Proceedings of the Institute of Acoustics

SYNTHETIC APERTURE SONAR SPECKLE REDUCTION USING PARTIAL DIFFERENTIAL EQUATION TECHNIQUES

Daniel A. Cook
Naval Surface Warfare Center Panama City
110 Vernon Ave.
Panama City, FL 32407

1. INTRODUCTION

The purpose of this work is to explore the application of partial differential equation (PDE) based image processing techniques to the problem of reducing speckle noise in synthetic aperture sonar imagery. Three techniques are compared. First is the linear heat diffusion. It is included as a reference, as it is equivalent to an ordinary symmetric two-dimensional Gaussian low-pass filter. Next is a geometric heat diffusion which is standard PDE-based technique used for edge-preserving noise reduction in imagery. The last, Speckle-Reducing Anisotropic Diffusion (SRAD), is a technique introduced by Yu and Acton[1] that is specifically designed for the purpose of speckle mitigation in synthetic aperture radar (SAR) imagery. We will first give a brief overview of the problem of speckle in coherent imaging systems. Following is a description of the PDE-based image processing techniques being considered. Lastly, the results of applying these techniques to field data are presented.

2. THE PROBLEM OF SPECKLE NOISE

Speckle is a type of correlated multiplicative noise characteristic of coherent imaging systems[2], be they optical, microwave, or acoustic. Speckle arises when the surface being imaged is rough compared to the scale of the wavelength of the incident energy. Under these conditions, the signal observed within any given resolution cell is the phasor sum of many randomly-distributed scattering centers. It has been shown[3] that the probability density function of speckle intensity, I, is described by the negative exponential distribution. An important property of this distribution is that the standard deviation is equal to the mean. Consequently, the contrast is equal to unity when defined as $K = \sigma_I / \bar{I}$.

A second quantity of interest is the average size of the 'cells' of speckle. It is possible to derive the distribution of the scale sizes of speckle[3]. However, Jakowatz et al.[4] give a more direct argument that is useful for the purpose at hand: The data collected by a synthetic aperture system occupy some finite region of support in the spatiospectral domain (i.e., the image results from the inverse 2D Fourier transform of the data). The size of this region of support determines the highest resolution achievable by the imaging system. In a given dimension, the spatial resolution is inversely proportional to the size of the region of spectral support (since the resolution is inversely proportional to bandwidth). Because the spectral support is always finite, the image of an ideal point scatterer will appear spread out. Consequently, the noise caused by the random rough scattering is correlated, and the correlation length is equal to the maximum resolution of the system. An interesting fact is that the scale distribution of speckle is independent of the quality of focus.

The author gratefully acknowledges the US Office of Naval Research for its support of this work and Prof. Anthony Yezzi of the Georgia Institute of Technology for his course notes which were used in the composition of Sections 3.1 and 3.2.

In summary, speckle is particularly troublesome because it is a high-contrast phenomenon whose characteristic size is the same as the resolution of the system it afflicts. Both of these cause speckle to have a very distracting effect upon the viewer. The latter effect makes speckle noise difficult to effectively counter without sacrificing an appreciable amount of resolution. Thus, the ability to preserve known features while reducing speckle noise is an important performance metric.

3. PDE-BASED TECHNIQUES FOR SPECKLE REDUCTION

In this section we present the salient features of the PDE-based techniques evaluated for speckle reduction. In the following, the image $I(\mathbf{x}, t)$ is defined for a finite 2D region of space, $\mathbf{x} = [x\ y]^{\mathrm{T}} \in \Omega$ and for positive time, $t \in [0, t]$. The image is a mapping of its domain onto the set of real numbers: $\Omega \times [0, t] \rightarrow \mathbb{R}$, where the value of I represents the image intensity.

3.1 The Linear Heat Equation

The linear heat diffusion equation is given by $I_t = \Delta I$ where Δ is the Laplacian operator $\left(\partial^2/\partial x^2 + \partial^2/\partial y^2\right)$, and I_t is the partial derivative with respect to time. This equation appears in a multitude of physical contexts. In image processing applications, the intensity distribution of the image is analogous to the distribution of heat in a uniform 2D planar substance. As time is allowed to progress, the heat (image intensity) diffuses into neighboring pixels.

If we consider the input image, $I_0(x, y, t = 0)$, as the initial condition, it can be shown[5] that allowing the heat equation to evolve for a duration t is equivalent to applying a low-pass Gaussian filter with variance $\sigma^2 = 2t$. Let us consider the 1D case:

$$I_t(x, t) = \frac{\partial^2}{\partial x^2} I(x, t), \tag{1}$$

and take its spatial Fourier transform:

$$I_t(k_x, t) = -k_x^2\ I(k_x, t). \tag{2}$$

The last equation is an ordinary differential equation (for fixed k_x) whose solution is given by:

$$I(k_x, t) = I_0(k_x)\ \exp\left(-k_x^2 t\right). \tag{3}$$

Taking the inverse spatial Fourier transform gives:

$$\begin{aligned} I_t(x, t) &= \frac{1}{\sqrt{2\pi}}\ \mathcal{F}^{-1}\left\{\exp\left(-k_x^2 t\right)\right\} \odot_x I_0(x) \\ &= \frac{1}{\sqrt{4\pi t}}\ \exp\left(-\frac{x^2}{4t}\right) \odot_x I_0(x), \end{aligned} \tag{4}$$

where $\odot_x$ indicates convolution with respect to x, and the last line is recognized as the convolution of the initial image with a Gaussian of variance $2t$.

This diffusion will serve as a reference that is representative of the traditional low-pass filtering approach to speckle reduction. The linear heat equation is related to the geometric heat equation described next. Thus for the sake of evaluating performance, the diffusion time t_{heat} will be set equal to that used for the geometric heat equation, t_{geo}.

3.2 Geometric Heat Diffusion

The drawback to using the linear heat equation is that it smoothes equally in all directions and thus causes blurring of edges. The geometric heat equation avoids this problem: For any given (x, y) location, the smoothing is 1D and applied in the direction orthogonal to the gradient $\nabla I(x, y, t)$. This has the effect of smoothing along edges, and not across them. To derive the geometric heat diffusion, we first note that the Laplacian is rotationally invariant. We are therefore free to write the linear heat equation as:

$$\begin{aligned} \Delta I &= I_{xx} + I_{yy} \\ &= I_{\eta\eta} + I_{\xi\xi}, \end{aligned} \tag{5}$$

in which we have recast the Laplacian operator in terms of a local coordinate system where $\eta(x, y) = \nabla I / \|\nabla I\|$ is the unit normal pointing in the direction of the image gradient at the point (x, y) and $\xi(x, y)$ is chosen to be the orthogonal complement of η. The next step is to isolate the term $I_{\xi\xi}$ that represents the contribution to the Laplacian from the edge at any given location. We first rewrite $I_{\eta\eta}$ as a directional derivative:

$$I_{\eta\eta} = \frac{\nabla I^T}{\|\nabla I\|} \nabla^2 I \frac{\nabla I}{\|\nabla I\|}, \tag{6}$$

where $\nabla^2 I$ is the Hessian matrix of I whose entries are given by $h_{ij} = \partial^2 I / \partial x_i \partial x_j$. Subtracting $I_{\eta\eta}$ from both sides of Equation (5) and setting the result equal to I_t gives:

$$\begin{aligned} I_t &= I_{\xi\xi} \\ &= \Delta I - I_{\eta\eta} \\ &= \frac{I_y^2 I_{xx} - 2 I_x I_y I_{xy} + I_x^2 I_{yy}}{I_x^2 + I_y^2}. \end{aligned} \tag{7}$$

The first line above indicates that we are diffusing only along edges and not across them. This diffusion is locally one-dimensional. The last line is the equation to be solved. Although the problem is formulated in terms of η and ξ, the PDE that we actually implement numerically has been expressed in terms of derivatives in x and y, which are easily obtained using simple finite difference schemes.

Further insight can be gained if we think of the image as being composed of a family of level-set curves and then consider only a single simply-connected curve, C, embedded in I. It can be shown that the level-set form of (7) is given by: $C_t = \kappa N$ in which κ is the curvature and N is the inward unit normal to C. This is known as *curvature flow* because the speed with which a point on C moves inward is proportional to the curvature at that point. We can derive from the curvature flow equation the rate at which area is lost from C as it shrinks to a point and disappears. The area enclosed by C is given by the integral inside C:

$$A = \int_C 1 \; dxdy. \tag{8}$$

Taking the time derivative gives:

$$A_t = -\int_C C_t \cdot N \; ds, \tag{9}$$

in which s is the arc length parameter. Into this, we substitute $C_t = \kappa N$ to obtain:

$$A_t = -\int_C \kappa \; ds = -2\pi, \tag{10}$$

where the integral has been carried out over a single complete traversal of C. Thus, the rate of change of area enclosed by the curve C is constant regardless of its shape. We use this fact to determine how long to run the diffusion PDE (7). As mentioned above, the average speckle size is known to be equal to the resolution of the synthetic aperture system. This suggests a diffusion time given by $t_{\text{geo}} = A_{\text{speckle}}/2\pi$.

3.3 Speckle-Reducing Anisotropic Diffusion

The third and final technique to be described is called Speckle-Reducing Anisotropic Diffusion (SRAD). It is a PDE-based method specifically designed for dealing with speckle in synthetic aperture radar imagery. Yu and Acton[1] review two popular conventional filters for speckle reduction that modify each image pixel based on the statistics within a window around that pixel. The drawback to these filters is that they are sensitive to the size and shape of the chosen window. The authors overcome this problem by formulating a PDE-based approach (SRAD) that is related to these filters, but does not use a local window for each pixel. In place of the statistics computed in a neighborhood around each pixel, Yu and Acton introduce the *instantaneous coefficient of variation*. This coefficient is recomputed at each iteration based either on an analytical expression or on the speckle statistics computed over a uniform region of the image. The SRAD technique has a form that is structurally similar to the Perona-Malik equation[6] and is given by:

$$\begin{cases} I_t = \nabla \cdot [c(q)\ \nabla I] \\ I(t=0) = I_0, \end{cases} \tag{11}$$

with Neumann boundary conditions imposed at the image border. In the above, the coefficient $c(q) = \left[1 + \left[q^2 - q_0^2\right] / \left[q_0^2\left(1+q_0^2\right)\right]\right]^{-1}$, where $q = q(x,y,t)$ and $q_0 = q_0(t)$ are given by:

$$q(x,y,t) = \sqrt{\frac{0.5\left(\|\nabla I\|/I\right)^2 - (0.25)^2\left(\Delta I/I\right)^2}{\left[1 + 0.25\left(\Delta I/I\right)\right]^2}} \quad \text{and} \quad q_0(t) = \frac{\sqrt{\text{var}\left[z(t)\right]}}{\overline{z(t)}}. \tag{12}$$

The $q(x,y,t)$ term is the instantaneous coefficient of variation. It serves as the edge detector, as it attains large values on edges or high-contrast features and small values in homogeneous regions. The quantity $q_0(t)$ is the speckle scale function. In the form given, q_0 is computed from a homogeneous region in the image (the same set of pixels are used to compute q_0 at each iteration). The quantity $z(t)$ indicates the set of pixels contained inside some window located in a region of the image that is expected to be homogeneous. The window location is fixed from one time step to the next, although the pixel values inside it clearly are not. In cases where selecting a homogeneous region is undesirable or impossible, an analytical approximation is $q_0(t) \approx q_0 \exp\{-\rho t\}$ where ρ and q_0 are constants. In the present work, they are taken to be equal to one. According to Yu and Acton, the instantaneous coefficient of variation enables SRAD to not only preserve edges, but actually enhance them. As mentioned before, SRAD is similar in form to the Perona-Malik equation, which also possesses an edge-enhancement property. A discussion of this effect is given by Pollak[7].

4. RESULTS

This section will present examples of applying the PDE techniques described above. The first is shown in Figure 1. Successful speckle reduction is measured in terms of the ability to smooth out the speckle in homogeneous regions while preserving sharp edges. A cursory look at Figure 1 reveals that the linear heat equation is not well-suited to this purpose. While it does a good job

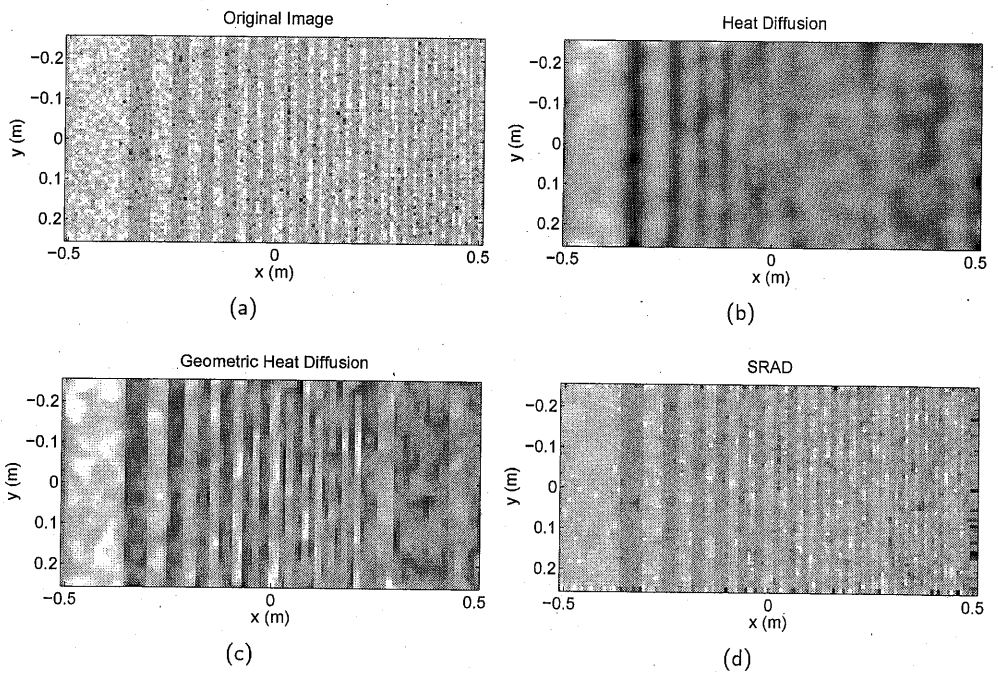

Figure 1: Comparison of PDE-based speckle reduction schemes.

of smoothing, is does not preserve edges. As a result, no edges are distinguishable to the right of $x = 0$ in Figure 1(b).

On the other hand, the geometric heat equation and SRAD succeed in both smoothing and preserving edges. For the spatial chirp shown, the geometric heat heat equation is better than SRAD at preserving edges. This is not surprising, as the geometric heat equation is derived specifically for that purpose. SRAD is designed to simply diffuse less strongly in the presence of edges. Indeed, the edges in Figure 1(c) are smooth and straight compared to those in 1(d) which are somewhat ragged due to the lack of diffusion. In terms of preserving the high-frequency edges (near the right edges of the images), SRAD outperforms the geometric heat equation. This is because the width of the high-frequency banding approaches the speckle correlation length. The geometric heat equation cannot distinguish between the band edge and the edges of speckle cells. In this region, SRAD simply does not apply much smoothing and the high-frequency vertical bands are preserved.

On the other hand, SRAD is better at smoothing out the speckle in the homogeneous regions in the image. This is clearly seen on the left side of Figure 1(d). The wide light and dark bands are almost uniform due to the SRAD smoothing– nearly as uniform as in the case of the linear heat equation. Contrast this to Figure 1(c) in which the geometric heat equation does not specifically work toward producing a uniformly smooth region. Rather, the cells of speckle smooth and coalesce slowly.

Next, we discuss the performance of the PDE methods on real SAS imagery. The imagery comes from the Small Synthetic Aperture Minehunter (SSAM) sonar system sponsored by the Office of Naval Research. This sonar has a theoretical resolution of 2.5cm by 2.5cm (1" by 1"). For the results shown, the PDE evolution time was chosen so as to eliminate an ideal piece of circular speckle 6cm in diameter. This is somewhat more than might be used in practice, but the extra smoothing helps to illustrate the differences between the methods. The diffusion time t_{geo} was

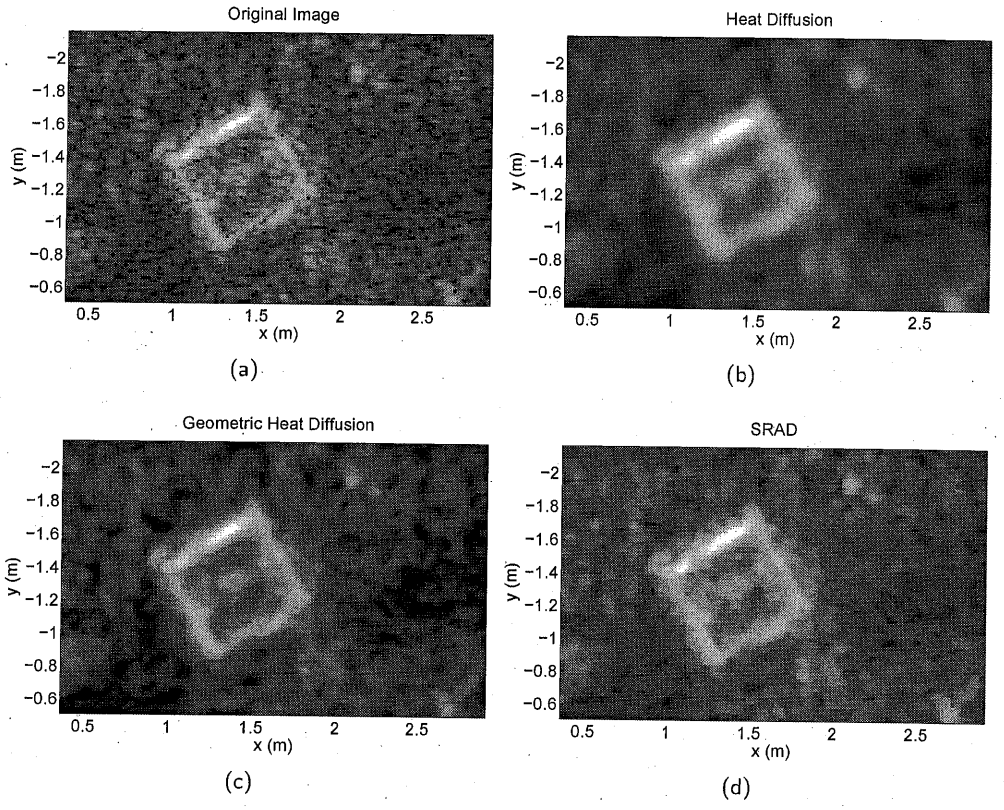

Figure 2: Comparison of PDE-based speckle reduction schemes.

computed as shown in equation (10). The times t_{heat} and t_{SRAD} were set equal to t_{geo}. Setting $t_{\text{heat}} = t_{\text{geo}}$ makes sense, as the geometric heat equation is closely related to the linear heat equation. The time used to evolve SRAD was set equal to t_{geo} for the sake of simplicity (*i.e.*, reducing the number of parameters in the experiment). Experience has shown that this choice makes for a fair comparison.

Figure 2 shows a lobster pot (a trap) situated next to a rock in the waters of Buzzard's Bay at Cape Cod, Massachusetts in the US. The effects of speckle can be seen in the unaltered imagery in the form of small-valued outliers (holes) which show up as isolated dark pixels. As is the case with nearly all imagery collected in the field, it is difficult to tell which small-scale features are real and which are speckle. Thus, the evaluation of the results is somewhat subjective and dependent on the experience of the person interpreting the images. This task is easier when known objects appear in the scene, such as the square trap. Its edges should be bright and sharp.

The PDE methods behave for real data much as they did in the previous artificial example. The linear heat equation smoothes out the speckle but does so at a significant cost of resolution. The geometric heat equation and SRAD produce very similar results due to the relatively short time of evolution. There are no significant differences between Figures 2(c) and 2(d).

All of the PDE techniques have succeeded in improving the image quality in terms of reducing the speckle noise. The linear heat equation, included primarily as a 'worst-case' reference isn't well-suited to general use since it causes a large loss of resolution for a given amount of speckle reduction. The geometric heat equation and SRAD produce similar results when run for only a

short time. These techniques successfully eliminate the very worst aspects of speckle noise while apparently preserving the resolution of the SAS image, as evidenced by noting the thickness and continuity of the four bright lines marking the outline of the trap. Additionally, the region inside the trap has been smoothed to more clearly reveal a circular feature in the center. This is the part of the trap used to hold the bait.

5. CONCLUSION

In this paper we have discussed three PDE-based image processing techniques and evaluated their usefulness for reducing speckle noise in SAS imagery. The goal is to find a successful technique whose evolution time can be determined without human intervention or relying heavily on heuristics. This criterion is very important, as the techniques chosen would be used to automatically process large amounts of imagery. The linear heat equation is not well suited to this purpose because of the resolution loss it induces. The other two methods perform similarly with respect to removing speckle noise when they are only run for a short time. However, the qualitative 'feel' or appearance of the resulting imagery differs in each case. This distinction may cause one technique to be preferred over another in certain situations. Neglecting this difference, however, the geometric heat equation is arguably better than SRAD in the sense that it is cheaper to implement numerically.

It should be noted that the results presented here are relatively narrow in scope. A number of other uses and refinements suggest themselves. The focus here was cleaning up the image while retaining the maximum amount of system resolution, and the geometric heat equation was found to be best for this purpose. If one wished instead to smooth extensively in an attempt to segment the image into a number of homogeneous regions, then SRAD would be the technique of choice. Running the geometric heat equation for a long time does not quickly produce homogeneous regions. Instead, the result is a swirling pattern reminiscent of Van Gogh's "Starry Night." This results from the fact that the geometric heat equation is a 1D diffusion. SRAD diffuses in 2D, and is therefore able to homogenize a region more effectively. Finally, we might also consider new techniques such as modifying the geometric heat equation so that it not only diffuses along edges, but is also allowed to diffuse across them. The extent of this lateral diffusion could be dictated by some metric like the local curvature, magnitude of the gradient, etc. As these results have shown, PDE image processing methods are indeed useful for eliminating speckle noise in real-world SAS imagery. The PDE-based approach is well-suited to this purpose because it allows the user to exploit local and nonlinear behavior with an ease and effectiveness that is impossible with conventional filtering.

References

1. Y. Yu and S. T. Acton, "Speckle reducing anisotropic diffusion," *IEEE Transactions on Image Processing*, vol. 11, no. 11, pp. 1260–1270, November 2002.

2. J. W. Goodman, *Statistical Optics*. John Wiley and Sons, Inc., 1985.

3. ——, "Some fundamental properties of speckle," *Journal of the Optical Society of America*, vol. 66, no. 11, pp. 1145–1150, November 1976.

4. C. V. Jakowatz, D. E. Wahl, P. H. Eichel, D. C. Ghiglia, and P. A. Thompson, *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*. Kluwer Academic Publishers, 1996.

5. J. C. Strikwerda, *Finite Difference Schemes and Partial Differential Equations.* Wadsworth and Brooks/Cole, 1989.

6. P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 7, pp. 629–637, July 1990.

7. I. Pollak, "Segmentation and restoration via nonlinear multiscale filtering," *IEEE Signal Processing Magazine*, September 2002.