EDGE DETECTION IN ULTRASOUND SPECKLE NOISE

Richard N. Czerwinski *Douglas L. Jones* *William D. O'Brien, Jr.*

Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, Illinois, 61801, USA *

ABSTRACT

In this paper, we present a statistical approach to edge detection in ultrasound speckle, and use actual noise statistics to derive an expression for an optimal detection rule. We compute the optimal detector for the special case of uncorrelated speckle, and an approximation to the optimal detector for the case when signal-to-noise ratio (SNR) is high. We also discuss a simple sub-optimal detection scheme, and show that its performance is close to that of the optimal uncorrelated speckle detection rule, and surpasses that of the high-SNR approximation except in the case of extremely high SNR, in which case they both perform equally well.

1. INTRODUCTION

The problem of automatically detecting regions of interest in an ultrasonogram is of fundamental importance in the design of medical diagnostic or non-destructive evaluation systems. However, segmentation algorithms require accurate edge maps for good performance and the highly signal dependent nature of ultrasound speckle makes these difficult to obtain.

In this paper, we derive the statistically optimal detection rule for tissue boundaries in ultrasound speckle. While implementation of the exactly optimal detector in arbitrarily correlated speckle noise is computationally infeasible, we derive an optimal detector for the special case of uncorrelated speckle and an approximation to the exactly optimal test in the case where the signal-to-noise ratio (SNR) is high.

Finally, we evaluate the performance of the optimal and approximate detection rules, and compare those schemes to the Sticks procedure discussed in [1]. We find that Sticks performs almost as well as exactly optimal detection, and since that procedure is not statistically motivated, it is relatively immune to noise modeling errors that can reduce the performance of the optimal detectors. Those properties, along with its computational efficiency, make it the most successful technique of any we describe here.

*This work was supported in part by the National Live Stock and Meat Board, and the United States Department of Agriculture.

2. ULTRASOUND BOUNDARY DETECTION

For many types of imagery, such as photography, the edges of interest are between regions of different brightness, so that edge detection is normally approached by looking for abrupt discontinuities in pixel intensity. These techniques typically compute the average pixel intensity in small regions on either side of a line through each pixel, vary the line orientation, and decide if an edge is present or not by comparing the maximum difference or ratio of the averages to a threshold. Examples of operators that are frequently used are the Roberts and Sobel operators [2]. A related technique by Bovik [3] evaluates a ratio of averages operator at those points where a Laplacian of a Gaussian filter produces a zero-crossing. This technique performs extremely well in synthetic aperture radar images, and is the method of choice for detecting step discontinuities in speckle.

In ultrasound imagery, however, the edges are caused by a fundamentally different process, and thus have a different appearance. The edges represent the places where the incoming sound wave encounters an acoustical impedance mismatch, for example at a tissue boundary. Thus, the edges appear as bright streaks against a dark background. In a sufficiently small region about each point on an edge, the edge looks like a "stick," a straight line at some orientation [1]. The optimal detection schemes derived in this paper assume the edges have this character and invoke the speckle model developed by Goodman [4], Burckhardt [5] and Wagner *et al.* [6] to derive the higher order statistics of the speckle.

3. MODEL FOR SPECKLE NOISE

Speckle noise is a phenomenon caused when a coherent imaging system such as ultrasound is used to image a surface which is rough on the scale of the wavelength used. The surface produces many reflections in each resolution cell which can add constructively or destructively to produce the speckle pattern.

At every point, an ultrasound speckle image is the magnitude of a complex Gaussian field with real and imaginary parts independent and identically distributed. In the case of "fully developed" speckle (*i.e.* an image region in which every resolution cell contains a large number of small scatterers and no coherently reflecting targets), the real and imaginary parts have zero mean; when a target is present,

0-8186-6950-0/94 $4.00 © 1994 IEEE

the field has a complex-valued mean but identical correlation structure to the noise only case [4, 5, 6]. For the simulations in this paper, we have taken the mean of the underlying Gaussian field under the ith hypothesis to be the ith orientation stick, and the correlation function to be a Gaussian with known correlation. In practice, the form of the autocorrelation function can be calculated from physical characteristics of the transducer and pulse shape.

For example, a fully developed speckle pattern obtained with a square transducer aperture and a Gaussian envelope pulse would be the magnitude of a complex Gaussian field with zero mean, independent real and imaginary parts, each with a correlation function with a Gaussian shape in the range direction and a sinc squared taper in the transverse direction.

4. THE LIKELIHOOD RATIO FUNCTION

Optimal M-ary hypothesis testing is the selection of the hypothesis which maximizes the likelihood ratio function [7],

$$\Lambda_i(x) = \frac{P[x|H_i]}{P[x|H_0]}, \qquad (1)$$

the ratio of the probability density of a particular realization x under hypothesis i (H_i) to its probability density under the null hypothesis (H_0). In the stick detection problem, we wish to distinguish between hypotheses H_1 through H_M, each of which corresponds to the presence of a straight line at some orientation, and the null hypothesis H_0. The different hypotheses are characterized statistically by

$$H_i : \mathbf{x} = |\mathbf{n}_c + \mu_i|^2 \qquad (2)$$

$$H_0 : \mathbf{x} = |\mathbf{n}_c|^2, \qquad (3)$$

where $\mathbf{n}_c = \mathbf{a} + j\,\mathbf{b}$ is a zero mean complex Gaussian random vector whose real and imaginary parts are Gaussian, independent and identically distributed, and μ_i is the stick at the ith orientation.

Note that we have "unrolled" the square $N \times N$ image regions into $N^2 \times 1$ vectors for ease of manipulation. This is done without loss of generality since all the spatial correlations of the two dimensional discrete Gaussian random field can be expressed in a correlation matrix. Note also that we will be performing the hypothesis test on the *squared* magnitude of the image rather than the magnitude. This invertible transformation is done to simplify computation.

The probability density function of the image vector $\mathbf{x}$ under hypothesis i can be computed from the underlying normal distribution on the constituents of $\mathbf{x}$:

$$P_{\mathbf{X}}[\mathbf{x}|H_i] = \lim_{\delta\to 0} \frac{1}{\prod_i 2\delta_i} P\left[\mathbf{x} - \delta \le X \le \mathbf{x} + \delta\right] \qquad (4)$$

$$= \lim_{\delta\to 0} \frac{1}{\prod_i 2\delta_i} P\left[\mathbf{x} - \delta \le (\mathbf{a} + \mu_i)^2 + \mathbf{b}^2 \le \mathbf{x} + \delta\right] \qquad (5)$$

$$= \int_{-\infty}^{\infty} p_{\mathbf{B}}(\mathbf{b}) \lim_{\delta\to 0} \frac{1}{\prod_i 2\delta_i} \left\{ P\left[\sqrt{\mathbf{x} - \delta - \mathbf{b}^2} - \mu_i \le \mathbf{a} \le \sqrt{\mathbf{x} + \delta - \mathbf{b}^2} - \mu_i\right] \right. \qquad (6)$$

$$\left. + P\left[\sqrt{\mathbf{x} - \delta - \mathbf{b}^2} - \mu_i \le \mathbf{a} \le \sqrt{\mathbf{x} + \delta - \mathbf{b}^2} + \mu_i\right] \right\} d\mathbf{b}$$

$$= \int_{-\infty}^{\infty} p_{\mathbf{B}}(\mathbf{b}) \frac{1}{\prod_i \sqrt{x^{(i)} - (b^{(i)})^2}} \left\{ p_{\mathbf{A}}\left(\sqrt{x^{(i)} - (b^{(i)})^2} - \mu_i\right) + p_{\mathbf{A}}\left(\sqrt{x^{(i)} - (b^{(i)})^2} + \mu_i\right) \right\} d\mathbf{b}, \qquad (7)$$

expressed in terms of an integral over $\mathbf{b}$, a nuisance parameter with known Gaussian distribution.

The density function under the null hypothesis is obtained from (7) by setting μ to the zero vector:

$$P_{\mathbf{X}}[\mathbf{x}|H_0] = \int_{-\infty}^{\infty} p_{\mathbf{B}}(\mathbf{b}) \frac{2}{\prod_i \sqrt{x^{(i)} - (b^{(i)})^2}} p_{\mathbf{A}}\left(\sqrt{x^{(i)} - (b^{(i)})^2}\right) d\mathbf{b}. \qquad (8)$$

The likelihood ratio function for each hypothesis is thus the ratio of two N^2-dimensional integrals, where N is the length of the stick used. The most likely hypothesis for any given image segment is the hypothesis which maximizes the likelihood ratio function. If all the likelihood ratio functions are less than some threshold, then the null hypothesis should be selected. So the optimal stick detection scheme is to evaluate a family of integrals for each point, compute the maximum result and compare to a threshold.

5. SPECIAL CASES

5.1. White Noise

The true likelihood ratio function requires tremendous calculation to evaluate because of all the interpixel correlations that make (7) and (8) iterated integrals rather than products of integrals. However, if the speckle is known to have uncorrelated pixels, optimal detection can be performed at more modest cost.

A speckle field tends to decorrelate rather quickly, so one way of dealing with speckle is to decimate the image to the point where the interpixel correlations are insignificant. In this case, the speckle field can be modeled as the magnitude squared of a white Gaussian field. Thus, the pixels are independent, and the problem reduces to a multidimensional Rayleigh/Rician detection problem, well known in non-coherent communications. Note that it is the *magnitude* not magnitude squared of a Gaussian random variable that has a Rayleigh or Rician distribution. Thus, the likelihood ratio below is expressed in terms of $\sqrt{\mathbf{x}}$, where $\mathbf{x}$ is the square of the image pixels, consistent with the derivation above.

For white speckle noise, the likelihood function is given by:

$$\Lambda_i(\mathbf{x}) = \prod_j \exp\left(-\frac{\mu_j^{(i)}}{2\sigma^2}\right) I_0\left(\frac{\sqrt{\mathbf{x}_j \mu_j^{(i)}}}{\sigma^2}\right), \qquad (9)$$

where $\mathrm{I}_0(\cdot)$ is a modified Bessel function, and $\mu_j^{(i)}$ is the jth component of the stick at the ith orientation. Collecting terms, we obtain

$$\Lambda_i(\mathbf{x}) = \exp\left(\sum_j -\frac{\mu_j^{(i)}}{2\sigma^2}\right)\prod_j \mathrm{I}_0\left(\frac{\sqrt{\mathbf{x}_j\mu_j^{(i)}}}{\sigma^2}\right) = K\prod_j \mathrm{I}_0\left(\frac{\sqrt{\mathbf{x}_j\mu_j^{(i)}}}{\sigma^2}\right), \quad (10)$$

where K is a constant with respect to $\mathbf{x}$ which can be incorporated into the threshold.

5.2. High SNR Noise

When the SNR is very high, the derivation of the likelihood ratio function admits a simplification. If $\sqrt{\mathbf{x}-\mathbf{b}^2} \approx \sqrt{\mathbf{x}}$, we can avoid the integral by canceling out the nuisance parameter $\mathbf{b}$. After manipulating, an equivalent optimal test statistic is given by

$$\Lambda_i(\mathbf{x}) = \exp\left(-\frac{1}{2}(\mu^{(j)})^T\mathbb{R}^{-1}\mu^{(j)}\right)\cosh\left((\mu^{(j)})^T\mathbb{R}^{-1}\sqrt{\mathbf{x}}\right). \quad (11)$$

An interesting special case occurs when the noise is white. In that case, since ${\mu^{(j)}}^T\mathbb{R}^{-1}\mu^{(j)}$ is constant with respect to i, Equation (11) simplifies to

$$K\cosh\left(\frac{1}{\sigma^2}(\mu^{(j)})^T\sqrt{\mathbf{x}}\right), \quad (12)$$

which can be equivalently implemented by

$${\mu^{(j)}}^T\sqrt{\mathbf{x}}, \quad (13)$$

which is the same as the Sticks technique developed in [1].

Sticks is the exactly optimal likelihood ratio test for the case of zero-mean additive Gaussian noise. It is essentially a matched filtering operation, in which we project the image onto the short line segments we are trying to detect. It is also closely related to the Hough transform [8] which searches images for simple patterns (such as straight lines) at various orientations and translations. For non-Gaussian noise, Sticks is a sub-optimal but useful technique which performs well in real images and requires comparatively little computation to implement. For these reasons, we evaluate its performance in the following sections along with that of the optimal techniques.

6. SIMULATIONS

To evaluate the performance of these various detection schemes, we set up a simulation of a realistic ultrasound imaging environment. The test image contains a square with edges a single pixel thick. The square is immersed in complex Gaussian noise either uncorrelated between pixels, or colored by a two-dimensional Gaussian filter. To simulate coherent detection, we take the magnitude squared of the resulting complex valued image.

The model for the observation is thus:

$$i(x,y) = |s(x,y) + n_c(x,y)|^2, \quad (14)$$

where the real and imaginary parts of $n_c(x,y)$ are uncorrelated with each other, and have autocorrelation given by $R_n(\Delta x, \Delta y) = \sigma^2\exp(-4\ln 2(\Delta x^2 + \Delta y^2)/W^2)$. Since SNR is an elusive quantity in a synthetic image, we will refer to different test cases by the σ^2 and W parameters of the correlation function. In the discussion of the results, however, it will be convenient to describe the noise as decreasing in SNR when the variance parameter σ^2 or the spread parameter W increases.

In all simulations, the operator dimension (the "stick length") was fixed at 9. That is, around every image point in the simulations we consider a 9 × 9 pixel region. This parameter must be optimized separately from the detection rule. A longer stick length will reduce noise more than a shorter one, but will also tend to smear out the edges if they bend too sharply. The correct stick length to use in any given application depends on the properties of the image. The issue of how to select the stick length is discussed in [1].

For each algorithm and each noise setting, we present an approximate receiver operating characteristic (ROC) curve. This is obtained by plotting the number of correct detections versus the number of false alarms as the threshold is varied parametrically. Each axis is normalized to probability one when either all pixels on the square are correctly detected or all pixels off the square are incorrectly detected as edges. The graphs' axes are labeled "Probability of Detection" and "Probability of False Alarm" even though the curve represents only an approximation to the true ROC curves; the data come from a single realization of the noise process, and may be biased due to the nature of the simulation — a square in noise to simulate the true imaging environment rather than a pure line detection scenario to test the capability of the detection routines.

7. RESULTS

7.1. White Noise Case

In the case of uncorrelated speckle, we can evaluate the performance of the exactly optimal detection rule (10). Figures 1 and 2 show the test images used to compare the optimal detection rule with the Sticks procedure and with simply thresholding the image. For Figure 1, the low SNR case, the noise variance was $\sigma^2 = 0.3$; for Figure 2, the high SNR case, the noise was $\sigma^2 = 0.2$. The ROCs for these test images are shown, respectively, in Figures 3 and 4. In the ROCs, the circles indicate the ROC curve for thresholding the image, the dotted curve indicates the performance of Sticks, and the solid line represents the performance of the optimal detector.

The optimal and Sticks curves are significantly above the null operator at both SNR levels. In both cases, the optimal curve is above the Sticks curve for almost all values of the threshold. However, the two curves are very close together, indicating that the two algorithms are close in

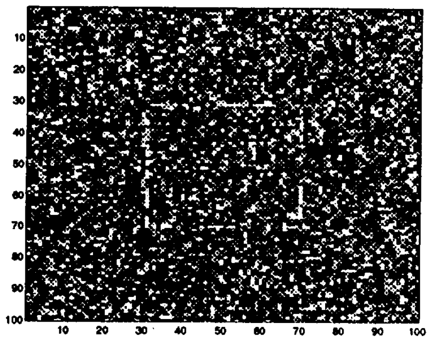

Figure 1. Test image for low SNR white noise case.

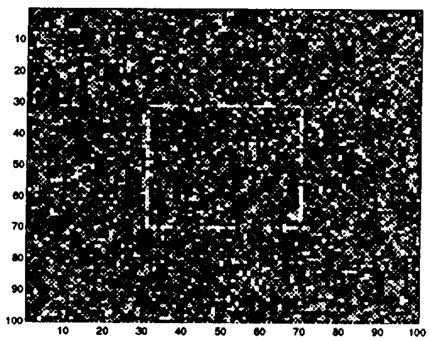

Figure 2. Test image for high SNR white noise case.

performance. This is significant, since the optimal detection rule is computed using actual noise statistics, while Sticks is not. Since Sticks can offer performance so closely approaching the optimal, it is hard to justify the additional computation required to estimate the statistics and implement the optimal test in the case of white speckle noise.

7.2. Colored Noise Case

Since Sticks is, in fact, the optimal detection rule for the case of additive white Gaussian noise, it is perhaps not surprising that it performs well in the white speckle case. In the colored speckle noise case, however, the optimal detection rule is impractical, so we must rely on the high SNR approximation (11). Figures 5 and 6 show the test images used to evaluate the approximate optimal detection rule, and Figures 7 and 8 show a performance comparison between the approximately optimal detector, Sticks and the null operator. In these cases, Sticks actually outperforms the approximately optimal detector, although the margin is slim and gets smaller as SNR improves. For lower SNRs, the approximately optimal detector does very poorly, which is to be expected since the approximation we use to derive the detector is valid only when SNR is high.

In the higher SNR cases, we can attribute the performance advantage of Sticks to the fact that the approximate optimal detector is only truly valid in the very high SNR case. Of course in that case the image features are obvious

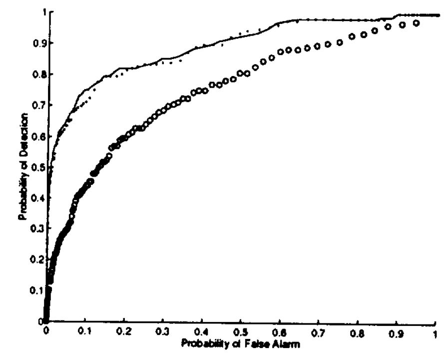

Figure 3. ROC curves for the low SNR white noise simulation. The circles represent the result of thresholding the original image, the dots represent the performance of Sticks, and the solid line represents the optimal detector.

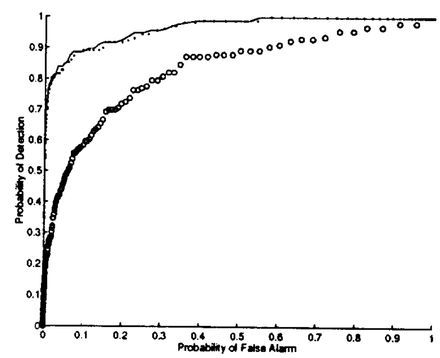

Figure 4. ROC curves for the high SNR white noise simulation. The circles represent the result of thresholding the original image, the dots represent the performance of Sticks, and the solid line represents the optimal detector.

without computer detection. This is consistent with the fact that the performance difference between Sticks and the approximately optimal detector decreases as SNR increases. Thus in the colored noise case, just as with white noise, case, Sticks is preferable to the statistically motivated detection rule, in this case with a performance advantage over the statistical technique.

8. CONCLUSIONS

We have shown that the simple Sticks technique presented in [1] is close in performance to the optimal detection scheme in the case of uncorrelated speckle when the statistics of the noise are exactly known. We have also given an approximation to the optimal colored speckle noise stick detector, valid in cases of high SNR. However, in all but the highest SNR simulations, the sub-optimal Sticks detector outperformed the approximation to the optimal detector in the sense that its approximate receiver operating curve is higher at almost every threshold value. Since the Sticks detection scheme is computationally much more efficient, and the statistically optimal schemes depend heavily on exact knowledge of the noise statistics which is generally unavailable, we conclude that Sticks is the best technique we have discussed for detecting thin edges in ultrasound noise.

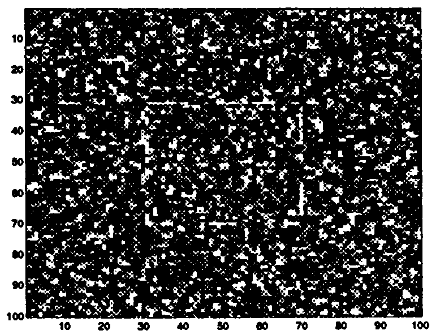

Figure 5. Test image for low SNR colored noise case.

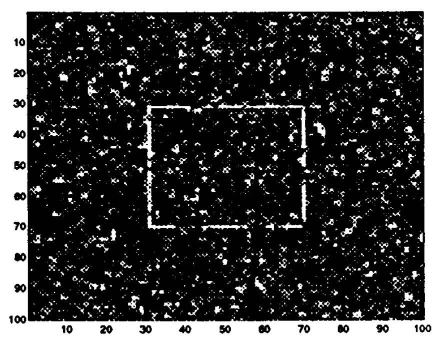

Figure 6. Test image for high SNR colored noise case.

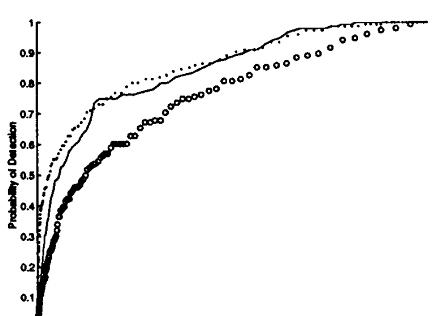

Figure 7. ROC curves for the low SNR colored noise simulation. The circles represent the result of thresholding the original image, the dots represent the performance of Sticks, and the solid line represents the optimal detector.

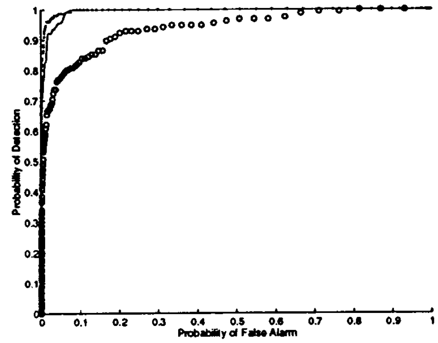

Figure 8. ROC curves for the high SNR colored noise simulation. The circles represent the result of thresholding the original image, the dots represent the performance of Sticks, and the solid line represents the optimal detector.

REFERENCES

[1] R. N. Czerwinski, D. L. Jones, and W. D. O'Brien, Jr., "An approach to boundary detection in ultrasound imaging", in *1993 IEEE Ultrasonics Symposium Proceedings*, Baltimore, MD, November 1993.

[2] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, Addison-Wesley, Reading, MA, 1992.

[3] A. C. Bovik, "On detecting edges in speckle imagery", *IEEE Transactions on Signal Processing*, vol. 36, no. 10, pp. 1618–1627, October 1988.

[4] J. W. Goodman, "Statistical properties of laser speckle patterns", in *Laser Speckle and Related Phenomena*, J. C. Dainty, Ed., number 9 in Topics in Applied Physics, chapter 2, pp. 9–77. Springer-Verlag, Berlin, 1977.

[5] C. B. Burckhardt, "Speckle in ultrasound B-mode scans", *IEEE Transactions on Sonics and Ultrasonics*, vol. SU-25, no. 1, pp. 1–6, January 1978.

[6] R. F. Wagner, S. W. Smith, J. M. Sandrik, and H. Lopez, "Statistics of speckle in ultrasound B-scans", *IEEE Transactions on Sonics and Ultrasonics*, vol. 30, no. 3, pp. 156–163, May 1983.

[7] H. V. Poor, *An Introduction to Signal Detection and Estimation*, Springer-Verlag, New York, 1988.

[8] R. O. Duda and P. E. Hart, "Use of the Hough transform to detect lines and curves in pictures", *Communications of the ACM*, vol. 15, no. 1, pp. 11–15, January 1972.