point from the centre of the window. The local mean is denoted by $\bar{x}$ and the local variance by σ^2.

In the adaptive mean-based filter the measure of similarity, p, given by the ratio of the local variance over the local mean, is used to control a constant k such that $0 \leq k \leq 1$ by

$$k = (gp - \bar{p}_s)/p$$

where $\bar{p}_s$ is the mean value of p for a local area of an image that is considered to consist of fully developed speckle. The constant g is a scaling factor that ensures k is always less than one. K adaptively varies the output $\hat{x}$ between the original value (x) and the local mean ($\bar{x}$), such that

$$\hat{x} = \bar{x} + k(x - \bar{x})$$

THE TRUNCATED MEDIAN FILTER

The truncated median filter was first proposed by Davies [5] and approximates to a modal filter. The most frequent value in any distribution is the mode, but the difficultly of determining the mode of a small population, such as that contained within a 9 × 9 mask has meant that whilst the median filter enjoys great popularity the modal filter has largely been ignored. Near edges the modal filter will suppress noise, where little is present, or enhance the edge in the case of high noise. When directly on an edge the mode can be either one of two peaks as the population is essentially bi-modal and this can result in rather coarse decisions being taken. The mode thus appears to offer a improvement over the median filter near to edges but may give erratic results when directly on the edge.

The order of the mode, median and mean can be used to estimate the mode for many distributions including those of ultrasound images. Figure 1 illustrates that if the median is moved away from the mean then it will approximate to the mode. To move the median, Davies proposed to truncate the original frequency distribution graph so that the median bisects the range of distributions and a new median value found that is closer to the mode. This process is iterative, each pass of the filter moving the truncated median output further towards the mode, although in practice one pass has been found to be sufficient. Additionally, if the filter is only applied once the output errs towards the median thus overcoming the problem of the mode being too harsh at edges.

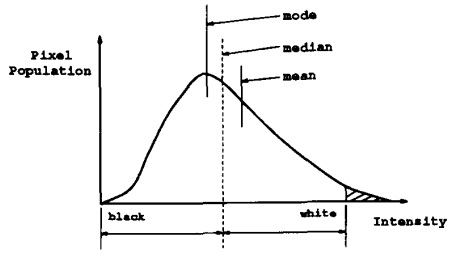

Figure 1: Normal Order of Mean, Mode and Median in a Unimodal Distribution and Method of Truncation

EXPERIMENTAL RESULTS

The four speckle reduction filters were applied to test images using a 9 × 9 mask. To allow accurate evaluation of the truncated median filter in comparison with contemporary techniques not only subjective but also objective differences need to be appraised. One such quantitative measure is the Contrast to Speckle Ratio (CSR) of Patterson and Foster [12], a measure of the observer's ability to observe a void or dark area against a background of speckle, defined as

$$CSR = \frac{\bar{x}_o - \bar{x}_i}{(\sigma_o^2 + \sigma_i^2)^{1/2}}$$

where $\bar{x}_i$ and σ_i^2 are the average signal value and the variance respectively inside the void and $\bar{x}_o$ and σ_o^2 are those outside the region. In addition to the CSR, local SNR measurements were also taken, defined as

$$SNR = \frac{\bar{x}}{\sigma}$$

where $\bar{x}$ and σ denote mean and standard deviation respectively. Both the CSR and SNR were calculated over a 9 × 9 region, matching the filter mask size. The effect of filtering on the position and strength of edges present was evaluated by the use of the Canny edge detector [4]. However in an *in vivo* image, correct position of the edges is not known and it is therefore impossible to determine whether the edges have been modified by the filters. To overcome this limitation a circular pipe in a water bath, modelling an artery, was used as a phantom providing a real B-scan image containing an object of known size and shape.The pipe used was hard plastic, 250mm long and 28mm in diameter. With an ideal imaging system the B-scan image would contain a complete circle corresponding to the pipe. The actual image only approximates to this but as the exact position of the edge is known this provides a standard for comparison of the filtered results.

Filter Type	Median	Truncated Median	Adaptive Weighted	Unsharp Masking
CSR	18.3	15.3	18.3	10.5
Improvement (dB)	21.0	19.4	21.0	16.2
SNR On Edge	1.63	1.26	1.37	2.31
Improvement (dB)	1.5	-0.7	-0.03	4.6
SNR Off Edge	73.7	62.9	73.7	44.6
Improvement (dB)	20.2	18.9	20.2	15.9

Figure 2: CSR and SNR of In Vivo Image

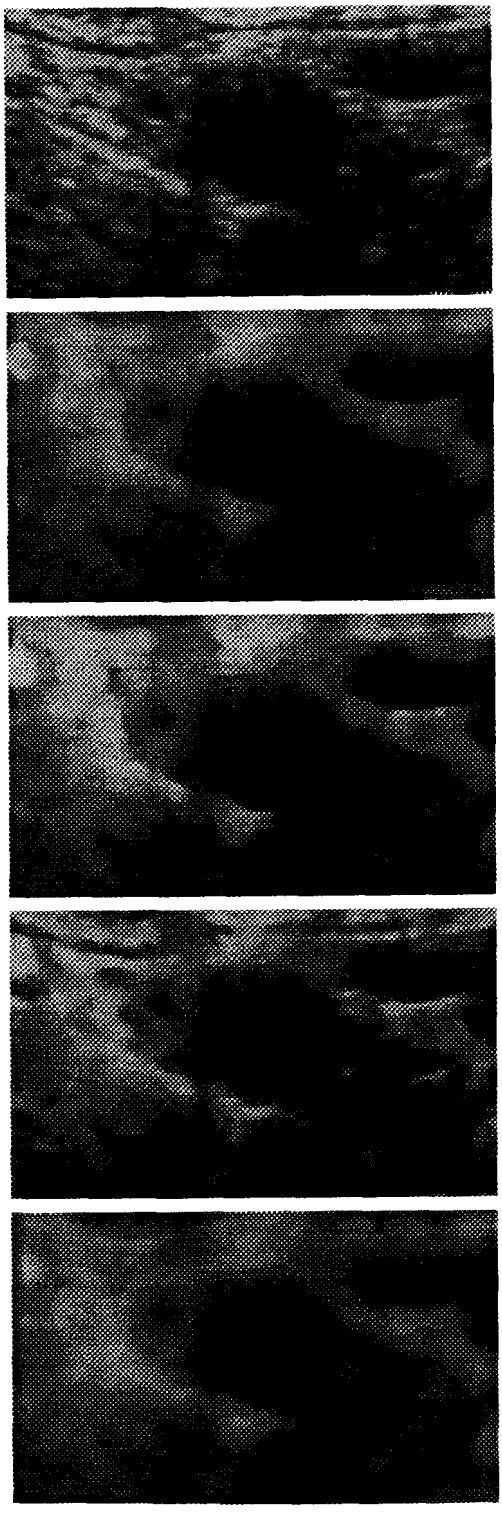

Figure 3 (opposite): In Vivo Images (from top to bottom): (a) original, (b) median filtered, (c) truncated median filtered, (d) adaptive weighted median filtered, (e) unsharp masking filtered

In Vivo Test Image

Figure 3 presents the original and filtered images of an *in vivo* scan of the carotid artery. Cosmetic improvement to the original image is achieved by all the filtering techniques although the edge blurring by the unsharp masking filter can be observed at the artery/tissue junction in figure 3 (e). In the median and adaptive median filtered images the edges around the artery are clearly defined with the strongest edge being displayed be the truncated median filter. The adaptive weighted filter clearly preserves features best and has been able to highlight small structures removed by other techniques, see for example the dark vein in the top left corner of figure 3 (d).

The Contrast to Speckle Ratio (CSR) and Signal to Noise Ratio (SNR) are used to interpret the results, figure 2. Here the CSR is calculated using the centre of the artery to calculate $\bar{x}_i$ and σ_i^2 and an area consisting solely of speckle to calculate $\bar{x}_o$ and σ_o^2; the SNR at an edge was found at the tissue/artery boundary. The unsharp masking filter consistently offered the smallest improvement in CSR and SNR away from an edge and also showed the largest improvement in SNR at an edge, indicating the most edge modification.

The CSR and SNR away from an edge are improved by $21.0dB$ and $20.3dB$ respectively by both the median and adaptive median filters, with the truncated median filter offering a slightly lower improvement in each case. At an edge the median filter preserves

Filter Type	$\frac{Number of Matching Edge Points}{Total Number of Edge Points}$
Median	25.8
Truncated Median	28.1
Adaptive Weighted	25.3
Unsharp Masking	12.6
Original	4.9

Figure 4: Ratio of Number of Edge Points Matching Ideal Edge to Total Number of Edge Points

least edge information (i.e. the greatest change in SNR), then the truncated median filter and the adaptive weighted filter and the appears to give the best edge preservation, with virtually no change to the SNR. The ideal edge model is not known but by observation of the edge detected images in figure 5, conclusions can be drawn about the edge preserving properties of the filters.

The adaptive weighted and truncated median filters have preserved more complete edges than the other techniques, for example the detail to the left of the artery in the figure 5 (b) and (c) which is not present in the other images. The adaptive median technique shows the greatest edge detail, as would be expected from its nature. Although the truncated median filter has less edge detail, those edges that are present are strongly defined and continuous, comparing favourably with the median filter. Fewest edges are determined by the unsharp masking filter, illustrated in figure 5 (d) and those that can be seen are rough and insubstantial.

Phantom Test Image

The edge detected images overlayed with the edges expected from an ideal imaging system in figure 7 show a close coupling between ideal and actual results although figure 7 (d) shows edges due to speckle that have been removed by the other techniques, illustrating the disadvantage of mean based filtering.

To quantify how close the resulting images are to the overlayed image, the ratio of the number of edge

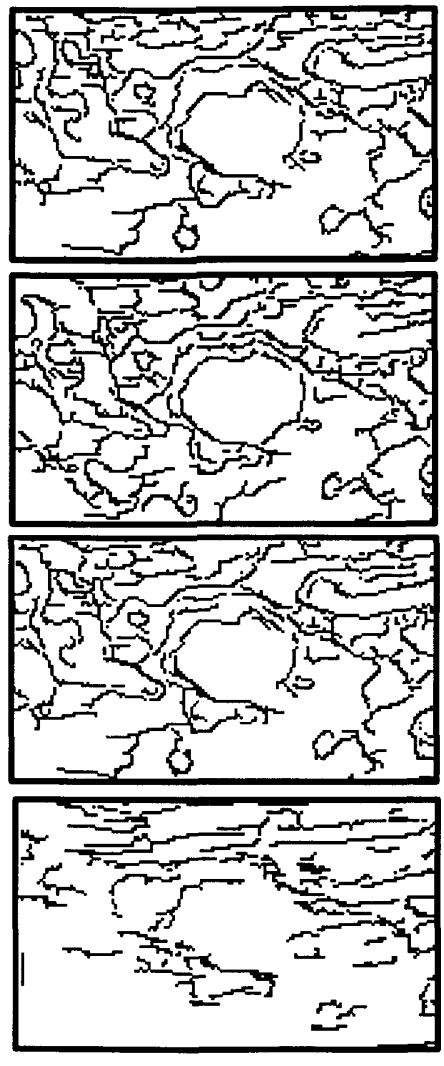

Figure 5: In Vivo Images Edge Detected (from top to bottom): (a) median filtered, (b) truncated median filtered, (c) adaptive weighted median filtered, (d) unsharp masking filtered

Filter Type	Median	Truncated median	Adaptive Weighted	Unsharp Masking
Relative Processing Time	2.3	2.9	5.4	1

Figure 6: CPU Times for Filtering Techniques

points that correspond to points in the ideal edge image to the total number of edge points was taken for each edge detected filtered image, the results being tabulated in figure 4. The truncated median technique shows the highest percentage, then the median filter which is just above that of the adaptive weighted median filter. The unsharp masking filter performs considerably worse than the other three methods.

Processing Times

The relative processing time for each technique, implemented is software, is used as a measure for comparison, see figure 6. Although these figures are obviously implementation dependent it can be seen that relative to the fastest technique the median and truncated median filters have similar values while the adaptive weighted median filter's complex nature is reflected in its long processing time.

CONCLUSIONS

The truncated median filter has been applied to ultrasound images and its performance evaluated using qualitative and quantitative methods, based on *in vivo* and phantom test images. The use of a phantom to model the carotid artery provided an exact basis for comparison of filtered edge data with edge data derived from the raw images.

It has been found that the truncated median filter improves upon the performance of the median filter, especially at edges, with little increase in computational cost and therefore appears very suited to become part of the stock of tools for speckle reduction for enhanced image analysis. The smoother, more continuous edges produced by the truncated median filter are particularly attractive if a further computer vision stage, such as feature extraction, is to be applied. The filter result at the arterial wall surpasses that of the median filter and the adaptive unsharp masking filter, highlighting the deficiencies of mean-based filtering, and is comparable with that of the adaptive weighted median filter without the computational expense afforded by adaptivity. In addition the truncated median filter can be used in a fully automated mode with no requirement for the user to preset a constant value, offering an advantage over the adaptive class of filters.

Acknowledgements

The authors would like to thank T.K. Hames of Southend Hospital for his advice and suggestions and SERC for the funding that made this work possible.

REFERENCES

1 Bamber, J., and Daft, C. 1986. Ultrasonics. 41–44.

2 Bamber, J., and Phelps, J. 1991. Ultrasonics. 218–224.

3 Burckhardt, C. 1978. IEEE Trans Son Ul. 25. 1–6.

4 Canny, J. 1986. IEEE Trans PAMI. 8. 679–698.

5 Davies, E. 1988. Pattern Recognition Letters. 7. 87–97.

6 Forsberg, F., Healey, A., Leeman, S., and Jensen, J. 1991. Phys Med Bio. 36. 1539–1549.

7 Goodman, J. 1976. J Opt Soc A. 66. 1145–1150.

8 Herlin, I., Ayache, N. 1992. Image and Vision Comp. 10 673–682.

9 Loupas, T. 1988. PhD thesis, University of Edinburgh.

10 Loupas, T., McDicken, W., and Allan, P. 1989. IEEE Trans Cir Sys. 36. 129–135.

11 Nixon, M., Hames, T., Martin, P., Powell, S., and de Paiva, M. 1992. In IEE Image Processing and its Applications, Maastricht, 373–376.

12 Patterson, M., and Foster, F. 1983. Ultrasonic Imaging. 5. 195–213.

13 Tuthill, T., Sperry, R., and Parker, K. 1988. Ultrasonic Imaging. 10. 81–89.

14 Waag, R., Demczar, B., and Case, T. 1991. IEEE Trans Biomed Eng. 38. 628–633.

15 Wagner, R., Smith, S., Sandrik, J., and Lopez, H. 1983. IEEE Trans Son Ul. 30. 156–163.

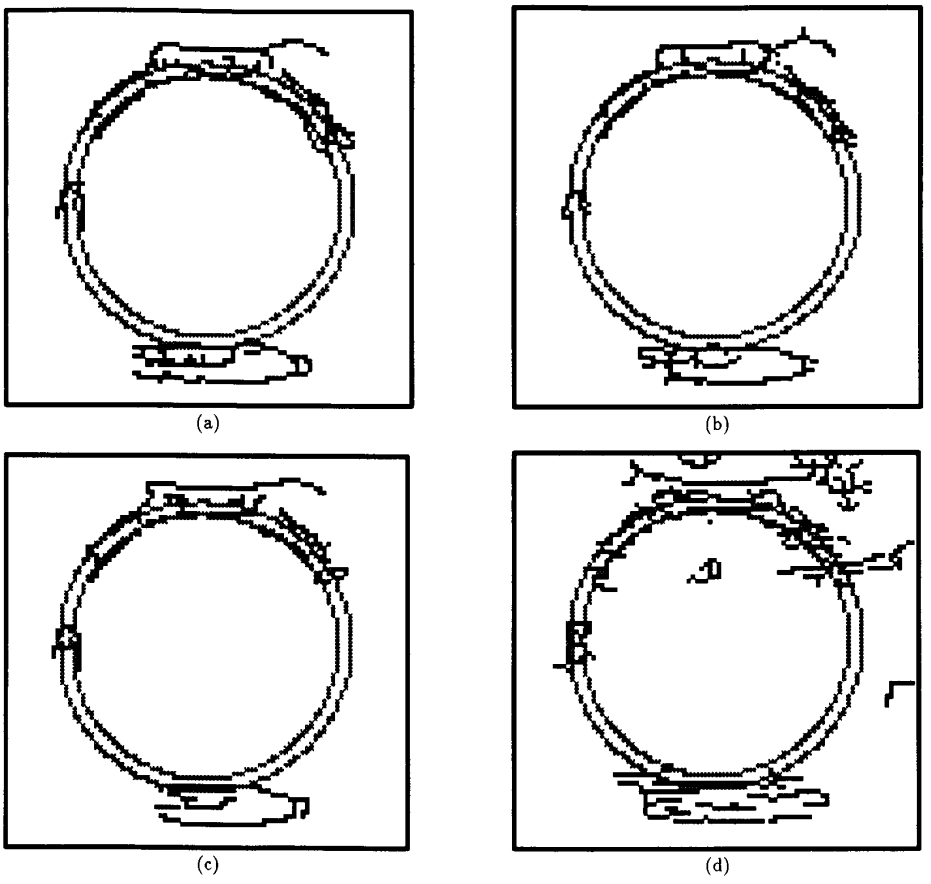

Figure 7: Filtered Phantom Images Edge Detected (black) and Ideal Image Edge Detected (grey): (a) median, (b) truncated median, (c) adaptive weighted and (d) unsharp masking