Use of anisotropic diffusion filters in Synthetic Aperture Radar (SAR) imagery

G. Barrie
DRDC – Ottawa Research Centre

Canada

IMPORTANT INFORMATIVE STATEMENTS

This document was reviewed for Controlled Goods by DRDC using the Schedule to the *Defence Production Act.*

Disclaimer: Her Majesty the Queen in right of Canada, as represented by the Minister of National Defence ("Canada"), makes no representations or warranties, express or implied, of any kind whatsoever, and assumes no liability for the accuracy, reliability, completeness, currency or usefulness of any information, product, process or material included in this document. Nothing in this document should be interpreted as an endorsement for the specific use of any tool, technique or process examined in it. Any reliance on, or use of, any information, product, process or material included in this document is at the sole risk of the person so using it or relying on it. Canada does not assume any liability in respect of any damages or losses arising out of or in connection with the use of, or reliance on, any information, product, process or material included in this document.

Endorsement statement: This publication has been peer-reviewed and published by the Editorial Office of Defence Research and Development Canada, an agency of the Department of National Defence of Canada. Inquiries can be sent to: Publications.DRDC-RDDC@drdc-rddc.gc.ca.

Abstract

The Defence R&D Canada (DRDC) Ottawa Research Center, ISR Data Exploitation Group, specializes in development of algorithms and tools enabling uptake of Synthetic Aperture Radar (SAR) data for Canadian Armed Forces (CAF) planning and operational purposes. This scientific report addresses the extraction of coastlines in SAR images exhibiting poor contrast between land and water backscatter. In place of traditional speckle filtering, we apply a nonlinear filter whose design principles are rooted in the physics of diffusion, where image blurring is treated in a manner analogous to diffusing particles. After a brief discussion of how the diffusion equation is applied to image processing, we introduce the Speckle Reducing Anisotropic Diffusion (SRAD) filter. SRAD is shown to be related to more commonly used speckle filters, which constitute a special case of isotropic diffusion. The SRAD filter is first used to process an artificial coastline consisting of two adjacent blocks of RADARSAT-2 (RS-2) data of water and land backscatter forming an idealized linear coastline. This construct allows us to establish easily quantifiable ground truth, and perform statistical analysis on both individual and combined land and water pixel distributions. Following this, SRAD is applied to large ($\sim$ 2 GB) Wide Ultra-Fine RS-2 data sets to examine its use in poor ($\lesssim$ 6 dB) contrast imagery. Results indicate this is an effective filter, and when used in conjunction with a traditional edge detector such as the Canny algorithm, can contribute to the delineation of edges in SAR images that would otherwise be difficult to process.

Significance for defence and security

DRDC Ottawa Research Centre develops application tools and procedures for use in Geographic Information Systems (GIS), allowing for the uptake of SAR data in more traditional cartography-based systems. This investigation addresses advanced image processing techniques to extract shorelines when contrast is too low to properly segment land-water masses. The technique can be applied in non-ideal imaging conditions, thereby filling collection and data exploitation gaps, enabling enhanced production of Geospatial Intelligence (GEOINT) products and Tactical Decision Aids (TDAs).

Résumé

Le Groupe d'exploitation des données RSR du Centre de recherche de R&D pour la défense Canada – Ottawa se spécialise dans la création d'algorithmes et d'outils permettant d'intégrer des données de radar à synthèse d'ouverture (RSO) aux fins de la planification et des opérations des Forces armées canadiennes. Dans le présent rapport scientifique, nous abordons l'extraction du tracé des côtes dans les images RSO dont le contraste entre la rétrodiffusion de la terre et celle de l'eau est faible. Plutôt que d'utiliser le filtrage classique du chatoiement, nous appliquons un filtre non linéaire dont les principes de conception proviennent de la physique de la diffusion et qui traite le flou de l'image comme si des particules avaient causé la diffusion. Après une brève discussion sur la manière dont nous appliquons l'équation de diffusion au traitement de l'image, nous présentons le "filtre de diffusion anisotrope réduisant le chatoiement" ou filtre SRAD (pour Speckle Reducing Anisotropic Diffusion). Nous montrons que le filtre SRAD s'apparente aux filtres à chatoiement régulièrement utilisés qui constituent un cas spécial de la diffusion isotrope. Nous l'avons d'abord utilisé pour traiter une côte artificielle, composée de deux blocs adjacents de données de rétrodiffusion par l'eau et la terre captées par RADARSAT 2 qui forment un littoral linéaire idéalisé. Cette construction permet de créer une vérification au sol facilement quantifiable et de réaliser une analyse statistique sur les distributions de pixels de terre et de pixels d'eau, individuellement et collectivement. Après cette opération, nous avons appliqué le filtre SRAD à de grands ensembles ($\sim$ 2 GB) de données ultrafines dans un champ large, captées par RADARSAT-2 pour étudier son utilisation sur des images à faibles contrastes (0,6 dB). Les résultats indiquent que le filtre est efficace et que s'il est utilisé en conjonction avec un détecteur de bords classique, comme l'algorithme de Canny, peut contribuer à définir les arêtes dans des images de radar à synthèse d'ouverture qui seraient autrement difficiles à traiter.

Importance pour la défense et la sécurité

Le Centre de recherche de RDDC Ottawa produit des outils et des procédures qui seront utilisés dans des systèmes d'information géographique (SIG), ce qui permet d'intégrer les données de radar à synthèse d'ouverture dans les systèmes de cartographie plus classique. Cette recherche présente des techniques de traitement d'image permettant d'extraire des côtes, lorsque le contraste est trop faible pour segmenter correctement les masses de terres de celles d'eau. On peut les appliquer sur des images prises en conditions non idéales, ce qui permet de combler les lacunes dans la collection de l'exploitation de données, et de générer de meilleurs produits de renseignements géospatiaux (GEOINT) et d'aide à la prise de décisions tactiques (TDA).

Table of contents

List of figures

Acknowledgements

1   Introduction

The Canadian Armed Forces (CAF) and the Department of National Defence (DND) are tasked with development and maintenance of an Arctic surveillance capability, a key element in exercising Canadian sovereignty. In cooperation with other federal departments and agencies, the CAF undertakes operations in the North, conducting regular patrols for Intelligence, Surveillance, & Reconnaissance (ISR) purposes.

Defence Research and Development Canada (DRDC) continually explores options for cost-effective Arctic monitoring systems, providing specialist scientific and technical advice and decision support to enhance CAF operations. DRDC delivers enablers for CAF capabilities in exploiting space-based synthetic aperture radar (SAR). The work in this scientific report (SR) took place under the Space-Based Radar Exploitation (SBRE) project, which is focused on developing specialized capabilities for the CAF, related to exploitation of SAR imagery, both for RADARSAT-2 (RS-2) data and the upcoming RADARSAT Constellation Mission (RCM).

SAR data is a powerful tool for both Geospatial Intelligence (GEOINT) and Geomatics Intelligence Preparation of the Battlefield (IPB). DRDC efforts are directed at concurrent development of specialized processing algorithms, and GIS-based software implementation for use by the CAF client. The desired outcome of this project is to enable extraction of features of interest from SAR data with minimal analyst input, and to subsequently display these features as map layers within geoprocessing software.

This work explores one aspect of the military's need for advanced GEOINT software tools to aid analysts in exploiting SAR imagery, a means to reliably and systematically extract shoreline boundaries in areas of interest (AOIs). A shoreline is intuitively understood to be the boundary between land and water, but as is well known in the remote sensing community, it is exceedingly difficult to situate a dynamic boundary precisely using a static medium such as a map or satellite image.

Shoreline extraction is typically enabled by sufficient contrast between land and water, allowing for proper classification, and thereby, separation, of the two regions. A number of circumstances can arise that reduce this contrast, either from unfavourable SAR imaging geometry and/or polarization states of transmitted and received waveforms, or due to environmental conditions such as wind-roughened sea surface, breaking surf or shoals. These challenges are exacerbated in Northern regions where the regular presence of ice cover further diminishes land-water contrast. This report focuses on the open water scenario, and will examine ice-covered regions at a later date.

1.1   Background

Algorithms are available that classify land-water areas based on a bimodal distribution of pixel intensity. Generating a histogram from image pixel values of a coastal region will show

two modes centered around each of the mean pixel values for land and water, resulting in a bimodal distribution. In general, radar returns from land masses are expected to be brighter than returns from the sea. The random nature of radar backscatter produces variation about these two means. If the mean values are not sufficiently separated relative to the variance of the individual distributions, proper segmentation suffers. One needs to determine an effective threshold value to obtain the best separation, and depending on overlap, this becomes progressively more difficult to classify the two regions.

QuadPol data is the most suitable for identifying segmentation thresholds, due to the much lower Noise Equivalent Sigma Zero (NESZ). In fact, for this beam mode, we have used the Otsu method [1] to automate the classification stage of shoreline extraction[1]. An automated process is desirable as it removes the need for specialist knowledge to set image processing parameters, and frees the analyst to concentrate on their main task.

Three full collects of RS-2 imagery have been completed, from August 2015, March 2016, and September 2016, of selected Arctic AOIs. Beam modes included Wide Multi-Look Fine, Wide Ultra-Fine and Wide Fine QuadPol. Details of beam modes can be found in MDA's RS-2 Product Description [2]. The ensuing exploitation focused on the HV polarization channel with incident angles from approximately $25°$ to $50°$. Wide Multi-Look Fine and Wide Ultra-Fine, as compared to the Wide QuadPol mode, exhibit lower land-water contrast. While a bimodal distribution can still be discerned, there is much greater overlap, leading to classification error. In some cases the second distribution appears as a shoulder, and in other cases, the two distributions completely merge. This is the situation we examine in this report. For the case of ice cover, the Otsu method breaks down entirely. Practical large-scale extraction of shorelines requires greater swath widths, which are accompanied by a degradation of image resolution and contrast.

To address these issues, alternative shoreline extraction methods based on edge detection have been developed. While the Canny algorithm is a standard technique [3], it requires three properly chosen parameters to run effectively, and in spite of so-called "automatic" Canny edge detection [4],[5], this is often a trial-and-error process, the quality of which is very much dependent on the particular image. Also, Canny's development specifically targeted optimal edge detection in additive noise, while SAR images are dominated by *speckle*, generally regarded as a form of multiplicative noise in coherent imaging systems [6],[7]. Some work has been carried out attempting to modify the Canny algorithm to incorporate realistic SAR image statistics, most notably a Gamma distribution [8].

Speckle is all-pervasive in masking features and boundaries within the image, so one typically performs some type of speckle reduction before applying the edge detector. The challenge lies in reducing speckle sufficiently but not distorting or shifting edge boundaries. Algorithms used in feature extraction tend to be based on local non-adaptive operators [9]. These are sampling kernels of fixed size, say a $n \times n$ window, where n is odd-valued, ensuring a center

---

[1] Otsu published an unsupervised method to obtain an optimum threshold value by stepping through the gray levels of the image, computing the partitioned statistics, and selecting the value that would minimize in-class variance between the two distributions.

pixel. The size of the kernel imposes a spatial context on the processed image, not always matching the scale of image features. Depending on image homogeneity and other texture characteristics, sampled statistics may not be sufficient for accurate estimates. The window also acts as a low pass convolution filter to smooth out image noise. Filtering however, softens edges, an effect that is exacerbated when increasing window dimensions or choosing larger values of the smoothing parameter. Overall, there exist conflicting demands in that larger windows provide better statistics, but lead to a blurring of boundaries and increased difficulty in defining a sharp edge.

An alternative is to employ the so-called Speckle Reducing Anisotropic Diffusion (SRAD) filter. SRAD was introduced in 2002 by Yu and Acton [10] for processing images containing multiplicative noise, and finds application in such technologies as ultrasound and SAR imaging. Their work was based on Perona and Malik [11], which introduced anisotropic diffusion (AD) as an approach to edge detection in additive noise situations. Both methods are considered hybrids, combining noise filtering with the simultaneous capacity to not only maintain, but actually *enhance* edges. The work contained herein is part of a study to better characterize the properties of both SRAD, and the broader AD family of filters. There is understandably a tendency for diffusion filters to obliterate terrain cover, but as we are using it specifically to delineate the land-water boundary, this is not a serious limitation.

2   The diffusion equation

Broadly speaking, AD filters are a family of partial differential equation (PDE) based algorithms that concurrently provide image smoothing and feature enhancement via nonlinear diffusion. The fundamental operating principle is that diffusion is regulated as a function of the gradients occurring within the evolving image. The physical motivation for this type of filter is the well-known phenomenon of diffusion.

Given a collection of particles, the tendency is for net displacement from areas of high concentration to areas of lower concentration. The resultant concentration gradient, $\nabla \rho$, drives particle motion as summarized in Fick's Law [12]

$$j = -D\nabla\rho(r, t), \tag{1}$$

where $\rho(r, t)$ is the density at position $r = (x, y, z)$ at time t, and D is the diffusion coefficient. The quantity j is called *flux*, and is defined as the rate of particle movement per unit area measured normal to the surface. Both j and $\nabla\rho$ are vector quantities, the negative sign indicating that flux is parallel to the gradient, directed away from regions of higher relative concentration.

The continuity equation is a conservation expression:

$$\frac{\partial\rho}{\partial t} + \nabla \cdot j = 0. \tag{2}$$

This equation states that the time rate of change of particles within an arbitrary volume element V is equal to the net flow into or out of that volume. This is determined by the flux across the volume boundary, ∂V. Substituting eq. (1) into eq. (2) yields the standard form of the diffusion equation:

$$\frac{\partial \rho(r,t)}{\partial t} = D\nabla^2 \rho(r,t). \tag{3}$$

2.1 Anisotropic diffusion

In practice, the diffusion coefficient is rarely a simple scalar. In many substances, diffusion occurs at different rates along different directions, and generally the flux vector does not point directly along the gradient. This situation is referred to as *anisotropic* diffusion, and the diffusion *tensor* replaces D in eq. (1). In $\mathbb{R}^3$ we have

$$\begin{bmatrix} j_x \\ j_y \\ j_z \end{bmatrix} = - \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix} \begin{bmatrix} \partial_x \rho \\ \partial_y \rho \\ \partial_z \rho \end{bmatrix}. \tag{4}$$

Consider, for example, the x-component of flux

$$j_x = -D_{xx}\frac{\partial \rho}{\partial x} - D_{xy}\frac{\partial \rho}{\partial y} - D_{xz}\frac{\partial \rho}{\partial z}, \tag{5}$$

where it is clear that flux in the x direction has contributions from gradients along all three coordinate axes. The second term, as an illustration, indicates that flux in the x direction is due to a concentration gradient directed along y with proportionality D_{xy}. Physical tensors are symmetric so that $D_{xy} = D_{yx}$ [13]. The isotropic form, eq. (1), is recovered when $D_{xx} = D_{yy} = D_{zz} = D$ and all off-diagonal elements vanish, leaving $D\mathbb{I}$, where $\mathbb{I}$ is the identity matrix.

2.2 Nonlinear diffusion

Another situation that arises is when the diffusion coefficient is dependent on local concentration. The diffusion coefficient can be a function of both position and the local concentration, $D(r; \rho(r,t))$. The result is a nonlinear relationship in which eq. (1) becomes

$$j = -D(r; \rho(r,t)) \, \nabla \rho(r,t) \tag{6}$$

Substituting eq. (6) into eq. (2), we have

$$\begin{aligned} \frac{\partial \rho(r,t)}{\partial t} &= \nabla \cdot [D(r; \rho(r,t)) \, \nabla \rho(r,t)] \\ &= \nabla D(r; \rho(r,t)) \cdot \nabla \rho(r,t) + D(r; \rho(r,t))\nabla^2 \rho(r,t) \end{aligned} \tag{7}$$

Any further elaboration of this expression requires a functional form for $D(r; \rho(r,t))$. One possibility has diffusion as a function of the concentration gradient, $D(|\nabla \rho|)$, where $|\cdot|$ indicates the ℓ_2 norm of the gradient

$$| \nabla \rho | = \sqrt{\left(\frac{\partial \rho}{\partial x}\right)^2 + \left(\frac{\partial \rho}{\partial y}\right)^2 + \left(\frac{\partial \rho}{\partial z}\right)^2}. \qquad (8)$$

Application of this in constructing diffusion filters will be analyzed in greater detail in Section 3.2 starting on page 7.

Note: A caution to the reader regarding diffusion filter nomenclature — The primary means of classification centers around anisotropy and nonlinearity. From this perspective, SRAD is mislabeled. It is more appropriately categorized as a nonlinear diffusion filter because it is only location dependent. To qualify as an anisotropic filter, we would need to recast the formulation with a diffusion *tensor* in order to achieve directional dependence. However, while the naming conventions are strictly speaking in error, these quirks remain in the literature and are likely to continue.

3 Diffusion in image processing

Discussion to this point has been focused on formulation of eq. (3), where we have briefly examined the physical basis for diffusion from the perspective of continuum mechanics[2]. We now turn to using diffusion as a filter to improve the interpretation of images. By analogy, we treat the gray levels within an image as diffusing particulate and the above formulation applies equally well to image processing. In this situation, $I(r,t)$ is used in place of $\rho(r,t)$ and the position vector $r = (x,y)$ is recast to denote an image in the plane $I \in \mathbb{R}^2$. In a sense, we are treating $I(r,t)$ as an *image density*. In place of say, mass per unit volume, which is typical for $\rho(r,t)$, we consider I to be a brightness function with units of gray level per unit area. Extending this analogy, we still use the shorthand notation ∇, but it now explicitly represents the 2×1 gradient operator:

$$\nabla = \left[\begin{array}{c} \partial_x \\ \partial_y \end{array} \right]. \qquad (9)$$

3.1 Green's function solution

A standard procedure in edge detection, certainly for differential-based filters, is to first convolve the image with a Gaussian filter thereby reducing noise-induced fluctuations. The filter width controls trade-offs between noise reduction and image blurring. Because this

---

[2] More detailed analysis can be found in texts on Statistical Mechanics and Nonlinear Dynamics [14] for their treatment of phase space, Fluid Dynamics [15] in developing the Navier-Stokes equation, or Electromagnetism [16] in deriving for example, Gauss' Law.

filtering causes object edges within the original image to spread out, one can intuitively think of this as a diffusive process. Researchers noticed ([17] and [18], for example) that the convolution integral employed in this smoothing operation, and correspondingly the source of image degradation, is also the Green's function solution to the diffusion equation.

This equivalence can be demonstrated by solving eq. (3), the isotropic diffusion equation for $I(r, t)$

$$\frac{\partial I}{\partial t} = D\nabla^2 I, \tag{10}$$

where we identify the original image as the initial condition $I(r, 0) = I(r)$. We choose here an unbounded domain to avoid considering specific boundary conditions. We also require the solution to be square integrable for all $t > \infty$ (which implies that the solution vanishes at infinity, guaranteeing uniqueness of solutions). This is a classical problem, with a well-characterized solution [13]. Applying the spectral method [19] to eq. (10), we obtain a solution in the form (for a 2-dimensional system):

$$G(r, r'; t) = \frac{1}{4\pi Dt} \, e^{-\frac{|r-r'|^2}{4Dt}}. \tag{11}$$

The solution is labeled $G(r, r'; t)$ to denote the Green's function for diffusion. If we consider the diffusion operator $L \equiv \partial_t - D\nabla^2$, a Green's function solves the related system $LG(r, r'; t) = \delta(r - r')$ for a point source located at r'. This is also commonly referred to as the *fundamental solution* [20]. The significance of this is that once we have a solution for a point source, then the solution for an arbitrary initial configuration is simply the convolution integral:

$$\underbrace{I(r, t)}_{\substack{\text{Observed} \\ \text{image}}} = \int_{-\infty}^{\infty} \underbrace{G(r, r'; t)}_{\substack{\text{Blurring} \\ \text{function}}} \underbrace{I(r')}_{\substack{\text{Original} \\ \text{image}}} \, dr', \tag{12}$$

While eq. (10) can be solved a number of ways, it is important to note that the fundamental solution exhibits a dimensionless scaling quantity $\sim x/\sqrt{Dt}$, alluding to the physical basis of diffusion being a random walk, where the time required to move over a given length scale is proportional to the square of the length. Note that the Green's function is readily identified as a Gaussian kernel, and that comparison with the standard Gaussian filter establishes the correspondence $\sigma = \sqrt{2Dt}$, allowing us to equate temporal evolution of diffusion to the process of image blurring.

Additional constraints are required to account for boundary effects. One needs to impose zero flux at the external edges of the image. This takes the form of von Neumann boundary conditions, with the directional derivative set to zero:

$$\left. \frac{\partial I}{\partial n} \right|_{\partial\Omega} = 0. \tag{13}$$

where $\partial\Omega$ denotes the image boundary and n is the outward normal unit vector. We haven't done this in the Green's function formalism above, but eq. (13) is incorporated within the numerical implementation of the filter. Note that the von Neumann condition only exists at the external boundary, not within the image itself. Internally, diffusion is controlled adaptively as described in the following section.

3.2 Nonlinear filtering

The challenge in diffusion filtering is to selectively turn diffusion on or off depending on proximity to edges contained within the image — otherwise all features will be smoothed out, yielding an indistinct image and no discernible information. This is addressed by modifying the diffusion parameter to be a function of image gradient. The idea here is to throttle diffusion so that it slows the flux and then shuts it off as the gradient magnitude increases. To achieve this, we define flux in the form of eq. (6), producing the nonlinear diffusion equation [10]:

$$\frac{\partial I}{\partial t} = \nabla \cdot [D(|\nabla I|)\nabla I]. \tag{14}$$

Two alternatives for $D(|\nabla I|)$ were introduced in [11]. For comparison, both are plotted in Fig. 1.

$$D(|\nabla I|) = \frac{1}{1 + \left(\frac{|\nabla I|}{K}\right)^2}, \tag{15}$$

or

$$D(|\nabla I|) = e^{-\left(\frac{|\nabla I|}{K}\right)^2}. \tag{16}$$

Either expression can be used to define gradient-dependent diffusion

$$\varphi = D(|\nabla I|)\nabla I. \tag{17}$$

There are minor differences, but the importance of the form taken by either eq. (15) or eq. (16) lies in how they respond to edges. The parameter K acts as a threshold for selecting image gradients. The "image flux" as used in diffusion filtering has the desired characteristic of being monotonically increasing for $|\nabla I| < K$, and monotonically decreasing for $|\nabla I| > K$.

For example, substituting eq. (15) into eq. (17) and differentiating, we find

$$\varphi'(|\nabla I|) = \frac{d}{d(\nabla I)}\left[\frac{|\nabla I|}{1 + \left(\frac{|\nabla I|}{K}\right)^2}\right] = \frac{K^2(K^2 - |\nabla I|^2)}{(K^2 + |\nabla I|^2)^2}. \tag{18}$$

Correspondingly, substituting eq. (16) into eq. (17), results in

$$\varphi'(|\nabla I|) = \frac{d}{d(\nabla I)}\left[|\nabla I| \, e^{-\left(\frac{|\nabla I|}{K}\right)^2}\right] = \left(1 - \frac{2|\nabla I|^2}{K^2}\right)e^{-\left(\frac{|\nabla I|}{K}\right)^2}. \tag{19}$$
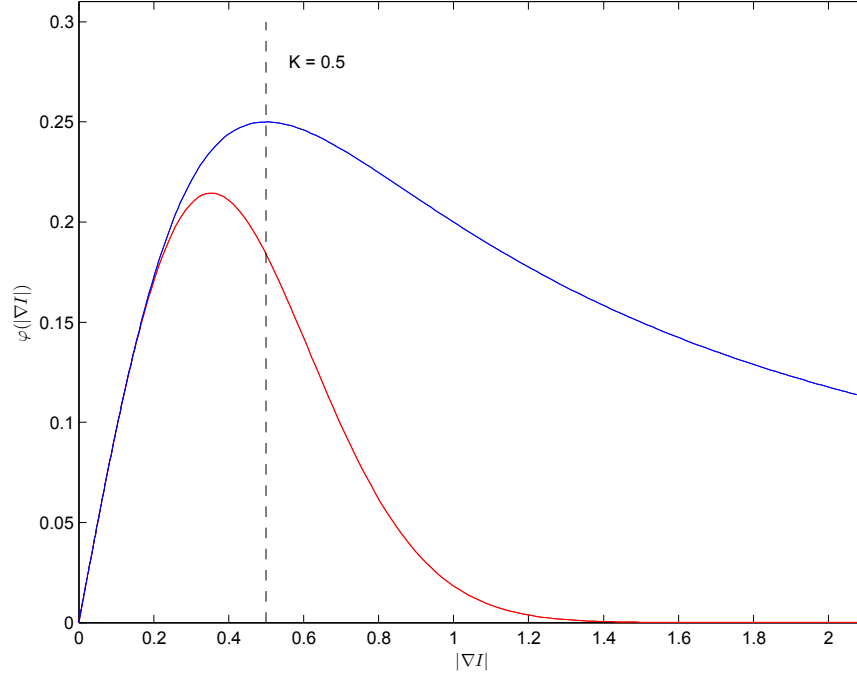
Figure 1: *Example for two alternatives of diffusion coefficients, with differing flux functions as a result. Red curve (lower) is the exponential expression. In both cases, K = 0.5.*

The resulting plots for derivatives are shown in Fig. 2. One noticeable difference between the two is that the maximum for the exponential-based function is shifted from $\mid \nabla I \mid = K$ to $\mid \nabla I \mid = K/\sqrt{2}$.

In summary, for images displaying large gradients (greater than the threshold K), the edge will be sharpened with time, and smaller gradients (below the threshold) will be blurred out. This behaviour runs counter to the physics of diffusion, where concentrated distributions disperse. However, the piecewise equations employed above ensure that stronger edges within an image are enhanced, and that the finite-difference implementation remains stable.

4 Connection to traditional adaptive speckle filters

The method most commonly applied to reduce speckle in SAR involves the use of an adaptive filter based on a multiplicative noise model. There are various formulations, and image processing software typically includes a speckle function option where the user can select from, for example, Lee [21], Enhanced Lee [22], Frost [23], or Kuan [24] filters, although this is far from being an exhaustive list [25].
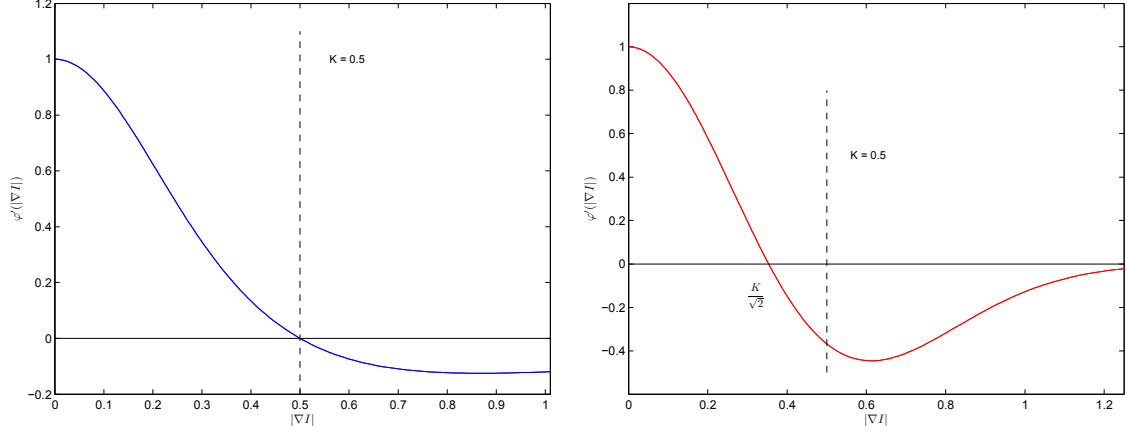
Figure 2: *First derivatives φ' for flux functions, showing extrema. The plot on the right (red line) corresponds to the exponential-based flux. For both plots the same edge threshold K is used.*

It is important to emphasize that metrics used in characterization of multiplicative (speckle) noise are different from those of additive noise. In [21], Lee examined both situations, using variance as the control parameter for adaptive filtering of additive noise, and the *coefficient of variation* for the multiplicative noise case.

This distinction results from analysis of the speckle model, where the observed image pixel I, is taken to be the product of the noise-free signal I_0 (i.e, ground reflectivity) and the noise, ξ:

$$I_{i,j} = I_{0;i,j}\xi_{i,j}, \tag{20}$$

where noise is assumed to be of unit mean. Subscripts are used to denote the pixel at coordinate (i,j), but to reduce notational clutter, will be suppressed in the following.

Image variance is

$$\begin{aligned} var(I) &= E[(I - \bar{I})^2] \\ &= E[I_0^2]E[\xi^2] - \bar{I}_0^2\bar{\xi}^2, \end{aligned} \tag{21}$$

which holds in general. However, for the special case where underlying ground reflectivity is homogeneous, we know the variance of I_0 will be zero: $var(I_0) = E[I_0^2] - \bar{I}_0^2 = 0$. We can then substitute $E[I_0^2] = \bar{I}_0^2$ into eq. (21) to obtain

$$\begin{aligned} var(I) &= \bar{I}_0^2 E[\xi^2] - \bar{I}_0^2\bar{\xi}^2 \\ &= \bar{I}_0^2(E[\xi^2] - \bar{\xi}^2) \\ &= \bar{I}_0^2\ var(\xi). \end{aligned} \tag{22}$$

This last line allows us to write out an expression for noise variance

$$var(\xi) = \frac{var(I)}{\bar{I}_0^2} = \frac{var(I)}{\bar{I}^2}, \tag{23}$$

where the last term has equated $\bar{I}_0$ with $\bar{I}$, which follows directly from eq. (20).

This result shows that noise variance is equal to the local image variance, scaled by the square of the local mean. Consequently, we see that within this model, the appropriate measure of speckle strength is indeed related to the coefficient of variation (CV) equal to $\sigma_I/\bar{I}$. This is why CV is such an important parameter in both speckle modeling and estimation—it effectively provides a measure of "normalized" or scaled variance, analogous to variance in systems with additive noise only [26].

From this analysis, Lee obtained an update equation for the speckle filter:

$$\hat{I} = \bar{I} + k(I - \bar{I}), \tag{24}$$

where $\hat{I}$ is the estimated value of I_0, $\bar{I}$ is the mean intensity within the filter window, and I is the original noise-corrupted pixel. The adaptive filter coefficient k can be written

$$k = 1 - \frac{C_\xi^2}{C_I^2}, \tag{25}$$

where C_ξ and C_I are the coefficients of variation for noise and image, respectively [27].

To better understand the properties of this filter (eqns. (24) and (25)), we consider the limiting behaviour of the coefficient k:

- In a homogeneous region, k is small so that $\hat{I} \simeq \bar{I}$, and the pixel takes on the value of the local mean.

- In high contrast areas (i.e., edges), k approaches unity and we obtain $\hat{I} \simeq I$, which is just the value of the pixel itself.

The effectiveness of this filter is therefore degraded when operating in high contrast areas. Consequently, a refined Lee filter [22] was developed to perform speckle filtering along edges.

4.1   Speckle filter implementation

Both the Lee and Frost formulations can be interpreted as discrete implementations of PDE-based *isotropic* diffusion filters. To illustrate this for the Lee filter, we begin with the isotropic diffusion equation

$$\begin{aligned}
\frac{\partial I}{\partial t} &= D\nabla^2 I, \\
&= D\left(\frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2}\right). \tag{26}
\end{aligned}$$

The finite difference, time-centered approximation is given by

$$I_{i,j}^{n+1} = I_{i,j}^n + \frac{D}{h^2}\left(I_{i,j+1}^n + I_{i,j-1}^n + I_{i+1,j}^n + I_{i-1,j}^n - 4I_{i,j}^n\right). \tag{27}$$

where the spatial increments $\Delta x = \Delta y = h$, and $t_n = n\Delta t$ with time increment $\Delta t = 1$. In such a discrete representation, we define a grid

$$\eta_s := \{(i+1, j), (i-1, j), (i, j+1), (i, j-1)\} \tag{28}$$

representing a discrete four-point lattice surrounding the pixel located at $s = (i, j)$.

With some algebraic manipulation, we find that eq. (24) can be rewritten as:

$$\hat{I} = I + (1-k)(\bar{I} - I), \tag{29}$$

and changing to index notation, the central pixel I becomes $I_{i,j}$ with mean $\bar{I}$ taken over the nearest neighbours. Here we choose a four-point sampling window corresponding to the neighbourhood η_s (eq. (28)). Expanding and then collecting terms, we find:

$$
\begin{aligned}
\hat{I}_{i,j} &= I_{i,j} + (1 - k_{i,j}) \left(\frac{1}{|\eta_s|} \sum_{p \in |\eta_s|} (I_p - I_{i,j}) \right) \\
&= I_{i,j} + (1 - k_{i,j}) \frac{1}{|\eta_s|} \left[(I_{i,j+1} - I_{i,j}) + (I_{i,j-1} - I_{i,j}) + (I_{i+1,j} - I_{i,j}) + (I_{i-1,j} - I_{i,j}) \right] \\
&= I_{i,j} + (1 - k_{i,j}) \frac{1}{|\eta_s|} \left[I_{i,j+1} + I_{i,j-1} + I_{i+1,j} + I_{i-1,j} - 4I_{i,j} \right] \\
&= I_{i,j} + (1 - k_{i,j}) \frac{1}{|\eta_s|} \nabla^2 I_{i,j}
\end{aligned}
\tag{30}
$$

This result is consistent with a constant diffusion parameter; the del operator (i.e., ∇) passes through to act on the gradient. The significance of this is that we see the Lee adaptive filter can be expressed in the form of a discrete isotropic filter, where the sampling window is the Laplacian operator:

$$\nabla^2 = \pm \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix} \tag{31}$$

4.2   The SRAD filter

For the formulation of the SRAD filter, CV again plays an important role, and in [10], diffusion is now based on an *instantaneous* CV measurement. Briefly, it combines a normalized gradient and a normalized Laplacian to better enable edge detection in both high and low-contrast areas. The diffusion equation developed earlier (eq. (14)) becomes

$$\frac{\partial I}{\partial t} = \nabla \cdot [D(q)\nabla I], \tag{32}$$

where

$$q = \sqrt{\frac{\frac{1}{2}\left(\frac{|\nabla I|}{I}\right)^2 - \frac{1}{4^2}\left(\frac{\nabla^2 I}{I}\right)^2}{\left[1 + \frac{1}{4}\left(\frac{\nabla^2 I}{I}\right)\right]^2}}. \tag{33}$$

Note that AD filters have been used in shoreline extraction [28], but Liu and Jezek's work did not include the instantaneous CV measurement. Strictly speaking, their implementation was appropriate for additive noise, not multiplicative, although their results are still impressive.

5   Test data

To examine SRAD filter performance, portions of an SLC Wide Ultra Fine beam mode (U5W2) RADARSAT-2 image were selected and analyzed. A typical SAR scene is shown in Fig. 3. Note that the land-water visual contrast apparent in the image is the result of spatial averaging carried out automatically within the image display software. This does not affect the underlying data; it is only used as part of the user screen interface. The actual contrast, which is much poorer, will be analyzed in the following pages.
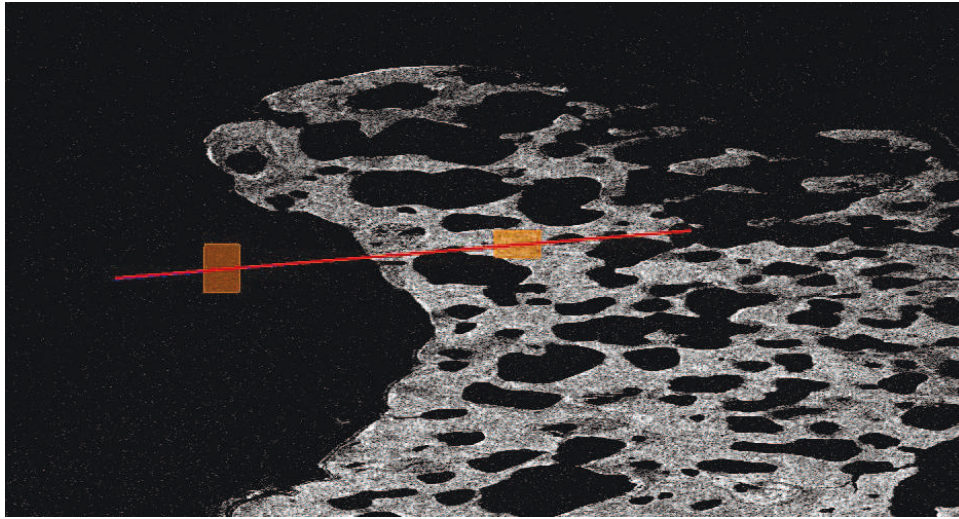
Figure 3: *RS-2 U5W2 image of Richards island. Areas marked highlight regions identified for statistical analysis. RADARSAT-2 Data and Products © Macdonald, Dettwiler and Associates Ltd. (2015) – all Rights Reserved. RADARSAT is an official trademark of the Canadian Space Agency.*

Superimposed on Fig. 3 is a line drawn from water, extending onto land. The graphic in Fig. 4 shows σ_0 in dB along this transect. The quantity σ_0, pronounced "sigma-nought" is a measure of the ground reflectivity. It measures the amount of RF energy backscatter returning to the SAR platform. Each pixel is represented by a digital number (DN), expressing the magnitude

of detected data, and look-up tables (LUTs) are used to relate this to calibrated values for radar backscatter [29]. Annex A provides an outline of the calculations. At the transition point between water and land (around pixel number 1600), we see the water backscatter is approximately -25 dB, while that of the land appears to be about -19 dB, yielding a contrast estimate of 6 dB. The dotted line shows the noise floor. The nominal NESZ for this beam mode is -22±7 dB [2]. This is one of the reasons we are considering beam modes such as this – the combination of low contrast between land and water, and the proximity to the noise floor make these types of images difficult to extract shorelines using standard techniques.
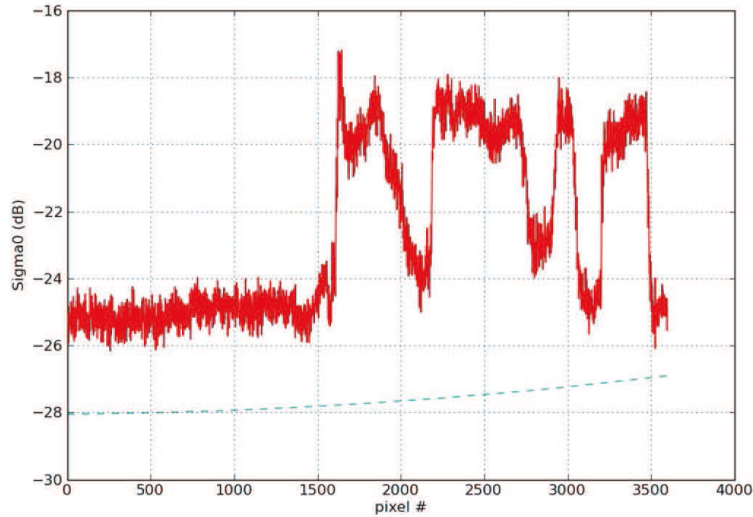
Figure 4: *Backscatter (sigma0, or σ_0) in dB obtained from transect indicated in Fig. 3. To reduce fluctuations, the profile shown here has been averaged over transect width, in this case 100 pixels.*

5.1   A model coastline

Fig. 3 also shows an overlay of two rectangular areas. These indicate regions where data has been chipped out to measure backscatter separately for water and land. We take the data from both areas to generate 512×256 sub-images of water pixels next to land pixels, resulting in a 512×512 image. This produces a vertically oriented "shoreline" (Fig. 5). While this may appear as an artificial construct, it allows us to run filter tests with real rather than synthesized data, knowing exactly where the edge is. It provides a more realistic depiction of the actual contrast between land and water masses. This also provides a more realistic visual depiction of the actual contrast between land and water masses – one can see that the edge is barely discernible in either the dB or linear plots of Fig. 5.
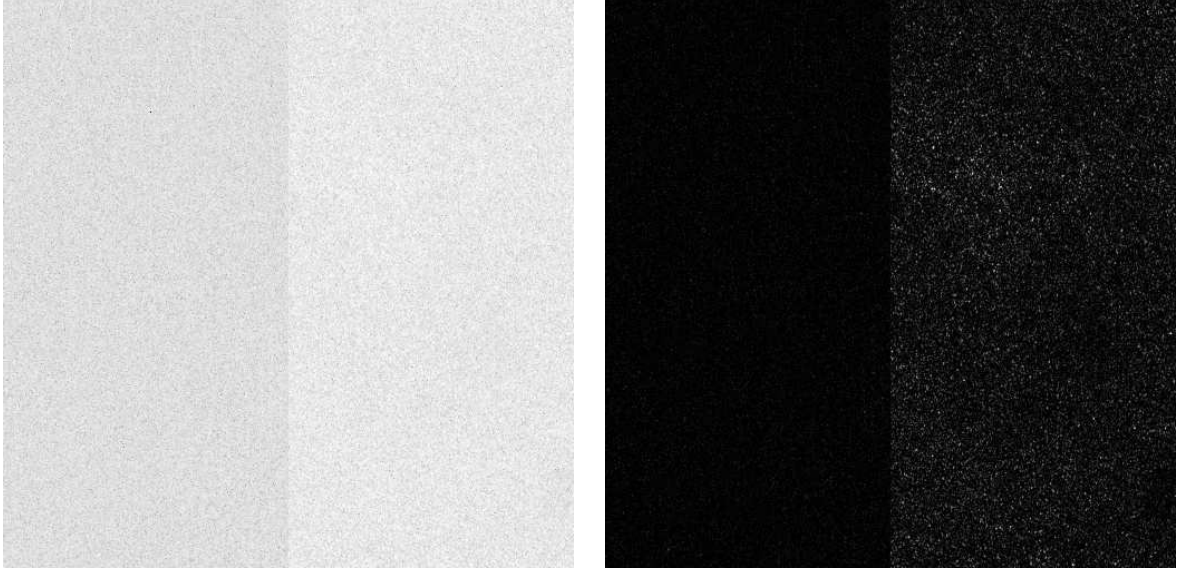
Figure 5: *Idealized shoreline. Image on the left is RS-2 data plotted as σ_0 in dB, image on the right as σ_0 in linear units. In both images the darker water backscatter is to the left, land to the right. This image was generated by combining data sets obtained from rectangles marked out in Fig. 3.*

For the separate regions, statistical analysis reveals negative exponential distributions[3] with means of $\langle\sigma_0\rangle_{water} = 0.0033$ and $\langle\sigma_0\rangle_{land} = 0.0121$. Converting the data to dB units, $10\log_{10}\langle\sigma_0\rangle$, we obtain, for water and land respectively, values of -24.81 dB and -19.17 dB. This results in a contrast of 5.64 dB, consistent with the plot obtained in Fig. 4. Histograms of the individual distributions are shown in Fig. 6. Superimposed (dotted line) is the combined distribution for land and water pixels. These are all plotted as normalized histograms using relative frequencies (Otherwise, combined and individual distributions would not plot on the same scale). The important point to note from Fig. 6 is that separation between the two distributions is insufficient to produce a bimodal distribution.

---

[3] A discussion of SAR image distributions is contained in the annex.
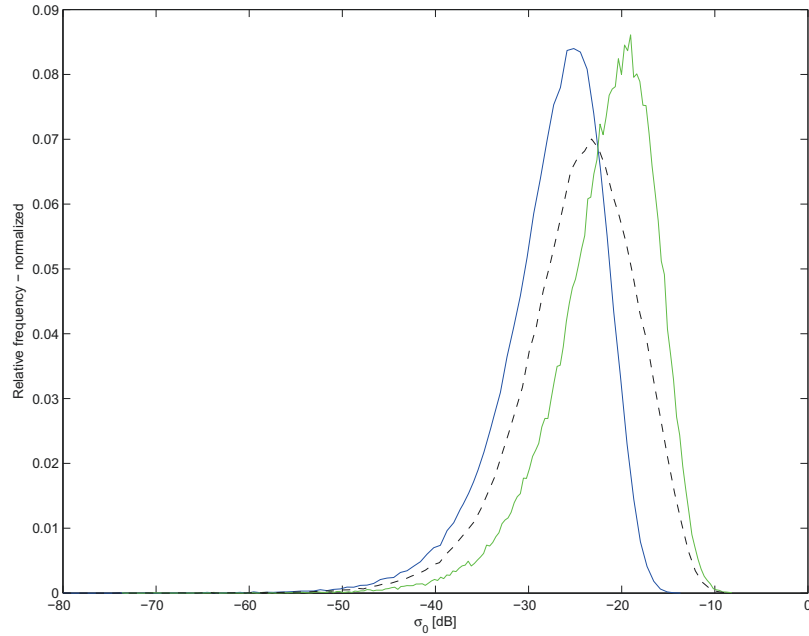
**Figure 6:** *Comparison of distributions for land (green) and water (blue) individually, and combined (dotted line). Histograms show relative frequency vs. sigma0 in dB.*

5.2   Processing and edge detection

Next, the combined distribution (that producing Fig. 5, or dotted line of Fig. 6), is input to the SRAD filter. As this is an iterative algorithm, we select N = 50 and N = 200 as iterations for the test runs. Output from this is shown in Fig. 7, which is a cross-section of the original and filtered image amplitudes oriented perpendicular to the edge profile. While there is an edge in the original (green) data, it is very hard to discern. For reference, the edge position at range index 256 is marked with a vertical dotted line. As the number of iterations increases, the smoothing on either side of the edge is apparent. Even at N = 50, the profile (blue) shows a marked improvement, with a small amplitude ripple running along the cross-section.

There are two points to note from this figure. The first is that SRAD is mean-preserving. While the variance has reduced considerably, the overall image intensity remains the same. Considering the conservation property of diffusion and the boundary condition eq. (13), we sum the original image pixels and compare with the sum after 200 iterations, yielding a difference of 0.5%. Second, the edge has not shifted. Many of the kernel-based filters tend to shift the edge a pixel or two, depending on the size of the mask. Once the SRAD algorithm identifies an edge from combined gradient and Laplacian operators, that position is fixed and the diffusive flow spreads out from there.
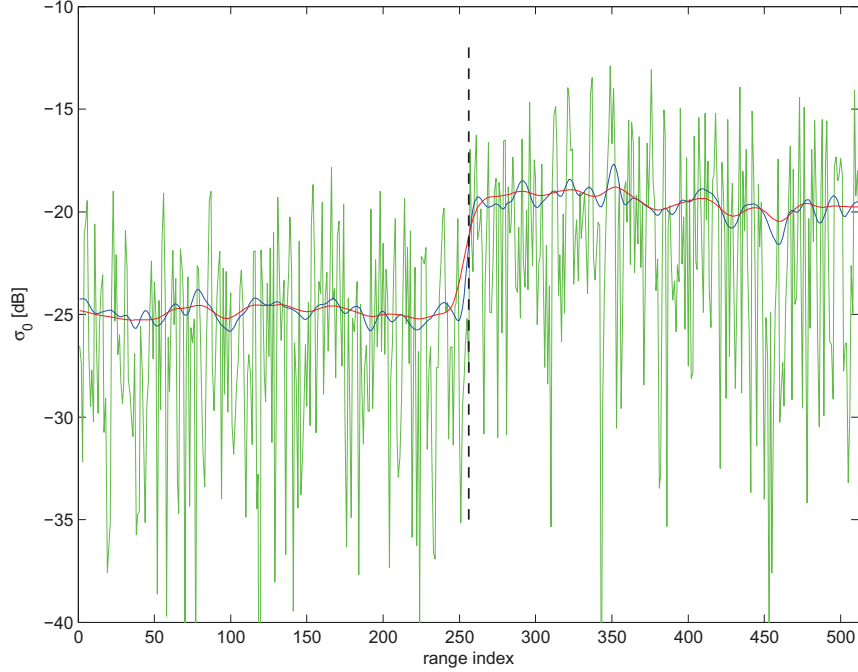
Figure 7: *Cross-sectional view of the model shoreline data set. The increasingly smoothed plots show i). the initial unprocessed data (green), ii). N = 50 output (blue), and N = 200 output (red). For reference, the water-land boundary is indicated by a vertical dotted line.*

How these results appear visually, or at least to an edge detector, are shown in Fig. 8. The Canny operator was applied in MATLAB using the "automatic" option, and edges overlaid in red on the image. The results are not optimized but still provide a good indicator for SRAD filter performance. The first image (original data set, Fig. 5) clearly has too many false positives to be of any practical use. The human eye can start to discern a vertical line in the center of the image, but computer vision applications would be severely challenged. For N=50, there is now a disjoint vertical line running along the image edge (i.e., "coastline"), but again, it is swamped by false positives. Finally, for N=200, we now have a more continuous line representing the edge, with perhaps a manageable number of false positives. This could likely be improved by setting up an optimized Canny operator at the start, or the false edges could be cleaned up using morphological operations [9].

The right-hand column of Fig. 8 shows how the histogram of the image evolves with the number of SRAD iterations. For reference, the mean values obtained from the individual distributions are superimposed on the plots. Again, we see the bimodal distribution is totally nonexistent in the original image, becoming quite apparent even at the N = 50 mark.
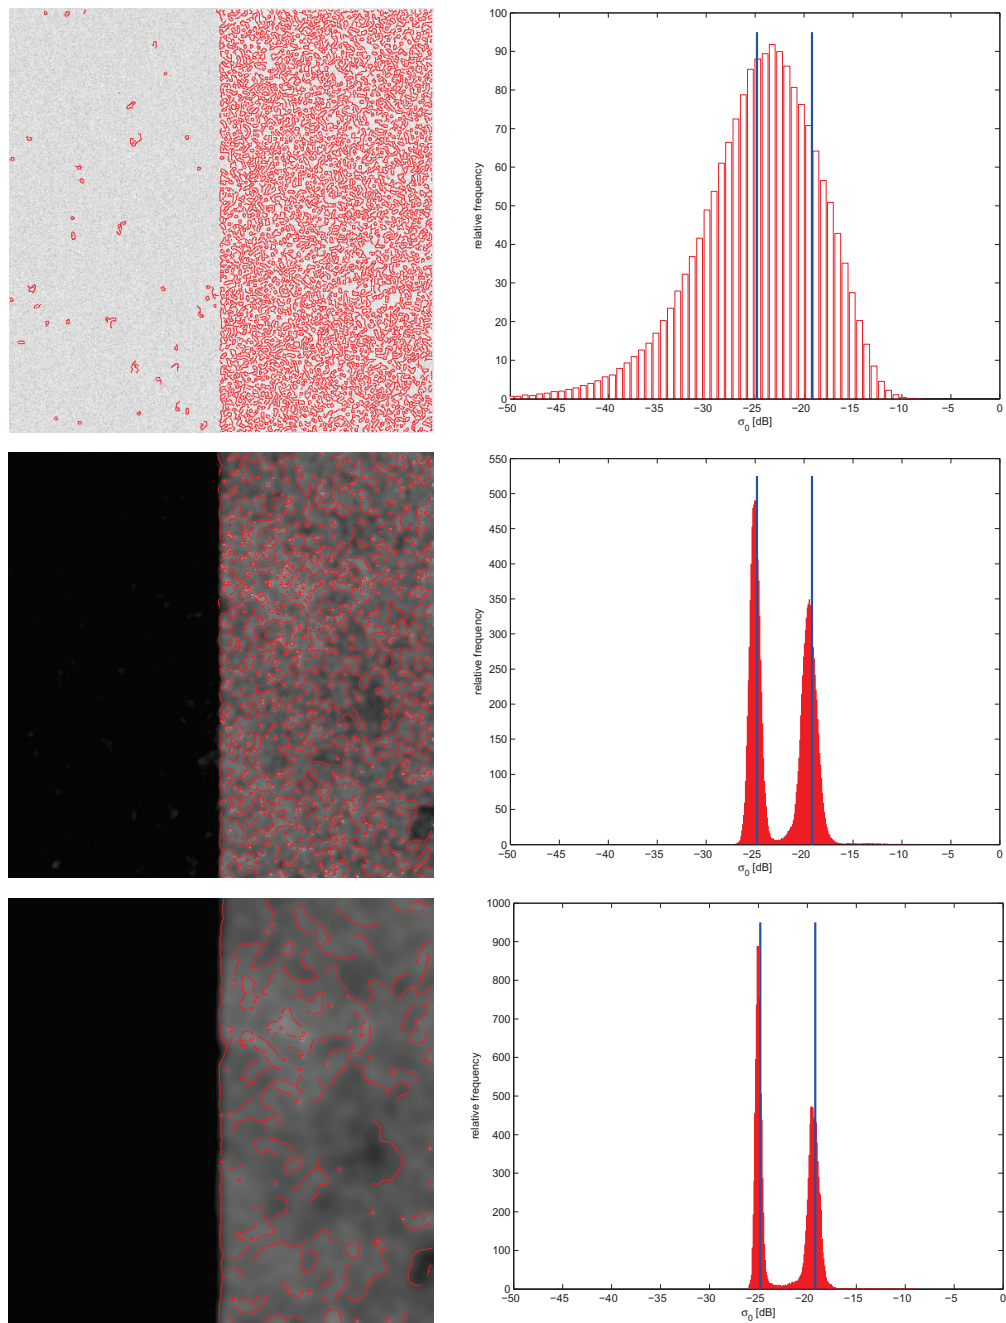
Figure 8: *Left: Results of the Canny edge detector applied to (from top to bottom) original image, N=50, and N=200 iterations of the SRAD filter. Right: Histograms derived from images. Blue vertical lines indicate means from the original individual pixel distributions.*

6 Application to RADARSAT-2 images

Looking at a less contrived situation, we apply the SRAD algorithm to a larger (∼2 GB) RS-2 data file. A SAR scene of Richards Island, NWT (60.69°N, 134.36°W) is analyzed. This is an SLC Wide Ultra Fine beam mode (U5W2) image[4]. It has a 1.6m × 2.8m nominal resolution, and swath width of approximately 50 km. The collection occured during an ascending pass, with an incidence angle of 34.2°, and HV polarization.

6.1 Visual and histogram analysis

We compare the original, unfiltered data (Fig. 9) with that of SRAD for 20 (Fig. 10) and 100 (Fig. 11) iterations. Since the algorithm is a simple finite-difference solver for the diffusion equation, the processing time is relatively long. The 20 iteration case took approximately 10 minutes, while 100 iterations were a little over 50 minutes[5].

In all three figures, the left hand side shows the original and filtered images, while the right hand side shows a histogram of the image ($\sim 2.2 \times 10^8$ pixels). Showing similar behaviour as the model coastline, there is no indication of two modes in the initial image because of the poor contrast. However, even after 20 iterations, the bimodal distribution is clearly visible. To relate back to the model problem, these plots can be compared with the results obtained in Fig. 8. There was no attempt to further process with the Canny edge detector at this point because the focus is on application of SRAD to enable separation of land and water masses via histogram analysis, which is what is shown here.
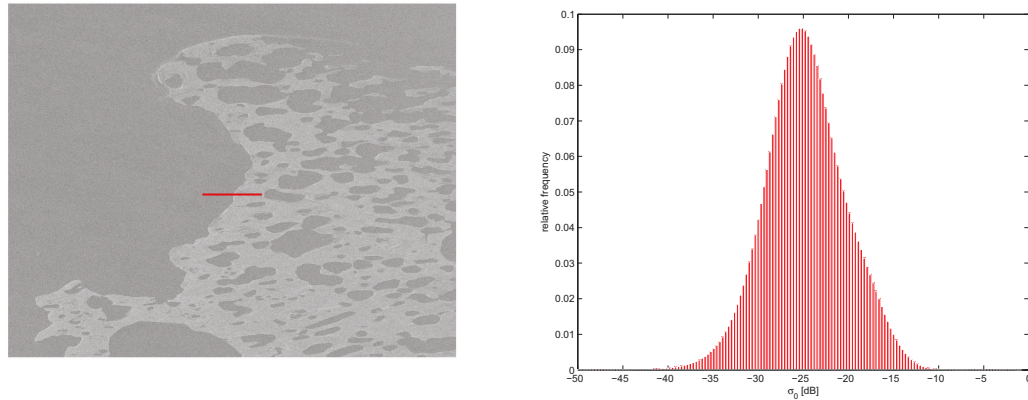
Figure 9: Unprocessed U5W2 magnitude image and histogram.

---

[4] $RS2\_OK66427\_IK414435\_PEK001488244\_U5W2\_20150812\_021510\_HV\_SLC$.
[5] IBM Z800 Workstation, 64-b OS, Intel Xeon Processor X5675, 3.06 GHz, 24 GB RAM.
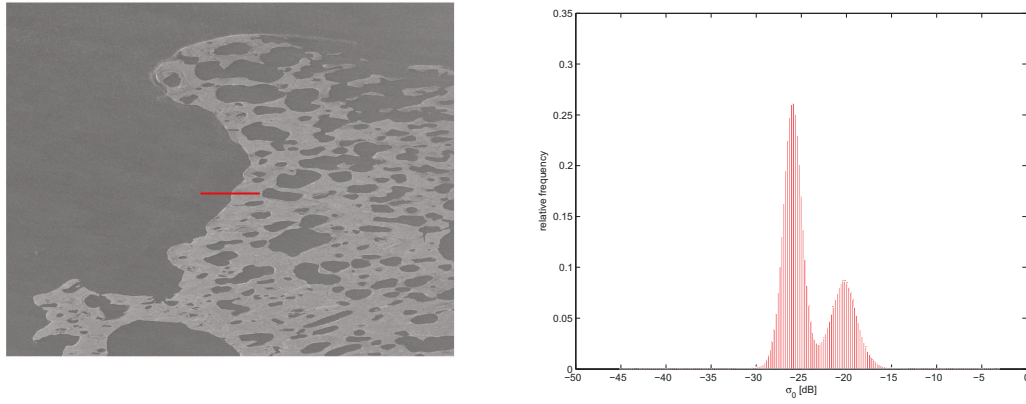
Figure 10: *SRAD (N=20) U5W2 magnitude image and histogram.*
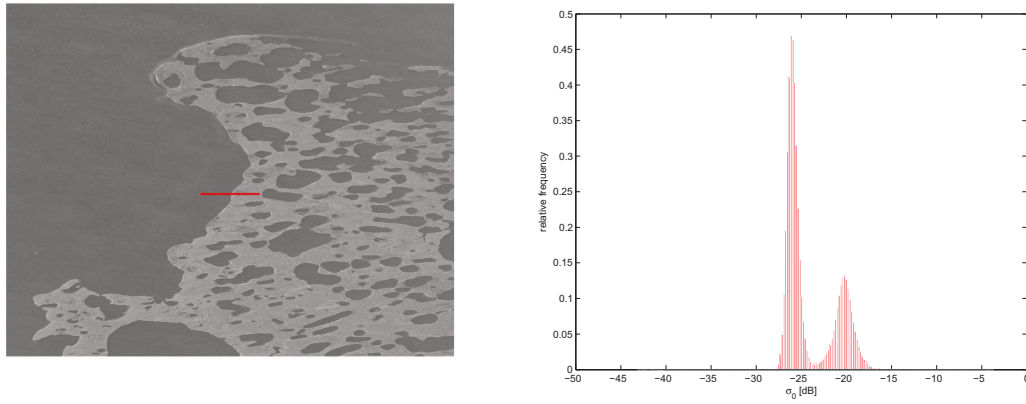
Figure 11: *SRAD (N=100) U5W2 magnitude image and histogram.*

6.2 Water-Land transects

For further comparison, cross-sections were extracted from the RS-2 processed data, showing the transition from water to land, as shown in Fig. 12. The transect used to obtain this is visible as a thick red line in Fig. 9 to Fig. 11. Qualitatively, the same results appear as for the model problem (Fig. 7) – reduction in overall variance, increasing smoothing with the number of iterations, and a fixed edge location.
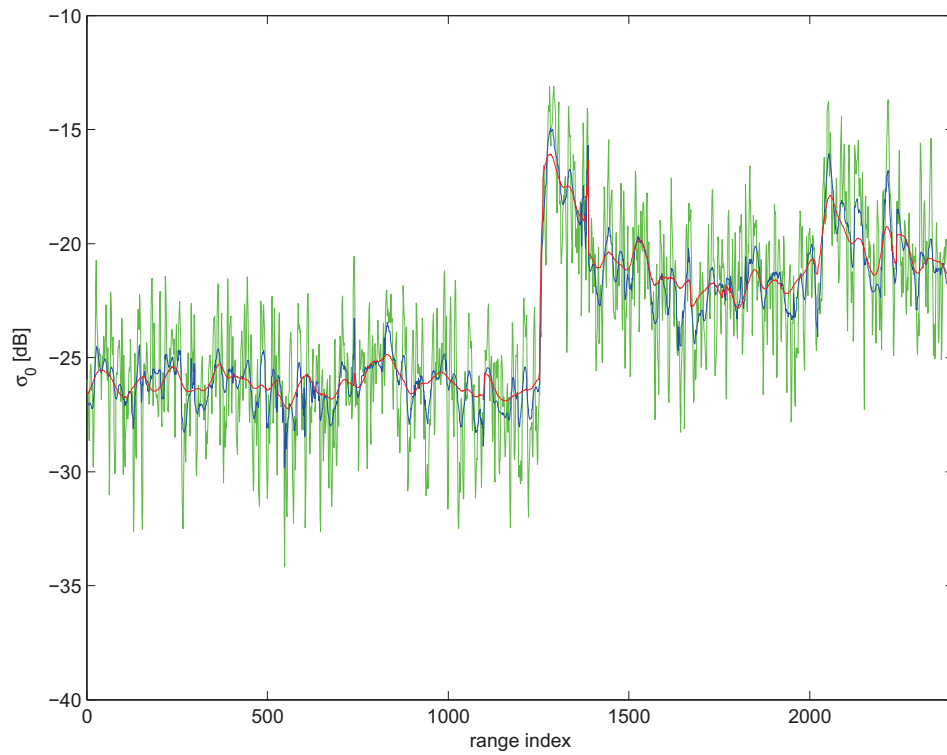
Figure 12: *Cross-sections obtained from RS-2 data. Unprocessed (green), N=20 (blue), N=100 (red).*

7 Discussion

This report demonstrated the potential effectiveness of using SRAD filters to extract shorelines in poor imaging conditions. By "poor" we mean low backscatter contrast between water and land. The standard practice, and most straight-forward technique, is to classify land and water masses based on their observed statistical distributions. Water pixels are generally darker than land pixels, leading to a bimodal distribution. Either by visual inspection, or other techniques such as Otsu's method, one can determine an appropriate threshold to achieve good separation of the two classes.

However, as illustrated in this work, differences of 5 to 6 dB render the land-water distributions inseparable. This situation can arise a number of ways, either through imaging geometry (most notably the incidence angle), or as a result of using a non-optimal beam mode for the particular imaging conditions. With respect to this work, RS-2 collects made for DRDC's shoreline investigation were specifically chosen to enhance shoreline extraction, but in practice, large scale work will require wider swaths and correspondingly less favourable conditions. Environmental conditions such as wind or ice cover will also reduce contrast.

In using SRAD to pre-process low contrast imagery, we found that bimodal distributions can be recovered. Moving forward with this promising result, we next want to incorporate this algorithm in a fully supported geospatial processing environment to test on a variety of images, including RCM. There are a variety of AD-related filters to examine, such as Detail Preserving Anosotropic Diffusion (DPAD) [30] and other methods that explicitly use diffusion tensors[31]. These approaches have the benefit of not only minimizing smoothing across edges in an image, but also maximizing smoothing along the edges. Equally of interest is applying this technique to delineation of ocean fronts by augmenting the edge detection currently employed [32].

Also, the perceived improvement in images shown in this work needs to be fully quantified. There are a number of metrics in the literature used to measure how well a speckle filter works vs. another [33], [34]. We will test AD filters against more commonly used methods such as the minimum mean-square based filter family (Lee, Enhanced Lee, Frost, Kuan, etc.) [25].

References

[1] N. Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1): 62–66, 1979.

[2] MDA. RADARSAT-2 Product Description. Technical Report RN-SP-52-1238, MacDonald, Dettwiler and Associates, 2016.

[3] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(6): 679–698, 1986.

[4] D. Sen and S.K. Pal. Thresholding for edge detection in SAR images. In *IEEE-International Conference on Signal Processing, Communications and Networking*, pages 311–316, 2008.

[5] M. Fang, G. Yue, and Q. Yu. The study on an application of Otsu method in Canny operator. In *Proceedings of the 2009 International Symposium on Information Processing (ISIP'09)*, pages 109–112, 2009.

[6] J.W. Goodman. Some fundamental properties of speckle. *Journal of the Optical Society of America*, 66(11): 1145–1150, 1976.

[7] J.S. Lee, et.al. Speckle filtering of synthetic aperture radar images: A review. *Remote Sensing Reviews*, 8: 313–340, 1994.

[8] L. Weibin, H. Mingyi, and Z. Shunli. A new Canny-based edge detector for SAR image. In *Congress on Image and Signal Processing (CISP) - IEEE Computer Society*, volume 5, pages 211–215, 2008.

[9] R.C. Gonzalez and R.E. Woods. *Digital Image Processing*. Pearson Prentice Hall, Third edition, 2018.

[10] Y. Yu and S.T. Acton. Speckle reducing anisotropic diffusion. *IEEE Transactions on Image Processing*, 11(11): 1260–1270, 2002.

[11] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7): 629–639, 1990.

[12] D.A. McQuarrie. *Statistical Mechanics*, chapter 17. Harper & Row, 1976.

[13] J. Crank. *The Mathematics of Diffusion*. Clarendon Press, Second edition, 1975.

[14] G. Nicolis. *Introduction to Nonlinear Science*. Cambridge University Press, 1995.

[15] D.J. Tritton. *Physical Fluid Dynamics*, chapter 5. Oxford University Press, Second edition, 1988.

[16] J.D. Jackson. *Classical Electrodynamics*. Wiley, 1999.

[17] A. Witkin. Scale–space filtering: A new approach to multi-scale description. In *IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP '84*, 1984.

[18] J.J. Koenderink. The structure of images. *Biological Cybernetics*, 50: 363–370, 1984.

[19] G.F.D. Duff and D. Naylor. *Differential Equations of Applied Mathematics*. John Wiley & Sons, 2018.

[20] W.A. Strauss. *Partial Differential Equations*. John Wiley & Sons, 2018.

[21] J.S. Lee. Digital image enhancement and noise filtering by use of local statistics. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-2(2): 165–168, 1980.

[22] J.S. Lee. Refined filtering of image noise using local statistics. *Computer Graphics and Image Processing*, 15: 380–389, 1981.

[23] V.S. Frost, et.al. A model for radar images and its application for adaptive digital filtering of multiplicative noise. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 4(2): 157–165, 1982.

[24] D.T. Kuan, et.al. Adaptive restoration of images with speckle. *IEEE Transactions on Acoustics*, ASSP-35:373–383, 1987.

[25] F. Argenti, et. al. A tutorial on speckle reduction in synthetic aperture radar images. *IEEE Geoscience and Remote Sensing Magazine*, 1(3): 6–35, 2013.

[26] S. Aja-Fernández, et.al. Automatic noise estimation in images using local statistics. *Image and Vision Computing*, 27: 756–770, 2009.

[27] R. Touzi A. Lopes and E. Nezry. Adaptive speckle filters and scene heterogeneity. *IEEE Transactions on Geoscience and Remote Sensing*, 28(6): 992–1000, 1990.

[28] H. Liu and K.C. Jezek. A complete high-resolution coastline of antarctica extracted from orthorectified radarsat sar imagery. *Photogrammetric Engineering & Remote Sensing*, 70(5): 605–616, 2004.

[29] K. El-Darymli, et. al. Understanding the significance of radiometric calibration for synthetic aperture radar imagery. In *2014 IEEE 27th Canadian Conference on Electrical and Computer Engineering (CCECE)*, 2014.

[30] S. Aja-Fernández and C. Alberola-López. On the estimation of the coefficient of variation for anisotropic diffusion speckle filtering. *IEEE Transactions on Image Processing*, 15(9), 2006.

[31] F. Benzarti and H. Amiri. Image denoising using non linear diffusion tensors. *Advances in Computing*, 2(1): 12–16, 2012.

[32] C.T. Jones, et. al. Automated discrimination of certain brightness fronts in RADARSAT-2 images of the ocean surface. *Journal of Atmospheric and Oceanic Technology*, 30: 2203–2215, 2013.

[33] Z. Wang, et. al. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 1–14, 2004.

[34] S.G. Dellepiane and E. Angiati. Quality assessment of despeckled SAR images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 7(2): 691–707, 2014.

[35] MDA. RADARSAT-2 Product Format Definition. Technical Report RN-RP-51-2713, MacDonald, Dettwiler and Associates, 2016.

[36] S. Bocquet. Calculation of radar probability of detection in k-distributed sea clutter and noise. Technical Report DSTO-TN-1000, Defence Science and Technology Organization (DSTO), Australia, 2011.

[37] G. Gao. Statistical modeling of SAR images: A survey. *Sensors*, 10: 775–795, 2010.

[38] J.S. Bendat and A.G. Piersol. *Random Data: Analysis and Measurement Procedures.* Wiley-Interscience, 1971.

[39] A. Papoulis. *Probability, Random Variables, and Stochastic Processes.* McGraw-Hill, 1965.

Annex A Representation of SAR backscatter

For amplitude (or intensity) products, one considers a greyscale image composed of pixels. Each pixel is represented by a digital number (DN), expressing the magnitude of detected data. In the case of complex data, there are two DNs, one for each of the in-phase and quadrature ($I\&Q$) components. DNs are in turn, converted to radiometrically calibrated values with the use of lookup tables (LUTs) provided with image data . Proper calibration relates digital values to physical quantities of the imaged scene, and allows one to compare targets imaged with different sensors, or with the same sensor under different collection conditions.

For example, RS-2 data is supplied with three LUTs, either β_0, σ_0, or γ_0 (pronounced beta-nought, sigma-nought or gamma-nought, respectively), each containing a fixed offset value, B, and a set of range dependent gain parameters, A_j. For most cases, $B = 0$. To convert DNs to calibrated values, we have the general form ([2], [35])

$$calibrated\ value = \frac{DN_j^2 + B}{A_j}, \tag{A.1}$$

where DN_j is the digital number for the j^{th} range pixel. For SLC products, the square of the modulus of the DN divided by the square of the gain parameter is used:

$$calibrated\ value = \frac{\mid DN_j \mid^2}{A_j^2}. \tag{A.2}$$

For specific backscatter coefficients, radar brightness, β_0 is a "natural" radiometric value as it expresses backscatter per unit area in the slant range and does not require the local incidence angle, θ_i.

$$\beta_{0,j} = 10 log_{10} \left[\frac{DN_j + B}{A_j} \right], \tag{A.3}$$

(units in dB). For single-look complex (SLC) products containing both $I\&Q$ components, this expression becomes

$$\beta_{0,j} = 20 log_{10} \left[\frac{DN_j}{A_j} \right]. \tag{A.4}$$

However, variable terrain within a SAR image causes geometric and radiometric distortion. Calibration of backscatter, accounting for variation in local incidence angles is carried out using the sigma-nought LUT, backscatter per unit area in ground range (in dB):

$$\sigma_{0,j} = \beta_{0,j} + 10 log_{10}(\sin \theta_i), \tag{A.5}$$

or γ_0 (gamma-nought), backscatter per unit area of the incident wavefront:

$$\gamma_{0,j} = \beta_{0,j} + 10 log_{10}(\tan \theta_i), \tag{A.6}$$

Calculations assume earth is a smooth ellipsoid at sea level (for example, GRS80 or WGS84) and platform-specific corrections are incorporated within the data furnished by tabulated data. These tables help to compensate for range dependent changes in the backscatter response, resulting in images that are visually consistent from near to far range. Additionally, LUTs are used to scale pixel values within the dynamic range of the SAR processor. For example, 16-bit unsigned integers with DNs ranging from 1 to 65535 result in a dynamic range of $20 log_{10}(2^{16} - 1) = 96.3$ dB. For SLC products, the pixel values are signed 16-bit integers ranging from -32768 to $+32767$.

Annex B Modeling SAR distributions

Much of the data analysis in this report depends on understanding distributions generated from radar backscatter. A major component of statistical fluctuation is due to *speckle*, appearing in SAR images as a consequence of the coherent nature of signals transmitted by the radar. Scattering from individual facets within a resolution cell occurs when imaged terrain is rough with respect to the scale of the radar wavelength. All of the individual scattering return paths (now with randomly distributed phase) are combined at the receiver and observed as pixel-to-pixel amplitude fluctuations, producing the aforementioned speckle. An image is said to be that of a distributed target in situations where there is no single return dominating all others. Further, if the number of individual scatters is large, the central limit theorem allows us to express the complex radar signal as $Ae^{i\theta} = x + iy$, where the I&Q components are iid zero-mean Gaussian RV's with variance σ. The phase θ is uniformly distributed on the interval $[0, 2\pi)$. Under these conditions, speckle is considered "fully-developed." The phenomenon of speckle is well-studied and characterized in the image processing community (See for example, [25] and references within). The notes in this appendix only skim the surface of a vast field.

As an illustration, refer to Fig. B.1, showing amplitude and intensity plots of RS-2 digital numbers (DNs). As discussed in Annex A, each pixel in SLC data is represented by two DNs in the form of I&Q data. By application of scaling data contained in look-up tables (LUTs), these can be converted to a backscatter coefficient, σ_0. The left hand side of fig. B.1 is amplitude, $A = \sqrt{x^2 + y^2}$, leading to a Rayleigh distribution, while the right hand side is intensity, $I = A^2 = x^2 + y^2$, producing an exponential distribution.
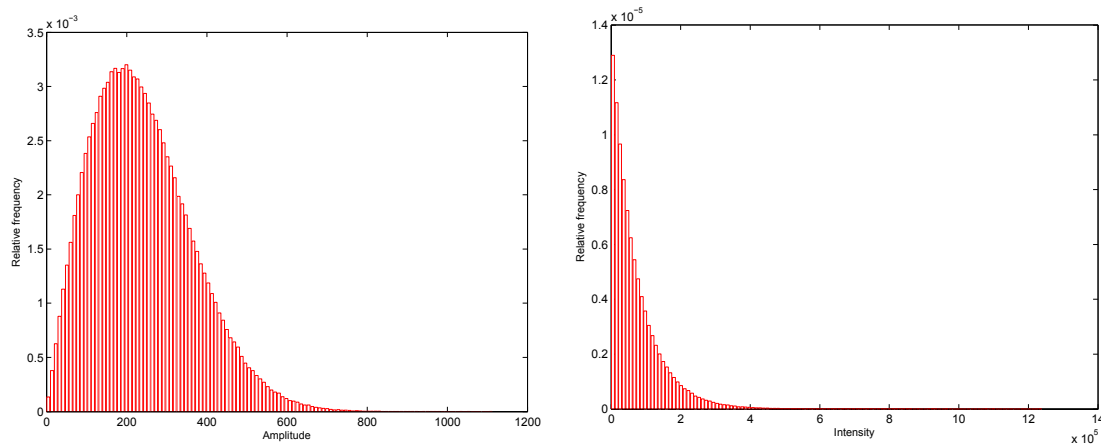
Figure B.1: *Amplitude (left) and intensity (right) histograms of RS-2 DNs. Amplidude data leads to a Rayleigh distribution, while intensity data produces an exponential distribution.*

We want to work with properly scaled data, say, σ_0 values. This is the case for fig. 3, from the main text. For a SAR system illuminating an area consisting of uniformly-distributed

scatterers, with backscattering cross-section σ_0, the single-look pixel intensity is described by an exponential probability distribution function (pdf) written as

$$f_X(x; \beta) = \frac{1}{\beta} \, e^{-x/\beta}, \tag{B.1}$$

where parameter β is the mean of the distribution (in other words, $\beta = \bar{\sigma}_0$). The cumulative distribution function (cdf) is $F_X(x) = 1 - e^{-x/\beta}$ and the two are related via a derivative operation:

$$f_X(x) = \frac{dF_X(x)}{dx}. \tag{B.2}$$

To describe a multi-look scenario, we consider a sum of random variables. In the case of two RV distributions, the convolution

$$(f * f)(z) = \int_{-\infty}^{\infty} f(x) f(z - x) dx \tag{B.3}$$

yields the following:

$$P(z) = \frac{1}{\beta^2} \int_0^z e^{-x/\beta} \, e^{-(z-x)/\beta} dx = \frac{z e^{-z/\beta}}{\beta^2}, \tag{B.4}$$

which can be recognized as the Gamma distribution for $\alpha = 2$. In turn, this can be generalized to higher order by simply repeating the process. For k convolved distributions, we obtain:

$$P(z; \alpha = k, \beta) = \frac{1}{1 \cdot 2 \cdots (k - 2)} \frac{e^{-z/\beta}}{\beta^k} \int_0^z x^{k-2} dx, \tag{B.5}$$

which results in the standard representation of the Gamma pdf

$$P(x; \alpha, \beta) = \frac{x^{\alpha-1} e^{-x/\beta}}{\beta^\alpha \, \Gamma(\alpha)}. \tag{B.6}$$

For single-look images, i.e., SLC, the parameter $\alpha = 1$ and eq. (B.6) reduces to the exponential distribution eq. (B.1). Additional complexity can be built into this model by considering the random nature of the terrain itself convolved with speckle, giving rise to the so-called K-Distribution [36],[37]. We have no need for this in our treatment, and it turns out that the exponential pdf matches the observed distributions quite nicely, as can be observed in Fig. B.2.

This plot shows a histogram of the data obtained from Fig. 3, where we have chipped out backscatter values for sea returns (rectangle on the left). This represents approximately 450k pixels with values ranging from $\sigma_0 = 0$ to 0.046, with mean 0.0033. Setting this as the β parameter, we plot the function $e^{-x/\beta}/\beta$, which appears superimposed on the histogram as a blue line. Histograms are properly normalized by dividing the individual bin counts by
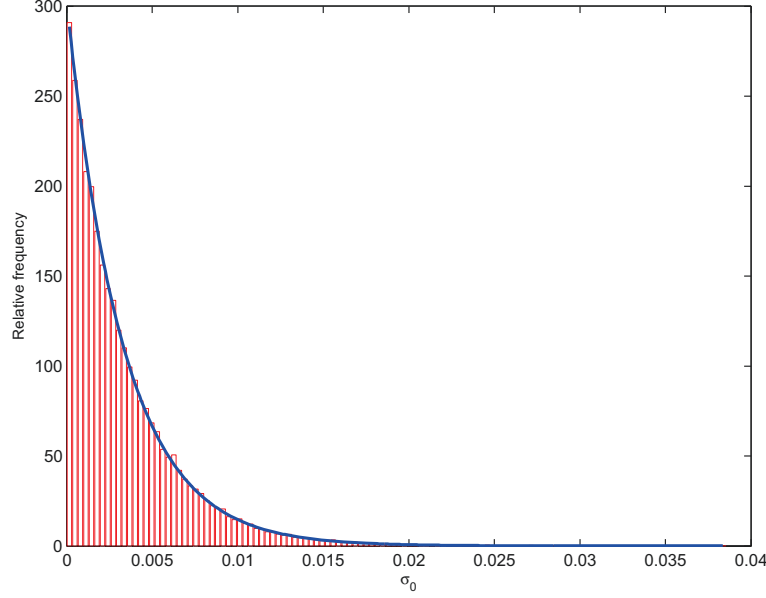
Figure B.2: *Exponential distribution obtained from the water segment of RS-2 data from Fig. 3. Blue overlay is the pdf for an exponential distribution with mean $\beta = 0.0033$.*

the overall area of the histogram (total count of all bins × bin width) [38]. This allows us to properly overlay pdfs on the histogram data.

However, the usual scaling used for SAR data is $\sigma_0 \rightarrow 10\log_{10}(\sigma_0)$, so we need to know what the corresponding distribution will look like. This is typically achieved with a random variable (RV) transform such as described in [39]. As an illustration, we consider $Y = LnX$, where X is the exponentially distributed RV, and we wish to perform a logarithmic transform.

$$
\begin{aligned}
F_Y(y) &= P[Y \leq y] \\
&= P[LnX \leq y] \\
&= P[X \leq e^y] \\
&= F_X(e^y) \\
&= 1 - e^{-\frac{e^y}{\beta}} \quad (B.7)
\end{aligned}
$$

This last step substitutes e^y into the original cdf for the exponential distribution, and we can then obtain the corresponding pdf.

$$
f_Y(y) = \frac{dF_Y(y)}{dy} = \frac{d}{dy}\left(1 - e^{-e^y/\beta}\right) = \frac{1}{\beta}e^{-e^y/\beta}e^y. \quad (B.8)
$$

This illustrates the technique, but to make a direct connection to SAR data, we want to focus on a base 10 log transform. Here, we write $F_Y(y) = P[10\log_{10} X \leq y]$, and following

the same steps as previous, the resulting cdf is

$$F_Y(y) = (1 - e^{-\frac{10^{y/10}}{\beta}}).$$ (B.9)

Once again, taking the derivative to find the pdf, we have

$$\begin{aligned}
f_Y(y) &= F_Y'(y) \\
&= \frac{1}{\beta} 10^{y/10-1} Ln(10) e^{-\frac{10^{y/10}}{\beta}} \\
&= \frac{0.23}{\beta} 10^{y/10} e^{-\frac{10^{y/10}}{\beta}}.
\end{aligned}$$ (B.10)

We have superimposed eq. (B.10) on the histogram in Fig. B.3. Again, we see there is very good agreement.
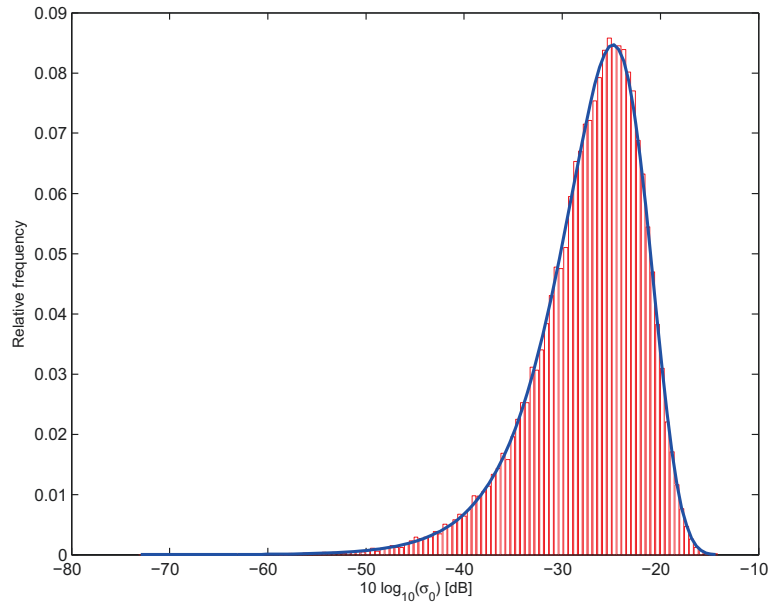
Figure B.3: *Log-transformed RS-2 data from Fig. B.2 with pdf overlay (see text for derivation).*

DOCUMENT CONTROL DATA

*Security markings for the title, authors, abstract and keywords must be entered when the document is sensitive

1. ORIGINATOR (Name and address of the organization preparing the document. A DRDC Centre sponsoring a contractor's report, or a tasking agency, is entered in Section 8.) DRDC – Ottawa Research Centre 3701 Carling Avenue, Ottawa ON K1A 0Z4, Canada	2a. SECURITY MARKING (Overall security marking of the document, including supplemental markings if applicable.) CAN UNCLASSIFIED
	2b. CONTROLLED GOODS NON-CONTROLLED GOODS DMC A

3. TITLE (The document title and sub-title as indicated on the title page.)

Use of anisotropic diffusion filters in Synthetic Aperture Radar (SAR) imagery

4. AUTHORS (Last name, followed by initials – ranks, titles, etc. not to be used. Use semi-colon as delimiter)

Barrie, G.

5. DATE OF PUBLICATION (Month and year of publication of document.) June 2018	6a. NO. OF PAGES (Total pages, including Annexes, excluding DCD, covering and verso pages.) 37	6b. NO. OF REFS (Total cited in document.) 39

7. DOCUMENT CATEGORY (e.g., Scientific Report, Contract Report, Scientific Letter)

Scientific Report

8. SPONSORING CENTRE (The name and address of the department project or laboratory sponsoring the research and development.)

DRDC – Ottawa Research Centre
3701 Carling Avenue, Ottawa ON K1A 0Z4, Canada

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.) JFD 5-1	9b. CONTRACT NO. (If appropriate, the applicable contract number under which the document was written.)

10a. DRDC DOCUMENT NUMBER DRDC-RDDC-2018-R171	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11a. FUTURE DISTRIBUTION WITHIN CANADA (Approval for further dissemination of the document. Security classification must also be considered.)

Public release

11b. FUTURE DISTRIBUTION OUTSIDE CANADA (Approval for further dissemination of the document. Security classification must also be considered.)

Public release

12. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Use semi-colon as a delimiter.)

Synthetic Aperture Radar (SAR); radar data exploitation; anisotropic diffusion; speckle filtering; shoreline extraction; edge detection

13. ABSTRACT/RÉSUMÉ (When available in the document, the French version of the abstract must be included here.)

The Defence R&D Canada (DRDC) Ottawa Research Center, ISR Data Exploitation Group, specializes in development of algorithms and tools enabling uptake of Synthetic Aperture Radar (SAR) data for Canadian Armed Forces (CAF) planning and operational purposes. This scientific report addresses the extraction of coastlines in SAR images exhibiting poor contrast between land and water backscatter. In place of traditional speckle filtering, we apply a nonlinear filter whose design principles are rooted in the physics of diffusion, where image blurring is treated in a manner analogous to diffusing particles. After a brief discussion of how the diffusion equation is applied to image processing, we introduce the Speckle Reducing Anisotropic Diffusion (SRAD) filter. SRAD is shown to be related to more commonly used speckle filters, which constitute a special case of isotropic diffusion. The SRAD filter is first used to process an artificial coastline consisting of two adjacent blocks of RADARSAT-2 (RS-2) data of water and land backscatter forming an idealized linear coastline. This construct allows us to establish easily quantifiable ground truth, and perform statistical analysis on both individual and combined land and water pixel distributions. Following this, SRAD is applied to large ($\sim$ 2 GB) Wide Ultra-Fine RS-2 data sets to examine its use in poor ($\lesssim$ 6 dB) contrast imagery. Results indicate this is an effective filter, and when used in conjunction with a traditional edge detector such as the Canny algorithm, can contribute to the delineation of edges in SAR images that would otherwise be difficult to process.

ABSTRACT/RÉSUMÉ (When available in the document, the French version of the abstract must be included here.)

Le Groupe d'exploitation des données RSR du Centre de recherche de R&D pour la défense Canada – Ottawa se spécialise dans la création d'algorithmes et d'outils permettant d'intégrer des données de radar à synthèse d'ouverture (RSO) aux fins de la planification et des opérations des Forces armées canadiennes. Dans le présent rapport scientifique, nous abordons l'extraction du tracé des côtes dans les images RSO dont le contraste entre la rétrodiffusion de la terre et celle de l'eau est faible. Plutôt que d'utiliser le filtrage classique du chatoiement, nous appliquons un filtre non linéaire dont les principes de conception proviennent de la physique de la diffusion et qui traite le flou de l'image comme si des particules avaient causé la diffusion. Après une brève discussion sur la manière dont nous appliquons l'équation de diffusion au traitement de l'image, nous présentons le "filtre de diffusion anisotrope réduisant le chatoiement" ou filtre SRAD (pour Speckle Reducing Anisotropic Diffusion). Nous montrons que le filtre SRAD s'apparente aux filtres à chatoiement régulièrement utilisés qui constituent un cas spécial de la diffusion isotrope. Nous l'avons d'abord utilisé pour traiter une côte artificielle, composée de deux blocs adjacents de données de rétrodiffusion par l'eau et la terre captées par RADARSAT 2 qui forment un littoral linéaire idéalisé. Cette construction permet de créer une vérification au sol facilement quantifiable et de réaliser une analyse statistique sur les distributions de pixels de terre et de pixels d'eau, individuellement et collectivement. Après cette opération, nous avons appliqué le filtre SRAD à de grands ensembles ($\sim$ 2 GB) de données ultrafines dans un champ large, captées par RADARSAT-2 pour étudier son utilisation sur des images à faibles contrastes (0,6 dB). Les résultats indiquent que le filtre est efficace et que s'il est utilisé en conjonction avec un détecteur de bords classique, comme l'algorithme de Canny, peut contribuer à définir les arêtes dans des images de radar à synthèse d'ouverture qui seraient autrement difficiles à traiter.