

General Adaptive Monte Carlo Bayesian Image Denoising

by

Wen Zhang

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2010

© Wen Zhang 2010

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Image noise reduction, or denoising, is an active area of research, although many of the techniques cited in the literature mainly target additive white noise. With an emphasis on signal-dependent noise, this thesis presents the **General Adaptive Monte Carlo Bayesian Image Denoising** (GAMBID) algorithm, a model-free approach based on random sampling. Testing is conducted on synthetic images with two different signal-dependent noise types as well as on real synthetic aperture radar and ultrasound images. Results show that GAMBID can achieve state-of-the-art performance, but suffers from some limitations in dealing with textures and fine low-contrast features. These aspects can be addressed in future iterations when GAMBID is expanded to become a versatile denoising framework.

Acknowledgements

This thesis would not have been possible without the assistance from a number of very talented and helpful individuals. I am grateful to the members of the Vision and Image Processing group for their support, especially Alexander Wong for his input on the mathematical theory and Peter Yu for providing synthetic test images.

I would like to thank my supervisors Professors John Zelek and David Clausi for their mentorship and their patience, and for providing an environment where I can pursue my research interests regardless of where they may lead.

Contents

List of Tables	vii
List of Figures	ix
1 Introduction	1
2 Background	4
2.1 Image Denoising Problem	4
2.2 Data Models	4
2.2.1 Electro-optical	6
2.2.2 Synthetic Aperture Radar (SAR)	6
2.2.3 Ultrasound	7
2.2.4 Magnetic resonance imaging (MRI)	7
2.3 State-of-the-Art Denoising Methods	7
2.3.1 Additive noise reduction	8
2.3.2 Signal-dependent noise reduction	8
2.4 Rationale	10
2.5 Objectives	11

3	General Adaptive Monte Carlo Bayesian Image Denoising	12
3.1	Overview	12
3.2	Bayesian Framework	14
3.3	Posterior Estimation via Sampling	14
3.4	Sample Weighting	15
3.5	Implementation	17
4	Experimental Results	21
4.1	Testing Methodology	21
4.1.1	Parameter selection	22
4.1.2	Synthetic images	23
4.1.3	Real images	26
4.2	Results	27
4.2.1	Synthetic images	27
4.2.2	Synthetic Aperture Radar Images	36
4.2.3	Ultrasound images	39
4.3	Discussion	39
5	Conclusions and Future Work	48
5.1	Conclusions	48
5.2	Recommendations for Future Work	49
	APPENDICES	50
	A Mathematical Derivations	51
	References	51

List of Tables

4.1	Tested denoising algorithms	22
4.2	Algorithm parameters	23
4.3	Tested noise levels	24
4.4	Image quality metrics for “Texture” image with Type I noise	27
4.5	Image quality metrics for “Structure” image with Type I noise . . .	29
4.6	Image quality metrics for “Texture” image with Type II noise . . .	32
4.7	Image quality metrics for “Structure” image with Type II noise . .	36

List of Figures

1.1	Examples of speckle	2
2.1	Problem formulation and notation	5
3.1	Overview of the GAMBID algorithm	13
3.2	Example posterior distribution estimates	18
3.3	Weighting as function of pixel location	19
4.1	Synthetic reference images	25
4.2	Image quality metrics for “Texture” image with Type I noise	28
4.3	“Texture” image with Type I noise $L = 2$ (N4)	30
4.4	Image quality metrics for “Structure” image with Type I noise	31
4.5	“Structure” image with Type I noise $L = 2$ (N4)	33
4.6	Image quality metrics for “Texture” image with Type II noise	34
4.7	“Texture” image with Type II noise with $\sigma_{add} = 0.2$ (N8)	35
4.8	Image quality metrics for “Structure” image with Type II noise	37
4.9	“Structure” image with Type II noise with $\sigma_{add} = 0.2$ (N8)	38
4.10	SAR1 processed with GAMBID	40
4.11	SAR2 processed with GAMBID	41
4.12	Detail of SAR1	42

4.13	Detail of SAR2	43
4.14	“Embryo” ultrasound image	44
4.15	“Prostate ultrasound images	45

Chapter 1


Introduction

Images are a natural way for humans to think about spatial information, and digital images are a natural representation of spatial data. Like all recorded signals, digital images are often corrupted by noise, increasing the difficulty with which human observers or computer algorithms are able to extract the useful underlying information. Although noise can be mitigated by improved image acquisition hardware, in some modalities, such as coherent imaging, the noise is an inherent part of the imaging process.

However, information content may be preserved even at high noise levels, as can be seen in Figure 1.1, which shows images corrupted by correlated speckle noise. With some effort, one can still discern the structural details in the presence of noise. Thus, the goal of noise filtering, or image *denoising*, is to exploit the available information in the observed image to obtain an estimate of the noise-free signal.


There are two main purposes for obtaining this estimate. First, the noise filtering can be performed as a pre-processing step for further machine analysis, such as scene segmentation, object detection, or visual tracking. Secondly, denoised images are easier to interpret by human observers, aiding in tasks such as classifying ice types in SAR images, or assessing arterial disease in ultrasound images.

The topic of image denoising is an active area of research in image processing



(a)

(b)



(c)

Figure 1.1: Examples of speckle. (a) Synthetic aperture radar image of sea ice (source: Canadian Ice Service). (b) Ultrasound image of fetus (source: DHD multimedia gallery [1]). (c) Confocal laser scan of human retina (source: Bueno et al. [2]).

and computer vision. Many techniques have been proposed, ranging from spatial filters [3, 4, 5, 6], frequency-domain filters [7, 8], to multiscale wavelet filters [9, 10]. Many of the popular methods work under the assumption of signal-independent additive noise. However, some imaging modalities give rise to noise models that are signal-dependent. This is especially true in coherent imaging systems, where the images are characterized by an optical phenomenon known as speckle [11]. Figure 1.1 shows examples of speckle images, produced by synthetic aperture radar [12], ultrasound [13], and confocal laser scanning ophthalmoscopy [2].

We propose a versatile statistical algorithm for image denoising technique called General Adaptive Monte Carlo Bayesian Image Denoising (GAMBID). The idea behind GAMBID is to produce a Bayesian least squares estimate of the signal, assuming that the posterior probability distribution of a pixel is conditioned around the observed values in its spatial neighbourhood. This conditional posterior density is approximated by Monte Carlo sampling of the other image pixels and assigning an appropriate weight to each sample. The advantage of this model-free approach is its flexibility to various noise models.

This thesis is structured as follows. First, the background and literature review are given in Chapter 2, where the image denoising problem is described, noise models for imaging systems are introduced, and state-of-the-art denoising methods are described with an emphasis on signal-dependent noise. Next, Chapter 3 describes a Markov model of the noisy image and the development of the GAMBID algorithm. Test results on synthetic noisy images and real images are described in Chapter 4. Finally, conclusions and recommendations for future work are given in Chapter 5.

Chapter 2

Background

2.1 Image Denoising Problem

The image denoising problem is formulated as follows. Let \mathcal{L} be a discrete rectangular lattice and S be the set of indices into the lattice. Individual pixel values are indexed by $\mathbf{x} = [i \ j]^T \in S$. Let the ideal image be r and the observed intensity be the random field I , both defined on the lattice \mathcal{L} . The goal of image denoising is to obtain an estimate \hat{r} of the ideal image r from I . This formulation is illustrated in Figure 2.1 and will be used in the remainder of thesis. The distribution of I depends on the specifics of the imaging system, and some models are described in Section 2.2. Current denoising methods, that is, methods to obtain an estimate \hat{r} , are described in Section 2.3.

2.2 Data Models

The characteristics of the image noise is a function of the imaging modality. Based on the underlying image formation process, each modality is described by a different data model that includes statistical properties of the noise.

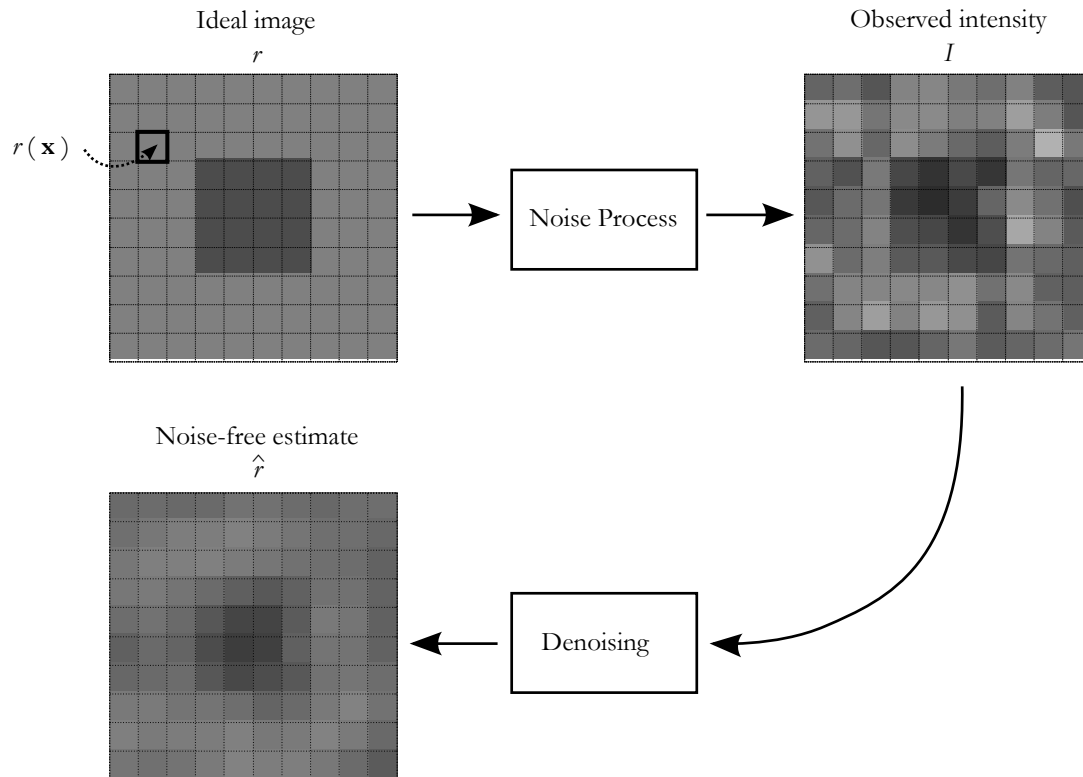


Figure 2.1: Problem formulation and notation. The goal is to obtain the denoised estimate \hat{r} from the observed image I , whose distribution is conditioned on the noise-free image r . All the images are defined on a lattice indexed by $\mathbf{x} = [i \ j]^T$.

2.2.1 Electro-optical

The most common type of electro-optical imaging technology is the charged-couple device (CCD) [14]. These sensors are governed by three principal noise sources:

1. shot noise: caused by a finite number of photons hitting the sensor,
2. dark/obscurity noise: spurious photons generated by the circuit, and
3. read-out noise: electronic noise.

These components are additive, and gives the model for image formation in CCDs [3]

$$I = r + c\sqrt{r}n_0 + n_1 + n_2, \quad (2.1)$$

where n_0 , n_1 , and n_2 are signal-independent noise processes and c is a constant. The overall noise is signal-dependent due to the $c\sqrt{r}n_0$ term. An alternative model for infrared clutter has been proposed based on both scene and noise statistics

$$I = r \cdot [n_m * h] + n_a \quad (2.2)$$

where the multiplicative noise term n_m is gamma-distributed and n_a is a zero-mean Gaussian process [15].

2.2.2 Synthetic Aperture Radar (SAR)

SAR images are formed from the backscattered signal from an electromagnetic wave. These images are characterized by a noise-like phenomenon known as speckle, which arises from the summation of random phases due to presence of multiple scatters in each resolution cell [12]. Although speckle is a real electromagnetic quality that is not technically noise, the observed intensity (power) of a SAR image can be modelled by a multiplicative noise model

$$I = [r \cdot n] * h \quad (2.3)$$

where r is a measure of the average reflectivity of the resolution cell, n is a gamma-distributed random process, and h is the point-spread function (PSF) of the SAR system.

2.2.3 Ultrasound

Sonograms are generated by measuring the backscattered echo from an ultrasound signal. Like SAR, ultrasound images exhibit a speckle pattern. The statistical model is actually identical to single-look SAR amplitude signals. Although some sources use a product model (Equation 2.3) [13], this is not strictly accurate, as the signal amplitude for fully-developed speckle follows a Rayleigh distribution given by the conditional density [16]

$$p(I|r) = \frac{I}{r} \exp\left(-\frac{I^2}{2r^2}\right), \quad (2.4)$$

with the noise variance is equal to $\frac{4-\pi}{2}r^2$.

2.2.4 Magnetic resonance imaging (MRI)

The noise characteristics of the MR signal are Gaussian. The complex MR image is generated by the Fourier transform of the data preserving the Gaussian characteristics. To avoid phase artifacts, the magnitude image is used and the phase is discarded. It can be shown that Gaussian white noise in the real and imaginary components give rise to a Rice distribution in the magnitude image I [17]

$$p(I|r) = \frac{I}{\sigma^2} \exp\left[-\frac{I+r^2}{2\sigma^2}\right] J_0\left[\frac{r \cdot I}{\sigma^2}\right] \quad (2.5)$$

where J_0 is the Bessel function of the first kind and σ is the standard deviation of the Gaussian noise in the real and imaginary components.

2.3 State-of-the-Art Denoising Methods

Image denoising is one of the fundamental challenges in image processing and computer vision, since many of the image noise sources, both intrinsic and extrinsic, cannot be avoided. Thus, many denoising algorithms have been developed to recover the noise-free image from a noisy input. The methods described in this

section can be categorized into two main types. The first type of denoising filters are obtained under the assumption that noise is additive and signal-independent (typically Gaussian). An overview of the methods in this category is given in the next section. The second type assume signal-dependent noise. The applications for the second type are of greater relevance to this work and will be described in more detail in the remainder of this section.

2.3.1 Additive noise reduction

Additive noise suppression methods can be further broken down into two classes. The first class, transform domain methods, perform noise reduction based on the transform domain coefficients. Methods in this class include Wiener filtering [8], collaborative Wiener filtering [7], Gaussian scale mixture denoising [10], and wavelet shrinkage [18]. The second class, spatial domain methods, utilize spatial information redundancy. Methods in this class include Gaussian filtering [19], anisotropic filtering [6], bilateral filtering [4], and nonlocal means [3].

2.3.2 Signal-dependent noise reduction

Adaptive local statistics filter

A widely-used class of speckle reduction filters perform estimation based on the local statistics within a sliding window. These include local linear minimum mean squared error (LLMMSE) and Bayesian filters.

The LLMMSE approach was employed in the filters proposed by Lee [20, 21], Frost [22] et al., and Kuan et al. [5]. The estimate is obtained from the first-order Taylor expansion of the local mean \bar{I} , that is

$$\hat{r}_{LLMMSE} = \bar{I} + k(I - \bar{I}), \quad (2.6)$$

where k is chosen based on the local statistics (mean and variance) to minimize the mean squared-error between the estimate and the noise-free image.

These LLMMSE filters are derived according to the scene and noise models. The Lee filter assumes that both the scene and the speckle are spatially uncorrelated. The method proposed by Frost incorporates correlation properties of the speckle based on the system PSF. The Kuan filter further incorporates correlation properties of both the scene and the speckle. Since only the local mean and variances are considered, all three filters make an implicit assumption that the image signal is Gaussian-distributed. To address the tradeoff between speckle attenuation and feature preservation, Lopes et al. [23] proposed an enhanced version of the Lee filter that classifies each pixel as belonging to a homogeneous region, textured region, or point target. Although motivated by signal-dependent noise, LLMMSE filters can also be formulated for additive noise or a combination of additive and signal-dependent noise.

The product model from Equation 2.3 can be explicitly used to derive local Bayesian estimators, where non-Gaussian forms for the scene and noise processes can be assumed. These include Γ -MAP [5, 24], suitable for image with gamma-distributed speckle (e.g. SAR), and Rayleigh maximum likelihood [13], suitable for signals image that can be modelled as Rayleigh processes (e.g. ultrasound signal in Equation 2.4).

Median filter

Median filters are another popular method for speckle suppression. However, since median filtering tends to suppress image details as well as speckle, adaptive median filtering methods based on local variation are usually employed [25].

Anisotropic diffusion

Anisotropic diffusion is an image smoothing and denoising approach based on the partial differential equation [26]

$$\begin{cases} \frac{\partial I}{\partial t} = \text{div} [c(\|\nabla I\|) \cdot \nabla I] \\ I(t = 0) = I_0 \end{cases} \quad (2.7)$$

where div denotes the divergence operator, $c(\cdot)$ is the diffusion coefficient, and I_0 is the initial image. The diffusion coefficient is chosen to promote smoothing within regions inhibit smoothing across edges. Multiple iterations of the PDE can be applied until the desired level of smoothing is achieved. Linking this approach with the LLMMSE filters of Lee and Frost, speckle-reducing anisotropic diffusion was developed by Yu and Acton [27].

Multiscale denoising

Since image features and noise are usually prominent in different scales, various wavelet domain despeckling methods have been proposed. Achim et al. modelled the wavelet coefficients of speckled images using heavy-tailed distributions, performed denoising using Bayesian shrinkage [9]. The multiscale method employed by Yue et al. is iterative and noise is reduced by diffusion of the wavelet coefficients [28]. Another related method performs classification based on wavelet coefficients to derive the statistics of the image (GenLik) [29], which has been shown to work on both speckle and Rician noise (with bias correction as a separate step).

2.4 Rationale

A large body of research currently exists concerning the filtering of additive white noise. These methods, for the most part, provide excellent results when the additive assumption holds.

The issue becomes more complex when discussing signal-dependent noise suppression, as the noise characteristics vary greatly depending on the image modality. Firstly, the data models for ultrasound (Rayleigh) and MR (Rice) have signal-dependent means. Consequently, the denoising methods described within this chapter will give a *biased* estimate of the noise-free image, with the exception of Rayleigh maximum likelihood, which explicitly takes the model into account. Secondly, while

some denoising methods employ specific data models (such as the aforementioned Rayleigh maximum likelihood), others depend on the signal being some form of multiplicative noise (e.g, SRAD). Other methods have been shown to be appropriate for a variety of different noise types (e.g. GenLik, Lee). The focus of this work will be on this latter class of signal-dependent noise because it allows greater opportunities for advancements in the field.

2.5 Objectives

This work proposes a versatile denoising method called General Adaptive Monte-Carlo Bayesian Image Denoising (GAMBID). The goal of GAMBID is to provide a denoising method that relaxes some of the modelling assumptions (e.g. Rayleigh or gamma-distributed noise, purely multiplicative or additive noise). Hence, we wish to show that GAMBID can achieve state-of-the-art denoising performance in a variety of signal-dependent noise cases as compared to other current methods. This performance is evaluated objectively on images with artificial noise of different types: (1) pure speckle following the model for SAR data in Equation 2.3; and (2) combined speckle and additive noise following the proposed model for infrared data in Equation 2.1. In addition, the performance will be evaluated subjectively on real SAR and ultrasound images. For the purposes of the the evaluation, the bias in the ultrasound estimates are ignored. In practice, the bias can be corrected as a separate step.

Chapter 3

General Adaptive Monte Carlo Bayesian Image Denoising

3.1 Overview

This thesis proposes a new non-local spatial denoising technique based on Monte Carlo estimation, called GAMBID. Since the method makes only weak assumptions about the underlying signal and noise models, it is well-suited to handle signal-dependent noise. The algorithm has three main steps, performed for each pixel of the image:

1. For the target pixel, sample pixels and their neighbourhoods are drawn at random.
2. Each sample is given an importance weight.
3. A posterior density estimate is generated from all the samples and their respective weights, and a Bayesian least-squares estimate for the target pixel is obtained through weighted averaging.

An overview of the GAMBID algorithm is shown in Figure 3.1. A detailed development of the algorithm is given in the following sections.

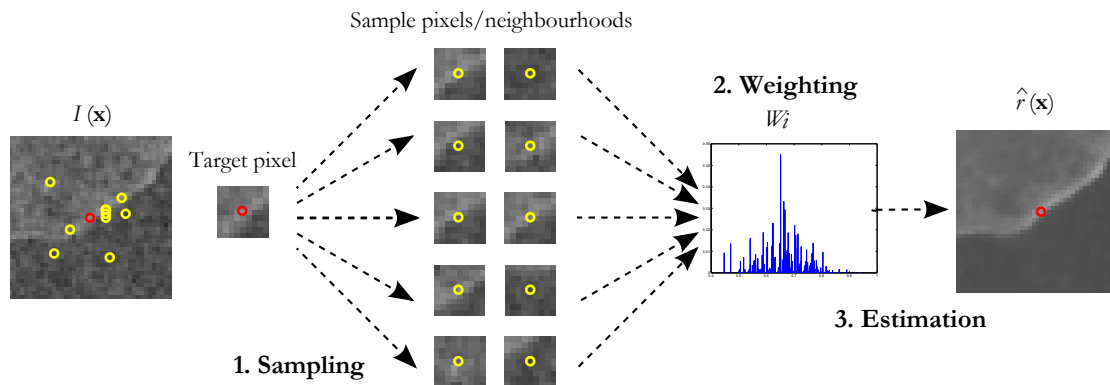


Figure 3.1: Overview of the GAMBID algorithm, where the intensity of the target pixel is to be estimated. 1. Sample pixels and their neighbourhoods are drawn at random. 2. Each sample is given an importance weight. 3. A posterior density estimate is generated from all the samples and their respective weights, and a Bayesian least-squares estimate is obtained through weighted averaging. (Although this figure illustrates the process for one pixel, the full noisy and denoised images are shown.)

3.2 Bayesian Framework

Let \mathbf{x}_c be the target pixel to be estimated and \mathcal{N} be the set of neighbouring indices. For notational simplicity we will use r and \hat{r} refer to the ideal and estimated values at \mathbf{x}_c respectively. Let $\mathbf{I} = \{I(\mathbf{x}), \mathbf{x} \in \mathcal{N}\}$ be the vector of intensity values in a neighbourhood around \mathbf{x}_c , where \mathcal{N} is assumed to be contiguous and symmetric about \mathbf{x}_c .

First, the image is modelled as a hidden Markov process, in which the unobserved state (the noise-free value) at a pixel is conditionally dependent on the observed intensity neighbourhood.

$$p(r) = p(r|\mathbf{I}). \quad (3.1)$$

Next, we formulate the noise filtering as a Bayesian least squares problem:

$$\hat{r} = \arg \min_{\hat{r}} \{E[(\hat{r} - r)^2|\mathbf{I}]\} \quad (3.2)$$

Performing this minimization (details given in Appendix A), we get

$$\hat{r} = \int r p(r|\mathbf{I}) dr = E[r|\mathbf{I}]. \quad (3.3)$$

Thus, the Bayesian least squares estimate is the expected intensity value given its neighbourhood intensities.

3.3 Posterior Estimation via Sampling

One possible approach is to assume the form of $p(r|\mathbf{I})$ (e.g. gamma-distributed neighbourhood model), obtain the maximum likelihood estimate of the parameters, and proceed with the integration [12]. However, we may not wish to assume a distribution given the non-stationary statistics of many images.

This thesis proposes a non-parametric estimator in which the form of the posterior $p(r|\mathbf{I})$ is unknown, and $E[r|\mathbf{I}]$ is estimated using an adaptive Monte Carlo

technique, made possible by exploiting the high level of information redundancy found in most images.

The first central assumption is that the maximum likelihood estimate of a single pixel's true value is the observed intensity at that location, so we let

$$\hat{r}^* = I. \tag{3.4}$$

The assumption holds in cases where $E[I] = r$, such as in SAR and CCD images. Otherwise, we will obtain a biased estimate. In the cases where the assumption breaks down (e.g. MRI), the bias must be corrected using an explicit noise model.

Next, we draw M random indices from the image lattice, ξ_1, \dots, ξ_M , with intensity values $\hat{r}_1^*, \dots, \hat{r}_M^*$ and neighbourhood intensities $\mathbf{I}_1, \dots, \mathbf{I}_M$. These values comprise samples drawn from a (rough) posterior estimate $p^*(r|\mathbf{I})$. The obvious limitation of this density estimate is that the sample neighbourhoods may not correspond to the target neighbourhood. Thus, to refine the density estimate, each sample is assigned a weight W_i proportional to the likelihood that it is drawn from the actual posterior $p(r|\mathbf{I})$. The method for determining the weights is discussed in detail the next section.

The Monte Carlo estimate of the posterior distribution is composed of a weighted Dirac delta function located at each intensity sample

$$\hat{p}_{MC}(r|\mathbf{I}) = \frac{\sum_{i=1}^M W_i \delta(r - \hat{r}_i^*)}{\sum_{i=1}^M W_i}. \tag{3.5}$$

Evaluating Equation 3.3 substituting $\hat{p}_{MC}(r|\mathbf{I})$ for $p(r|\mathbf{I})$, we get

$$\hat{r}_{MC} = \int r \hat{p}_{MC}(r|\mathbf{I}) dr = \frac{\sum_{i=1}^M W_i \hat{r}_i^*}{\sum_{i=1}^M W_i}, \tag{3.6}$$

using the fact that $\delta(r - \hat{r}^*) = 1$ only when $r = \hat{r}^*$.

3.4 Sample Weighting

We wish to weight the sample given by the centre intensity \hat{r}_i^* given the neighbourhood intensities \mathbf{I}_i proportional to the likelihood that it belongs to the same

distribution as the target pixel

$$W_i \propto P(\hat{r}_i^* \sim p(r|\mathbf{I})). \quad (3.7)$$

We use the squared difference between two intensity values $\epsilon = |I_1 - I_2|^2$ as the similarity measure. Although $p_\epsilon(\epsilon)$ depends on the distribution of the intensity values, we make the second major assumption that ϵ follows a negative exponential distribution

$$p_\epsilon(\epsilon) = \frac{1}{\beta} e^{-\beta\epsilon}, \quad \epsilon \geq 0 \quad (3.8)$$

which can be shown empirically to give reasonable results.

Extending upon this, for two neighbourhoods \mathbf{I}_1 and \mathbf{I}_2 each containing N pixels, the likelihood for the set of pairwise intensity differences is given by

$$\begin{aligned} p(\epsilon_1, \epsilon_2, \dots, \epsilon_N) &= \prod_{t=1}^N \frac{1}{\beta} \exp\{-\beta\epsilon_t\} \\ &= \frac{1}{\beta^N} \exp\left\{\sum_{t=1}^N -\beta\epsilon_t\right\} \\ &= \frac{1}{\beta^N} \exp\left\{-\beta \sum_{t=1}^N |\mathbf{I}_1(t) - \mathbf{I}_2(t)|^2\right\}. \end{aligned} \quad (3.9)$$

Setting the weight of each sample W_i to be equal to the above likelihood for the target and sample neighbourhoods, it becomes a Gibbs measure [30]

$$W_i = \frac{1}{Z} \exp\{-\beta\Phi(\mathbf{I}, \mathbf{I}_i)\}$$

or

$$W_i = \frac{1}{Z} \exp\left\{-\frac{\Phi(\mathbf{I}, \mathbf{I}_i)}{\eta^2}\right\}, \quad (3.10)$$

where η is a smoothing factor and Z is a normalizing factor. As was shown in the posterior estimation step in Equation 3.5, the value of Z does not affect the outcome because it is cancelled out. The “energy” function

$$\Phi(\mathbf{I}, \mathbf{I}_i) = \sum_{t=1}^N |\mathbf{I}(t) - \mathbf{I}_i(t)|^2 \quad (3.11)$$

is a sum of squared differences (SSD) of pixel intensities between the target neighbourhood and the sample neighbourhood.

Figure 3.2 gives an example of the posterior estimates using this weighting scheme. The synthetic image in the centre has intensities modelled as, clockwise from top-left, gamma, Gaussian, exponential, and uniform random processes. The yellow squares denote 100 randomly sampled pixels, and the red squares denote the locations for which the estimates were obtained. The four plots show the estimated posterior density $\hat{p}_{MC}(r|\mathbf{I})$ (vertical bars) overlaid with the ideal pdf at each of these locations. While the estimates are not perfect, they manage to capture the general characteristics of the ideal pdfs, such as the mean and overall shape.

Figure 3.3 demonstrates how the weighting scheme favours pixels with similar neighbourhoods, thereby allowing for nonlocal sampling from larger spaces while preserving the features. The yellow squares denote the randomly sampled pixels, and the red squares denote the locations at which the weights were calculated. The blue bars in the bottom two portions show the weights W_i at each sample pixel location. In the left case, the sample pixel resides in a homogeneous area, so all other samples in that area are given relatively high weights. In the right case, the sample pixel is on an edge, so other pixel samples on the same edge are given the highest weights.

3.5 Implementation

The GAMBID algorithm is relatively straightforward to implement directly from Equations 3.6 and 3.10. The smoothing parameter η is chosen based on the desired amount of smoothing, and should be proportional to the noise level. In practice, the noise standard deviation (for constant SNR) or the coefficient of variation (for multiplicative noise) are found to be effective choices for η . The neighbourhoods \mathcal{N} are 9×9 centred at the target pixel. The samples are generated from a uniform distribution in a 21×21 pixel sample space around each pixel, with a maximum of

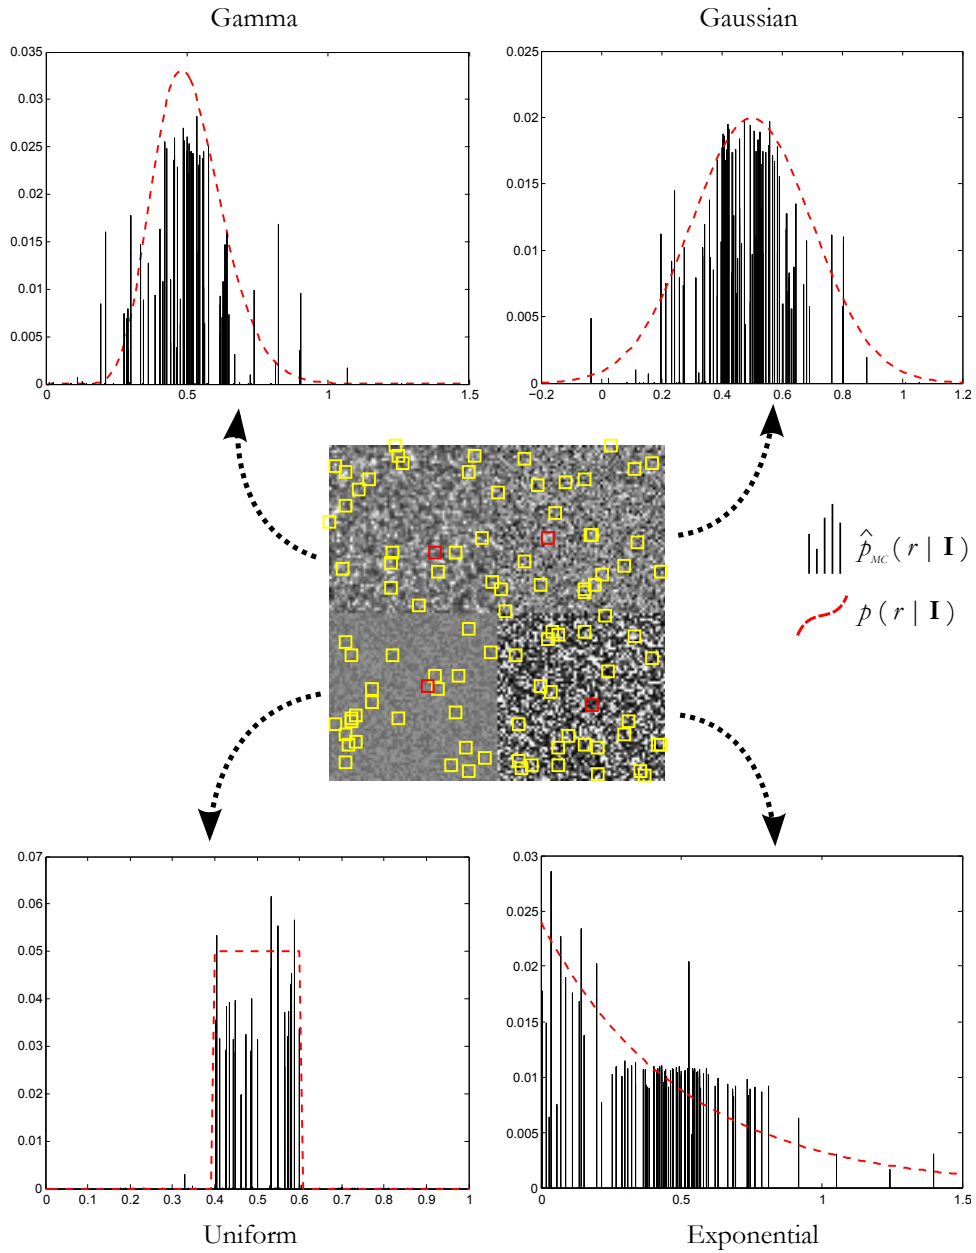


Figure 3.2: Example of posterior distribution estimates. For each of the four intensity regions, the set of sample pixels indicated by the yellow markers are used to generate impulse-based estimated posterior density functions. The dotted red lines plot the ideal continuous distribution functions.

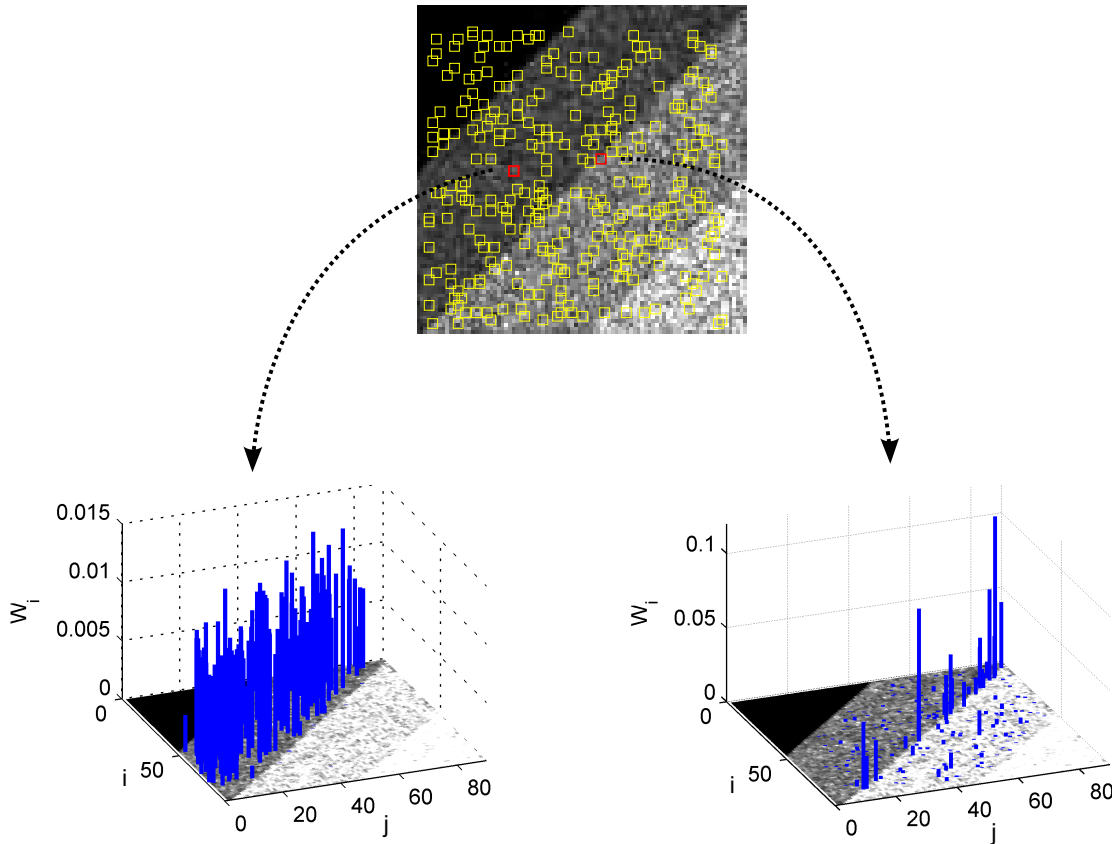


Figure 3.3: Sample weights as function of pixel location. Top: Synthetic image with sampled pixels in yellow. Bottom row: sample weights W_i . In the left case, the sample pixel (red) resides in a homogeneous area, so all other samples in that area are given relatively high weights. In the right case, the sample pixel is on an edge, so other pixel samples on the same edge are given the highest weights.

$M \leq 200$ unique samples.

With a square image of size $p \times p$, and using square $n \times n$ neighbourhoods and a maximum of M samples per pixel, the complexity of GAMBID is $O(p^2 n^2 M)$. This is significantly more complex than the popular LLMMSE filters, which are $O(p^2 n)$.

Chapter 4

Experimental Results

4.1 Testing Methodology

The proposed GAMBID is demonstrated to be a versatile denoising method that can achieve state-of-the-art performance in comparison to existing methods under different signal-dependent noise conditions. Firstly, the performance is evaluated objectively using synthetic images. Two types of simulated noise are tested:

- Type I: Multiplicative gamma-distributed noise, which simulates cases where the image is dominated by speckle-like noise sources (e.g. SAR, ultrasound, and other coherent imaging).
- Type II: Combined multiplicative and additive noise, which simulates cases where multiple different noise sources are present (e.g. electro-optical measurements).

Secondly, real SAR and ultrasound images are evaluated. The difficulty lies in a lack of ground truth, so objective measures are impossible. Thus, these images are evaluated in a subjective visual manner based on the level of noise reduction and preservation of perceptual details.

Table 4.1: Tested denoising algorithms

Name	Type	Data model
Lee filter [20, 24]	Local statistics	Multiplicative or combined additive and multiplicative noise (Gaussian white)
Γ -MAP [5]	Local statistics	Gamma distributed noise
SRAD [27]	Partial differential equation	Multiplicative noise
BM3D [7]	Transform domain filtering	Additive white noise
GenLik [29]	Wavelet domain filtering	General
GAMBID	Stochastic Bayesian least squares	General

The denoising algorithms tested comprise a cross-section of current methods, listed in Table 4.1. The state-of-the-art additive noise suppression algorithm is known to be collaborative Wiener filtering (BM3D) [7]. Although BM3D is not theoretically justified for signal-dependent noise filtering, it has been applied successfully in practical applications to speckled images [31]. From the LLMMSE filters, the enhanced Lee filter [23] and the Lee filter [20] for combined additive and multiplicative noise were chosen. Other tested methods for reducing signal-dependent noise are Γ -MAP [5] and anisotropic diffusion (SRAD) [27]. Lastly, the wavelet domain general likelihood method (GenLik) [29] was tested.

4.1.1 Parameter selection

To mimic operational conditions where the noise properties of the image are unknown, some filter parameters are estimated from the noisy image. The Lee, Gamma-MAP, and SRAD filters require as a parameter the number of looks $L = \bar{I}^2/\sigma_I^2$, where \bar{I} is the local mean and σ_I^2 is the local variance. These values are estimated by manually selecting a relatively homogeneous region in the image. The BM3D method requires the noise standard deviation, so the local standard deviation in the aforementioned region is used. GAMBID’s smoothing parameter η is

Table 4.2: Algorithm parameters

	Parameters	Selection
Lee filter	window size	5×5
	number of looks	\bar{I}^2/σ_I^2 from homogeneous region
Γ -MAP	window size	5×5
	number of looks	\bar{I}^2/σ_I^2 from homogeneous region
SRAD	number of iterations	variable (until convergence)
	timestep	0.5
	number of looks	\bar{I}^2/σ_I^2 from homogeneous region
BM3D	noise standard deviation	σ_I from homogeneous region
GenLik	window size	5×5
GAMBID	smoothing factor β	CV^2 in homogeneous region

set to the squared coefficient of variation $CV^2 = 1/L = \sigma_I^2/\bar{I}^2$. These and any additional parameters are listed in Table 4.2.

4.1.2 Synthetic images

Test images

Synthetic noisy images are generated according to the noise model containing correlated multiplicative noise and white Gaussian additive noise

$$I = [r \cdot n_{mult}] * h + n_{add}, \quad (4.1)$$

where r is a noise-free reference image, and n_{mult} is an uncorrelated gamma process

$$p(n_{mult}) = \frac{L^L n_{mult}^{L-1}}{(L-1)!} \exp[-Ln_{mult}]. \quad (4.2)$$

The convolution kernel h used is a 2D Gaussian function simulating the system point-spread function (PSF). The level of multiplicative noise is controlled by the number of looks L , where a higher number of looks produces lower noise variance.

Table 4.3: Tested noise levels

		Looks L	Additive noise σ_{add}
Type I	N1	16	0
	N2	8	0
	N3	4	0
	N4	2	0
Type II	N5	4	0.01
	N6	4	0.05
	N7	4	0.1
	N8	4	0.2

n_{add} is a white zero-mean Gaussian process with standard deviation σ_{add} . Using this model, two types of noisy images are generated. Type I contains multiplicative noise only, simulating images dominated by speckle (where $\sigma_{add} = 0$) for $L = 16, 8, 4, 2$. Type II contains a combination of multiplicative and additive noise where $L = 4$ and $\sigma_{add} = 0.01, 0.05, 0.1, 0.2$. These are listed in Table 4.3.

The reference images are chosen to test two scenarios present in remote sensing, shown in Figure 4.1. The first image, “Texture”, is a mosaic of four textures previously used as the segmentation benchmark T4 by Qin and Clausi [32]. The textures are originals or composites of images from the Brodatz set [33]. This image is generated using, clockwise from top-left, D5, $0.5 \cdot D8 + 0.5 \cdot D84$, D84 and D92¹. The second image, “Structure”, consists of homogenous areas of constant reflectivity, where the features of interest are structural details in the form of boundaries (edges) and thin elements. The dimensions of both reference images are 512×512 pixels.

The reference images are grey level images in the range $[0, 255]$. However, the long tail of the gamma-distributed noise will result in many values exceeding 255. The noisy test images will have increased dynamic range compared to the original

¹The D* refers to the numbering system used in the Brodatz album.




Figure 4.1: Reference images (512×512) used in synthetic noise tests. Left: “Texture”. Right: “Structure”.

and are processed without clipping. To make visual comparison to the reference images easier and to maintain the contrast, the grey value mapping in the printed images is in the original $[0, 255]$, where saturating values are clipped.

Evaluation metrics

The synthetic images provide ground truth with which we can evaluate denoising performance under two main criteria:

1. Fidelity to the original signal as measured by the peak signal-to-noise ratio (PSNR)

$$PSNR = 10 \log_{10} \left(\frac{MAX_I^2}{MSE} \right), \quad (4.3)$$

where

$$MSE = \frac{1}{|S|} \sum_{\mathbf{x} \in S} [r(\mathbf{x}) - \hat{r}(\mathbf{x})]^2, \quad (4.4)$$

$MAX_I = 255$ for 8-bit greyscale images, and $|S|$ is the total number of pixels in the image.

2. Visual quality as measured by the structural similarity measure (SSIM) [34], which is sensitive to the distortions that are significant to human perception.

4.1.3 Real images

Two sets of real speckled images are used. The first consists of RADARSAT-2 synthetic aperture radar images of arctic sea ice. These two images, denoted by SAR1 and SAR2, measure 1209×865 and 599×635 respectively and are provided by the Canadian Ice Service. These images are assumed to contain some form of speckle noise like all SAR images, although the actual noise model is unknown, and the clutter is most likely a combination of fluctuations in the reflectivity (an underlying texture) as well as noise [12]. The second set consists of two publicly-available medical ultrasound images obtained from the DHD multimedia gallery [1], denoted by “Embryo” and “Prostate”, and measure 394×454 and 430×288 respectively. The ultrasound images contain highly-correlated noise due to the log-polar transform used to rectify them for display.

Since there are no noise-free reference images, the metrics listed in Section 4.1.2 cannot be used. Thus, these are evaluated on a subjective visual basis. The subjective visual quality is based on the amount of noise attenuation and preservation of structural information. Since the SAR images follow a multiplicative noise model, we can also examine the noise field estimate $\hat{n}_{mult} = I/\hat{r}$, which is ideally independent of the underlying reflectivity, meaning that it should appear to be a stationary process with no structural information. To quantify the tradeoff between structure preservation and smoothing, we also calculate a variant of the contrast-to-noise ratio (CNR) [35]

$$CNR = \frac{\sigma_{edge}}{\sigma_{noise}}, \quad (4.5)$$

where σ_{edge} is the local standard deviation averaged over 10 manually-selected edge pixels and σ_{noise} is standard deviation at a manually-selected homogeneous region.

Table 4.4: Image quality metrics for “Texture” image with Type I noise (see Figure 4.2)

Noise level	L	PSNR					
		Lee	Γ -MAP	SRAD	BM3D	GenLik	GAMBID
N1	16	18.48	21.67	18.80	22.24	19.02	22.33
N2	8	18.21	20.17	18.28	21.28	18.27	21.10
N3	4	17.60	18.60	17.95	19.90	17.59	19.66
N4	2	17.29	16.63	17.51	18.73	17.01	17.91
Noise level	L	SSIM					
		Lee	Γ -MAP	SRAD	BM3D	GenLik	GAMBID
N1	16	0.68	0.86	0.71	0.92	0.71	0.90
N2	8	0.66	0.83	0.67	0.89	0.65	0.85
N3	4	0.63	0.75	0.66	0.82	0.59	0.78
N4	2	0.61	0.72	0.62	0.77	0.54	0.70

4.2 Results

4.2.1 Synthetic images

Type I: Pure multiplicative noise

The metrics for the “Texture” image tests under Type I noise are given in Figure 4.2 and Table 4.4. The BM3D algorithm, despite being designed for additive white noise, performs very well even on an image contaminated by multiplicative noise. This is a trend that can be seen in the rest of the results shown in this chapter. Save BM3D, the metrics show that GAMBID outperforms all the other methods, except in the high-noise $L = 2$ case, where the SSIM is comparable to Γ -MAP.

The test images and the denoised results for “Texture” under the high noise case N4 ($L = 2$) are shown in Figure 4.3. Despite the high PSNR and SSIM measures, GAMBID appears to over-flatten the texture details. This is especially

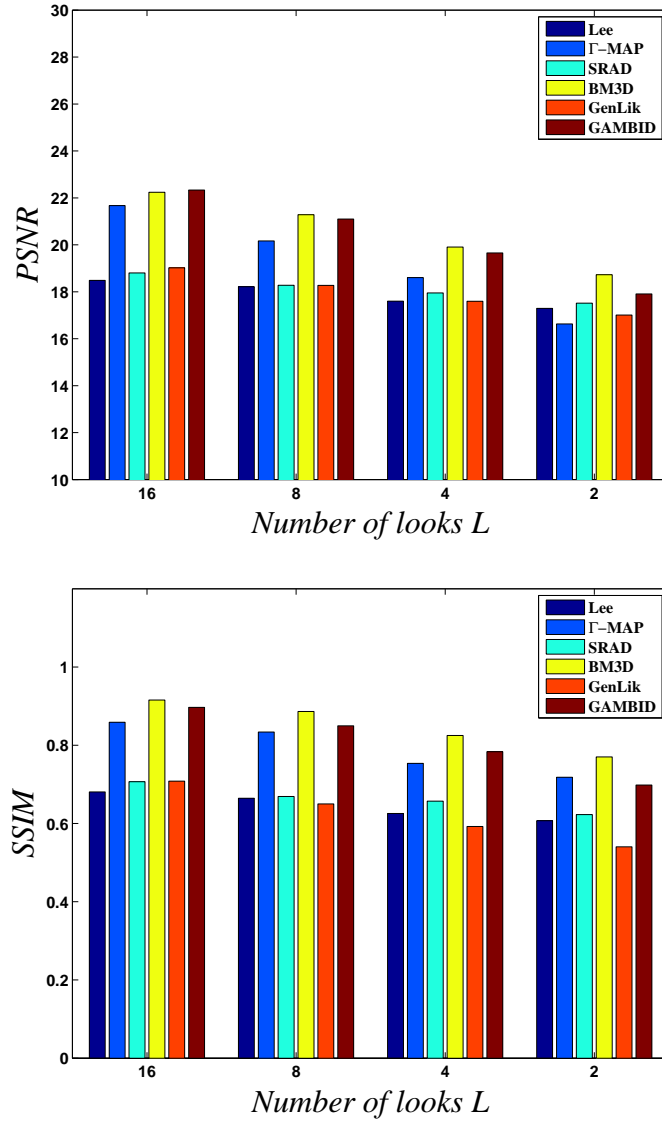


Figure 4.2: Image quality metrics for “Texture” image with Type I noise; top: PSNR, bottom: SSIM. BM3D, despite being designed for additive white noise, performs very well even on image contaminated by multiplicative noise. This is a trend that can be seen in the rest of the results shown in this chapter. Save BM3D, the metrics show that GAMBID outperforms all the other methods, except in the high-noise $L = 2$ case, where the SSIM is comparable to Γ -MAP.


Table 4.5: Image quality metrics for “Structure” image with Type I noise (see Figure 4.4)

Noise level	L	PSNR					
		Lee	Γ -MAP	SRAD	BM3D	GenLik	GAMBID
N1	16	25.63	26.61	29.32	28.82	26.59	28.69
N2	8	24.65	24.48	26.73	26.22	24.71	26.94
N3	4	23.63	21.77	24.58	23.19	23.07	24.74
N4	2	22.18	19.61	22.45	22.04	21.68	21.46
Noise level	L	SSIM					
		Lee	Γ -MAP	SRAD	BM3D	GenLik	GAMBID
N1	16	0.86	0.89	0.87	0.89	0.90	0.94
N2	8	0.81	0.83	0.82	0.84	0.84	0.90
N3	4	0.74	0.73	0.76	0.75	0.76	0.84
N4	2	0.66	0.63	0.67	0.72	0.68	0.75

apparent in the top-right section. Although BM3D performs quite well visually, it does introduce subtle distortion artifacts which are most noticeable in the top-left region. The Γ -MAP result has strong noise in the edge regions which is not heavily penalized in the SSIM measure, but does decrease the PSNR.


The metrics for the “Structure” image under Type I noise are given in Figure 4.4 and Table 4.5. In terms of PSNR, the GAMBID is comparable, but not necessarily superior to the other methods. GAMBID is particularly weak in the $L = 2$ case for PSNR. However, it does outperform all the competing methods in terms of perceptual quality as given by SSIM.

The test images and the denoised results for “Structure” under the high noise case N4 ($L = 2$) are shown in Figure 4.5. Visually, GAMBID applies a high level of smoothing in the homogeneous regions while maintaining edge contrast. However, some of the finer low-contrast features are blurred out. This is obvious in the wavy cross-hatch pattern near the top-right, which some of the competing methods are



(a) Original


(b) Noisy



(c) Lee

(d) Γ -MAP

(e) SRAD



(f) BM3D

(g) GenLik

(h) GAMBID

Figure 4.3: “Texture” image with Type I noise $L = 2$ (N4). Despite the high PSNR and SSIM measures, GAMBID appears to over-flatten the texture details. This is especially apparent in the top-right section. Although BM3D also performs quite well visually, it does introduce subtle distortion artifacts which are most noticeable in the top-left region. The Γ -MAP result has strong noise in the edge regions which is not heavily penalized in the SSIM measure, but does decrease the PSNR.

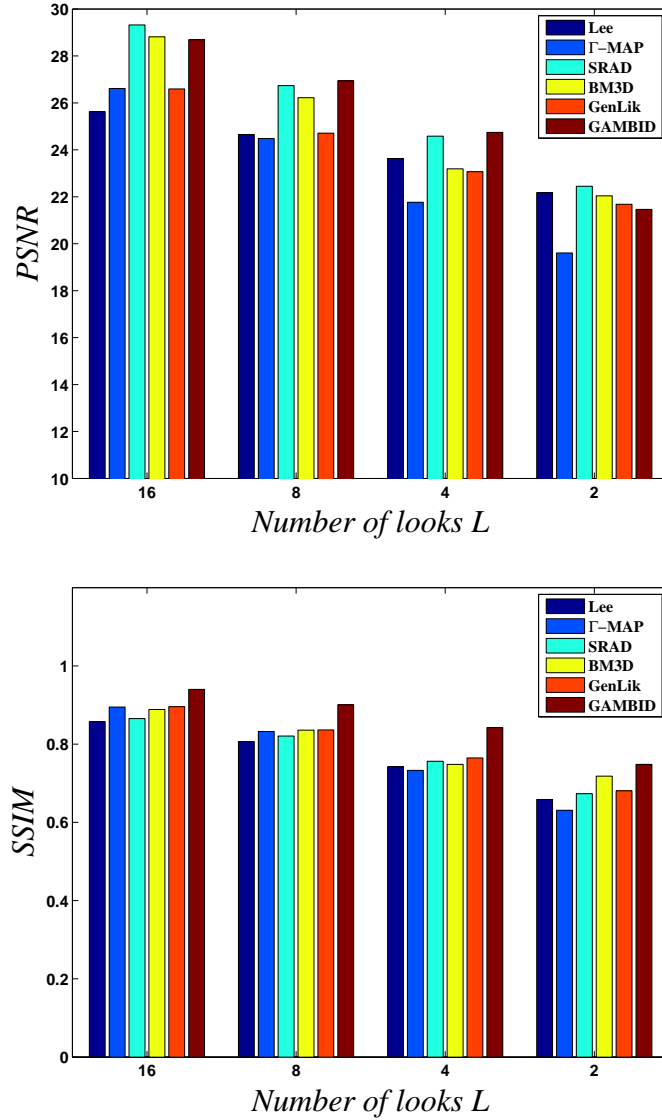


Figure 4.4: Image quality metrics for “Structure” image with Type I noise; top: PSNR, bottom: SSIM. GAMBID is comparable, but not necessarily superior to the other methods. GAMBID is particularly weak in the $L = 2$ case for PSNR. However, it does outperform all the competing methods in terms of perceptual quality as given by SSIM.

Table 4.6: Image quality metrics for “Texture” image with Type II noise (see Figure 4.6)


Noise level	σ_{add}	PSNR					
		Lee	Γ -MAP	SRAD	BM3D	GenLik	GAMBID
N5	0.01	17.23	18.00	17.89	20.23	17.61	18.47
N6	0.05	17.23	18.23	17.92	20.04	17.39	18.69
N7	0.1	17.18	17.33	17.80	19.80	17.05	18.69
N8	0.2	17.05	14.86	17.34	19.16	16.78	17.97
Noise level	σ_{add}	SSIM					
		Lee	Γ -MAP	SRAD	BM3D	GenLik	GAMBID
N5	0.01	0.60	0.80	0.64	0.85	0.59	0.81
N6	0.05	0.60	0.78	0.65	0.84	0.57	0.80
N7	0.1	0.60	0.77	0.64	0.82	0.54	0.77
N8	0.2	0.59	0.72	0.60	0.79	0.51	0.68

better at preserving.

Type II: Combined additive and multiplicative noise


The metrics for the “Texture” image tests for Type II noise are given in Figure 4.6 and Table 4.6. Increasing additive noise level up to 0.1 does not significantly reduce the performance of the tested algorithms. Here, GAMBID is able to outperform all the methods except BM3D in terms of PSNR. BM3D, GAMBID, and Γ -MAP show the strongest SSIM numbers.

The test images and the denoised results for “Texture” under noise level N8 ($\sigma_{add} = 0.2$) are shown in Figure 4.7. Visually, these results are comparable to those with only multiplicative noise applied (seen in Figure 4.3). Again, GAMBID results in a flattening of the texture detail. Although Γ -MAP is comparable to GAMBID in terms of SSIM, the edge region noise residual is quite pronounced (more so here than the Type I case in Figure 4.3).



(a) Original

(b) Noisy



(c) Lee

(d) Γ -MAP

(e) SRAD



(f) BM3D

(g) GenLik

(h) GAMBID

Figure 4.5: “Structure” image with Type I noise $L = 2$ (N4). Visually, GAMBID applies a high level of smoothing in the homogeneous regions while maintaining edge contrast. However, some of the finer low-contrast features are blurred out. This is obvious in the wavy cross-hatch pattern near the top-right, which some of the competing methods are better at preserving.






Figure 4.6: Image quality metrics for “Texture” image with Type II noise, top: PSNR, bottom: SSIM. Increasing additive noise level up to 0.1 does not significantly reduce the performance of the tested algorithms. GAMBID is able to outperform all the methods except BM3D in terms of PSNR. BM3D, GAMBID, and Γ -MAP show the strongest SSIM numbers.



(a) Original


(b) Noisy



(c) Lee

(d) Γ -MAP

(e) SRAD



(f) BM3D

(g) GenLik

(h) GAMBID

Figure 4.7: “Texture” image with Type II with $\sigma_{add} = 0.2$ (N8). Visually, these results are comparable to those with only multiplicative noise applied (Figure 4.3). Again, GAMBID results in a flattening of the texture detail. Although Γ -MAP is comparable to GAMBID in terms of SSIM and slightly better when $\sigma_{add} = 0.2$, the edge region noise residual is more pronounced here than the Type I case in Figure 4.3

Table 4.7: Image quality metrics for “Structure” image with Type II noise (see Figure 4.8)

Noise level	σ_{add}	PSNR					
		Lee	Γ -MAP	SRAD	BM3D	GenLik	GAMBID
N5	0.01	21.33	22.05	24.62	23.93	23.15	24.67
N6	0.05	21.29	22.41	23.91	24.43	22.66	24.57
N7	0.1	21.21	21.02	22.88	24.64	22.15	24.20
N8	0.2	20.85	16.09	20.50	23.66	21.37	22.77
Noise level	σ_{add}	SSIM					
		Lee	Γ -MAP	SRAD	BM3D	GenLik	GAMBID
N5	0.01	0.69	0.74	0.75	0.78	0.77	0.84
N6	0.05	0.67	0.73	0.72	0.79	0.75	0.84
N7	0.1	0.65	0.65	0.69	0.80	0.72	0.83
N8	0.2	0.60	0.47	0.63	0.77	0.66	0.77

The metrics for the “Structure” image under Type II noise are given in Figure 4.8 and Table 4.7. GAMBID produces favourable PSNR values compared to the other methods, with BM3D producing better PSNR for $\sigma_{add} = 0.1, 0.2$. Unlike the “Texture” image results, combined noise in the homogeneous regions poses problems with the competing methods, as can be seen in Figure 4.7.

The test images and the denoised results for “Structure” under noise level N8 ($\sigma_{add} = 0.2$) are shown in Figure 4.9. In this case, Γ -MAP and SRAD do not sufficiently suppress the noise. This is especially obvious in the black region, where the noise is essentially intact. Although GAMBID successfully removes the noise in these regions, it doesn’t preserve the fine structures.

4.2.2 Synthetic Aperture Radar Images

The full images and the GAMBID results, with comparisons showing 128×128 extracted regions, are given in Figures 4.10 to 4.13. The grey levels in the noise






Figure 4.8: Image quality metrics for “Structure image” with Type II noise. GAMBID produces favourable PSNR values compared to the other methods, with BM3D producing better PSNR for $\sigma_{add} = 0.1, 0.2$. Unlike the “Texture” image results, combined noise in the homogeneous regions poses problems with the competing methods, as can be seen in Figure 4.7.



(a) Original


(b) Noisy



(c) Lee

(d) Γ -MAP

(e) SRAD



(f) BM3D

(g) GenLik

(h) GAMBID

Figure 4.9: “Structure” image with Type II noise at noise level N8 (with $\sigma_{add} = 0.2$). Γ -MAP and SRAD do not sufficiently suppress the noise. This is especially obvious in the black region, where the noise is essentially intact. Although GAMBID successfully removes the noise in these regions, it doesn’t preserve the fine structures.

field estimates are mapped to $[0.5, 1.5]$.

As shown in the detail of SAR1 in Figure 4.12, GAMBID and BM3D are able to remove speckle while retaining structural features better than Lee, SRAD, and GenLik, while Γ -MAP leaves very noisy edges. Beside each denoised image, we also show the noise signal estimates \hat{n}_{mult} . Some structure is evident in the GAMBID \hat{n}_{mult} , though not to the extent of Lee, SRAD, GenLik, and Γ -MAP. The BM3D \hat{n}_{mult} exhibits some subtle nonstationarities across different regions.

The results from SAR2 in Figure 4.13 are similar to those from SAR1, in that GAMBID and BM3D are able to remove speckle while retaining structural features better than Lee, SRAD, and GenLik. Here, some structural details, the fine fissures, appear in noise estimate for GAMBID. In this case, as it is with SAR1, choosing between the BM3D and GAMBID results based on purely visual inspection is difficult.


4.2.3 Ultrasound images

In the “Embryo” image in Figure 4.14, none of the methods save SRAD and GAMBID provides any significant noise attenuation, but GAMBID produces sharper edges than SRAD.


In the “Prostate” image in Figure 4.15, the Lee and SRAD filters are visually inferior as they oversmooth the image. Of the remaining filters, GAMBID provides the highest amount of speckle attenuation while retaining all the features visible in the Γ -MAP, BM3D and GenLik results. A significant amount of correlated noise remains in all the results.

4.3 Discussion

The first main testing objective is to demonstrate the viability of GAMBID using synthetic images using two noise types at varying noise levels. In general, GAMBID




(a) Original




(b) GAMBID

Figure 4.10: SAR1 processed with GAMBID (1209×865). A comparison of results for this image is given in Figure 4.12




(a) Original




(b) GAMBID

Figure 4.11: SAR2 processed with GAMBID (599×635). A comparison of results for this image is given in Figure 4.13




(a) Original CNR = 4.40




(b) Lee CNR = 3.61

(c) Γ -MAP CNR = 13.16



(d) SRAD CNR = 5.92


(e) BM3D CNR = 13.12




(f) GenLik CNR = 5.93

(g) GAMBID CNR = 29.25

Figure 4.12: Detail of SAR1. GAMBID and BM3D are able to remove speckle while retaining structural features better than Lee, SRAD, and GenLik, while Γ -MAP leaves very noisy edges. Beside each denoised image, we also show the noise signal estimates $\hat{n}_{mult} = I/\hat{r}$. Some structure is evident in the GAMBID noise estimate, though not to the extent of Lee, SRAD, GenLik, and Γ -MAP. The BM3D noise estimate exhibits some subtle nonstationarities across different regions.




(a) Original CNR = 1.96




(b) Lee CNR = 1.67

(c) Γ -MAP CNR = 3.41



(d) SRAD CNR = 2.00

(e) BM3D CNR = 3.63



(f) GenLik CNR = 2.58

(g) GAMBID CNR = 5.56

Figure 4.13: Detail of SAR2. GAMBID and BM3D are able to remove speckle while retaining structural features better than Lee, SRAD, and GenLik. Here, some structural details, the fine fissures, appear in noise estimate for GAMBID. In this case, as it is with SAR1, choosing between the BM3D and GAMBID results based on purely visual inspection is difficult.

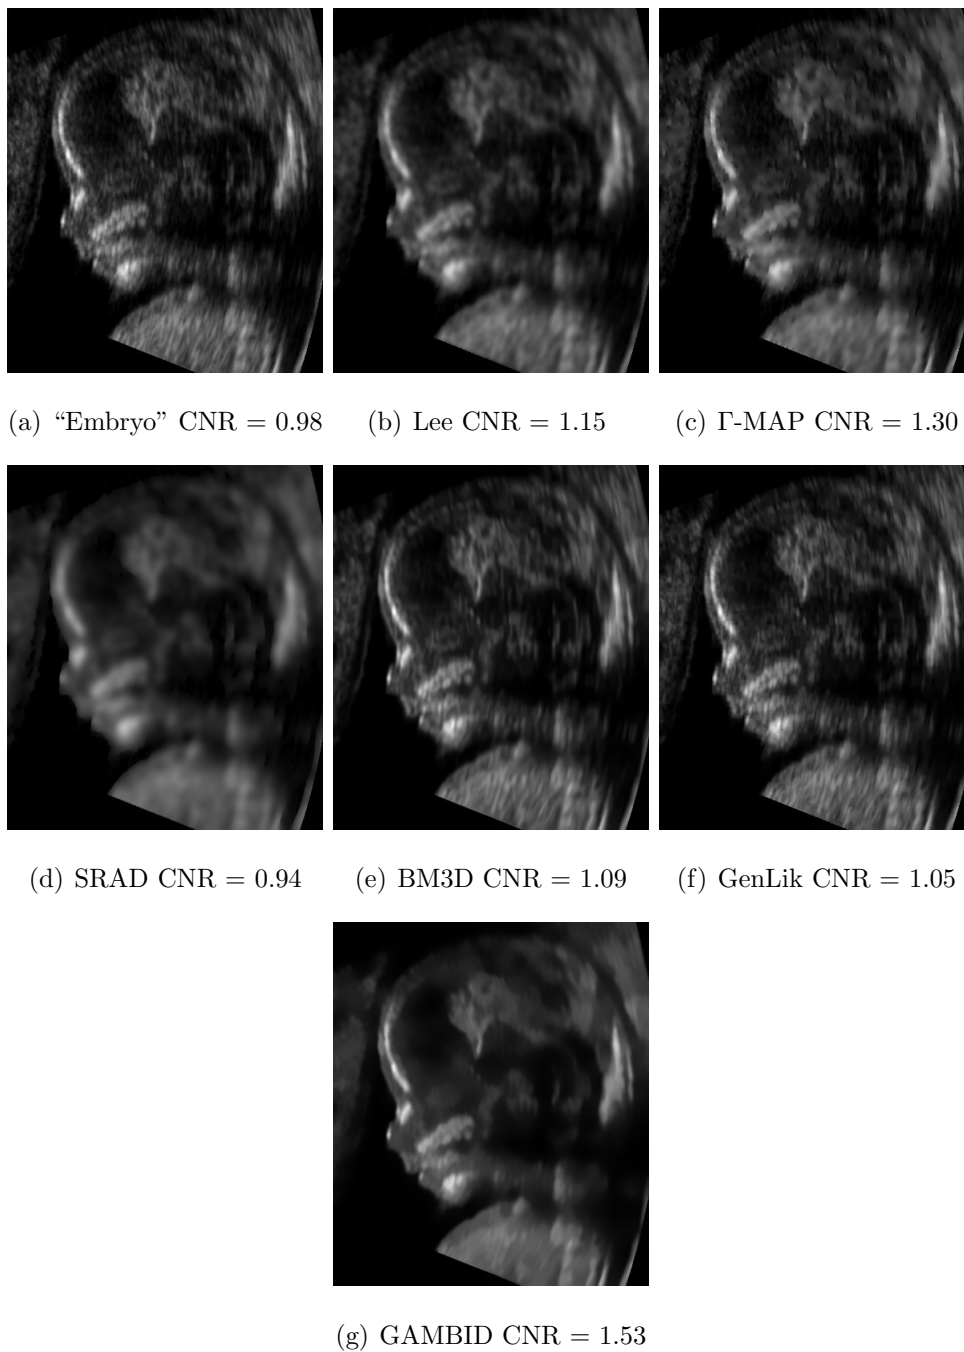




Figure 4.14: “Embryo” ultrasound image. None of the methods save SRAD and GAMBID provides any significant noise attenuation, but GAMBID produces sharper edges than SRAD. Unlike with SAR data, the noise signal estimates are not shown because they cannot be obtained by a simple ratio because ultrasound images do not follow the product model.




(a) “Prostate” CNR = 0.80

(b) Lee CNR = 0.75




(c) Γ -MAP CNR = 1.06

(d) SRAD CNR = 0.88



(e) BM3D CNR = 0.82

(f) GenLik CNR = 0.82



(g) GAMBID CNR = 1.20

Figure 4.15: “Prostate” ultrasound image. The Lee and SRAD filters are visually inferior as they oversmooth the image. Of the remaining filters, GAMBID provides the highest amount of speckle attenuation while retaining all the features visible in the Γ -MAP, BM3D and GenLik results. A significant amount of correlated noise remains in all the results.

provides consistent performance across the two noise types, which is a quality shared with Lee, GenLik and BM3D (where the Lee filter was adapted depending on the type of noise). Of these methods, GAMBID and BM3D consistently outperform Lee and GenLik. Although BM3D was developed for additive white noise, these tests show that it can be used to denoise multiplicative noise as well, often providing the strongest results. However, the metrics generated do not support arguments for which algorithm is the “best”, as the PSNR and SSIM can give conflicting numbers, where a method can score high in one metric and low in the other. Also, the metrics do not adequately capture the cases where GAMBID provides less than satisfactory results.

The synthetic test results show that GAMBID suffers from two limitations. In the “Texture” test, there is a loss of texture contrast, resulting in flat-looking images. This seems to indicate that the design of GAMBID favours homogeneous regions due to the estimation being based on intensity only. Although the weights calculated from neighbourhood differences could account for some of the texture information, this aspect will need to be improved, perhaps by incorporating texture characteristics such as entropy. The other limitation can be seen the “Structure” tests, where some of the fine details are lost in the GAMBID results, such as the wavy cross-hatch in near the top-right of the image, although other methods suffer from a similar limitation. For some reason, fine, low-contrast features are not captured by the estimator.

The second main testing objective is to demonstrate the viability of GAMBID with real signal-dependent noise, in this case SAR and ultrasound speckle. For the four tested images, GAMBID and BM3D provided the best results in terms of edge preservation and noise attenuation, as evaluated subjectively. However, none of the tested methods can fully restore ultrasound images due to the high level of spatially-varying correlation. Although GAMBID performs well for these images, decorrelation is required for proper restoration of ultrasound images, which is beyond the scope of this thesis.

From the tests, we can conclude that the proposed GAMBID algorithm is competitive with current methods in terms of the main objectives: denoising of (1) pure artificial speckle, (2) combined multiplicative and additive noise, and (3) real SAR and ultrasound speckle. However, the current implementation over-smooths textures and fine structural details. These problems can be addressed in future iterations, which are discussed in Chapter 5.

Chapter 5

Conclusions and Future Work

5.1 Conclusions

In this thesis, we presented GAMBID, a general adaptive Monte Carlo Bayesian image denoising algorithm. GAMBID works by building a posterior distribution estimate for each pixel via random sampling. This approach does not depend on any particular image model, and thus can be applied to different noise types, both signal-dependent and signal-independent.

The testing demonstrates that GAMBID is viable under two types of signal-dependent noise and with two types of synthetic images, as well as with real SAR and ultrasound imagery. In the synthetic image tests, GAMBID produced competitive and consistent (albeit not always superior) results in terms of objective quality metrics, but over-smooths textures and fine details. In the SAR and ultrasound image tests, GAMBID outperforms current methods in terms subjective evaluation of edge preservation and noise attenuation. However, some correlated speckle remains in the ultrasound results for all methods, indicating that the ultrasound data needs to be decorrelated and denoising should be performed in the log-normal domain where the correlation does not vary spatially.

Because it is a model-free approach, GAMBID can potentially be applied in

a wide variety of applications. These include remote sensing, medical imaging, computer vision, and video denoising.

5.2 Recommendations for Future Work

An advantage offered by GAMBID method is that it can be expanded to become a general framework instead of a fixed algorithm. While still maintaining the Monte Carlo Bayesian least squares approach, GAMBID can be adapted in two main ways. Firstly, it is not necessary to perform the estimation purely in the spatial domain. GAMBID can be applied in the wavelet domain, for example, to better capture the statistics that exist at different scales and spatial frequencies. Secondly, the weights used in the posterior density estimation do not have to be based on neighbourhood intensity differences. Other image features, such as texture features, can be used in the weighting. Also, heuristics can be incorporated into the weighting in a manner similar to the Enhanced Lee Filter (i.e. classification of the pixels into homogeneous, point-target, and textured types). With these adaptations, we can address the limitations in texture and small-scale detail preservation in the current implementation. Furthermore, the adaptability of GAMBID offers additional versatility for other denoising applications not discussed here.

The GAMBID framework is currently formulated for single-channel images. However, multichannel images are common in many applications, such as colour images and polarimetric data. GAMBID can be extended to take advantage of any correlation in the noise between channels by considering the each pixel as a vector-valued observation. Additionally, it can also be extended to three-dimensional data such as video or volumes by sampling along the third axis.

APPENDICES

Appendix A

Mathematical Derivations

In the Bayesian least squares problem, the goal is to minimize the expected value of the squared error (difference between the estimate and the true value), given the observed neighbourhood intensities \mathbf{I}

$$\hat{r} = \arg \min_{\hat{r}} \{E[(\hat{r} - r)^2 | \mathbf{I}]\}. \quad (\text{A.1})$$

Differentiating the objective function with respect to the estimate \hat{r} ,

$$\begin{aligned} \frac{\partial}{\partial \hat{r}} E[(\hat{r} - r)^2 | \mathbf{I}] &= \frac{\partial}{\partial \hat{r}} \int (\hat{r} - r)^2 p(r | \mathbf{I}) dr \\ &= \int \frac{\partial}{\partial \hat{r}} (\hat{r} - r)^2 p(r | \mathbf{I}) dr \\ &= \int \frac{1}{2} (\hat{r} - r) p(r | \mathbf{I}) dr \\ &= \frac{1}{2} \int \hat{r} p(r | \mathbf{I}) dr - \frac{1}{2} \int r p(r | \mathbf{I}) dr \\ &= \frac{1}{2} \hat{r} - \frac{1}{2} \int r p(r | \mathbf{I}) dr. \end{aligned}$$

To find the minimum, we set the derivative of the objective function to zero

$$\begin{aligned} \frac{1}{2} \hat{r} - \frac{1}{2} \int r p(r | \mathbf{I}) dr &= 0 \\ \hat{r} &= \int r p(r | \mathbf{I}) dr = E[r | \mathbf{I}]. \end{aligned} \quad (\text{A.2})$$

References

- [1] “DHD multimedia gallery.” http://gallery.hd.org/_virtual/ByCategory/medicine/sonogram/. 2, 26
- [2] J. M. Bueno, J. J. Hunter, C. J. Cookson, M. L. Kisilak, and M. C. W. Campbell, “Improved scanning laser fundus imaging using polarimetry,” *Journal of the Optical Society of America A*, vol. 24, pp. 1337–1348, May 2007. 2, 3
- [3] A. Buades, B. Coll, and J. Morel, “Nonlocal image and movie denoising,” *Int. J. Comput. Vision*, vol. 76, no. 2, pp. 123–139, 2008. 3, 6, 8
- [4] C. Tomasi and R. Manduchi, “Bilateral filtering for gray and color images,” in *Computer Vision, 1998. Sixth International Conference on*, pp. 839–846, 1998. 3, 8
- [5] D. T. Kuan, A. A. Sawchuk, T. C. Strand, and P. Chavel, “Adaptive restoration of images with speckle,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 35, no. 3, pp. 373–383, 1987. 3, 8, 9, 22
- [6] S. Greenberg and D. Kogan, “Improved structure-adaptive anisotropic filter,” *Pattern Recognition Letters*, vol. 27, no. 1, pp. 59–65, 2006. 3, 8
- [7] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, “Image denoising by sparse 3-D Transform-Domain collaborative filtering,” *Image Processing, IEEE Transactions on*, vol. 16, no. 8, pp. 2080–2095, 2007. 3, 8, 22
- [8] N. Wiener, *Extrapolation, interpolation, and smoothing of stationary time series*. New York: Wiley, 1950. 3, 8

- [9] A. Achim, P. Tsakalides, and A. Bezerianos, “SAR image denoising via bayesian wavelet shrinkage based on heavy-tailed modeling,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, pp. 1773–1784, Aug. 2003. 3, 10
- [10] J. Portilla, V. Strela, M. Wainwright, and E. Simoncelli, “Image denoising using scale mixtures of gaussians in the wavelet domain,” *Image Processing, IEEE Transactions on*, vol. 12, no. 11, pp. 1338–1351, 2003. 3, 8
- [11] J. W. Goodman, “Statistical properties of laser speckle patterns,” in *Laser Speckle and Related Phenomena* (J. C. Dainty, ed.), pp. 9–75, New York: Springer-Verlag, 1984. 3
- [12] C. Oliver and S. Quegan, *Understanding Synthetic Aperture Radar Images*. Boston: Artech House, 1998. 3, 6, 14, 26
- [13] T. Aysal and K. Barner, “Rayleigh-Maximum-Likelihood filtering for speckle reduction of ultrasound images,” *Medical Imaging, IEEE Transactions on*, vol. 26, no. 5, pp. 712–727, 2007. 3, 7, 9
- [14] R. G. Driggers, P. Cox, and T. Edwards, *Introduction to infrared and electro-optical systems*. Artech House, 1999. 6
- [15] V. Anastassopoulos, G. A. Lampropoulos, M. Strojnik, and B. F. Andresen, “Statistical infrared image analysis,” in *Infrared Spaceborne Remote Sensing III*, vol. 2553, (San Diego, CA, USA), pp. 171–181, SPIE, 1995. 6
- [16] T. Eltoft, “Modeling the amplitude statistics of ultrasonic images,” *Medical Imaging, IEEE Transactions on*, vol. 25, no. 2, pp. 229–240, 2006. 7
- [17] H. Gudbjartsson and S. Patz, “The rician distribution of noisy MRI data,” *Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine*, vol. 34, pp. 910–914, Dec. 1995. PMID: PMC2254141. 7

- [18] F. Luisier, T. Blu, and M. Unser, “A new SURE approach to image denoising: Interscale orthonormal wavelet thresholding,” *Image Processing, IEEE Transactions on*, vol. 16, no. 3, pp. 593–606, 2007. 8
- [19] M. Lindenbaum, M. Fischer, and A. Bruckstein, “On gabor’s contribution to image enhancement,” *Pattern Recognition*, vol. 27, no. 1, pp. 1–8, 1994. 8
- [20] J. Lee, “Digital image enhancement and noise filtering by use of local statistics,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. PAMI-2, no. 2, pp. 165–168, 1980. 8, 22
- [21] J. Lee, “Speckle suppression and analysis for synthetic aperture radar images,” *Optical Engineering*, vol. 25, pp. 636–643, May 1986. 8
- [22] V. S. Frost, J. A. Stiles, K. S. Shanmugan, and J. C. Holtzman, “A model for radar images and its application to adaptive digital filtering of multiplicative noise,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. PAMI-4, no. 2, pp. 157–166, 1982. 8
- [23] A. Lopes, R. Touzi, and E. Nezry, “Adaptive speckle filters and scene heterogeneity,” *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 28, no. 6, pp. 992–1000, 1990. 9, 22
- [24] A. Lopes, E. Nezry, R. Touzi, and H. Laur, “Structure detection and statistical adaptive speckle filtering in SAR - images,” *International Journal of Remote Sensing*, vol. 14, no. 9, p. 1735, 1993. 9, 22
- [25] T. Loupas, W. McDicken, and P. Allan, “An adaptive weighted median filter for speckle suppression in medical ultrasonic images,” *Circuits and Systems, IEEE Transactions on*, vol. 36, no. 1, pp. 129–135, 1989. 9
- [26] P. Perona and J. Malik, “Scale-space and edge detection using anisotropic diffusion,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 12, no. 7, pp. 629–639, 1990. 9

- [27] Y. Yu and S. T. Acton, “Speckle reducing anisotropic diffusion,” *Image Processing, IEEE Transactions on*, vol. 11, no. 11, pp. 1260–1270, 2002. 10, 22
- [28] Y. Yue, M. Croitoru, A. Bidani, J. Zwischenberger, and J. Clark, “Nonlinear multiscale wavelet diffusion for speckle suppression and edge enhancement in ultrasound images,” *Medical Imaging, IEEE Transactions on*, vol. 25, no. 3, pp. 297–311, 2006. 10
- [29] A. Pizurica, W. Philips, I. Lemahieu, and M. Acheroy, “A versatile wavelet domain noise filtration technique for medical imaging,” *Medical Imaging, IEEE Transactions on*, vol. 22, no. 3, pp. 323–331, 2003. 10, 22
- [30] S. Geman, “Stochastic relaxation, gibbs distributions, and the bayesian restoration of images,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. PAMI-6, no. 6, pp. 721–741, 1984. 16
- [31] H. Yu, Z. Lin, G. L. Chai, C. H. Yan, and N. Raman, “A DCT-based change detection method for Multi-Temporal SAR images,” in *Geoscience and Remote Sensing Symposium, 2008. IGARSS 2008. IEEE International*, vol. 3, pp. III – 166–III – 169, 2008. 22
- [32] A. K. Qin and D. A. Clausi, “Multivariate image segmentation based on semantic region growing with adaptive edge penalty,” *submitted to IEEE Transactions on Image Processing*, 2009. 24
- [33] P. Brodatz, *Textures: A Photographic Album for Artists and Designers*. New York: Dover Publications, 1966. 24
- [34] Z. Wang, E. Simoncelli, and A. Bovik, “Multiscale structural similarity for image quality assessment,” in *Signals, Systems and Computers, 2003. Conference Record of the Thirty-Seventh Asilomar Conference on*, vol. 2, pp. 1398–1402 Vol.2, 2003. 25
- [35] X. Song, B. W. Pogue, S. Jiang, M. M. Doyley, H. Dehghani, T. D. Tosteson, and K. D. Paulsen, “Automated region detection based on the contrast-to-noise

ratio in near-infrared tomography,” *Applied Optics*, vol. 43, pp. 1053–1062,
Feb. 2004. 26