

Automatic matching of features in Synthetic Aperture
Radar data to digital map data

Ronald George Caves

Thesis submitted for the degree of Doctor of Philosophy
Department of Applied and Computational Mathematics
University of Sheffield

June 1993

Summary

Automatic matching of features in Synthetic Aperture Radar data to digital map data

Ronald George Caves

The large amounts of Synthetic Aperture Radar (SAR) data now being generated demand automatic tools for image interpretation. Where available, map data provides a valuable aid for visual interpretation and it should aid automatic interpretation. Automatic map based interpretation will be heavily dependent on methods for matching image and map features, both for defining the initial registration and for comparing image and map. This thesis investigates methods for carrying out this matching.

Before beginning to develop image map matching methods, a full understanding of the nature of SAR data is first required. The general theory of SAR imaging, the effects of speckle and texture on image statistics, multi-look image statistics, and parameter estimation, are all discussed before addressing the main subject matter.

Initially the feasibility of directly matching map features to SAR image features is investigated. Simulations based on a simple image model produce promising results. However, the results of matching features in real images are disappointing. This is due to the limitations of the image model on which matching is based. Possible extensions to include texture and correlation are considered to be computationally too expensive. Rather, it is concluded that pre-processing is needed to structure the image prior to matching.

Structuring using edge detection and segmentation are investigated. Among operators for detecting edges in SAR an operator based on intensity ratios is identified as the most suitable. Its performance is fully analysed. Segmentation using an iterative edge detection/segment growing algorithm developed at the Royal Signals and Radar Establishment is investigated and various improvements are suggested. The output of segmentation is structured to a higher level than the output of edge detection. Thus the former is the more suitable candidate for map matching. Approaches to matching segmentations to map data are discussed.

Contents

Acknowledgement	5
1 Introduction	6
1.1 Outline	7
1.2 Synthetic aperture radar data	8
2 SAR theory	10
2.1 Antenna theory	11
2.2 Azimuth resolution	14
2.3 Range resolution	17
2.4 Summary	20
3 Image statistics	21
3.1 Imaging model	22
3.2 Statistics of the phase and complex components	23
3.3 Statistics of the detected intensity	25
3.3.1 General case	25
3.3.1.1 The exponential distribution	29
3.3.1.2 The multiplicative imaging model	29
3.3.2 Exponentially distributed intensity	30
3.3.3 K-distributed intensity	32
3.3.3.1 Gamma distributed imaged SCS	34
3.3.3.2 Justification for a gamma distributed imaged SCS	37
3.3.3.3 K-distributed field statistics	40
3.4 Statistics of the amplitude of the detected field	43
3.4.1 The Rayleigh distribution	43
3.4.2 Constant imaged SCS	44
3.4.3 Gamma distributed SCS	45
3.4.3.1 Square root gamma distribution	45
3.4.3.2 K-distributed amplitude	46
3.5 Transforming speckle to being additive	47
3.5.1 Constant imaged SCS	49
3.5.2 Gamma distributed SCS	49
3.5.2.1 K-distributed density	50
3.6 Second order intensity statistics	53
3.7 Summary	55

4	Parameter Estimation and Multi-looking	57
4.1	Averages of pixel values	58
4.1.1	Exponential intensity	60
4.1.2	Rayleigh distributed amplitude	61
4.1.3	K-distributed intensity	62
4.1.4	Effect of correlation on averaging	63
4.1.4.1	Correlated speckle	64
4.1.4.2	Correlated imaged SCS	66
4.2	Texture estimation	67
4.2.1	ML estimation of the mean and order parameter of K-distributed intensity	67
4.2.2	Non-optimal estimation of the order parameter of K-distributed intensity	68
4.2.2.1	Second normalised intensity moment	68
4.2.2.2	Mean normalised density	70
4.3	Summary	72
5	Template matching	75
5.1	Correlation measures	76
5.1.1	Cross-correlation	76
5.1.2	The Kolmogorov-Smirnov correlation measure	80
5.2	Comparison of correlation measures	80
5.2.1	Behaviour of the measures at a match	81
5.2.2	Behaviour as a template moves away from a match	81
5.3	Performance tests	83
5.3.1	Simulated data	83
5.3.2	Real Data	84
5.4	Discussion	87
5.5	Summary	90
6	Statistics along edges	92
6.1	Detected field along an edge	93
6.1.1	Terrain effects	94
6.1.1.1	Changes in surface slope	95
6.1.1.2	Layover	95
6.1.1.3	Radar shadow	97
6.2	Intensity statistics along edges	98
6.2.1	Characteristic function of the intensity	98
6.2.2	Intensity moments	99
6.2.3	The effect of system noise	100
6.2.4	Statistics of the imaged SCS	101
6.3	Measurement of intensity statistics along edges	102
6.3.1	Measurements along linear targets	105
6.3.2	Comparison to measurements over extended targets	106
6.4	Summary	107

7	Edge detection	108
7.1	The local operator approach to edge detection	109
7.2	Operators for detecting edges in SAR images	110
7.2.1	Disparity measures	110
7.2.1.1	ML estimates of the ratio and difference of mean values	111
7.2.1.2	Distribution of the normalised ratio	113
7.2.1.3	Distribution of the absolute difference	116
7.2.2	Heterogeneity measures	120
7.2.2.1	Coefficient of variation based operators	120
7.2.2.2	The Marr Hildreth operator	122
7.3	Performance of the normalised ratio operator	124
7.3.1	Conditions for a constant false alarm rate	125
7.3.2	Threshold selection and edge detectability	126
7.3.3	Correlated and K-distributed intensity	128
7.3.4	Measured False Alarm Rates	130
7.4	Edge scale size and detectability	139
7.4.1	Edge thickening	139
7.4.2	Multiple edges occurring within the window	140
7.5	Summary	142
8	Point target detection	144
8.1	Statistics of speckle peaks	144
8.1.1	Distributions of maximum and minimum values	145
8.1.2	Maximum and minimum intensity	145
8.1.3	Distribution of maximum and minimum ratios	146
8.1.4	Intensity fluctuations	148
8.2	Approaches to point feature detection	150
8.2.1	Point feature detection based on the magnitude of intensity fluctuations	152
8.3	Results of point target detection	152
8.4	Summary	154
9	Segmentation	158
9.1	Segmentation algorithm	159
9.1.1	Segmentation accuracy and halting	159
9.1.2	Edge Detection	160
9.1.2.1	The normalised gradient operator	160
9.1.2.2	Edge thresholding	163
9.1.2.3	Point target detection	163
9.1.2.4	Alternative edge detectors	164
9.1.2.5	Comparison of segmentations	166
9.1.2.6	Intensity segmentation	166
9.1.3	Segment growing and merging	169
9.1.3.1	Segment growing	169
9.1.3.2	Merging based on moments	170

9.1.3.3	Merging based on edge strength	170
9.2	Segmentation of ERS-1 data	173
9.2.1	The ERS-1 dataset	173
9.2.2	Segmentation of SLC data	181
9.2.3	Segmentation of PRI data	185
9.3	Summary	190
10	Conclusion and future work	195
10.1	Background theory	195
10.2	Direct matching	196
10.3	Edge detection	197
10.4	Segmentation	198
10.5	Image map matching	200
10.6	Future work	202
A	PDF's of binary operators	203
B	Rayleigh characteristic function	205
C	Pdf of the sum of Γ random variables	207
D	Moments of circularly symmetric RV's	209
E	Selection of map data	211
F	Glossary	213
	Bibliography	224

Acknowledgement

The research described in this thesis could not have been carried out without the help and guidance of a number of people all of whom I am grateful to. Most of all I wish to thank my supervisors Dr Shaun Quegan and Dr Peter Harley, first for offering the opportunity to carry out the research, and second for being very generous both with their time and knowledge when discussing progress and when answering queries. I also wish to thank research colleagues for helpful discussion; Dr David Blacknell, Dr Luciano Dutra and Kevin Grover, all of whom have shared an office with me at one stage, and Dr Ian Rhodes. Professor Chris Oliver and Dr Richard White from the SAR group at DRA Malvern have been very helpful in supplying data and segmentation software, but even more so in *their critical appraisal of my results presented at half yearly meetings*. These meetings have always generated tangential but highly beneficial discussion of fundamentals, usually with myself as a bystander to Shaun, Chris and Richard as the main protagonists. This, along with never being allowed to get beyond the opening statement of a presentation without being queried, has inspired me to gain a fuller understanding of many topics. Both the Jet Propulsion Laboratory and the European Space Agency also need to be thanked for supplying data. I am also grateful to SERC for funding me as a research assistant on grants GR/F14505 and GR/H90636 over the time this thesis was written. Finally, I wish to thank family and friends for general encouragement, particularly when the light at the end of the tunnel was still far off.

Chapter 1

Introduction

Synthetic Aperture Radar (SAR) is a high-resolution, active microwave imaging device which has advantages over passive optical imaging in that it can image through cloud and at night. This, along with the sensitivity of radar wavelengths to surface roughness and moisture content [59], means that SAR is potentially an ideal remote sensing tool for distinguishing land cover types and for detecting temporal change, e.g., in agriculture or forestry.

Prior to 1991 SAR images were only readily available from the short lived Seasat mission in 1978, brief Space Shuttle missions in 1981 and 1984 and various airborne SAR missions which at most gave only regional coverage [10, 27]. With this sparse spatial and temporal coverage the scope for change detection using SAR was limited and there was no urgent requirement for automatic methods for interpreting images, nor for detecting changes between successive images of the same scene. However, in July 1991 the ERS-1 satellite was launched by the European Space Agency (ESA). ERS-1 carries a SAR as part of its payload and is the first in a series of missions which will supply SAR images on a regular basis (the Japanese JERS-1, ERS-2, the Canadian Radarsat, and the ESA polar platform POEM-1).

To utilize the large number of SAR images now being generated, automatic tools for carrying out interpretation and change detection are urgently required. When these tasks are carried out manually, map data of the area covered by an image is invaluable both as an interpretative aid and as a framework for change detection between successive images. Automatic methods may benefit from an analogous use of map data [27, 51]. However, before map data can be used in an automatic manner for interpretation or change detection, the SAR and map data have to be registered. This is carried out by identifying matching image and map features from which a geometric mapping function [23] between SAR and map data can be defined. This matching process is fundamental to the whole automatic interpretation scheme as methods for automatically matching image and map features are also required for identifying change once registration has taken place. This thesis addresses how to carry out this central task. The problem can be broken into a number of sub-questions:

- Can map features be matched directly to SAR image features ?

- If not, what kind of pre-processing is needed prior to matching ?
- What algorithms are available to carry out this pre-processing and how do they perform ?
- How does matching proceed after pre-processing ?

These questions comprise the main aspects of the image map matching problem addressed in following chapters. As rough topography introduces additional problems, for the moment we will only consider the problem of matching over regions of flat topography.

In the remainder of this introductory chapter an outline plan of the thesis is presented and a brief description is given of the SAR data that will be used in later chapters.

1.1 Outline

Prior to being able to address any of the above questions, a full understanding of the nature of SAR data is needed, i.e., what information it represents and what its noise limitations are. This is a general prerequisite for developing tools for interpreting any type of data. Hence Chapters 2 to 4 deal with the background theory needed for interpreting SAR images. Few of the results in these chapters are new; rather, their intended contribution is to combine results from a large number of sources into a coherent discussion of models for SAR image statistics. These models will supply the theoretical basis for image analysis in later chapters. It is especially hoped that the discussion of first order image statistics is as complete as possible.

In Chapter 2 the general theory of SAR imaging is discussed; in particular image resolution and the response of the imaging system to a point target. This information will be required in later chapters to model spatial correlation in images. Because SAR is a coherent imaging system, images are corrupted by speckle noise. Due to the nature of speckle SAR images cannot be interpreted on a per pixel basis. Rather, interpretation has to be based on parameters estimated over regions of the image. Obvious questions arise as to what parameters to estimate, what the optimal estimates of these parameters are, and what error is present in the estimates. To answer these questions models for the first and second order statistics of a homogeneous region in a full resolution SAR image are discussed in Chapter 3. Estimation of the parameters describing these models are then discussed in Chapter 4. This includes a discussion of the statistics of multi-look, i.e., incoherently averaged, SAR images.

In Chapters 5 to 9 we go on to address the different aspects of the image map matching problem covered by the questions highlighted above. The first two are addressed in Chapter 5 where an attempt is made to directly match templates representing map features to corresponding features in SAR images. Although promising results are produced using simulated data, the methods do not perform successfully on real data. This failure is partially due to the matching process being based on too simple an image model but is more generally

due to the SAR image and map template being mismatched data types. Some structuring of the SAR image into objects that can be more easily matched to map features, i.e., edges or homogeneous segments, is needed prior to matching taking place. Before these topics are discussed in Chapters 7 to 9 (in answer to the third question above), the statistics of pixel value along the types of linear feature suitable for matching are investigated in more detail in Chapter 6 to see how they differ from the statistics over homogeneous regions already discussed in Chapter 3.

Methods of detecting edges in SAR images are discussed in Chapter 7. Edge detection using a local operator to measure the intensity ratio between adjacent regions is concentrated upon. The performance of this operator under the image models described in earlier chapters, is determined. A discussion of the problems involved in detecting edges at different scales leads on to a discussion of the need for point features to be detected prior to edges. Various approaches to point feature detection are described in Chapter 8. Segmentation is discussed in Chapter 9, primarily in terms of a segmentation algorithm originally developed at the Royal Signals and Radar Establishment (RSRE; now DRA Malvern). The operation of this algorithm and various modifications and weaknesses are described. The algorithm is used to investigate agricultural change detection in a series of ERS-1 images taken during the growing season of 1992.

Chapter 10 contains conclusions on the state of the image map matching problem given the results in Chapters 5 to 9, including pointers as to how an edge map or segmentation is to be matched to map data (in answer to the fourth and final question above).

A number of appendices are included. Appendix A lists the output distributions of basic binary operators in terms of the distributions of their inputs; these results are required at various stages in the main body of the thesis. The next three appendices (B-D) contain extended proofs of results referred to in the main body of the thesis. The selection of digital map data suitable for matching features in SAR images is discussed in Appendix E. Appendix F is a glossary which briefly defines various mathematical functions and relationships referred to in the main body of the thesis. For ease of reference topics listed in the glossary are printed in **bold** in the main body of the thesis wherever the reader may require further explication of the topic (e.g., where it is first introduced).

1.2 Synthetic aperture radar data

The results presented in this thesis are based on a number of datasets from both airborne and spaceborne SAR's:

1. Single-look complex data taken by the RSRE X-band airborne system near Amesbury, Wiltshire, England.
2. Quad-polarised single-look complex data taken by the Jet Propulsion Laboratory P, L and C band airborne system (AIRSAR) over the Goldstone calibration site in California, USA.

Band	Allocation	
	Frequency GHz	Wavelength cm
P	0.428-0.448	70.1-67.0
L	1.215-1.3	24.7-23.1
S	3.1-3.3	9.68-9.09
C	5.25-5.35	5.71-5.62
X	9.50-9.80	3.16-3.06

Table 1.1: Frequency band allocations for radar remote sensing [58]. Amongst past spaceborne SAR's NASA's Seasat (1978) and Shuttle imaging radar missions (SIR-A 1981, SIR-B 1984) operated at L-band while the Russian Almaz-1 (1992) operated at S-band. The current JERS-1 SAR operates at L-band. The SIR-C mission planned for 1994 will carry L and C band SAR's as well as a German X-band SAR. The SAR's on both Radarsat (1995) and POEM-1 (1998) will operate at C-band.

3. Quad-polarised single-look complex data taken by AIRSAR over the Feltwell agricultural test site in Norfolk, England.
4. Three-look intensity fast delivery (FDP) data taken by the C-band SAR on ERS-1 over Flevoland, Netherlands.
5. Single-look complex (SLC) and 3-look intensity precision image (PRI) data taken by ERS-1 over the Feltwell agricultural test site in Norfolk, England.

Seven-look amplitude data taken by the Canadian Centre for Remote Sensing (CCRS) airborne system from over Feltwell was also available, but unfortunately the data was heavily saturated due to it only having an 8-bit dynamic range and little use was made of it.

The frequency allocations corresponding to the radar frequency bands referred to above are listed in Table 1.1. Other relevant imaging parameters will be listed when the need arises.

Chapter 2

SAR theory

In this chapter the theory of how a side looking radar, flown on an airborne or spaceborne platform, uses a synthetic aperture and pulse compression to generate high resolution imagery is described. The resolution is derived in terms of the system response to a point target. This information will be required in later chapters for determining the degree of spatial correlation in a SAR image.

Before the properties of a radar antenna with a synthetic aperture can be discussed, the far-field pattern and resolution of the real aperture of a one-dimensional antenna needs to be determined; this is dealt with in Section 2.1. This information is then used in Section 2.2 to derive the resolution and system response in azimuth when a synthetic aperture is formed. It is shown that when a synthetic aperture is fully focused the azimuth resolution is given by half the real aperture length, and that the system response is described by a *sinc* function. The range resolution of a side looking radar is discussed in Section 2.3. The range resolution is inversely proportional to the bandwidth of the radar pulse. Frequency modulation is used to achieve the large bandwidth required for high range resolution. When linear modulation (i.e., a chirp pulse) is used the system response in range, like the system response in azimuth of a fully focused synthetic aperture, is described by a *sinc* function. In this chapter azimuth resolution is dealt with prior to range resolution because a discussion of the former follows on from the properties of real apertures discussed in Section 2.1. However, in practice range compression is applied to a received pulse prior to azimuth processing.

The imaging geometry will assume a flat Earth model and an imaging platform with a straight flight path. This is valid for considering the general theory of SAR and when dealing with airborne SAR's operating over narrow swath widths. However, for spaceborne SAR the imaging geometry must take into account the Earth's curvature and rotation, and the satellite's orbit [15, 54]. Because these factors affect the details of both azimuth and range processing, rather than the general theory, they are not considered further. The effects of range curvature [59] will also be ignored.

2.1 Antenna theory

By treating an antenna as an array of infinitesimal point sources, which are in phase, the field at a point target due to illumination by a one-dimensional antenna with real aperture D may be expressed

$$E = \kappa \int_{-\frac{D}{2}}^{\frac{D}{2}} a(x) e^{jk\phi(x)} dx \quad (2.1)$$

where κ is a proportionality factor, $a(x)$ is the weighting of the field across the antenna, $k = 2\pi/\lambda$ is the wavenumber of the illuminating wavelength λ , and $\phi(x)$ is the path length from position x along the antenna to the target¹. The proportionality factor κ must take into account the transmit power of the antenna and losses in power due to propagation; it is determined using the radar equation [34].

If the target is in the far-field of the antenna, i.e., the range R to the target satisfies the far-field condition $R \geq 2D^2/\lambda$ [58], rays converging on the target from different points along the antenna can be considered to be parallel at the antenna. The only difference in the length of the path different rays have to travel to the target will be due to the orientation of the antenna relative to the target. This is illustrated in Figure 2.1. When the target is oriented at an angle θ relative to being broadside of the antenna, the path length at position x along the antenna is given by

$$\phi(x) = x \sin \theta + R \quad (2.2)$$

where R is the path length from the centre of the antenna to the target, i.e., the range to the target. The far-field at orientation θ is then given by

$$E(\theta) = \kappa \int_{-\frac{D}{2}}^{\frac{D}{2}} a(x) e^{jkx \sin \theta} dx \quad (2.3)$$

Thus the far-field is proportional to the **Fourier transform** $A(\omega)$ of the antenna illumination at $\omega = k \sin \theta$ [7, 15]; this corresponds to Fraunhofer diffraction [5, 58].

Because the target is in the far-field of the antenna, when the field scattered back from the target reaches the antenna it can be considered to be a plane wavefront oriented at an angle θ relative to the antenna, and the path differences across the antenna will be the same as those when transmitting. Thus (2.3) also describes the field received by the antenna due to scattering from a point target in the far-field. The proportionality factor must now take into account losses in power due to two-way propagation and the radar cross-section of the point target.

When the weighting across the antenna is uniform ($a(x) = 1$), the far-field pattern is given by

$$E(\theta) = \kappa \int_{-\frac{D}{2}}^{\frac{D}{2}} e^{jkx \sin \theta} dx = \kappa \text{sinc} \left(\frac{D \sin \theta}{\lambda} \right) \quad (2.4)$$

¹ $j = \sqrt{-1}$

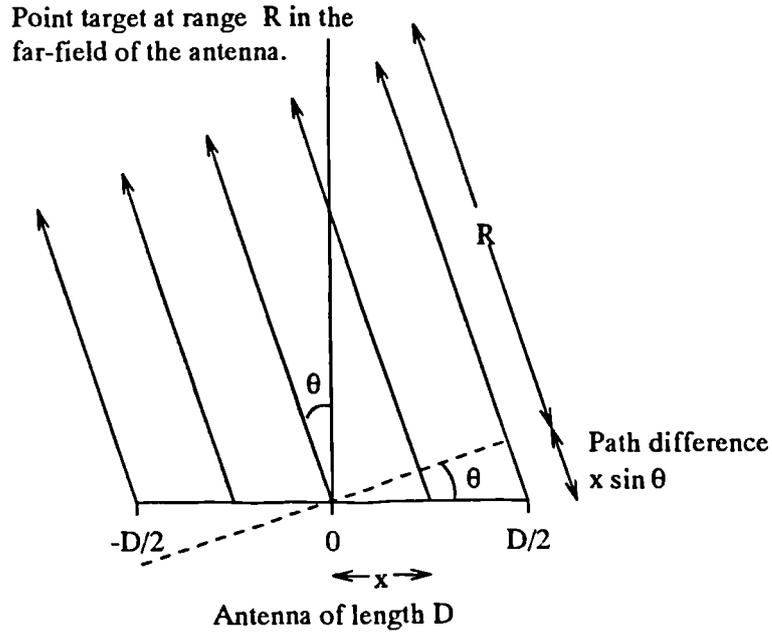


Figure 2.1: Geometry of the illumination of a point target in the far-field of a one-dimensional antenna.

where

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \quad (2.5)$$

The field pattern is shown in Figure 2.2(a). The zeros indicate where interference has completely cancelled the field. The first zeros occur at $\theta = \sin^{-1}(\lambda/D)$. Assuming $\lambda \ll D$ the width of the main beam between nulls $\theta_w = 2\lambda/D$. Of more interest in terms of detection is the half-power beamwidth θ_h defined by

$$|E(\theta_h/2)|^2 = \frac{1}{2} |E(0)|^2 \quad (2.6)$$

Solving this numerically $\theta_h = 0.88\lambda/D$. The field pattern has sidelobes which decrease in size moving away from the main beam. The peaks of the sidelobes occur where $E'(\theta) = 0$, i.e., where

$$\tan\left(\frac{\pi D \sin \theta}{\lambda}\right) = \frac{\pi D \sin \theta}{\lambda} \quad (2.7)$$

The first and largest sidelobe peaks occur where $\theta = \pm 1.43\lambda/D$, at -13.26dB below peak power.

The effects of triangular and cosine squared weightings on the far-field pattern are shown in Figures 2.2(b)-(c), and Table 2.1 compares the characteristics of the different weightings in terms of a reference angle defined by the ratio of the wavelength and aperture, $\theta_0 = \lambda/D$. It can be seen that triangular and cosine squared weightings produce successively wider beams than uniform weighting, but lead to a corresponding reduction in the level of sidelobes.

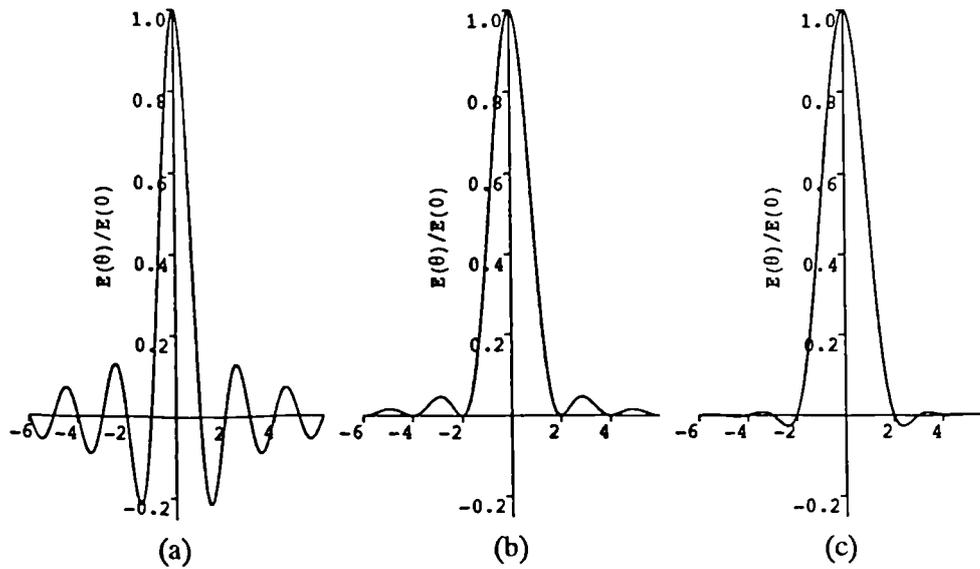


Figure 2.2: Far-field patterns resulting from different antenna weightings, (a) uniform, (b) triangular and (c) cosine squared. The x-axis is calibrated in units of $\sin \theta / \theta_0$, where θ is the orientation relative to broadside of the antenna, and $\theta_0 = \lambda / D$. The proportionality factor has been set to one.

$a(x)$	$E(\theta) / E(0)$	$E(0)$	Beamwidth		1st sidelobes	
			θ_w	θ_h	level (dB)	position
1	$\text{sinc} \left(\frac{\sin \theta}{\theta_0} \right)$	D	$2\theta_0$	$0.88\theta_0$	-6.63	$\pm 1.43\theta_0$
$2 \left(1 - \frac{2 x }{D} \right)$	$\text{sinc}^2 \left(\frac{\sin \theta}{2\theta_0} \right)$	D	$4\theta_0$	$1.28\theta_0$	-13.3	$\pm 2.86\theta_0$
$2 \cos^2 \left(\frac{\pi x}{D} \right)$	$\frac{1}{2} \text{sinc} \left(\frac{\sin \theta}{\theta_0} + 1 \right) + \text{sinc} \left(\frac{\sin \theta}{\theta_0} \right) + \frac{1}{2} \text{sinc} \left(\frac{\sin \theta}{\theta_0} - 1 \right)$	D	$4\theta_0$	$1.45\theta_0$	-15.7	$\pm 2.36\theta_0$

Table 2.1: Characteristics of the far-field pattern of an antenna with uniform, triangular and cosine squared weighting. It is assumed that the aperture D is much greater than the illuminating wavelength λ . Antenna characteristics are given in terms of the ratio, $\theta_0 = \lambda / D$. The table lists; the antenna weighting $a(x)$, the far-field pattern as a function of target orientation θ , the beamwidth between the first nulls θ_w , the half-power beamwidth θ_h , and the level and position of the first sidelobe peaks. The proportionality factor has been set to one.

In the following sections $\theta_0 = \lambda/D$ will be used as an approximate value for the beamwidth of the antenna. The ground track at a range R illuminated by such a beam will have length

$$R\theta_0 = \frac{R\lambda}{D} \quad (2.8)$$

This will be used as our working definition of the ground resolution of the antenna at range R .

2.2 Azimuth resolution

The azimuth resolution of a side-looking radar flown on an airborne or spaceborne platform can be improved by summing returns at different positions along the line of flight to form a synthetic aperture longer than the real aperture. The geometry of the illumination of a point target along a straight synthetic aperture of length L is shown in Figure 2.3. The phase of the detected field at each point along the synthetic aperture, relative to the phase at transmit, is a function of the two-way range to the target and any phase change ϱ induced by scattering

$$\psi(\theta) = 2kR(\theta) + \varrho(\theta) \quad (2.9)$$

Assuming that the phase change induced by scattering is constant over the narrow range of values taken by θ , the phase difference $\delta\psi(\theta)$ between when the target is at an orientation θ and when it is broadside of the antenna is given by

$$\delta\psi(\theta) = \psi(\theta) - \psi(0) = 2k(R(\theta) - R(0)) \quad (2.10)$$

The range $R(\theta)$ to the target can be expressed in terms of its range when it is broadside to the radar platform $R_0 = R(0)$, and its azimuth lag $x = R_0 \tan \theta$, along the ground track of the radar beam away from being broadside

$$R(x) = \sqrt{R_0^2 + x^2} \quad (2.11)$$

The phase difference is then given by

$$\delta\psi(x) = 2k \left(\sqrt{R_0^2 + x^2} - R_0 \right) \quad (2.12)$$

Normally $x \ll R_0$ when the target lies in the main lobe of the real beam. Therefore to a good approximation the phase difference is related to the azimuth lag by

$$\delta\psi(x) \approx 2kR_0 \left(\left(1 + \frac{x^2}{2R_0^2} \right) - 1 \right) = \frac{kx^2}{R_0} \quad (2.13)$$

The returns may be summed without any focusing when this quadratic phase difference term is less than a fraction of a cycle (a nominal value of $\pi/4$ is usually chosen [10, 15, 45, 59]). This is just the far-field condition applied to the synthetic aperture [58]. This requires that at the extremes of the synthetic aperture where $x = \pm L/2$

$$\frac{kx^2}{R_0} \leq \pi/4 \quad (2.14)$$

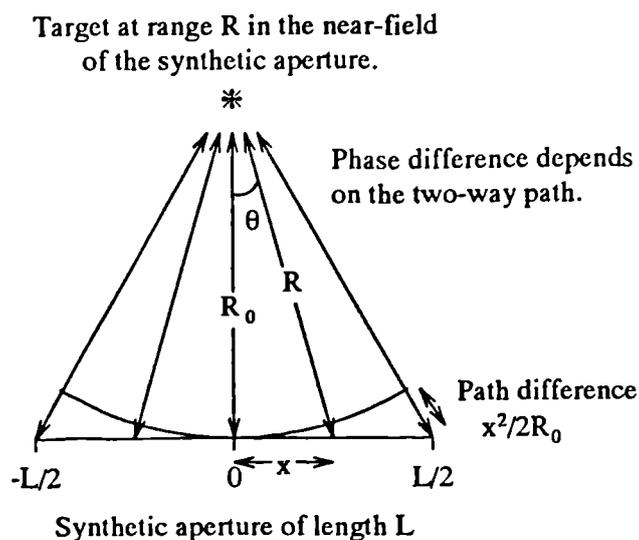


Figure 2.3: Geometry of the illumination of a point target across a synthetic aperture of length L .

Hence the optimum synthetic aperture that can be formed without any focusing is given by

$$L = \sqrt{\frac{\lambda R_0}{2}} \quad (2.15)$$

The maximum available synthetic aperture is determined by the distance over which the target remains within the main beam of the antenna, i.e., the ground resolution of the main beam in the direction of flight. Thus from (2.8) the maximum synthetic aperture in terms of the beamwidth is given by

$$L = R_0 \theta_0 = \frac{\lambda R_0}{D_a} \quad (2.16)$$

where D_a is the length of the antenna aperture in the direction of flight.

The field pattern of the synthetic aperture can be determined in the same way as that of the real aperture except that phase differences must be measured in terms of the two-way as opposed to the one-way path to the target. This results in the synthetic aperture having a beamwidth

$$\theta_0 = \frac{\lambda}{2L} \quad (2.17)$$

and corresponding ground resolution

$$d_a = \frac{\lambda R_0}{2L} \quad (2.18)$$

With the optimum unfocused synthetic aperture (2.15) the ground resolution in azimuth is

$$d_a = \sqrt{\frac{\lambda R_0}{2}} \quad (2.19)$$

Note that the ground resolution then equals the length of the synthetic aperture [59], i.e., (2.15) gives the maximum synthetic aperture length over which the synthetic apertures corresponding to adjacent azimuth resolution cells do not overlap.

Focusing over the maximum available aperture (2.16) the ground resolution in azimuth is

$$d_a = \frac{D_a}{2} \quad (2.20)$$

This means that whilst the azimuth resolution of the optimum unfocused SAR increases with range and illuminating wavelength, the azimuth resolution of a fully focused SAR is determined solely by the azimuthal length of the real aperture. The reason why the latter is independent of range is that when the range is increased the length of the synthetic aperture is increased by the same factor. This does not mean that the azimuthal length of the real aperture can be reduced indefinitely to improve resolution because, (a) this would also reduce the power of the antenna and thus its detection performance, and (b) the limits put on the pulse repetition frequency (PRF) by azimuth sampling (the Nyquist criterion demands that the PRF be greater than twice the maximum Doppler shift) and range ambiguity (the inverse of the PRF must be greater than the time delay between returns from near and far range) constrain the minimum dimensions of the real aperture [15, 59]. This second constraint is particularly severe for spaceborne SAR systems [59].

To focus on a target, the field detected from it along the synthetic aperture must be phase corrected. Let $E(\theta)$ be the field detected by a radar antenna from a point target at range R_0 in the far-field of its real aperture, and oriented at an angle θ away from broadside of the antenna. At each point x along a synthetic aperture of length L the range to the target is given by (2.11). Thus, the field $F(x)$ detected at x from the target is given by $E(\theta)$, where $\theta = \tan^{-1} x/R_0$, times a term accounting for the phase difference $\delta\psi(\theta)$ due to the target's increased range (2.13). This may be written

$$F(x) = \begin{cases} E\left(\tan^{-1} \frac{x}{R_0}\right) e^{\frac{jkx^2}{R_0}} & |x| \leq \frac{L}{2} \\ 0 & |x| > \frac{L}{2} \end{cases} \quad (2.21)$$

Because the variation in the target's range across the synthetic aperture will be small relative to the range itself, variations in the magnitude of the detected field across the synthetic aperture due to additional propagation losses are ignored in (2.21).

Phase correction involves convolving the detected field along the synthetic aperture with the matched filter of the phase difference term. The output, i.e., the azimuth response to a point target, is given by

$$h_a(x) = F(x) * e^{-j\delta\psi(-x)} = \int_{-L/2}^{L/2-|x|} F(|x|+y) e^{-j\delta\psi(y)} dy \quad (2.22)$$

for $|x| \leq L$ otherwise it is zero¹. Here we have used the fact that both the phase difference and the detected field are even functions and that they are defined

¹When the symbol $*$ is used as a unary operator as in x^* , it represents the complex conjugate

only over the interval $-L/2 \geq x \geq L/2$. Assuming $E(\theta)$ is a constant equalling E_0 over the same interval, when $|x| \leq L$ the azimuth response is given by a distorted sinc function

$$\begin{aligned} h_a(x) &= E_0 \int_{-L/2}^{L/2-|x|} e^{jk \frac{(|x|+y)^2}{R_0}} e^{-jk \frac{y^2}{R_0}} dy \\ &= E_0 L \left(1 - \frac{|x|}{L}\right) \operatorname{sinc} \left(\frac{2Lx}{\lambda R_0} \left(1 - \frac{|x|}{L}\right) \right) \end{aligned} \quad (2.23)$$

otherwise it is zero. The first nulls in the azimuth response occur where

$$x \left(1 - \frac{|x|}{L}\right) = \frac{\lambda R_0}{2L} = d_a \quad (2.24)$$

The distorted nature of the *sinc* function giving the azimuth response arises from the detected field, and thus the matched filter of the phase difference, only being defined over a finite aperture. In some references [45, 59] the matched filter of the phase difference is defined beyond this aperture. This results in the azimuth response being given by an undistorted *sinc* function but with an additional phase term. In fact within certain limits the distorted *sinc* response given by (2.23) can be approximated by an undistorted *sinc* function. As the target is in the far field of the real antenna $2D_a^2 \leq \lambda R_0$, thus from (2.16) and (2.18), $d_a/L \leq 1/4$. Therefore for $|x| \ll d_a$ (2.23) can be approximated by an undistorted *sinc* function with a width between its main lobe and first null equal to the azimuth resolution

$$h_a(x) \approx E_0 L \operatorname{sinc} \left(\frac{2Lx}{\lambda R_0} \right) = E_0 L \operatorname{sinc} \left(\frac{x}{d_a} \right) \quad (2.25)$$

In fact for many focused SAR systems, particularly spaceborne ones, $d_a/L \ll 1/4$, and the approximation given by (2.25) may be used over a wide range of values of the azimuth lag x .

2.3 Range resolution

Two components are required to form any two-dimensional array or image. When forming an image using a side looking radar the azimuth component of a target is generated by flying past it. The other component is given by the time it takes a radar pulse to return from the target when it is broadside; this gives its slant range. The range geometry of side looking radar imaging is shown in Figure 2.4.

As the length D_a of the antenna aperture determines the beamwidth in the direction of flight, the width D_r of the antenna aperture determines the beamwidth perpendicular to the direction of flight, i.e., in range. The width of the ground segment covered by the beam perpendicular to the direction of flight defines the maximum attainable 'swath' width of the image in the range

of x , whilst when it is used as a binary operator as in $x * y$, it represents the convolution of x and y .

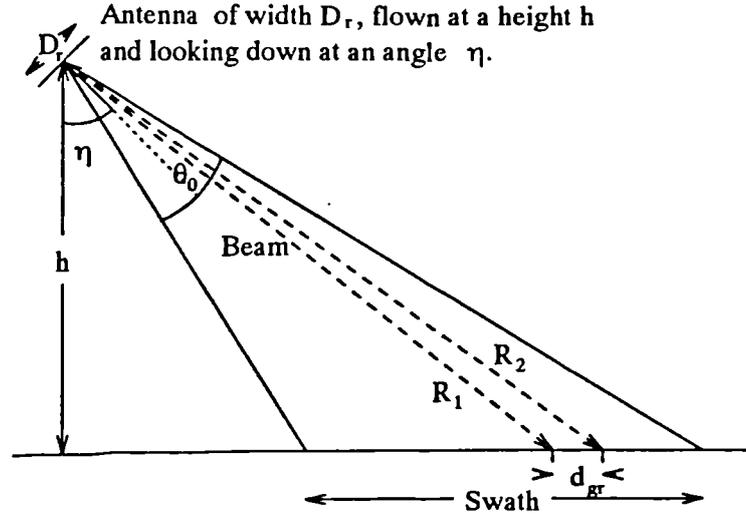


Figure 2.4: The range geometry of side looking radar imaging.

direction. The beam of a side looking radar flying at a height h and looking down at an angle η will cover a maximum swath width of [15]

$$\begin{aligned} S &= h \left(\tan \left(\eta + \frac{\theta_0}{2} \right) - \tan \left(\eta - \frac{\theta_0}{2} \right) \right) \\ &= \frac{2h \sin \theta_0}{\cos \theta_0 + \cos 2\eta} = \frac{h\theta_0}{\cos^2 \eta} \end{aligned} \quad (2.26)$$

where $\theta_0 = \lambda/D_r$ now defines the beamwidth perpendicular to the direction of flight; it is assumed that $\lambda \ll D_r$.

Using a pulse of bandwidth B , returns spaced $1/B$ or more in time apart may be resolved. In terms of the two-way path difference to targets this means that targets at slant ranges R_1 and R_2 may be resolved if

$$\frac{c}{2B} \leq |R_1 - R_2| \quad (2.27)$$

where c is the speed of electromagnetic propagation through the atmosphere ($c \approx 3 \times 10^8$ m/s). Thus the slant range resolution of the radar is given by

$$d_{sr} = \frac{c}{2B} \quad (2.28)$$

Transforming to ground range, the resolution is given by

$$d_{gr} = \frac{d_{sr}}{\sin \vartheta} \quad (2.29)$$

where ϑ is the local incidence angle of the beam with the surface. Assuming a flat surface this angle lies between $\eta - \theta_0/2$ in the near range and $\eta + \theta_0/2$ in the far range. This increase in incidence angle going across the swath from

near range to far range, causes a corresponding decrease in the ground range resolution.

To achieve a slant range resolution on the order of 15 m requires a 10MHz bandwidth (with a look angle of 30 degrees and a beamwidth of 6 degrees, the corresponding ground range resolution would vary between 33 m in near range and 27.5 m in far range). If a simple sinusoidal pulse were used it would have to be very short, $T = 1/B = 0.1 \mu\text{s}$ in length. Generating a detectable pulse of this length would require a transmitter with a very large peak power. Thus longer frequency modulated pulses are used to achieve the required bandwidth. This results in the system response in range being a *sinc* function [34, 45, 59].

Given a linear modulated pulse (i.e., a chirp) of bandwidth B and length $T' \gg 1/B$, where the frequency $f(t)$ at time t is given by

$$f(t) = \begin{cases} f_0 + \frac{Bt}{T'} & |t| \leq \frac{T'}{2} \\ 0 & |t| > \frac{T'}{2} \end{cases} \quad (2.30)$$

the pulse received from a point target at range R will be proportional to

$$s(t) = \begin{cases} e^{j2\pi(f(t-\tau)+f_d)(t-\tau)} & |t| \leq T'/2 \\ 0 & |t| > T'/2 \end{cases} \quad (2.31)$$

where $\tau = 2R/c$ is the time delay to and from the target and f_d is the Doppler shift of the target due to its velocity relative to the imaging platform [45]. The received pulse has the same length T' as the transmitted pulse. To meet the desired temporal resolution of $T = 1/B$ the received pulse has to be compressed. This is achieved by convolving it with the matched filter of the transmitted pulse [45]. The output is given by

$$h_r(t) = s(t) * e^{j2\pi f(-t)t} = \int_{\max(-\frac{T'}{2}, -\frac{T'}{2}-t)}^{\min(\frac{T'}{2}, \frac{T'}{2}-t)} s(\tau+t) e^{-j2\pi f(\tau)r} d\tau \quad (2.32)$$

when $|t - \tau| \leq T'$ otherwise it is zero. This gives the system response in range to a point target. Integrating we find that like the system response in azimuth, the system response in range is described by a distorted *sinc* function

$$h_r(t) = T' \left(1 - \frac{|t - \tau|}{T'}\right) \text{sinc} \left((2B(t - \tau) + f_d T') \left(1 - \frac{|t - \tau|}{T'}\right) \right) e^{j2\pi \left(f_0 + \frac{f_d}{2}\right)(t - \tau)} \quad (2.33)$$

For $|t - \tau| \ll T'$ the system response in range given by (2.33) may be approximated by an undistorted *sinc* function centred at $t = \tau - f_d T' / 2B$ with a width between the main lobe and first null of $1/2B = T/2$

$$h_r(t) = T' \text{sinc} (2B(t - \tau) + f_d T') e^{j2\pi \left(f_0 + \frac{f_d}{2}\right)(t - \tau)} \quad (2.34)$$

In terms of distance the system response in range is then given by

$$h_r(x) = T' \text{sinc} \left(\frac{x - 2R}{d_{sr}} + f_d T' \right) e^{j2\pi \left(f_0 + \frac{f_d}{2}\right) \left(\frac{x - 2R}{c}\right)} \quad (2.35)$$

When imaging a stationary scene in the far-field of the real aperture the maximum expected Doppler shift is v/D_a where v is the speed of the imaging platform. For a typical airborne SAR flying at a speed of 720 kph, with a real aperture of length 2 m and using a pulse length of $10 \mu\text{s}$, the centre of the system response in range will be offset at most a distance $d_{sr} \times 10^{-3}$ away from $x = 2R$ [45]. Thus for most practical purposes the offset due to the Doppler shift may be ignored.

2.4 Summary

In this chapter:

- The theory of how a side looking radar flown on an airborne or spaceborne platform uses a synthetic aperture and pulse compression, to generate high resolution imagery, was described.
- The properties of the radar beam generated by one-dimensional antenna were discussed. In particular, how the beamwidth can be approximated by the ratio of the illuminating wavelength and the length of the real aperture, and how this can be used to give a measure of resolution.
- The formation of a synthetic aperture to improve azimuth resolution was described, and the resulting resolution when the synthetic aperture is focused and unfocused.
- The system response in azimuth and range to a point target were both shown to be given by *sinc* functions when the synthetic aperture is fully focused and a chirp pulse is used to achieve high range resolution.

Chapter 3

Image statistics

Any image generated by a coherent imaging system such as a SAR, is affected by coherent interference between scatterers [22]. Over many types of extended targets (e.g., agricultural fields) this interference causes the detected intensity to fluctuate from resolution cell to resolution cell. This effect is termed speckle as it gives the image an overall speckled appearance (see Figure 3.1). Due to the nature of speckle SAR images cannot be analysed on a single pixel basis, rather analysis needs to be based on statistics estimated over regions of the image [27, 51]; this averages out the effect of speckle.

For designing optimal feature detection and matching algorithms for SAR a full understanding of image statistics is first required; in particular knowledge of the effect of speckle on image statistics. This information is needed to define the most suitable parameters upon which to base image analysis decisions, to determine how best to estimate these parameters from data, and to derive the sampling distributions of the resulting estimates for purposes of setting decision thresholds. Models for the first and second order statistics of pixel values over extended targets in full resolution (i.e., single-look) SAR imagery are described in this chapter. Parameter estimation and the effect of averaging (i.e., multi-look) on image statistics are dealt with separately in Chapter 4. These statistical models will supply the theoretical basis for the image analysis algorithms discussed in later chapters.

To make this overview of SAR image statistics as complete as possible, whenever a new statistical distribution is introduced (where they are known) its **probability density function (pdf)**, **cumulative distribution function (cdf)**, **characteristic function**, **moments**, **normalised moments**, **mean**, **variance** and **coefficient of variation**, are listed.

As far as possible the statistics of pixel values in a full resolution SAR image will be discussed in terms of a general model for the electromagnetic field detected by the imaging system; this model is presented in Section 3.1. Because most 'information' is contained in the intensity as opposed to the phase of the detected field, the statistics of the phase and complex components of the detected field are briefly described in Sections 3.2, prior to a more extensive discussion of the statistics of the intensity (and amplitude) of the detected field in Section 3.3. Assuming there are a large number of scattering centres within

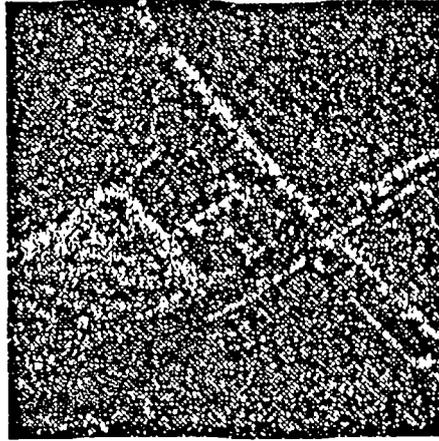


Figure 3.1: A RSRE X-band single-look SAR image of an agricultural area shown in intensity. The area contains several extended natural targets; fields which appear as dark regions, and woodland which appears as bright regions. Intensity fluctuations are clearly apparent within each extended target. These are mostly induced by speckle but are also due to fluctuations in the radar cross-section of the surface being detected by the imaging system, particularly over the woodland.

a resolution cell, it is shown how the correlation length of scatterers relative to the resolution of the imaging system influences the distribution of the detected intensity. The statistics of the amplitude of the detected field are then discussed in Section 3.4. The result of taking the log of SAR intensity data to transform multiplicative speckle to being additive is described in Section 3.5. Finally in this chapter, the second order statistics of the detected intensity are derived in Section 3.6.

3.1 Imaging model

The total electromagnetic field detected by a monochromatic, single polarised imaging system at a point $\bar{x} = (x_a, x_r)$, where x_a and x_r are azimuth and slant range coordinates, can be represented as a sum of contributions from discrete elementary scatterers [22, 31, 43, 46]

$$Z(\bar{x}) = \sum_{k=-\infty}^{\infty} a_k e^{j\theta_k} h(\bar{x}, \bar{x}_k) \quad (3.1)$$

where a_k , θ_k and \bar{x}_k are the amplitude, phase and position of the k 'th scatterer, and $h(\bar{x}, \bar{x}_k)$ is the system response at \bar{x} to a target at \bar{x}_k . Although this discrete model does not represent the physical scattering mechanism in the exact formulation of Maxwell's equations [58, 59], it offers a representation which allows the imaging process to be analysed in a manner which agrees with observed phenomena [31, 43].

So long as a target is rough relative to the illuminating wavelength and incidence angle, or the depth of the resolution cell is much greater than the wavelength, the phase of scatterers within the resolution cell will be randomly distributed over many cycles. This situation is referred to as fully developed speckle. The phase of each scatterer will then be uniformly distributed over the interval $-\pi$ to π and will be statistically independent of its amplitude [22]. Thus in fully developed speckle an elementary scatterer z can be modelled as a circularly symmetric complex random variable whose amplitude and phase have a joint pdf

$$g_z(a, \theta) = \frac{f_a(a)}{2\pi} \quad (3.2)$$

where $f_a(a)$ is the marginal pdf of the amplitude of the scatterer. It will be assumed that when speckle is fully developed each of the elementary scatterers in (3.1) are independent in phase though not necessarily in amplitude. This involves the assumption that scatterers decorrelate in phase faster than in amplitude.

Speckle will not be fully developed when a target is smooth relative to the illuminating wavelength and incidence angle, and the resolution cell is only a few wavelengths in depth. At X and C-band (3 and 6 cm wavelengths) this is likely to occur with man made targets but not with natural targets. However, at P-band (68 cm wavelength), and to a lesser extent at L-band (24 cm wavelength), many natural targets will appear smooth, and at resolutions of a few metres the detected field may be the result of scattering from only a few targets. It is also possible that speckle will not be fully developed when a smooth target which produces strong backscatter, e.g., a corner reflector, is immersed in a resolution cell which otherwise generates fully developed speckle. The field detected from such a resolution cell can be modelled as the sum of contributions from scatterers which on their own generate fully developed speckle, and a single constant phasor which represents the contribution of the smooth target [17, 22]. The resulting Rician statistics will depart significantly from those of fully developed speckle when the amplitude of the contribution from the smooth target is much greater than that from all the other contributing scatterers.

Because we are primarily interested in natural targets and because we will mainly work on C and X-band data, it will be assumed that speckle is fully developed, unless stated otherwise. A full analysis of the effect of speckle not being fully developed is still needed but is beyond the scope of this thesis.

3.2 Statistics of the phase and complex components of the detected field

The detected field (3.1) can be viewed as arising from a n -step random walk in the complex plane [22]

$$Z(\bar{x}) = \sum_{i=1}^n b_{k_i} e^{j\Theta_{k_i}} \quad (3.3)$$

where the k 'th step length b_k , is the product of the amplitudes of the k 'th scatterer and the system response to it

$$b_k = a_k |h(\bar{x}, \bar{x}_k)| \quad (3.4)$$

the k 'th step orientation Θ_k , is the sum of the phases of the k 'th scatterer and the system response to it

$$\Theta_k = \theta_k + \varphi_k \quad (3.5)$$

where¹ $\varphi_k = \angle h(\bar{x}, \bar{x}_k)$, and n is the number of scatterer positions k_1, \dots, k_n yielding a non-zero system response.

When speckle is fully developed the amplitude and phase of scatterers will be independent of each other and the system response. Therefore the step orientation will be uniformly distributed and independent of the step length and other step orientations. The phase of the detected field will then be uniformly distributed and it will be independent of the amplitude of the detected field [29], i.e., the detected field will also be a circularly symmetric complex random variable.

In itself the phase of the detected field contains no information when it is uniformly distributed. However, information is contained in phase differences between different polarised images [52], and pairs of images taken from slightly offset flight paths for use in interferometry [65]. Because polarimetry and interferometry are beyond our present scope the phase relationships between images will not be dealt with here.

When the phase of the detected field is uniformly distributed, the real and imaginary components of the detected field given by

$$Z_{\Re}(\bar{x}) = \sum_{k=-\infty}^{\infty} b_k \cos \Theta_k \quad (3.6)$$

$$Z_{\Im}(\bar{x}) = \sum_{k=-\infty}^{\infty} b_k \sin \Theta_k \quad (3.7)$$

will both have mean zero²

$$\langle Z_{\Re}(\bar{x}) \rangle = \sum_{k=-\infty}^{\infty} \langle b_k \rangle \langle \cos \Theta_k \rangle = 0 \quad (3.8)$$

$$\langle Z_{\Im}(\bar{x}) \rangle = \sum_{k=-\infty}^{\infty} \langle b_k \rangle \langle \sin \Theta_k \rangle = 0 \quad (3.9)$$

equal variance, which depends only on the step length

$$\langle Z_{\Re}(\bar{x})^2 \rangle = \sum_{i=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \langle b_i b_k \rangle \langle \cos \Theta_i \cos \Theta_k \rangle = \frac{1}{2} \sum_{k=-\infty}^{\infty} \langle b_k^2 \rangle \quad (3.10)$$

$$\langle Z_{\Im}(\bar{x})^2 \rangle = \sum_{i=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \langle b_i b_k \rangle \langle \sin \Theta_i \sin \Theta_k \rangle = \frac{1}{2} \sum_{k=-\infty}^{\infty} \langle b_k^2 \rangle \quad (3.11)$$

¹ \angle denotes the argument of a complex variable.

² The operator $\langle \cdot \rangle$ indicates ensemble averaging.

and they will be uncorrelated [22]. This is because

$$\langle Z_{\Re}(\bar{x}) Z_{\Im}(\bar{x}) \rangle = \sum_{i=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} \langle b_i b_k \rangle \langle \cos \Theta_i \sin \Theta_k \rangle = 0 = \langle Z_{\Re}(\bar{x}) \rangle \langle Z_{\Im}(\bar{x}) \rangle \quad (3.12)$$

Equations (3.8) to (3.12) rely on the fact that when Θ_k is an independent random variable uniformly distributed on the interval $(-\pi, \pi]$

$$\langle \cos \Theta_k \rangle = \langle \sin \Theta_k \rangle = 0 \quad (3.13)$$

$$\langle \cos \Theta_i \cos \Theta_k \rangle = \langle \sin \Theta_i \sin \Theta_k \rangle = \begin{cases} \frac{1}{2} & \text{if } i = k \\ 0 & \text{if } i \neq k \end{cases} \quad (3.14)$$

and

$$\langle \cos \Theta_i \sin \Theta_k \rangle = 0 \quad (3.15)$$

for all i and k .

Figure 3.2 shows the phase and complex components of the SAR image whose intensity is shown in Figure 3.1. Structural features present in the intensity image are barely discernible in the complex components. No structure is apparent in the phase image.

3.3 Statistics of the detected intensity

In this section the statistics of the detected intensity are discussed; first in Section 3.3.1 in terms of the general imaging model. It is shown that when a resolution cell contains a large number of scatterers the detected intensity may be modelled as the product of a unit mean exponential speckle process and the result of positive imaging of the surface cross-section. The exponential and K-distributed intensity statistics which arise from this model are then described in Sections 3.3.2 and 3.3.3 respectively.

3.3.1 General case

The intensity of the detected field defined by

$$\begin{aligned} I(\bar{x}) &= |Z(\bar{x})|^2 = Z_{\Re}(\bar{x})^2 + Z_{\Im}(\bar{x})^2 \\ &= \sum_{i=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} a_i a_k e^{j(\theta_i - \theta_k)} h(\bar{x}, \bar{x}_i) h^*(\bar{x}, \bar{x}_k) \end{aligned} \quad (3.16)$$

has mean, from (3.4), (3.10) and (3.11)

$$\mu_I = \langle I(\bar{x}) \rangle = \sum_{i=-\infty}^{\infty} \langle a_i^2 \rangle |h(\bar{x}, \bar{x}_i)|^2 \quad (3.17)$$

The mean intensity depends only on the surface cross-section (SCS) of the elementary scatterers defined as $\rho_i = a_i^2$ [31], and not on their phase [43]¹. If the

¹The symbol ρ is used for the surface cross-section rather than the more normal σ so as not to confuse σ 's indicating the surface cross-section with those indicating variances.

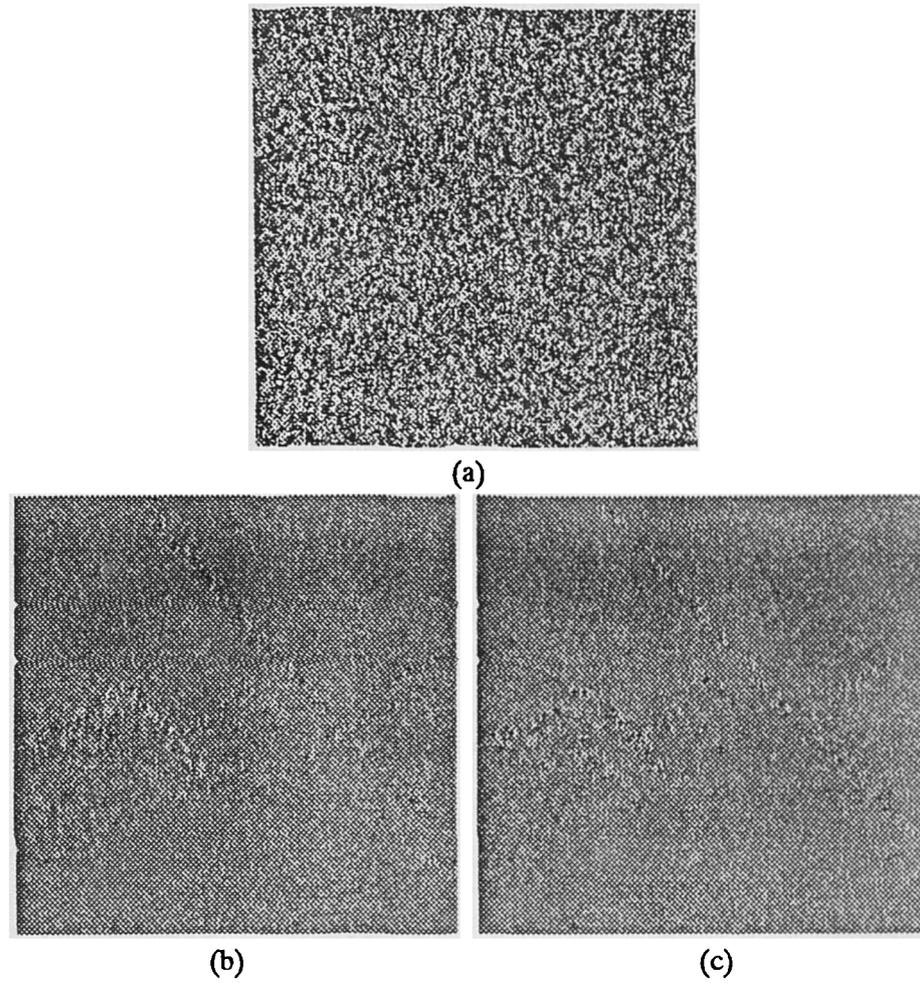


Figure 3.2: (a) the phase and, (b) and (c) the real and imaginary components of the SAR image whose intensity is shown in Figure 3.1.

SCS is spatially stationary the mean intensity is given by the product of its mean value, and the total power of the system response

$$\mu_I = \langle \rho \rangle \sum_{i=-\infty}^{\infty} |h(\bar{x}, \bar{x}_i)|^2 \quad (3.18)$$

The moments of the intensity are given by

$$\begin{aligned} \langle I(\bar{x})^m \rangle = & \sum_{i_1=-\infty}^{\infty} \sum_{k_1=-\infty}^{\infty} \dots \sum_{i_m=-\infty}^{\infty} \sum_{k_m=-\infty}^{\infty} \langle a_{i_1} a_{k_1} \dots a_{i_m} a_{k_m} \rangle \times \\ & \langle e^{j(\theta_{i_1} - \theta_{k_1} + \dots + \theta_{i_m} - \theta_{k_m})} \rangle h(\bar{x}, \bar{x}_{i_1}) h^*(\bar{x}, \bar{x}_{k_1}) \dots h(\bar{x}, \bar{x}_{i_m}) h^*(\bar{x}, \bar{x}_{k_m}) \end{aligned} \quad (3.19)$$

The ensemble average of the phase terms is unity when they all cancel, otherwise it is zero. Considering all possible ways in which cancellation can occur gives

$$\begin{aligned} \langle I(\bar{x})^m \rangle = & \sum_{\gamma_1, \dots, \gamma_m} \frac{m!}{\gamma_1! \dots \gamma_m!} \sum_{i_1=-\infty}^{\infty} \dots \sum_{\substack{i_k=-\infty \\ i_k \neq i_1, \dots, i_{k-1}}}^{\infty} \langle \rho_{i_1}^{\gamma_1} \dots \rho_{i_k}^{\gamma_k} \rangle \times \\ & |h(\bar{x}, \bar{x}_{i_1})|^{2\gamma_1} \dots |h(\bar{x}, \bar{x}_{i_k})|^{2\gamma_k} \end{aligned} \quad (3.20)$$

where the first summation in (3.20) represents summation over all possible sequences of m non-negative integers γ_i , which decrease monotonically and whose sum equals m . In each sequence γ_k is the last non-zero exponent; i.e., for each sequence $\gamma_i \geq \gamma_{i+1} \geq 0$, $\gamma_1 + \dots + \gamma_m = m$ and $\gamma_{i>k} = 0$. As with the mean intensity the higher moments depend only on the SCS of the elementary scatterers and not on their phase. Thus the distribution of the intensity will only depend on the SCS and the system response [43]. The second, third and fourth moments are given by

$$\begin{aligned} \langle I(\bar{x})^2 \rangle = & \quad (3.21) \\ & 2 \sum_{i_1=-\infty}^{\infty} \sum_{\substack{i_2=-\infty \\ i_2 \neq i_1}}^{\infty} \langle \rho_{i_1} \rho_{i_2} \rangle |h(\bar{x}, \bar{x}_{i_1})|^2 |h(\bar{x}, \bar{x}_{i_2})|^2 + \sum_{i_1=-\infty}^{\infty} \langle \rho_{i_1}^2 \rangle |h(\bar{x}, \bar{x}_{i_1})|^4 \end{aligned}$$

$$\begin{aligned} \langle I(\bar{x})^3 \rangle = & \quad (3.22) \\ & 6 \sum_{i_1=-\infty}^{\infty} \sum_{\substack{i_2=-\infty \\ i_2 \neq i_1}}^{\infty} \sum_{\substack{i_3=-\infty \\ i_3 \neq i_1, i_2}}^{\infty} \langle \rho_{i_1} \rho_{i_2} \rho_{i_3} \rangle |h(\bar{x}, \bar{x}_{i_1})|^2 |h(\bar{x}, \bar{x}_{i_2})|^2 |h(\bar{x}, \bar{x}_{i_3})|^2 \\ & + 3 \sum_{i_1=-\infty}^{\infty} \sum_{\substack{i_2=-\infty \\ i_2 \neq i_1}}^{\infty} \langle \rho_{i_1}^2 \rho_{i_2} \rangle |h(\bar{x}, \bar{x}_{i_1})|^4 |h(\bar{x}, \bar{x}_{i_2})|^2 + \sum_{i_1=-\infty}^{\infty} \langle \rho_{i_1}^3 \rangle |h(\bar{x}, \bar{x}_{i_1})|^6 \end{aligned}$$

$$\begin{aligned} \langle I(\bar{x})^4 \rangle = & \quad (3.23) \\ & 24 \sum_{i_1=-\infty}^{\infty} \sum_{\substack{i_2=-\infty \\ i_2 \neq i_1}}^{\infty} \sum_{\substack{i_3=-\infty \\ i_3 \neq i_1, i_2}}^{\infty} \sum_{\substack{i_4=-\infty \\ i_4 \neq i_1, i_2, i_3}}^{\infty} \langle \rho_{i_1} \rho_{i_2} \rho_{i_3} \rho_{i_4} \rangle |h(\bar{x}, \bar{x}_{i_1})|^2 \dots |h(\bar{x}, \bar{x}_{i_4})|^2 \end{aligned}$$

$$\begin{aligned}
& + 12 \sum_{i_1=-\infty}^{\infty} \sum_{\substack{i_2=-\infty \\ i_2 \neq i_1}}^{\infty} \sum_{\substack{i_3=-\infty \\ i_3 \neq i_1, i_2}}^{\infty} \langle \rho_{i_1}^2 \rho_{i_2} \rho_{i_3} \rangle |h(\bar{x}, \bar{x}_{i_1})|^4 |h(\bar{x}, \bar{x}_{i_2})|^2 |h(\bar{x}, \bar{x}_{i_3})|^2 \\
& + 4 \sum_{i_1=-\infty}^{\infty} \sum_{\substack{i_2=-\infty \\ i_2 \neq i_1}}^{\infty} \langle \rho_{i_1}^3 \rho_{i_2} \rangle |h(\bar{x}, \bar{x}_{i_1})|^6 |h(\bar{x}, \bar{x}_{i_2})|^2 \\
& + 6 \sum_{i_1=-\infty}^{\infty} \sum_{\substack{i_2=-\infty \\ i_2 \neq i_1}}^{\infty} \langle \rho_{i_1}^2 \rho_{i_2}^2 \rangle |h(\bar{x}, \bar{x}_{i_1})|^4 |h(\bar{x}, \bar{x}_{i_2})|^4 + \sum_{i_1=-\infty}^{\infty} \langle \rho_{i_1}^4 \rangle |h(\bar{x}, \bar{x}_{i_1})|^8
\end{aligned}$$

Rearranging terms so as to replace the diagonal term in each of the inner summations gives

$$\langle I(\bar{x})^2 \rangle = \quad (3.24)$$

$$2 \sum_{i_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \langle \rho_{i_1} \rho_{i_2} \rangle |h(\bar{x}, \bar{x}_{i_1})|^2 |h(\bar{x}, \bar{x}_{i_2})|^2 - \sum_{i_1=-\infty}^{\infty} \langle \rho_{i_1}^2 \rangle |h(\bar{x}, \bar{x}_{i_1})|^4$$

$$\langle I(\bar{x})^3 \rangle = \quad (3.25)$$

$$\begin{aligned}
& 6 \sum_{i_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \sum_{i_3=-\infty}^{\infty} \langle \rho_{i_1} \rho_{i_2} \rho_{i_3} \rangle |h(\bar{x}, \bar{x}_{i_1})|^2 |h(\bar{x}, \bar{x}_{i_2})|^2 |h(\bar{x}, \bar{x}_{i_3})|^2 \\
& - 15 \sum_{i_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \langle \rho_{i_1}^2 \rho_{i_2} \rangle |h(\bar{x}, \bar{x}_{i_1})|^4 |h(\bar{x}, \bar{x}_{i_2})|^2 + 10 \sum_{i_1=-\infty}^{\infty} \langle \rho_{i_1}^3 \rangle |h(\bar{x}, \bar{x}_{i_1})|^6
\end{aligned}$$

$$\langle I(\bar{x})^4 \rangle = \quad (3.26)$$

$$\begin{aligned}
& 24 \sum_{i_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \sum_{i_3=-\infty}^{\infty} \sum_{i_4=-\infty}^{\infty} \langle \rho_{i_1} \rho_{i_2} \rho_{i_3} \rho_{i_4} \rangle |h(\bar{x}, \bar{x}_{i_1})|^2 \dots |h(\bar{x}, \bar{x}_{i_4})|^2 \\
& - 132 \sum_{i_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \sum_{i_3=-\infty}^{\infty} \langle \rho_{i_1}^2 \rho_{i_2} \rho_{i_3} \rangle |h(\bar{x}, \bar{x}_{i_1})|^4 |h(\bar{x}, \bar{x}_{i_2})|^2 |h(\bar{x}, \bar{x}_{i_3})|^2 \\
& + 148 \sum_{i_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \langle \rho_{i_1}^3 \rho_{i_2} \rangle |h(\bar{x}, \bar{x}_{i_1})|^6 |h(\bar{x}, \bar{x}_{i_2})|^2 \\
& + 90 \sum_{i_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \langle \rho_{i_1}^2 \rho_{i_2}^2 \rangle |h(\bar{x}, \bar{x}_{i_1})|^4 |h(\bar{x}, \bar{x}_{i_2})|^4 - 129 \sum_{i_1=-\infty}^{\infty} \langle \rho_{i_1}^4 \rangle |h(\bar{x}, \bar{x}_{i_1})|^8
\end{aligned}$$

The first term in each of (3.24) to (3.26) dominates when the system response is non-zero for a large number of scattering elements, i.e., when each resolution cell contains a large number of scatterers with independent phase. This result can be generalised to all moments to give

$$\begin{aligned}
\langle I(\bar{x})^m \rangle & = m! \sum_{i_1=-\infty}^{\infty} \dots \sum_{i_m=-\infty}^{\infty} \langle \rho_{i_1} \dots \rho_{i_m} \rangle |h(\bar{x}, \bar{x}_{i_1})|^2 \dots |h(\bar{x}, \bar{x}_{i_m})|^2 \\
& = m! \left\langle \left(\sum_{i=-\infty}^{\infty} \rho_i |h(\bar{x}, \bar{x}_i)|^2 \right)^m \right\rangle \quad (3.27)
\end{aligned}$$

The m 'th moment of the detected intensity given by (3.27) is the product of the m 'th moment of the result of imaging the surface cross-section with a system response $|h(\bar{x}, \bar{x}_i)|^2$ (note this is always real and positive), and $m!$ which is the m 'th moment of a unit mean exponential random variable. Thus the m 'th moment of the detected intensity may be expressed by

$$\langle I(\bar{x})^m \rangle = \langle \xi(\bar{x})^m \rangle \langle \Upsilon(\bar{x})^m \rangle \quad (3.28)$$

where ξ is an exponential unit mean random variable and Υ is the imaged SCS defined as

$$\Upsilon(\bar{x}) = \sum_{i=-\infty}^{\infty} \rho_i |h(\bar{x}, \bar{x}_i)|^2 \quad (3.29)$$

3.3.1.1 The exponential distribution

An exponentially distributed random variable x with mean μ_x has pdf

$$f_x(x) = \frac{1}{\mu_x} e^{-\frac{x}{\mu_x}} \quad (3.30)$$

cumulative distribution function

$$p_x(x) = 1 - e^{-\frac{x}{\mu_x}} \quad (3.31)$$

characteristic function

$$\Phi_x(\omega) = (1 - j\omega\mu_x)^{-1} \quad (3.32)$$

moments²

$$\langle x^m \rangle = \Gamma(m+1) \mu_x^m \quad (3.33)$$

variance

$$\sigma_x^2 = \mu_x^2 \quad (3.34)$$

normalised moments

$$x^{(m)} = \Gamma(m+1) = m! \quad (3.35)$$

and coefficient of variation

$$v_x^2 = 1 \quad (3.36)$$

3.3.1.2 The multiplicative imaging model

Because (3.28) applies for all moments we may assume that the detected intensity is given by the product of the imaged SCS and an *independent* unit mean exponentially distributed random variable

$$I(\bar{x}) = \xi(\bar{x}) \Upsilon(\bar{x}) \quad (3.37)$$

The exponential random variable ξ in (3.28) and (3.37) is the result of cancellation of phase terms, or rather coherent interference, and it causes the detected intensity to fluctuate, i.e., it is the process which causes speckle.

² $\Gamma(z)$ is the gamma function.

The representation of the detected intensity as the product of the imaged SCS and an independent exponential unit mean speckle process is often referred to as the multiplicative model for radar imaging [19, 60]. The imaged SCS is itself a random variable, being the product of the mean imaged SCS and a unit mean random process which is due to fluctuations (textures) in the SCS being detected by the imaging system. Fluctuations in intensity are thus the result of both speckle and imaged surface texture and their combined effect is referred to as *clutter*. So defined clutter contains both surface texture *information* and speckle *noise*. Thus, the multiplicative model given by (3.37) can be written more fully as

$$I(\bar{x}) = \xi(\bar{x}) \chi(\bar{x}) \mu_{\Upsilon}(\bar{x}) \quad (3.38)$$

where χ is the unit mean imaged surface texture and μ_{Υ} is the mean imaged SCS which equals the mean intensity μ_I . The clutter Δ is then given by the intensity normalised by its mean

$$\Delta(\bar{x}) = \frac{I(\bar{x})}{\mu_I} = \xi(\bar{x}) \chi(\bar{x}) \quad (3.39)$$

Following from (3.39) the moments of the clutter (the normalised intensity moments) are given by

$$\begin{aligned} \langle \Delta(\bar{x})^m \rangle &= \langle \xi(\bar{x})^m \rangle \langle \chi(\bar{x})^m \rangle = \xi(\bar{x})^{(m)} \Upsilon(\bar{x})^{(m)} \\ &= m! \left\langle \left(\sum_{i=-\infty}^{\infty} \rho_i |h(\bar{x}, \bar{x}_i)|^2 \right)^m \right\rangle / \left(\sum_{i=-\infty}^{\infty} \langle \rho_i \rangle |h(\bar{x}, \bar{x}_i)|^2 \right)^m \end{aligned} \quad (3.40)$$

and are dependent on the moments of the imaged surface texture (normalised moments of the imaged SCS), which in turn are dependent on the correlation properties of the SCS. When the SCS is stationary the moments of the imaged surface texture are given by

$$\begin{aligned} \langle \chi(\bar{x})^m \rangle &= \Upsilon(\bar{x})^{(m)} = 1 + \sum_{i_1=-\infty}^{\infty} \dots \sum_{i_m=-\infty}^{\infty} \left(\frac{\langle \rho_{i_1} \dots \rho_{i_m} \rangle}{\langle \rho \rangle^m} - 1 \right) \times \\ &\quad |h(\bar{x}, \bar{x}_{i_1})|^2 \dots |h(\bar{x}, \bar{x}_{i_m})|^2 / \left(\sum_{i=-\infty}^{\infty} |h(\bar{x}, \bar{x}_i)|^2 \right)^m \end{aligned} \quad (3.41)$$

3.3.2 Exponentially distributed intensity

When the correlation length of the SCS is much less than the width of the system response (resolution of the imaging system), the SCS of a large number of scatterers with non-zero system response will be independent. As the number of independent scatterers with non-zero system response tends to infinity the second term in (3.41) will tend to zero. Thus, over stationary surfaces with a correlation length much less than the resolution of the imaging system the imaged SCS will be constant, i.e., the imaging system detects no surface texture. All clutter will then be purely due to speckle, and the intensity will be

exponentially distributed with mean μ_I and normalised moments given by

$$I^{(m)} = m! \quad (3.42)$$

Because the contrast ($\sqrt{\text{coefficient of variation}}$) within a homogeneous region of exponentially distributed intensity equals one (3.36), the standard deviation of the intensity will equal its mean value. This means that the intensity will on average fluctuate away from its mean value by an amount equal to its mean value. This is often interpreted as a signal to noise ratio of one.

The result that the detected intensity is exponentially distributed when the number of independent scatterers with non-zero system response tends to infinity, may also be reached by applying the **central limit theorem** to the complex components of the detected field [22]. It follows from the results given in Section 3.2, that as the number of independent scatterers with non-zero system response tends to infinity the complex components will become **Gaussian** distributed with zero mean and the same variance. In addition as they are already uncorrelated the complex components will now be independent³. Thus, they will have a joint pdf

$$f_Z(Z_{\Re}, Z_{\Im}) = \frac{1}{2\pi\sigma_Z^2} e^{-\frac{(Z_{\Re}^2 + Z_{\Im}^2)}{2\sigma_Z^2}} \quad (3.43)$$

where σ_Z^2 is the variance of both the real and imaginary components (3.10,3.11). Transforming to polar coordinates the joint pdf of the intensity and phase of the detected field is given by

$$f_Z(I, \theta) = f_Z(\sqrt{I} \cos \theta, \sqrt{I} \sin \theta) |J| = \frac{1}{4\pi\sigma_Z^2} e^{-\frac{I}{2\sigma_Z^2}} \quad (3.44)$$

where $J = 1/2$ is the **Jacobian** of the transform. Integrating over the phase the marginal distribution of the intensity is exponentially distributed as expected

$$f_I(I) = \int_{-\pi}^{\pi} f_Z(I, \theta) d\theta = \frac{1}{\mu_I} e^{-\frac{I}{\mu_I}} \quad (3.45)$$

where $\mu_I = 2\sigma_Z^2$. In addition it can be confirmed that the marginal distribution of the phase is uniformly distributed by integrating over the intensity

$$f_{\theta}(\theta) = \int_0^{\infty} f_Z(I, \theta) dI = \frac{1}{2\pi} \quad (3.46)$$

and that the intensity and phase are independent as the product of their marginal distributions gives their joint distribution

$$f_Z(I, \theta) = f_I(I) f_{\theta}(\theta) \quad (3.47)$$

As illustration of these results Figure 3.3 shows histograms of the intensity, the phase, and the complex component, of pixel values from a homogeneous

³The complex components of a circularly symmetric complex random variable can only be independent if they are Gaussian [48].

region in a SAR image (the agricultural field in the top right hand corner of the image in Figure 3.1). The expected distributions of the data values are represented by solid curves. In each case there is good fit between histogram and expected distribution. This implies that speckle is fully developed in the region and that all clutter is purely due to it with no surface texture contribution.

3.3.3 K-distributed intensity

At the opposite extreme of there being a large number of independent scatterers within a resolution cell, when the correlation length of the surface cross-section is much greater than the system resolution, i.e., when

$$\langle \rho_{i_1} \dots \rho_{i_m} \rangle = \langle \rho^m \rangle \quad (3.48)$$

whenever

$$\langle |h(\bar{x}, \bar{x}_{i_1})|^2 \dots |h(\bar{x}, \bar{x}_{i_m})|^2 \rangle \geq 0 \quad (3.49)$$

the system response terms in (3.41) cancel and the normalised moments of the imaged SCS equal those of the SCS. In this case all surface textures are fully detected by the imaging system and the intensity is proportional to the SCS multiplied by speckle, and has normalised moments

$$I^{(m)} = m! \rho^{(m)} \quad (3.50)$$

In between the extremes described by (3.42) and (3.50) surface textures will be partially detected by the imaging system and the normalised moments of the imaged SCS will lie in the interval

$$1 \leq \Upsilon^{(m)} \leq \rho^{(m)} \quad (3.51)$$

Observation has shown that when surface textures are detected (partially or fully), the detected intensity fits the two parameter K-distribution better than the exponential distribution [31, 43]; i.e., clutter is K-distributed with unit mean. Figure 3.4 shows how the intensity in a region of woodland is better fitted by a K-distribution of order 10 than an exponential distribution.

The intensity K-distribution is defined by its mean μ_I and order parameter ν , and has pdf⁴

$$f_I(I) = \frac{2}{\Gamma(\nu)\beta} \left(\sqrt{\frac{I}{\beta}} \right)^{\nu-1} K_{\nu-1} \left(2\sqrt{\frac{I}{\beta}} \right) \quad (3.52)$$

where $\beta = \mu_I/\nu$, cumulative distribution function [31] (from (F.4))

$$P_I(I) = 1 - \frac{2}{\Gamma(\nu)} \left(\sqrt{\frac{I}{\beta}} \right)^{\nu} K_{\nu} \left(2\sqrt{\frac{I}{\beta}} \right) \quad (3.53)$$

moments (from (F.5))

$$\langle I^m \rangle = m! \frac{\Gamma(m+\nu)}{\Gamma(\nu)} \beta^m \quad (3.54)$$

⁴ $K_{\nu}(x)$ is the ν 'th-order modified Bessel function of the second kind.

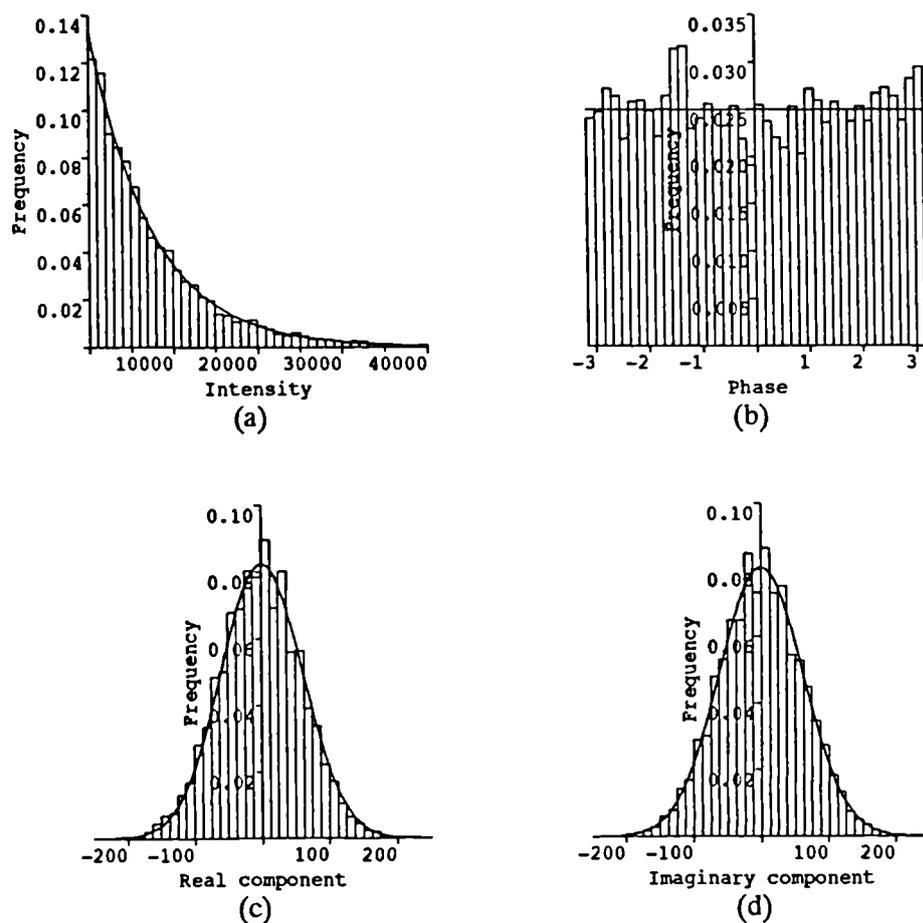


Figure 3.3: Histograms showing the distribution of the (a) intensity, (b) phase, (c) real component and (d) imaginary component, of pixel values from a homogeneous region (80×80 pixels) in the image shown in Figure 3.1. The curves represent expected distributions fitted to the data, (a) exponential, (b) uniform, and (c) and (d) zero mean Gaussian.

normalised moments

$$I^{(m)} = m! \frac{\Gamma(m + \nu)}{\Gamma(\nu) \nu^m} = m! \left(1 + \frac{1}{\nu}\right) \dots \left(1 + \frac{m-1}{\nu}\right) \quad (3.55)$$

variance

$$\sigma_I^2 = \mu_I^2 \left(1 + \frac{2}{\nu}\right) \quad (3.56)$$

and coefficient of variation

$$\nu_I^2 = 1 + \frac{2}{\nu} \quad (3.57)$$

Figure 3.5 shows how the shape of a K-distribution changes with order parameter. The mode of a K-distribution occurs at zero and is given by $f_I(0) = (\mu(1 - 1/\nu))^{-1}$; the mode is infinite when $\nu = 1$. It has already been pointed out that when intensity is exponentially distributed it will on average fluctuate away from its mean value by an amount equal to its mean value. From (3.57) it is evident that the intensity will on average fluctuate away from its mean value by an amount significantly greater than its mean value when it is K-distributed with a small order parameter, i.e., image contrast increases as the amount of imaged surface texture increases.

3.3.3.1 Gamma distributed imaged SCS

It is clear from the multiplicative model (3.37) and the moments of the K-distribution (3.54), that the detected intensity will be K-distributed with mean μ_I and order parameter ν when the imaged SCS has m 'th moment

$$\langle \Upsilon^m \rangle = \frac{\Gamma(m + \nu)}{\Gamma(\nu)} \beta^m \quad (3.58)$$

where $\beta = \mu_I/\nu$. This is the m 'th moment of a gamma distributed random variable with the same mean and order parameter as the K-distribution. Thus, the detected intensity will be K-distributed when the imaged surface texture and thus the imaged SCS is gamma distributed.

A gamma distributed random variable x with mean μ_x and order parameter ν has pdf

$$f_x(x) = \frac{e^{-\frac{x}{\beta}} x^{\nu-1}}{\Gamma(\nu) \beta^\nu} \quad (3.59)$$

where $\beta = \mu_x/\nu$, cumulative distribution function

$$p_x(x) = 1 - e^{-\frac{x}{\beta}} \sum_{k=0}^{\nu-1} \frac{\left(\frac{x}{\beta}\right)^k}{k!} \quad (3.60)$$

characteristic function

$$\Phi_x(\omega) = (1 - j\omega\beta)^{-\nu} \quad (3.61)$$

moments

$$\langle x^m \rangle = \frac{\Gamma(m + \nu)}{\Gamma(\nu)} \beta^m \quad (3.62)$$

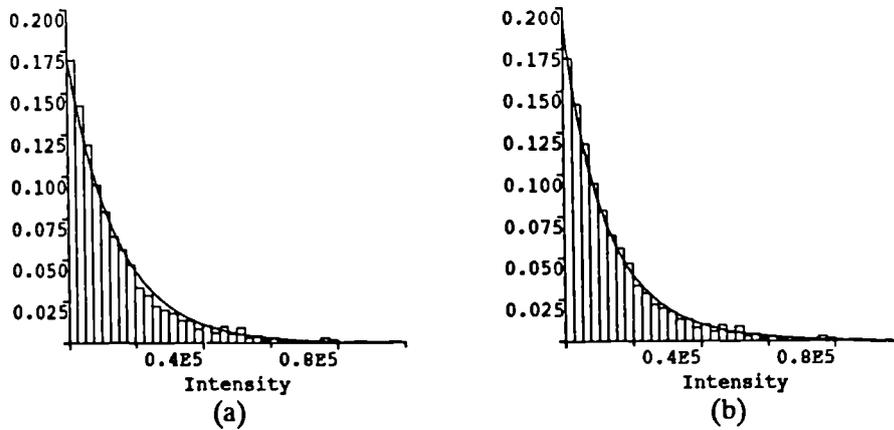


Figure 3.4: Histogram of the distribution of the intensity in a homogeneous region of woodland (bottom left of the image shown in Figure 3.1) fitted to (a) an exponential distribution, and (b) a K-distribution of order 10. The K-distribution gives the better fit even though there is only a small difference between the two fitted distributions. Both distributions have the same mean as the data but the variance of the K-distribution is 1.2 times greater than that of the exponential distribution. The goodness of fit was confirmed using the **Kolmogorov-Smirnov test** which gave values of 0.36 for the exponential distribution and 0.84 for the K-distribution.

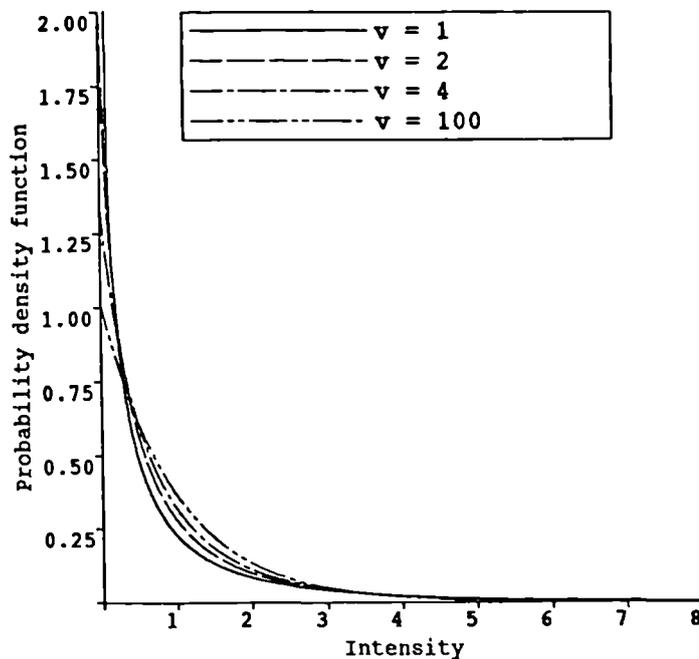


Figure 3.5: Pdf's of K-distributions with unit mean and order parameters $\nu = 1, 2, 4$ and 100 (i.e., variances 3, 2, 1.5 and 1.02 respectively). On this linear scale a K-distribution with order 100 or higher is virtually indistinguishable from an exponential distribution with the same mean value.

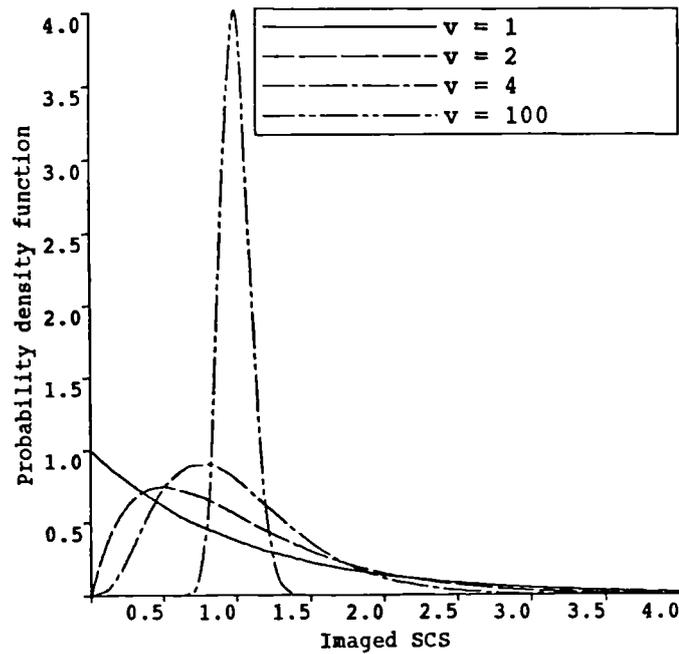


Figure 3.6: Pdf's of gamma distributions for the imaged SCS corresponding to the K-distributions for the intensity shown in Figure 3.5.

normalised moments

$$x^{(m)} = \frac{\Gamma(m + \nu)}{\Gamma(\nu) \nu^m} \quad (3.63)$$

variance

$$\sigma_x^2 = \frac{\mu^2}{\nu} \quad (3.64)$$

and coefficient of variation

$$v_x^2 = \frac{1}{\nu} \quad (3.65)$$

When $\nu = 1$ the gamma distribution reduces to the exponential distribution. Thus the product of two exponential random variables gives a K-distributed random variable of order one. As the order parameter tends to infinity the gamma distribution tends to a Gaussian distribution with mean μ_x and zero variance. Therefore a K-distribution is equivalent to an exponential distribution when its order parameter is large. This accords with saying that the detected intensity is exponential when the imaged SCS is constant (gamma distributed with large order parameter). Figure 3.6 shows the gamma distributions for the imaged SCS corresponding to the K-distributions for the intensity shown in Figure 3.5.

Most proofs that the product of an exponentially distributed random variable x_1 with unit mean, and an independent gamma distributed random variable x_2 with mean μ and order parameter ν , is K-distributed with the same mean and order parameter as x_2 , are based on comparing moments [29, 31, 44]. It is

possible to prove this directly using pdf's. Substituting the pdf's of x_1 and x_2 given by (3.30) and (3.59) into (A.5) their product y has pdf

$$f_y(y) = \int_0^\infty e^{-\frac{y}{x}} \frac{e^{-\frac{x}{\beta}} x^{\nu-1}}{\Gamma(\nu)\beta^\nu} \frac{1}{x} dx \quad (3.66)$$

where $\beta = \mu/\nu$. Letting $z = \ln \frac{x}{\sqrt{y\beta}}$

$$f_y(y) = \frac{2}{\Gamma(\nu)\beta} \left(\sqrt{\frac{y}{\beta}}\right)^{\nu-1} \int_0^\infty e^{-2\sqrt{\frac{y}{\beta}} \cosh z} \cosh((\nu-1)z) dz \quad (3.67)$$

which by (F.2) gives

$$f_y(y) = \frac{2}{\Gamma(\nu)\beta} \left(\sqrt{\frac{y}{\beta}}\right)^{\nu-1} K_{\nu-1} \left(2\sqrt{\frac{y}{\beta}}\right) \quad (3.68)$$

the required pdf.

3.3.3.2 Justification for a gamma distributed imaged SCS

To determine why the imaged SCS should be gamma distributed we need to examine how such distributions arise. Gamma distributions have been shown to arise as limiting cases of two random walk mechanisms [29, 30]

1. A one dimensional random walk with a fluctuating number of steps.
2. A multi-dimensional random walk with a constant number of steps.

The conditions under which these random walks generate gamma distributed outputs will be described before going on to show how they may be used to justify treating the imaged SCS as being gamma distributed.

The output of the one dimensional random walk

$$X = \sum_{k=1}^n x_k \quad (3.69)$$

will be gamma distributed with order parameter ν when the contributions x_k are statistically identical and independent, and the number of steps n has a negative binomial distribution given by

$$f_n = \binom{n+\nu-1}{n} \frac{\left(\frac{\mu_n}{\nu}\right)^n}{\left(1+\frac{\mu_n}{\nu}\right)^{n+\nu}} \quad (3.70)$$

where μ_n the mean number of steps is large. Such a distribution for n will arise when the number of steps is a Markov process controlled by the birth-death-immigration equation [3, 29]

$$\frac{df_n}{dt} = D(n+1)f_{n+1} - ((B+D)n+I)f_n + (B(n-1)+I)f_{n-1} \quad (3.71)$$

where \mathcal{B} is the birth rate, \mathcal{D} is the death rate and \mathcal{I} is the spontaneous immigration rate. An equilibrium solution to (3.71) exist when the death rate is greater than the birth rate, in which case f_n tends to a negative binomial distribution with mean $\mu_n = \mathcal{I}/(\mathcal{D} - \mathcal{B})$ and order parameter $\nu = \mathcal{I}/\mathcal{B}$ as $t \rightarrow \infty$. The number of steps will then have variance

$$\sigma_n^2 = \mu_n^2 \left(\frac{1}{\mu_n} + \frac{1}{\nu} \right) \quad (3.72)$$

Thus when the mean number of steps is large (i.e., $\mathcal{I} \gg \mathcal{D} - \mathcal{B}$) the order parameter ν is inversely proportional to the variance in the number of steps and the variance tends to zero as $\nu \rightarrow \infty$ (i.e., $\mathcal{I} \gg \mathcal{B}$).

It follows from the convolution theorem that the output of the one-dimensional random walk defined by (3.69) has characteristic function

$$\Phi_X(\omega) = \Phi_x(\omega)^n \quad (3.73)$$

where $\Phi_x(\omega) = \langle e^{j\omega x} \rangle$ is the characteristic function of the contributions x_k . When n has a negative binomial distribution with mean μ_n and order parameter ν (3.70), on average the characteristic function of the output is given by

$$\Phi_X(\omega) = \left[1 + \frac{\mu_n}{\nu} (1 - \langle e^{j\omega x} \rangle) \right]^{-\nu} \quad (3.74)$$

since for arbitrary u

$$\langle u^n \rangle = \sum_{n=0}^{\infty} u^n f_n = \left[1 + \frac{\mu_n}{\nu} (1 - u) \right]^{-\nu} \quad (3.75)$$

Normalising x by μ_n and substituting the series expansion for $e^{j\omega x}$ (F.7) into (3.74) gives

$$\Phi_X(\omega) = \left[1 - \frac{j\omega \langle x \rangle}{\nu} - \sum_{k=2}^{\infty} \frac{j^k \omega^k \langle x^k \rangle}{\nu \mu_n^{k-1} k!} \right]^{-\nu} \quad (3.76)$$

As the mean number of steps μ_n tends to infinity the characteristic function tends to

$$\Phi_X(b) \rightarrow \left[1 - \frac{j\omega \langle x \rangle}{\nu} \right]^{-\nu} \quad (3.77)$$

This is the characteristic function of a gamma distributed random variable with mean $\langle x \rangle$ and order parameter ν (3.61).

The negative binomial distribution is in fact the discrete analogue of the gamma distribution, and it has been shown that the output of a *continuous* birth-death-immigration process, with a death rate greater than the birth rate, tends to being gamma distributed as $t \rightarrow \infty$ [30]. The mean and order parameter of the gamma distribution are then defined in the same way as the mean and order parameter of the negative binomial distribution in the discrete case.

We now turn to the second random walk mechanism. Let z_1, \dots, z_n be 2ν dimensional vectors each with statistically identical but uncorrelated zero mean

components. The inner product $Y = W \cdot W$ of the output of the random walk in 2ν dimensions

$$W = \sum_{k=1}^n z_k \quad (3.78)$$

will be gamma distributed with order parameter ν when n is large but not fluctuating. It follows from the central limit theorem that the output W of the random walk will have a joint-Gaussian distribution when n is constantly large relative to the number of correlated contributions in each dimension. Each of the Gaussian components will then be independent with zero mean and the same variance. The inner product Y is given by the sum of the squares of these components. The square of a Gaussian distributed random variable with zero mean and variance σ^2 is gamma distributed with mean σ^2 and order parameter one half (3.59), and has a characteristic function given by (3.61). Therefore, from the convolution theorem the characteristic function of Y , the sum of 2ν such variables, is also gamma distributed but with mean $2\nu\sigma^2$ and order parameter ν .

The higher order statistics of the two gamma distributed outputs X and Y will be the same if their autocorrelation functions (ACF's) are equal [30]. The underlying random walk mechanisms will then be indistinguishable in terms of the statistics of X and Y . When the number of steps contributing to the one-dimensional random walk is controlled by the birth-death-immigration process defined by (3.71) its output X will have a Lorentzian spectrum and corresponding ACF given by [29]

$$\frac{\langle X(t_1) X(t_2) \rangle}{\langle X \rangle^2} = 1 + \frac{B}{I} e^{-(D-B)|t_1-t_2|} \quad (3.79)$$

The correlation length of the output may then be defined as $1/(D-B)$ and the correlation coefficient will tend to zero when $|t_1 - t_2| \gg 1/(D-B)$ or when $I \gg B$ (i.e., $\nu \rightarrow \infty$).

Both of the above random walk mechanisms have been used to justify treating the imaged SCS as being gamma distributed. Approximating the system response $|h(\bar{x}, \bar{x}_i)|^2$ to the SCS by a rectangular function in both azimuth and range, the imaged SCS (3.29) can be represented as a sum of n statistically identical and independent SCS contributions. This rectangular function may be taken as representing the main lobe of the system response. When the number of contributions is controlled by the birth-death-immigration process (3.71), the imaged SCS over an extended target will be gamma distributed. It is also possible using the same approach to prove that the intensity resulting from n elementary scatterers is K-distributed without having to represent the intensity as the product of an exponential speckle process and the imaged SCS [29]. This involves expressing the characteristic function of the sum of the circularly symmetric scatterer contributions in terms of the Hankel transform.

Alternatively, each scattering contributions can be viewed as the result of a random walk in 2ν real dimensions (ν complex dimensions) [43]. The SCS will then be gamma distributed and the imaged (i.e., integrated) SCS will be

constant or gamma distributed according to whether the correlation length of the SCS is much less than or on a par with the width of the system response [30, 43].

It should be noted that the above random walk mechanisms provide a purely phenomenological justification for a gamma distributed imaged SCS. The mechanisms still have to be related to the physical world. This is particularly problematic when the correlations between multi-channel datasets are considered, as all the physics is pushed onto the single elementary scatterers [64].

3.3.3.3 K-distributed field statistics

We have not yet considered the distribution of the complex components of the detected field when the intensity is K-distributed. The complex components are of course zero mean Gaussian distributed when the intensity is exponentially distributed (see Section 3.3.2). If the field Z is a circularly symmetric complex random variable which is K-distributed in intensity with mean μ and order parameter ν , the joint pdf of the field in terms of the intensity and phase, which are independent (see Section 3.2), is given by

$$f_Z(I, \theta) = f_I(I) f_\theta(\theta) = \frac{1}{\pi\beta\Gamma(\nu)} \left(\sqrt{\frac{I}{\beta}} \right)^{\nu-1} K_{\nu-1} \left(2\sqrt{\frac{I}{\beta}} \right) \quad (3.80)$$

where $\beta = \mu/\nu$. Transforming to Cartesian coordinates, in terms of its real and imaginary components Z_{\Re} and Z_{\Im} the joint pdf of the field is given by

$$f_Z(Z_{\Re}, Z_{\Im}) = \frac{2}{\pi\beta\Gamma(\nu)} \left(\sqrt{\frac{Z_{\Re}^2 + Z_{\Im}^2}{\beta}} \right)^{\nu-1} K_{\nu-1} \left(2\sqrt{\frac{Z_{\Re}^2 + Z_{\Im}^2}{\beta}} \right) \quad (3.81)$$

Integrating over the imaginary component the marginal pdf of the real component is given by⁵

$$f_{Z_{\Re}}(Z_{\Re}) = \int_{-\infty}^{\infty} f_Z(Z_{\Re}, Z_{\Im}) dZ_{\Im} = \frac{2}{\sqrt{\pi}\beta\Gamma(\nu)} \left(\frac{|Z_{\Re}|}{\sqrt{\beta}} \right)^{\nu-\frac{1}{2}} K_{\nu-\frac{1}{2}} \left(\frac{2|Z_{\Re}|}{\sqrt{\beta}} \right) \quad (3.82)$$

This pdf is plotted for various values of ν in Figure 3.7. The pdf is symmetric about zero and becomes sharper as the order parameter decreases. The marginal pdf of the imaginary component given by integrating the joint pdf over the real component, is identical.

Figure 3.8 shows how the real and imaginary components in the region of woodland whose intensity distribution was better fitted by a K-distribution than an exponential distribution in Figure 3.4, is fitted by a zero mean Gaussian distribution and the distribution given by (3.82). Whilst the K-distribution clearly gave the better fit to the intensity distribution it is unclear whether (3.82) gives a better fit to the complex components than the zero mean Gaussian distribution. Using the Kolmogorov-Smirnov test the real and imaginary components gave respective fits of 0.16 and 0.46 to (3.82) whilst they gave respective fits of 0.15

⁵This integral was initially solved using the symbolic computation program Mathematica.

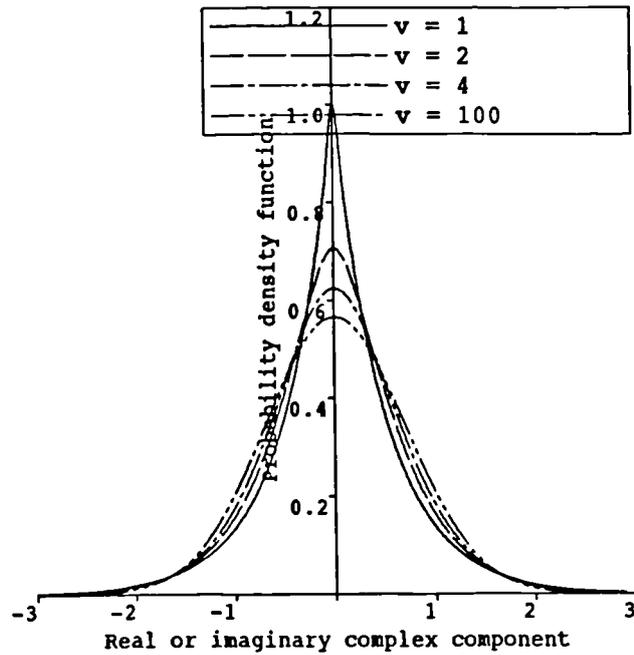


Figure 3.7: Pdf's of the real or imaginary component of a circularly symmetric complex random variable which is K-distributed in intensity, as shown in Figure 3.5, with unit mean and order parameter $\nu = 1, 2, 4, 100$.

and 0.32 to the zero mean Gaussian distribution. This implies that the complex components are better fitted by (3.82) but only slightly so. However, a full series of measurements is needed to properly check the validity of (3.82) as the pdf of the real and imaginary components of data whose intensity distribution fits a K-distribution.

Note that for integer order parameters the pdf (3.82) is most easily calculated using the series expansion for $K_{n-\frac{1}{2}}(z)$ given by (F.6) where n is an integer. The pdf is then given by

$$f_{Z_{\Re}}(Z_{\Re}) = \frac{1}{\sqrt{\beta}\Gamma(\nu)} e^{-\frac{2|Z_{\Re}|}{\sqrt{\beta}}} \sum_{k=0}^{\nu-1} \frac{\Gamma(\nu+k)}{\Gamma(\nu-k)\Gamma(k+1)} \frac{1}{4^k} \left(\frac{|Z_{\Re}|}{\sqrt{\beta}}\right)^{\nu-1-k} \quad (3.83)$$

Written in this form it can be seen that the pdf is exponential when $\nu = 1$

$$f_{Z_{\Re}}(Z_{\Re}) = \frac{1}{\sqrt{\mu}} e^{-\frac{2|Z_{\Re}|}{\sqrt{\mu}}} \quad (3.84)$$

and it can also be seen that the mode of the pdf which occurs at zero equals

$$f_{Z_{\Re}}(0) = \frac{1}{\sqrt{\beta}} \frac{\Gamma(2\nu-1)}{\Gamma(\nu)^2 4^{\nu-1}} \quad (3.85)$$

When the pdf of the real component is given by (3.82) its moments may be

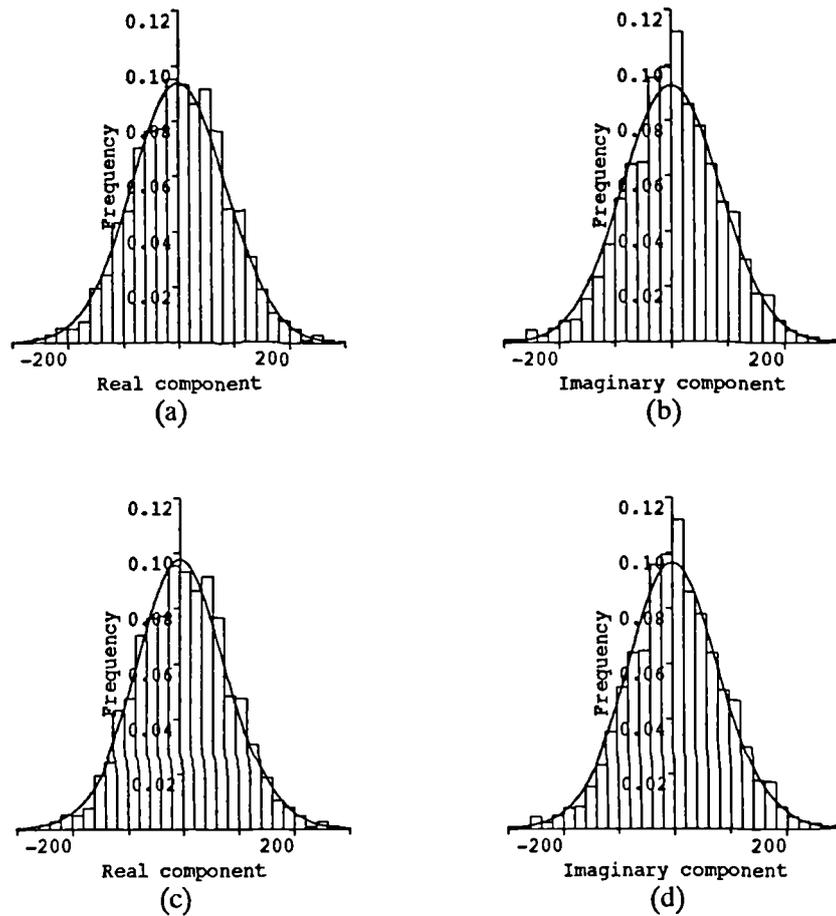


Figure 3.8: Histograms showing the distribution of the real and imaginary components of pixel values in a homogeneous region of woodland (the same as that in Figure 3.4) fitted to (a) and (b) a zero mean Gaussian distribution, and (c) and (d) the distribution given by (3.82) with order parameter $\nu = 10$.

calculated using (F.5) and are given by

$$\langle Z_{\mathfrak{R}}^m \rangle = \begin{cases} \frac{\Gamma(\frac{m+1}{2}) \Gamma(\frac{m}{2} + \nu)}{\Gamma(\frac{1}{2}) \Gamma(\nu)} \beta^{\frac{m}{2}} & m \text{ even} \\ 0 & m \text{ odd} \end{cases} \quad (3.86)$$

These are the moments of the product of a Gaussian distributed random variable with zero mean and variance one half (F.9), and the square root of an independent gamma distributed random variable with order parameter ν (3.100). Substituting the pdf's of these random variables given by (F.8) and (3.98) into (A.5) and using (F.2) it may be shown that (3.82) defines the pdf of such a product. From this we can infer that when the intensity is K-distributed the detected field is the product of a jointly Gaussian speckle process and the square root of the gamma distributed imaged SCS. As we would expect from Section 3.3.2 the moments of the real and imaginary components tend to those of a zero mean Gaussian distribution with variance $\mu/2$ as the imaged SCS tends to a constant ($\nu \rightarrow \infty$).

The real and imaginary components of the detected field are uncorrelated (3.12). However, they will only be independent if the imaged SCS is constant. This can be seen from the higher joint moments given by

$$\langle Z_{\mathfrak{R}}^m Z_{\mathfrak{I}}^n \rangle = \begin{cases} \frac{\Gamma(\frac{m+1}{2}) \Gamma(\frac{n+1}{2}) \Gamma(\frac{m+n}{2} + \nu)}{\Gamma(\frac{1}{2}) \Gamma(\frac{1}{2}) \Gamma(\nu)} \beta^{\frac{m+n}{2}} & m \text{ and } n \text{ both even} \\ 0 & \text{otherwise} \end{cases} \quad (3.87)$$

For independence $\langle Z_{\mathfrak{R}}^m Z_{\mathfrak{I}}^n \rangle$ must equal $\langle Z_{\mathfrak{R}}^m \rangle \langle Z_{\mathfrak{I}}^n \rangle$. This will only occur when the order parameter tends to infinity.

3.4 Statistics of the amplitude of the detected field

SAR image analysis is often performed on amplitude (square root intensity) data rather than intensity data. When the intensity of the detected field accords with the multiplicative model (3.38), it follows that the amplitude of the detected field is given by the product of the square root of a unit mean exponential speckle process and the square root of the imaged SCS

$$A = \sqrt{I} = \sqrt{\xi} \sqrt{\Upsilon} = \sqrt{\xi} \sqrt{\chi} \sqrt{\mu_{\Upsilon}} \quad (3.88)$$

As the intensity speckle process ξ is exponentially distributed with unit mean, the amplitude speckle process $\sqrt{\xi}$ will be Rayleigh distributed with mean $\sqrt{\pi/4}$.

3.4.1 The Rayleigh distribution

If x has pdf $f_x(x)$ its square root $y = \sqrt{x}$ has pdf

$$f_y(y) = f_x(y^2) \frac{dx}{dy} \quad (3.89)$$

Thus the square root y of an exponentially distributed random variable x with mean μ_x has pdf

$$f_y(y) = \frac{2ye^{-\frac{y^2}{\mu_x}}}{\mu_x} \quad (3.90)$$

This is the Rayleigh distribution. When y is Rayleigh distributed it has cumulative distribution function

$$p_y(y) = 1 - e^{-\frac{y^2}{\mu_x}} \quad (3.91)$$

characteristic function (see Appendix B for proof)

$$\Phi_y(\omega) = 1 + j\frac{\omega\sqrt{\pi\mu_x}}{2}e^{-\frac{\omega^2\mu_x}{4}} - \frac{\omega\mu_x}{2}e^{-\frac{\omega^2\mu_x}{4}} \int_0^\omega e^{\frac{z^2\mu_x}{4}} dz \quad (3.92)$$

moments

$$\langle y^m \rangle = \Gamma\left(\frac{m}{2} + 1\right) \mu_x^{\frac{m}{2}} \quad (3.93)$$

mean¹

$$\mu_y = \sqrt{\frac{\pi\mu_x}{4}} \approx 0.886\sqrt{\mu_x} \quad (3.94)$$

variance

$$\sigma_y^2 = \left(1 - \frac{\pi}{4}\right) \mu_x \approx 0.215\sigma_x \quad (3.95)$$

normalised moments

$$y^{(m)} = \frac{\Gamma\left(\frac{m}{2} + 1\right)}{\Gamma\left(\frac{3}{2}\right)^m} \quad (3.96)$$

and coefficient of variation

$$v_y^2 = \frac{4}{\pi} - 1 \approx 0.273 \quad (3.97)$$

3.4.2 Constant imaged SCS: Rayleigh distributed amplitude

When the imaged SCS is constant, from (3.88) the amplitude of the detected field will be Rayleigh distributed with mean $\sqrt{\pi\mu_I/4}$. An amplitude image and the histogram of pixel amplitude values in a homogeneous area, along with expected Rayleigh distribution, are shown in Figure 3.9. The histogram and distribution give a good fit to each other.

The coefficient of variation of a Rayleigh distributed random variable (3.97), is approximately four times smaller than that of an exponential random variable (3.36). Thus the contrast of SAR amplitude data will be approximately half that of the corresponding intensity data. This is why amplitude data is often preferred to intensity data for display purposes as the former will have a smaller dynamic range. This can be seen in the images displayed in Figures 3.1 and 3.9. Also, comparing the plots of the intensity and amplitude distribution in Figures 3.3 and 3.9, it can be seen that unlike the exponential distribution,

¹ $\Gamma\left(\frac{1}{2}\right) = \sqrt{\pi}$

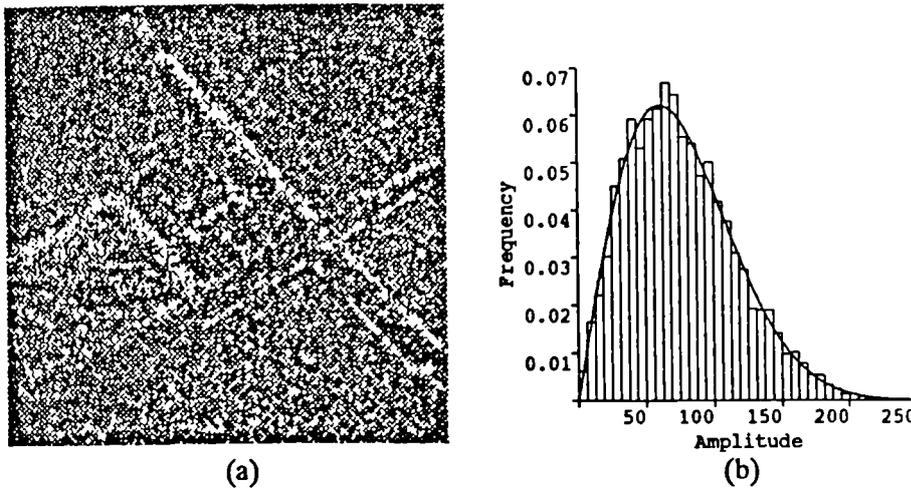


Figure 3.9: (a) the amplitude of the SAR image whose intensity is shown in Figure 3.1, and (b) the histogram of pixel amplitude values in a homogeneous area (the same as that in Figure 3.3) with expected Rayleigh distribution shown as a solid curve.

the amplitude distribution is bell shaped, although asymmetrically so. Thus image analysis tools which assume data to have a symmetric bell shaped distribution, e.g., a Gaussian distribution, often give better, though not necessarily acceptable results, when applied to amplitude as compared to intensity data.

3.4.3 Gamma distributed SCS

If the imaged SCS is gamma distributed the amplitude will be given by the product of a Rayleigh distributed speckle process and the square root of the gamma distributed imaged SCS.

3.4.3.1 Square root gamma distribution

The pdf of the square root y of a gamma distributed random variable x with mean μ_x and order parameter ν , is given by substituting its pdf given by (3.59) into the transform for the pdf of the square root of a random variable (3.89). y then has pdf

$$f_y(y) = \frac{2e^{-\frac{y^2}{\beta}} y^{2\nu-1}}{\Gamma(\nu) \beta^\nu} \quad (3.98)$$

where $\beta = \mu_x/\nu$, cumulative distribution function

$$p_y(y) = 1 - e^{-\frac{y^2}{\beta}} \sum_{k=0}^{\nu-1} \frac{\left(\frac{y^2}{\beta}\right)^k}{k!} \quad (3.99)$$

moments

$$\langle y^m \rangle = \frac{\Gamma(\nu + \frac{m}{2})}{\Gamma(\nu)} \beta^{\frac{m}{2}} \quad (3.100)$$

mean

$$\mu_y = \sqrt{\frac{\mu_x \Gamma(\nu + \frac{1}{2})^2}{\Gamma(\nu + 1) \Gamma(\nu)}} \quad (3.101)$$

variance

$$\sigma_y^2 = \mu_x \left(1 - \frac{\Gamma(\nu + \frac{1}{2})^2}{\Gamma(\nu + 1) \Gamma(\nu)} \right) \quad (3.102)$$

normalised moments

$$y^{(m)} = \frac{\Gamma(\nu + \frac{m}{2}) \Gamma(\nu)^{m-1}}{\Gamma(\nu + \frac{1}{2})^m} \quad (3.103)$$

and coefficient of variation

$$v_y^2 = \frac{\Gamma(\nu + 1) \Gamma(\nu)}{\Gamma(\nu + \frac{1}{2})^2} - 1 \quad (3.104)$$

As setting $\nu = 1$ in a gamma distribution gives the exponential distribution, setting $\nu = 1$ in the square root gamma distribution gives the Rayleigh distribution.

3.4.3.2 K-distributed amplitude

The distribution of the amplitude of data which has a gamma distributed imaged SCS, may be determined by substituting the pdf of its K-distributed intensity (3.52) into (3.89), or by substituting the pdf's of the Rayleigh speckle process (3.90) and the square root of the gamma distributed imaged SCS (3.98) into the integral for the pdf of the product of two random variables (A.5) and solving. Either way, data which is K-distributed in intensity with mean μ_I and order parameter ν has amplitude pdf

$$f_A(A) = \frac{4}{\Gamma(\nu) \sqrt{\beta}} \left(\frac{A}{\sqrt{\beta}} \right)^\nu K_{\nu-1} \left(\frac{2A}{\sqrt{\beta}} \right) \quad (3.105)$$

where $\beta = \mu_I/\nu$, cumulative distribution function

$$p_A(A) = 1 - \frac{2}{\Gamma(\nu)} \left(\frac{A}{\sqrt{\beta}} \right)^\nu K_\nu \left(\frac{2A}{\sqrt{\beta}} \right) \quad (3.106)$$

moments

$$\langle A^m \rangle = \frac{\Gamma(1 + \frac{m}{2}) \Gamma(\nu + \frac{m}{2})}{\Gamma(\nu)} \beta^{\frac{m}{2}} \quad (3.107)$$

mean

$$\mu_A = \sqrt{\frac{\pi \mu_I}{4} \frac{\Gamma(\nu + \frac{1}{2})^2}{\Gamma(\nu + 1) \Gamma(\nu)}} \quad (3.108)$$

variance

$$\sigma_A^2 = \mu_I \left(1 - \frac{\pi}{4} \frac{\Gamma(\nu + \frac{1}{2})^2}{\Gamma(\nu + 1)\Gamma(\nu)} \right) \quad (3.109)$$

normalised moments

$$A^{(m)} = \frac{\Gamma(1 + \frac{m}{2}) \Gamma(\nu + \frac{m}{2}) \Gamma(\nu)^{m-1}}{\Gamma(\frac{3}{2})^m \Gamma(\nu + \frac{1}{2})^m} \quad (3.110)$$

and coefficient of variation

$$v_A^2 = \frac{4}{\pi} \frac{\Gamma(\nu + 1)\Gamma(\nu)}{\Gamma(\nu + \frac{1}{2})^2} - 1 \quad (3.111)$$

Figure 3.10 shows how the amplitude in the region of woodland whose intensity distribution was better fitted by a K-distribution than an exponential distribution in Figure 3.4, is better fitted by an amplitude K-distribution than a Rayleigh distribution. Examples of amplitude K-distributions with different order parameters are shown in Figure 3.11. The asymmetry of the pdf increases as the order parameter decreases.

3.5 Transforming multiplicative speckle to being additive

As already mentioned the speckle found in SAR intensity and amplitude images is a multiplicative process (3.38,3.88). However, many image processing operations assume that any noise in an image is additive and perform poorly otherwise, e.g., edge detection using the Sobel operator (see Chapter 7). The problems this entails for analysing SAR intensity and amplitude images containing multiplicative speckle, can be circumvented by taking the log of the image prior to applying other operations, this transforms the speckle to being additive. Taking the log of the intensity described by (3.38) gives

$$\log I = \log \xi + \log \Upsilon = \log \xi + \log \chi + \log \mu_{\Upsilon} \quad (3.112)$$

where ξ is a unit mean exponential speckle process, Υ is the imaged SCS and χ is the imaged surface texture.

The natural log of the intensity of the detected field $D = \ln I$, is sometimes referred to as the density of the detected field [2]. Because speckle density is additive the shape of the density distribution is the same regardless of its mean value, i.e., the central moments of the distribution are independent of the mean. The only effect of the mean on the distribution is to shift the mode by a factor equal to the mean.

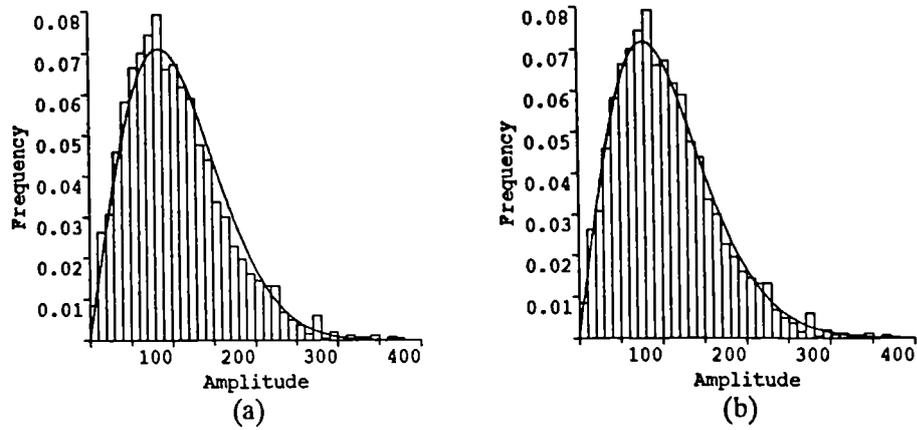


Figure 3.10: Histogram of the distribution of the amplitude in a homogeneous region of woodland (the same as that in Figure 3.4) fitted to (a) a Rayleigh distribution, and (b) an amplitude K-distribution of order 10.

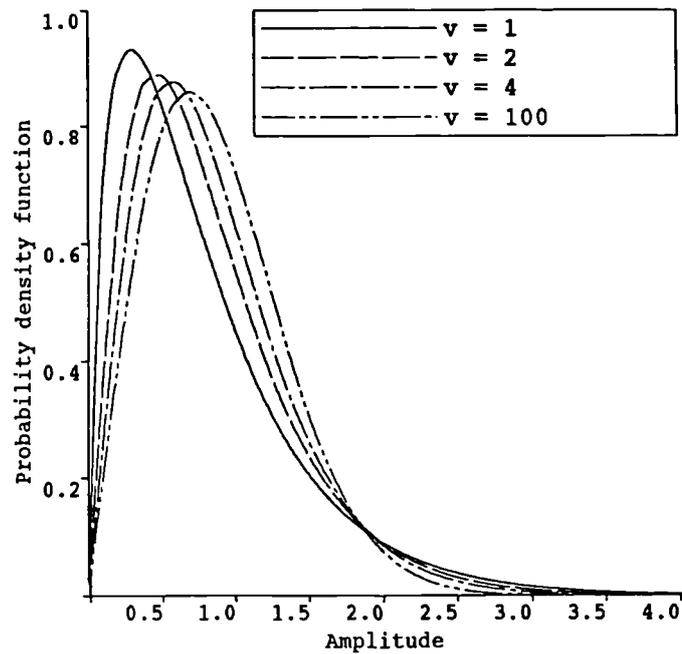


Figure 3.11: Amplitude pdf's corresponding to the K-distributed intensity pdf's shown in Figure 3.5.

3.5.1 Constant imaged SCS: Fischer-Tippett distributed density

If a random variable x has pdf $f_x(x)$ its natural log $y = \ln x$ will have pdf

$$f_y(y) = f_x(e^y) e^y \quad (3.113)$$

Therefore, when the intensity is exponentially distributed (3.30) with mean μ_I its density will have a pdf

$$f_D(D) = e^{-e^{D-\ln \mu_I}} e^{D-\ln \mu_I} \quad (3.114)$$

This is the Fischer-Tippett distribution [1, 2]. The density then has cumulative distribution function

$$p_D(D) = 1 - e^{-e^{D-\ln \mu_I}} \quad (3.115)$$

characteristic function

$$\Phi_D(\omega) = \Gamma(j\omega + 1) \mu_I^{j\omega} \quad (3.116)$$

mean¹

$$\mu_D = \ln \mu_I - \gamma_E \quad (3.117)$$

variance

$$\sigma_D^2 = \frac{\pi^2}{6} \quad (3.118)$$

and coefficient of variation

$$v_D^2 = \frac{\pi^2}{6 (\ln \mu_I - \gamma_E)^2} \quad (3.119)$$

The distribution of the density speckle process $\ln \xi$ is given by setting $\ln \mu_I = 0$ in (3.114). A density image and the histogram of pixel density values in a homogeneous area, along with expected Fischer-Tippett distribution are shown in Figure 3.12.

3.5.2 Gamma distributed SCS

When the SCS is gamma distributed (3.59) with mean $\mu_\Gamma = \mu_I$ and order parameter ν , from (3.113) its natural log will have pdf [2]

$$f_z(z) = \frac{e^{-e^{z-\alpha}} e^{(z-\alpha)\nu}}{\Gamma(\nu)} \quad (3.120)$$

where $\alpha = \ln \beta$ and $\beta = \mu_I/\nu$, cumulative distribution function (from (3.60))

$$p_z(z) = 1 - e^{-e^{z-\alpha}} \sum_{k=0}^{\nu-1} \frac{e^{k(z-\alpha)}}{k!} \quad (3.121)$$

characteristic function

$$\Phi_z(\omega) = \frac{\Gamma(j\omega + \nu)}{\Gamma(\nu)} \beta^{j\omega} \quad (3.122)$$

¹ $\gamma_E = 0.57722\dots$ is Euler's constant.

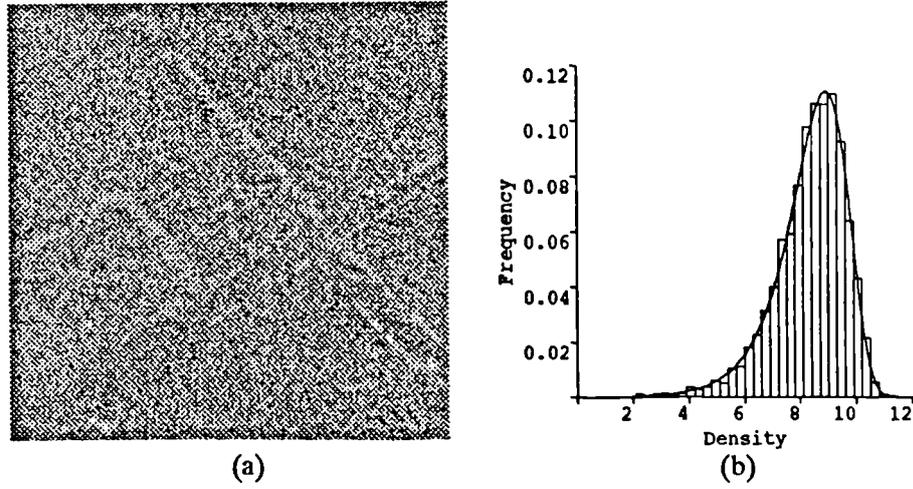


Figure 3.12: (a) the density (log) of the SAR image shown in Figure 3.1, and (b) the histogram of pixel density values in a homogeneous area (the same as that in Figure 3.3) with expected Fischer-Tippett distribution (3.114) shown as a solid curve.

mean²

$$\mu_z = \ln \mu_I - \ln \nu + \psi(\nu) \quad (3.123)$$

variance

$$\sigma_z^2 = \psi'(\nu) \quad (3.124)$$

and coefficient of variation

$$v_z^2 = \frac{\psi'(\nu)}{(\ln \mu_I - \ln \nu + \psi(\nu))^2} \quad (3.125)$$

(3.120) reduces to the Fisher-Tippett distribution when $\nu = 1$.

3.5.2.1 K-distributed density

When the intensity of the detected field is K-distributed with mean μ_I and order parameter ν , from (3.52) and (3.113) its density D will have pdf

$$f_D(D) = \frac{2}{\Gamma(\nu)} e^{\frac{(D-\alpha)(\nu+1)}{2}} K_{\nu-1} \left(2e^{\frac{(D-\alpha)}{2}} \right) \quad (3.126)$$

cumulative distribution function

$$p_D(D) = 1 - \frac{2}{\Gamma(\nu)} e^{\frac{(D-\alpha)\nu}{2}} K_\nu \left(2e^{\frac{(D-\alpha)}{2}} \right) \quad (3.127)$$

and characteristic function

$$\Phi_D(\omega) = \frac{\Gamma(j\omega + \nu) \Gamma(j\omega + 1)}{\Gamma(\nu)} \beta^{j\omega} \quad (3.128)$$

² $\psi(z)$ is the Digamma function.

Also, from (3.112) the mean density will be given by the mean of the speckle density $-\gamma_E$ plus the mean of the natural log of the imaged SCS (3.123)

$$\mu_D = \ln \mu_I - \gamma_E - \ln \nu + \psi(\nu) \quad (3.129)$$

and its variance will be given by the variance of the speckle density (3.118) plus the variance of the natural log of the imaged SCS (3.123)

$$\sigma_D^2 = \frac{\pi^2}{6} + \psi'(\nu) \quad (3.130)$$

The coefficient of variation will then be given by

$$v_D^2 = \frac{\frac{\pi^2}{6} + \psi'(\nu)}{(\ln \mu_I - \gamma_E - \ln \nu + \psi(\nu))^2} \quad (3.131)$$

Note that while the coefficient of variation of the K-distributed intensity (3.57) is independent of the mean intensity the coefficient of variation of its density does depend on it (3.131); linearly scaling the intensity does not alter the coefficient of variation of the intensity but it does alter the coefficient of variation of the density. This arises from the standard deviation of the density not being proportional to its mean value; as is the case with the intensity. The coefficients of variation of the natural log of exponential (3.119) and gamma distributed (3.125) random variables are also dependent on their mean values. This implies that the coefficient of variation is not a well defined measure of contrast (i.e., the relative width of distributions) in images degraded by additive noise; as it is in images degraded by multiplicative noise. More suitable measures of 'contrast' are given by the density variance on its own, and the expected bias of the density away from the natural log of the mean intensity. Both of these measures consist of a constant part due to speckle and a part dependent solely on the order parameter which increases as the order parameter decreases but which tends to zero as the order parameter tends to infinity. The expected bias of the density away from the natural log of the mean intensity equals the mean of the natural log of the intensity after it has been normalised to have unit mean, i.e., the expected value of

$$D_1 = \ln I - \ln \langle I \rangle = \ln \frac{I}{\langle I \rangle} \quad (3.132)$$

For the case of the K-distributed intensity the mean normalised density is given by

$$\langle D_1 \rangle = \psi(\nu) - \ln \nu - \gamma_E \quad (3.133)$$

This is the sum of the expected values of the natural logs of the unit mean exponential speckle and the unit mean gamma distributed imaged surface texture.

Figure 3.13 shows how the density in the region of woodland whose intensity distribution was better fitted by a K-distribution than an exponential distribution in Figure 3.4, is better fitted by a density K-distribution than a Fischer-Tippett distribution. Examples of density K-distributions with different order parameters are shown in Figure 3.14. The figure clearly illustrates both the

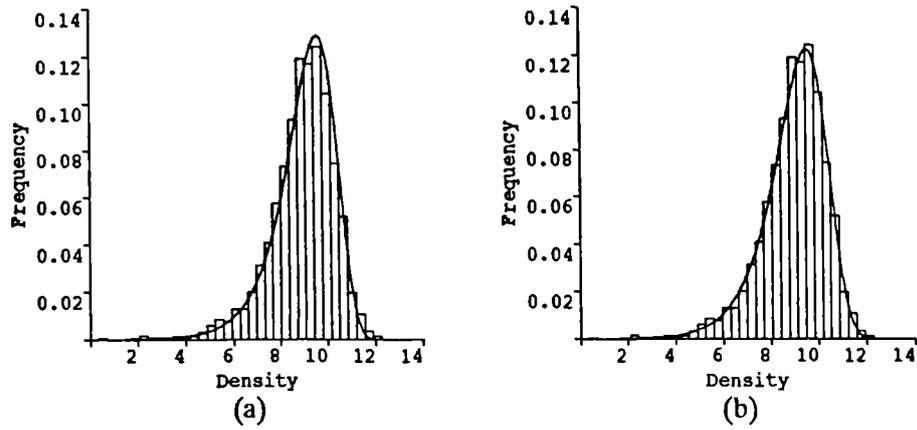


Figure 3.13: Histogram of the distribution of the density in a homogeneous region of woodland (the same as that in Figure 3.4) fitted to (a) a Fischer-Tippett distribution, and (b) a density K-distribution of order 10.

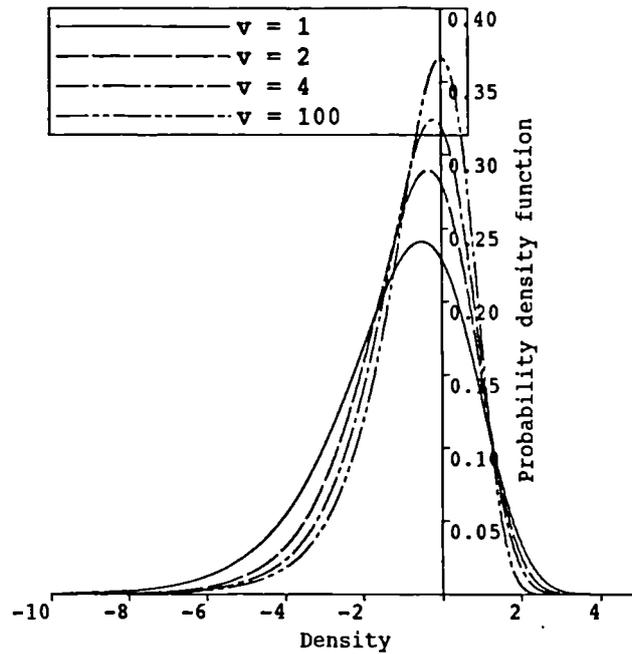


Figure 3.14: Density pdf's corresponding to the K-distributed intensity pdf's shown in Figure 3.5.

increase in the bias of the mean density away from the natural log of the mean intensity (zero in this case), and the increase in the density variance as the order parameter decreases.

Speckle can also be transformed to being additive by taking the log of the amplitude of the detected field (or any power of the intensity for that matter). The statistics of the natural log amplitude L_A of the detected field are easily derived from those of its density D using the relation

$$L_A = \ln A = \ln \sqrt{I} = \frac{\ln I}{2} = \frac{D}{2} \quad (3.134)$$

For example, the log amplitude has pdf

$$f_{L_A}(L_A) = 2f_D(2L_A) \quad (3.135)$$

mean

$$\mu_{L_A} = \frac{\mu_D}{2} \quad (3.136)$$

and variance

$$\sigma_{L_A}^2 = \frac{\sigma_D^2}{4} \quad (3.137)$$

3.6 Second order intensity statistics

In this section the second order statistics of the detected intensity are derived. In particular the results given in Chapter 2 regarding the response of a SAR system to a point target are used to derive the autocorrelation function of the detected intensity.

Assuming the SCS is stationary, from (3.16) the normalised autocorrelation function of the detected intensity is given by

$$\begin{aligned} \frac{\langle I(\bar{x}) I(\bar{y}) \rangle}{\langle I(\bar{x}) \rangle^2} = & \quad (3.138) \\ & \sum_{i_1=-\infty}^{\infty} \sum_{k_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \sum_{k_2=-\infty}^{\infty} \langle a_{i_1} a_{k_1} a_{i_2} a_{k_2} \rangle \langle e^{j(\theta_{i_1} - \theta_{k_1} + \theta_{i_2} - \theta_{k_2})} \rangle \times \\ & h(\bar{x}, \bar{x}_{i_1}) h^*(\bar{x}, \bar{x}_{k_1}) h(\bar{y}, \bar{x}_{i_2}) h^*(\bar{y}, \bar{x}_{k_2}) \left/ \left(\langle \rho \rangle \sum_{i_1=-\infty}^{\infty} |h(\bar{x}, \bar{x}_{i_1})|^2 \right)^2 \right. \end{aligned}$$

Cancelling phase terms in the same way as was used to derive the moments of the detected intensity in Section (3.3.1)

$$\begin{aligned} \frac{\langle I(\bar{x}) I(\bar{y}) \rangle}{\langle I(\bar{x}) \rangle^2} = & \\ & \sum_{i_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \frac{\langle \rho_{i_1} \rho_{i_2} \rangle}{\langle \rho \rangle^2} [|h(\bar{x}, \bar{x}_{i_1})|^2 |h(\bar{y}, \bar{x}_{i_2})|^2 + \\ & h(\bar{x}, \bar{x}_{i_1}) h^*(\bar{y}, \bar{x}_{i_1}) h(\bar{y}, \bar{x}_{i_2}) h^*(\bar{x}, \bar{x}_{i_2})] \left/ \left(\sum_{i_1=-\infty}^{\infty} |h(\bar{x}, \bar{x}_{i_1})|^2 \right)^2 \right. - \end{aligned}$$

$$\begin{aligned}
& \sum_{i_1=-\infty}^{\infty} \frac{\langle \rho_{i_1}^2 \rangle}{\langle \rho \rangle^2} |h(\bar{x}, \bar{x}_{i_1})|^2 |h(\bar{y}, \bar{x}_{i_1})|^2 / \left(\sum_{i_1=-\infty}^{\infty} |h(\bar{x}, \bar{x}_{i_1})|^2 \right)^2 \quad (3.139) \\
& = 1 + \left| \sum_{i_1=-\infty}^{\infty} h(\bar{x}, \bar{x}_{i_1}) h^*(\bar{y}, \bar{x}_{i_1}) \right|^2 / \left(\sum_{i_1=-\infty}^{\infty} |h(\bar{x}, \bar{x}_{i_1})|^2 \right)^2 + \\
& \quad \sum_{i_1=-\infty}^{\infty} \sum_{i_2=-\infty}^{\infty} \left(\frac{\langle \rho_{i_1} \rho_{i_2} \rangle}{\langle \rho \rangle^2} - 1 \right) [|h(\bar{x}, \bar{x}_{i_1})|^2 |h(\bar{y}, \bar{x}_{i_2})|^2 + \\
& \quad h(\bar{x}, \bar{x}_{i_1}) h^*(\bar{y}, \bar{x}_{i_1}) h(\bar{y}, \bar{x}_{i_2}) h^*(\bar{x}, \bar{x}_{i_2})] / \left(\sum_{i_1=-\infty}^{\infty} |h(\bar{x}, \bar{x}_{i_1})|^2 \right)^2 - \\
& \quad \sum_{i_1=-\infty}^{\infty} \frac{\langle \rho_{i_1}^2 \rangle}{\langle \rho \rangle^2} |h(\bar{x}, \bar{x}_{i_1})|^2 |h(\bar{y}, \bar{x}_{i_1})|^2 / \left(\sum_{i_1=-\infty}^{\infty} |h(\bar{x}, \bar{x}_{i_1})|^2 \right)^2 \quad (3.140)
\end{aligned}$$

The first component in (3.140) is the uncorrelated background term, the second component is caused by coherent interference between scatterers within the resolution cell (speckle), and the third and fourth components relate to the normalised ACF of the surface cross-section [46]. The fourth term disappears when the resolution cell contains a large number of scattering centres; as was the case with terms in the moments of the detected intensity (3.27). The third term will then also disappear if the correlation length of the SCS is much less than the system resolution, i.e., when the imaged SCS is constant.

From Chapter 2, for a fully focused SAR using a chirp pulse the system response to a point target is stationary and is separable into azimuth and slant range components which are both described by *sinc* functions (2.25,2.35). Thus the system response is given by

$$h(\bar{x}, \bar{y}) = \kappa h_a(z_a) h_r(z_r) \quad (3.141)$$

where κ is a proportionality factor expressing the peak power received according to the radar equation [34], z_a and z_r represent the vector $\bar{z} = \bar{x} - \bar{y}$ in azimuth and slant range coordinates, and

$$h_a(z_a) = \text{sinc} \left(\frac{z_a}{d_a} \right) \quad (3.142)$$

$$h_r(z_r) = \text{sinc} \left(\frac{z_r}{d_r} \right) \quad (3.143)$$

where d_a and d_r are the system resolution in azimuth and slant range.

The second term in (3.140) is just the square of the magnitude of the normalised ACF of the system response. Thus when the third and fourth terms in (3.140) are zero the **autocovariance** of the intensity is just given by the square of the autocovariance of the system response. This corresponds to the Siegert relation which says that the intensity autocovariance of a jointly Gaussian zero mean field with independent real and imaginary components is given by the square of the field autocovariance [30]. When the system response is

described by (3.141), its Fourier transform is given by

$$H(u, v) = \begin{cases} \frac{kd_a d_r}{16\pi^2} & |u| < \frac{2\pi}{d_a} \text{ and } |v| < \frac{2\pi}{d_r} \\ 0 & \text{otherwise} \end{cases} \quad (3.144)$$

and its power spectrum is given by

$$|H(u, v)|^2 = \begin{cases} \left(\frac{kd_a d_r}{16\pi^2}\right)^2 & |u| < \frac{2\pi}{d_a} \text{ and } |v| < \frac{2\pi}{d_r} \\ 0 & \text{otherwise} \end{cases} \quad (3.145)$$

As the power spectrum is just the Fourier transform of the ACF of the system response and as (3.145) is the Fourier transform of a *sinc* function, the ACF of the system response will itself be a *sinc* function. Therefore when the imaged SCS is constant the normalised ACF of the detected intensity is described by a *sinc*² function with a width between its peak and first null equal to the resolution of the imaging system

$$\frac{\langle I(\bar{x}) I(\bar{x} + \bar{z}) \rangle}{\langle I(\bar{x}) \rangle^2} = 1 + \text{sinc}^2\left(\frac{z_a}{d_a}\right) \text{sinc}^2\left(\frac{z_r}{d_r}\right) \quad (3.146)$$

In addition, assuming the multiplicative model applies, the detected field is then given by the product of a jointly Gaussian speckle process with independent real and imaginary components each with variance one half, and the square root of the gamma distributed imaged SCS (see Section 3.3.3), this results in the intensity ACF equalling

$$\frac{\langle I(\bar{x}) I(\bar{x} + \bar{z}) \rangle}{\langle I(\bar{x}) \rangle^2} = \frac{\langle \Upsilon(\bar{x}) \Upsilon(\bar{x} + \bar{z}) \rangle}{\langle \Upsilon(\bar{x}) \rangle^2} \left(1 + \text{sinc}^2\left(\frac{z_a}{d_a}\right) \text{sinc}^2\left(\frac{z_r}{d_r}\right) \right) \quad (3.147)$$

This equates to (3.146) when the imaged SCS is constant. Alternatively, when the imaged SCS is 100% correlated the intensity ACF is given by

$$\frac{\langle I(\bar{x}) I(\bar{x} + \bar{z}) \rangle}{\langle I(\bar{x}) \rangle^2} = \left(1 + \frac{1}{\nu} \right) \left(1 + \text{sinc}^2\left(\frac{z_a}{d_a}\right) \text{sinc}^2\left(\frac{z_r}{d_r}\right) \right) \quad (3.148)$$

Figure 3.15 shows an example of the intensity autocovariance in a homogeneous region normalised by the square of the mean intensity. The *sinc*² shape of the autocovariance is clearly apparent. Pixel intensity values are uncorrelated in azimuth but there is significant correlation in range.

3.7 Summary

In this chapter:

- It has been described how the field detected from a resolution cell may be represented as a sum of contributions from discrete elementary scatterers.
- The conditions for speckle being fully developed were discussed.

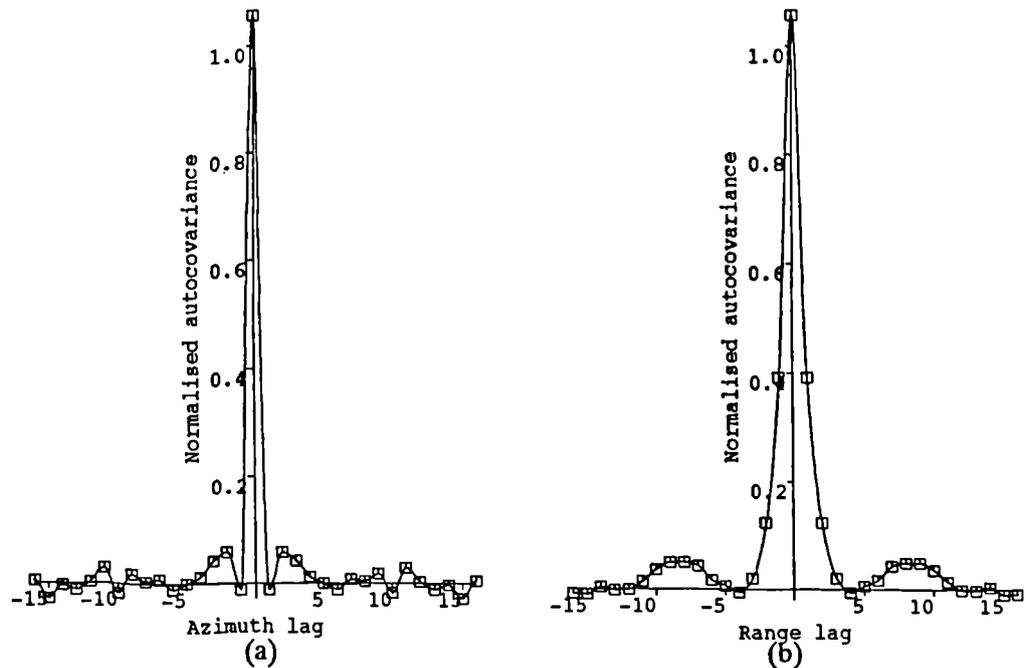


Figure 3.15: Normalised intensity autocovariance in the homogeneous region shown in Figure 3.3.

- The statistics of the phase and complex components of the detected field resulting from fully developed speckle were derived
- It was shown how the intensity of the detected field may be modelled as the product of an exponential speckle process and positive imaging of the surface cross-section.
- It was shown how the correlation length of the surface cross-section relative to the system resolution determines whether the intensity is (a) exponentially distributed, or (b) K-distributed.
- Assuming the intensity is exponentially or K-distributed, the statistics of the amplitude of the detected field were described.
- It was shown how speckle can be transformed to being additive by taking its log and the resulting distributions assuming the intensity is exponentially or K-distributed were derived.
- The autocorrelation function of the intensity was derived for the cases of constant and 100% correlated imaged SCS.

Chapter 4

Parameter Estimation and Multi-looking

A large proportion of image analysis involves estimating parameters within different regions of an image. These parameters supply the basis for local decision making. This is particularly the case with SAR where the presence of speckle precludes image analysis being carried out on a per pixel basis. For whichever application they are used, e.g., speckle reduction, feature detection or image classification, the sampling distributions of estimates need to be known for determining detection and false alarms rates.

The parameters required for image analysis are estimated by applying simple operations (e.g., averaging) to sets of neighbouring pixels. To determine the distribution of an estimate output by an operator the distribution of the variables input to it, and the relationship between these distributions and the distribution of the random variable output, need to be known. The general form of the relationship between the input and output distributions of basic binary operators used in this and other chapters are listed in Appendix A.

The parameter most commonly required for image analysis is the local mean. For example, estimates of the local mean are required by the edge detection and segmentation algorithms discussed in Chapters 7 and 9. A simple estimate of the local mean is the arithmetic average of pixel values. The result of averaging, including multi-looking, on image statistics is discussed in Section 4.1; the effects of correlation, K-distributed speckle and heterogeneous populations are all considered.

The distribution of pixel values over an extended target is fully described by the mean intensity whenever the imaged SCS is constant. However, when surface textures are imaged additional parameters are needed to define the intensity distribution. Texture estimation is addressed in Section 4.2. Assuming the imaged surface texture is gamma distributed with unit mean and unknown order parameter, methods of estimating the order parameter from the K-distributed intensity are derived and compared.

4.1 Averages of pixel values

Arithmetic averaging is used to estimate mean values because given samples from a stationary random process their average value gives an unbiased estimate of their mean, and so long as the autocovariance of the process is finite the estimate will be consistent.

The importance of knowing the statistics of averages which are used as estimates of local mean values in image analysis, has already been alluded to. Knowledge of the effect of averaging on image statistics is also needed when the whole of an image is incoherently averaged; so as to be able to determine the statistics of the resultant image. Incoherent averaging (or multi-looking) is used to reduce speckle noise and/or resolution. When n independent samples, or looks, in azimuth (or range) are averaged image contrast is reduced by a factor \sqrt{n} , and the azimuth (range) resolution of the image is reduced by a factor n . Multi-looking is performed by either spatially averaging discrete blocks of n pixels in the full resolution image, or during initial processing of the raw data, by splitting the image bandwidth into n sections and processing each section to produce n lower resolution images of the same scene, which are then averaged [59].

The average intensity is of particular importance because over a homogeneous region where the imaged SCS is constant (i.e., exponentially distributed intensity), it gives the maximum likelihood (ML) estimate of the mean intensity, and via (3.94) the ML estimate of the mean amplitude. Given n statistically identical independent exponentially distributed intensity values I_1, \dots, I_n with unknown mean, the likelihood F that the intensity values have mean μ_I is given by

$$F(\mu_I) = \prod_{i=1}^n \frac{1}{\mu_I} e^{-\frac{I_i}{\mu_I}} \quad (4.1)$$

and the log-likelihood $L = \ln F$ is given by

$$L(\mu_I) = -n \ln \mu_I - \frac{n}{\mu_I} \hat{I} \quad (4.2)$$

where \hat{I} is the average intensity defined by

$$\hat{I} = \frac{1}{n} \sum_{i=1}^n I_i \quad (4.3)$$

Because \ln is a monotonically increasing function the maximum log-likelihood will occur at the same position as the maximum likelihood. Hence the ML estimate of the mean may be found by solving where the first derivative of the log-likelihood with respect to μ_I given by

$$\frac{dL}{d\mu_I} = \frac{n}{\mu_I^2} \hat{I} - \frac{n}{\mu_I} = \frac{n}{\mu_I^2} (\hat{I} - \mu_I) \quad (4.4)$$

equals zero. This will occur when the mean is estimated by the average intensity.

Likewise, given n statistically identical independent Rayleigh distributed amplitude values A_1, \dots, A_n with unknown mean, from (3.90) the likelihood F that the amplitude values have mean μ_A is given by

$$F(\mu_A) = \prod_{i=1}^n \frac{\pi A_i}{2\mu_A^2} e^{-\frac{\pi A_i^2}{4\mu_A^2}} \quad (4.5)$$

and the first derivative of the log-likelihood with respect to μ_A is given by

$$\frac{dL}{d\mu_A} = \frac{2n}{\mu_A^3} \left(\frac{\pi \widehat{A}^2}{4} - \mu_A^2 \right) \quad (4.6)$$

This will equal zero when the mean amplitude is estimated by $\sqrt{\pi \widehat{A}^2}/2$. Since \widehat{A}^2 is just the average intensity, the ML estimate of the mean amplitude is given by calculating the average intensity, *rather than* the average amplitude, and then relating it to the mean amplitude using (3.94). Therefore in terms of estimating mean values over regions where the imaged SCS is constant, averaging is best carried out on intensity as opposed to amplitude data. It should also be noted that as the exponential distribution is fully defined by its mean value, the mean intensity is the only parameter that needs to be estimated to determine the distribution of data in such regions.

The effect of averaging pixel values with a variety of distributions are considered below. In Section 4.1.1 it is shown that the average of n exponential pixel values from a homogeneous region is gamma distributed with order n . The effect of heterogeneous populations on the distribution of the average are then discussed. As knowledge of the statistics of the average amplitude may sometimes be need, this is considered in Section 4.1.2, for instance when multi-looking has been applied to amplitude as opposed to intensity (e.g., CCRS airborne SAR data [24]). It will be seen that the statistics are analytically more tractable in the case of intensity. This, along with the fact that *the average intensity gives* the ML estimate of the mean intensity when the imaged SCS is constant, means that we will often prefer to work on intensity as opposed to amplitude data. In Section 4.1.3 the average of K-distributed intensity is examined. This together with correlation which is considered in Section 4.1.4, have the effect of increasing the variance of the resulting average.

Normally the distribution of averages of random variables with known distributions are most easily determined using characteristic functions. The characteristic function of the average \widehat{x} of n independent random variables x_i with characteristic functions $\Phi_{x_i}(\omega)$ is given by

$$\Phi_{\widehat{x}}(\omega) = \prod_{i=1}^n \Phi_{x_i} \left(\frac{\omega}{n} \right) \quad (4.7)$$

This comes from the **convolution theorem** and the fact that the characteristic function of x/n is related to the characteristic function of x by

$$\Phi_{\frac{x}{n}}(\omega) = \Phi_x \left(\frac{\omega}{n} \right) \quad (4.8)$$

4.1.1 Exponential intensity

From (3.32) and (4.7) the characteristic function of the average \hat{I} of n independent exponentially distributed intensity values with means μ_1, \dots, μ_n is given by

$$\Phi_{\hat{I}}(\omega) = \prod_{i=1}^n \left(1 - \frac{j\omega\mu_i}{n}\right)^{-1} \quad (4.9)$$

When the intensity values are statistically identical with mean μ the characteristic function is given by

$$\Phi_{\hat{I}}(\omega) = \left(1 - \frac{j\omega\mu}{n}\right)^{-n} \quad (4.10)$$

This is the characteristic function of a gamma distributed random variable with mean μ and order parameter n (3.61). Thus it is normally assumed that when the imaged SCS is constant, multi-look intensity data and averages of intensity values are gamma distributed [22]. When this is true, it can also be seen that the average of m statistically identical independent n -look intensity values (i.e., the average of mn statistically identical independent single-look intensity values), will be gamma distributed with order mn , by substituting the characteristic function of the gamma distributed intensity values given by (3.61) into the characteristic function of their average (4.7). As in the single-look case, such an average gives the ML estimate of the mean intensity.

Over a heterogeneous area where pixel intensity values are sampled from exponential populations with different mean values, in general, the distribution of the average cannot be determined analytically from its characteristic function. In the case of n values being sampled from M different populations such that n_i values have mean μ_i ($i = 1, \dots, M$), the average has mean

$$\mu_{\hat{I}} = \sum_{i=1}^M p_i \mu_i \quad (4.11)$$

where $p_i = n_i/n$ and $n_1 + \dots + n_M = n$, and has the same distribution as the sum of M gamma distributed random variables with means $p_i \mu_i$ and order parameters n_i . The characteristic function of the average is then given by

$$\Phi_{\hat{I}}(\omega) = \prod_{i=1}^M (1 - j\omega\beta_i)^{-n_i} \quad (4.12)$$

where $\beta_i = p_i \mu_i / n_i$. The pdf of the average has been determined for the case of $M = 2$, and is given by

$$\begin{aligned} f_{\hat{I}}(\hat{I}) = & \frac{1}{\beta_1^{n_1} \beta_2^{n_2}} \left\{ e^{\frac{\hat{I}}{\beta_1}} \alpha^{n_2} \sum_{k=0}^{n_1-1} \binom{n_2-1+k}{k} \frac{\hat{I}^{n_1-1-k} (-\alpha)^k}{\Gamma(n_1-k)} \right. \\ & \left. + e^{\frac{\hat{I}}{\beta_2}} (-\alpha)^{n_1} \sum_{k=0}^{n_2-1} \binom{n_1-1+k}{k} \frac{\hat{I}^{n_2-1-k} \alpha^k}{\Gamma(n_2-k)} \right\} \quad (4.13) \end{aligned}$$

where $1/\alpha = 1/\beta_2 - 1/\beta_1$. See Appendix C for the proof.

For $M > 2$ it has not been possible to obtain an analytic expression for the pdf of the average from its characteristic function given by (4.12). However, the moments of the average may be generated by differentiating its characteristic function. Normalising these by the mean of the average intensity gives

$$\hat{I}^{(2)} = 1 + \frac{1}{n}q_2 \quad (4.14)$$

$$\hat{I}^{(3)} = 1 + \frac{3}{n}q_2 + \frac{2}{n^2}q_3 \quad (4.15)$$

$$\hat{I}^{(4)} = 1 + \frac{6}{n}q_2 + \frac{3}{n^2}q_2^2 + \frac{8}{n^2}q_3 + \frac{6}{n^3}q_4 \quad (4.16)$$

etc, where

$$q_k = \sum_{i=1}^M p_i \left(\frac{\mu_i}{\mu_{\hat{I}}} \right)^k \quad (4.17)$$

These normalised moments lie between those of gamma distributed random variables with order parameters n and $n\mu_{\hat{I}}/\mu_{max}$, this is because

$$1 \leq q_k \leq \left(\frac{\mu_{max}}{\mu_{\hat{I}}} \right)^{k-1} \quad (4.18)$$

where $\mu_{max} = \max(\mu_i)$. The normalised moments of the average will tend to those of a gamma distribution with order parameter n as the μ_i 's tend to being equal, and they will tend to those of a gamma distribution with order parameter n_i when the i 'th term in q_k dominates all the other terms. This will occur when either most of the pixels are sampled from the i 'th population or when its mean value is much greater than the mean value of all the other populations ($\mu_i \gg \mu_{j \neq i}$). The coefficient of variation of the average intensity is given by q_2/n and will always be greater than that over a homogeneous area given by $1/n$ (3.65).

After multi-looking is carried out on intensity its square root is often taken to further reduce image contrast (e.g., ERS-1 multi-look SAR image products [16]). Assuming average intensity is gamma distributed its square root (or amplitude) is distributed according to (3.98). Also, its natural log (or density) is distributed according to (3.120).

4.1.2 Rayleigh distributed amplitude

The distribution of the average \hat{A} of n statistically identical independent Rayleigh distributed amplitude values A_1, \dots, A_n with mean intensity μ_I , cannot be determined in as straightforward a manner as the average intensity can. Obviously from (3.94) the average amplitude has mean

$$\mu_{\hat{A}} = \mu_A = \sqrt{\frac{\pi\mu_I}{4}} \quad (4.19)$$

and from (3.95) and (3.97) it is easily shown that the average has variance

$$\sigma_A^2 = \frac{\sigma_A^2}{n} = \frac{1}{n} \left(1 - \frac{\pi}{4}\right) \mu_I \quad (4.20)$$

and coefficient of variation

$$v_A^2 = \frac{v_A^2}{n} = \frac{1}{n} \left(\frac{4}{\pi} - 1\right) \quad (4.21)$$

The coefficient of variation of the average of independent Rayleigh amplitude values given by (4.21) is greater than that of the square root of the average of the corresponding exponential intensity values (3.104). This accords with the reasons given previously for averaging in intensity rather than amplitude.

The characteristic function of the average can be defined in terms of the characteristic function of the Rayleigh distribution (3.92) but it has not been possible to deduce the analytical form of its pdf from this. For generating the moments of the average by differentiating its characteristic function, it is better to express the characteristic function in terms of the moments of the Rayleigh distributed amplitude values (3.93)

$$\Phi_{\hat{A}}(\omega) = \Phi_A \left(\frac{\omega}{n}\right)^n = \left(\sum_{k=0}^{\infty} \frac{\left(\frac{\omega}{n}\right)^k}{k!} \langle A^k \rangle\right)^n = \left(\sum_{k=0}^{\infty} \Gamma\left(\frac{k}{2} + 1\right) \frac{\left(\frac{\omega}{n}\right)^k}{k!} \mu_I^{\frac{k}{2}}\right)^n \quad (4.22)$$

The normalised moments are then given by

$$\hat{A}^{(2)} = \left(1 - \frac{1}{n}\right) + \frac{4}{n\pi} \quad (4.23)$$

$$\hat{A}^{(3)} = \left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) + \frac{12}{n\pi} \left(1 - \frac{1}{n}\right) + \frac{6}{n^2\pi} \quad (4.24)$$

$$\begin{aligned} \hat{A}^{(4)} = & \left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) \left(1 - \frac{3}{n}\right) + \frac{24}{n\pi} \left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) \\ & + \frac{48}{n^2\pi^2} \left(1 - \frac{1}{n}\right) \left(1 + \frac{\pi}{2}\right) + \frac{32}{n^3\pi^2} \end{aligned} \quad (4.25)$$

etc.

4.1.3 K-distributed intensity

The average of n statistically identical independent K-distributed intensity values with mean μ_I and order parameter ν has mean μ_I and it is easily shown from (3.56) and (3.57) that the average has variance

$$\sigma_I^2 = \frac{\mu_I^2}{n} \left(1 + \frac{2}{\nu}\right) \quad (4.26)$$

and coefficient of variation

$$v_I^2 = \frac{1}{n} \left(1 + \frac{2}{\nu}\right) \quad (4.27)$$

Thus, the average intensity over a homogeneous region will have greater variance when the intensity values are K-distributed than if they were exponentially distributed with the same mean value.

As with the average Rayleigh distributed amplitude over a homogeneous region the characteristic function of the average K-distributed intensity over a homogeneous region can be expressed in terms of the moments of the averaged values (3.54)

$$\Phi_{\hat{I}}(\omega) = \left(\sum_{k=0}^{\infty} \frac{\Gamma(\nu + k)}{\Gamma(\nu)} \left(\frac{j\omega\mu_I}{n\nu} \right)^k \right)^n \quad (4.28)$$

from which the normalised moments can be generated

$$\hat{I}^{(2)} = \left(1 - \frac{1}{n}\right) + \frac{2}{n} \left(1 + \frac{1}{\nu}\right) \quad (4.29)$$

$$\begin{aligned} \hat{I}^{(3)} &= \left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) + \frac{6}{n} \left(1 - \frac{1}{n}\right) \left(1 + \frac{1}{\nu}\right) \\ &\quad + \frac{6}{n^2} \left(1 + \frac{1}{\nu}\right) \left(1 + \frac{2}{\nu}\right) \end{aligned} \quad (4.30)$$

$$\begin{aligned} \hat{I}^{(4)} &= \left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) \left(1 - \frac{3}{n}\right) + \frac{12}{n} \left(1 - \frac{1}{n}\right) \left(1 - \frac{2}{n}\right) \left(1 + \frac{1}{\nu}\right) \\ &\quad + \frac{12}{n^2} \left(1 - \frac{1}{n}\right) \left(\left(1 + \frac{1}{\nu}\right)^2 + 2 \left(1 + \frac{1}{\nu}\right) \left(1 + \frac{2}{\nu}\right) \right) \\ &\quad + \frac{24}{n^3} \left(1 + \frac{1}{\nu}\right) \left(1 + \frac{2}{\nu}\right) \left(1 + \frac{3}{\nu}\right) \end{aligned} \quad (4.31)$$

etc. The normalised moments tend to those of a gamma distribution as $\nu \rightarrow \infty$.

4.1.4 Effect of correlation on averaging

So far it has been assumed the pixel values being averaged are statistically independent, in practice this will generally not be the case. The average intensity \hat{I} over an area of size n_a pixels in azimuth by n_r pixels in slant range, is given by

$$\hat{I}(x_a, x_r) = \frac{1}{n_a n_r} \sum_{w_a=0}^{n_a-1} \sum_{w_r=0}^{n_r-1} I(x_a + w_a, x_r + w_r) \quad (4.32)$$

Assuming the autocovariance of the intensity is separable into azimuth and range components it may be expressed

$$C_I(\bar{z}) = \langle I(\bar{x}) I(\bar{x} + \bar{z}) \rangle - \langle I(\bar{x}) \rangle \langle I(\bar{x} + \bar{z}) \rangle = C_a(z_a) C_r(z_r) \quad (4.33)$$

the average intensity will then have autocovariance

$$\begin{aligned} C_{\hat{I}}(z_a, z_r) &= \langle \hat{I}(\bar{x}) \hat{I}(\bar{x} + \bar{z}) \rangle - \langle \hat{I}(\bar{x}) \rangle \langle \hat{I}(\bar{x} + \bar{z}) \rangle \\ &= \frac{1}{n_a n_r} \left\{ C_a(z_a) + \sum_{w_a=1}^{n_a-1} (C_a(z_a + w_a) + C_a(z_a - w_a)) \left(1 - \frac{w_a}{n_a}\right) \right\} \times \\ &\quad \left\{ C_r(z_r) + \sum_{w_r=1}^{n_r-1} (C_r(z_r + w_r) + C_r(z_r - w_r)) \left(1 - \frac{w_r}{n_r}\right) \right\} \end{aligned} \quad (4.34)$$

and the variance of the average intensity will be given by

$$\sigma_{\hat{I}}^2 = C_{\hat{I}}(0,0) = \frac{\sigma_I^2}{n_a n_r} \left\{ 1 + 2 \sum_{w_a=1}^{n_a-1} c_a(w_a) \left(1 - \frac{w_a}{n_a}\right) \right\} \times \left\{ 1 + 2 \sum_{w_r=1}^{n_r-1} c_r(w_r) \left(1 - \frac{w_r}{n_r}\right) \right\} \quad (4.35)$$

where $c_a(w_a) = \frac{C_a(w_a)}{C_a(0)}$ and $c_r(w_r) = \frac{C_r(w_r)}{C_r(0)}$. Thus the autocovariance of the average intensity is separable into azimuth and range components if the intensity autocovariance is likewise separable.

As pixel intensity values are the product of a speckle component and an imaged SCS component correlation may arise from either one of these components or both being correlated. As the combined effect of these two sources of correlation is somewhat complex their effect on averaging will only be dealt with separately.

4.1.4.1 Correlated speckle

All correlation will be purely due to speckle when the imaged SCS is constant. The intensity autocovariance will then just be determined by the system response to a point target. For a fully focused SAR using a chirp pulse the azimuth and slant range components of the intensity autocovariance are both *sinc*² functions (3.146). Figure 4.1 shows plots of the azimuth component of the autocovariance of the average intensity for different azimuth resolutions d_a when $n_a = 4$, and the azimuth component of the intensity autocovariance is given by

$$c_a(z_a) = \text{sinc}^2\left(\frac{z_a}{d_a}\right) \quad (4.36)$$

The slant range component of the autocovariance of the average intensity will have a similar form. The variance and correlation length of the average intensity both increase as d_a increases.

Figure 4.2 shows the plot corresponding Figure 4.1 when the intensity autocovariance in azimuth is approximated by a rectangular function of width d_a given by

$$c_a(z_a) = \begin{cases} 1 & |z_a| < d_a \\ 0 & \text{otherwise} \end{cases} \quad (4.37)$$

This may be taken to represent the main lobe of a *sinc*² function. The variance and correlation length of the average are both greater than those in Figure 4.1. Thus when the imaged SCS is constant the variance of the average intensity approximated from a rectangular intensity autocovariance may be used as a conservative estimate of the actual variance of the average. When the intensity autocovariance is given by (4.37) the variance of the average is given by

$$\sigma_{\hat{I}}^2 = \frac{\sigma_I^2}{n_a n_r} \left\{ 2d_a - 1 + \frac{d_a(d_a - 1)}{n_a} \right\} \left\{ 2d_r - 1 + \frac{d_r(d_r - 1)}{n_r} \right\} \quad (4.38)$$

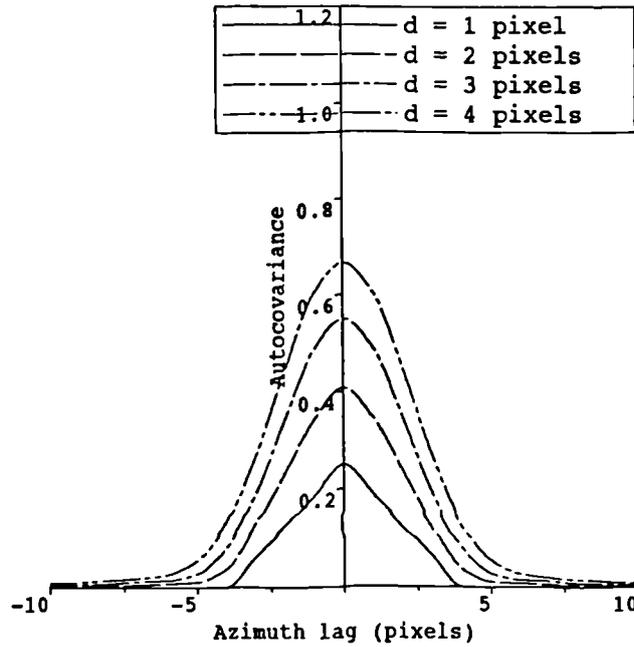


Figure 4.1: Azimuth component of the autocovariance of the average intensity for different azimuth resolutions when the intensity has unit variance, $n_a = 4$, and the azimuth component of the intensity autocovariance is given by $\text{sinc}^2(z_a/d_a)$.

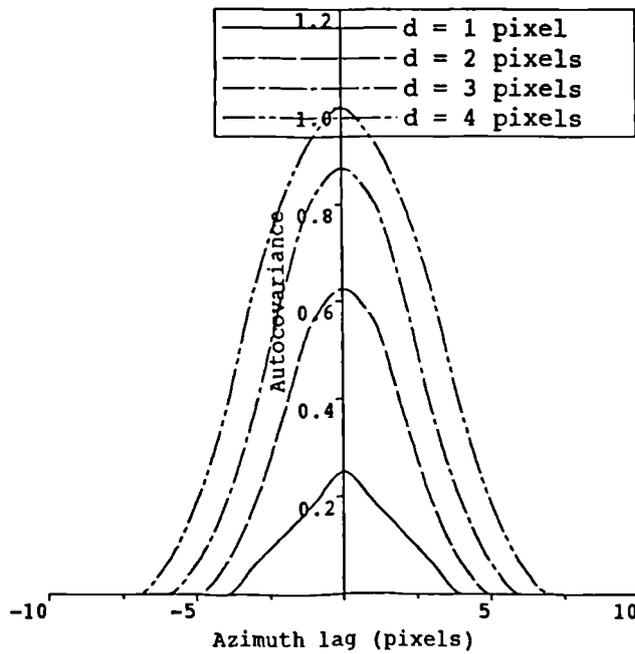


Figure 4.2: Azimuth component of the autocovariance of the average intensity when the intensity has unit variance, $n_a = 4$ and the azimuth component of the intensity autocovariance is approximated by a rectangular function of width d_a and unit height.

when $d_a \leq n_a$ and $d_r \leq n_r$, otherwise it is given by

$$\sigma_{\hat{I}}^2 = \sigma_I^2 \quad (4.39)$$

This corresponds to the intensity being 100% correlated over the averaging window. When the intensity is uncorrelated ($d_a \leq 1$ and $d_r \leq 1$) the variance is given by the variance of the intensity divided by the number of values averaged

$$\sigma_{\hat{I}}^2 = \frac{\sigma_I^2}{n_a n_r} \quad (4.40)$$

When $d_a \ll n_a$ and $d_r \ll n_r$ the variance may be approximated by

$$\sigma_{\hat{I}}^2 \approx \frac{\sigma_I^2}{n_a n_r} \{2d_a - 1\} \{2d_r - 1\} \quad (4.41)$$

4.1.4.2 Correlated imaged SCS: the multi-look K-distribution

When intensity values which have independent speckle components but 100% correlated imaged SCS components are averaged, the averaging only affects the speckle. This situation will arise when either multi-look is carried out by splitting the bandwidth or when spatial averaging is carried out over a region where the imaged SCS has a correlation length much greater than the pixel sampling rates in range and azimuth. The average of n such intensity values will be the product of the imaged cross section and a unit mean gamma distribution of order n (n -look speckle). If the imaged SCS is gamma distributed with mean μ and order parameter ν the average intensity \hat{I} will be given by the product of two independent gamma distributed random variables one with unit mean and order parameter n , and one with mean μ and order parameter ν . The average intensity will obviously have mean μ and from (3.63) it will have normalised moments

$$\hat{I}^{(m)} = \frac{\Gamma(n+m) \Gamma(\nu+m)}{\Gamma(n) n^m \Gamma(\nu) \nu^m} \quad (4.42)$$

and coefficient of variation

$$v_{\hat{I}}^2 = \frac{1}{n} \left(1 + \frac{n+1}{\nu} \right) \quad (4.43)$$

Because only the speckle and not the intensity fluctuations are averaged, the normalised moments and coefficient of variation are greater than those of the sum of n independent K-distributed random variables with mean μ and order parameter ν (see Section 4.1.3).

The pdf of the average intensity may be derived by substituting the pdf's of the gamma distributed speckle and imaged SCS (3.59) into the integral for the pdf of the product of two independent random variables (A.5) and solving using (F.2). This gives

$$f_{\hat{I}}(\hat{I}) = \frac{2}{\Gamma(n) \Gamma(\nu) \beta} \left(\sqrt{\frac{\hat{I}}{\beta}} \right)^{n+\nu-2} K_{\nu-n} \left(2\sqrt{\frac{\hat{I}}{\beta}} \right) \quad (4.44)$$

where $\beta = \mu/n\nu$. This distribution has been termed the multi-look K-distribution [36], and has strong similarities to the single-look K-distribution (3.52). When $\nu = 1$ the n -look K-distribution has the same form as the single-look K-distribution with order parameter n . As n or ν tends to infinity the multi-look K-distribution tends to a gamma distribution.

4.2 Texture estimation

The intensity distribution over an extended target is described by the speckle, the mean imaged SCS (the mean intensity) and the distribution of the imaged surface texture. Estimation of the mean intensity using the average intensity has already been discussed in Section 4.1. In this section we consider how best to estimate the parameters describing the texture distribution, i.e., the non-speckle part of the clutter distribution. Assuming imaged surface texture is gamma distributed with unit mean and order parameter ν , the intensity will be K-distributed with the same order parameter (see Section 3.3.3). Thus the distribution of the texture may be determined by estimating the order parameter of the K-distributed intensity. Several possible estimates of the order parameter are considered. First the ML estimate of both the mean and order parameter of n statistically identical independent K-distributed intensity values is derived but is found to be somewhat impracticable. Therefore non-optimal methods have to be used to estimate the order parameter [47]. Two statistics, each one of which the order parameter can be estimated from, are then discussed, namely the second normalised intensity moment $I^{(2)}$ and the mean normalised density $\langle \ln I \rangle - \ln \langle I \rangle$. The latter is shown to give the more reliable estimate of the order parameter.

4.2.1 ML estimation of the mean and order parameter of K-distributed intensity

Given n statistically identical independent K-distributed intensity values I_1, \dots, I_n with unknown mean and order parameter, the log-likelihood that they have mean μ and order parameter ν is given by

$$L(\mu, \nu) = n \ln 2 + \frac{n(\nu + 1)}{2} (\ln \nu - \ln \mu) - n \ln \Gamma(\nu) + \frac{\nu - 1}{2} \sum_{k=1}^n \ln I_k + \sum_{k=1}^n \ln K_{\nu-1} \left(2\sqrt{\frac{I_k \nu}{\mu}} \right) \quad (4.45)$$

The ML estimates of the mean and order parameter of the K-distributed intensity are given by the values of μ and ν which maximize the log-likelihood. Solving for where the derivatives of the log-likelihood with respect to the mean and order parameter both equal zero and using (F.3) gives [47]

$$\sqrt{\frac{\mu}{\nu}} (\nu + 1) = \frac{1}{n} \sum_{k=1}^n \sqrt{I_k} \frac{K_{\nu} \left(2\sqrt{\frac{I_k \nu}{\mu}} \right) + K_{\nu-2} \left(2\sqrt{\frac{I_k \nu}{\mu}} \right)}{K_{\nu-1} \left(2\sqrt{\frac{I_k \nu}{\mu}} \right)} \quad (4.46)$$

and

$$\ln \psi(\nu) + \ln \sqrt{\frac{\mu}{\nu}} = \frac{1}{n} \sum_{k=1}^n \ln \sqrt{I_k} + \frac{1}{n} \sum_{k=1}^n \frac{\frac{\partial}{\partial \nu} K_{\nu-1} \left(2\sqrt{\frac{I_k \nu}{\mu}} \right)}{K_{\nu-1} \left(2\sqrt{\frac{I_k \nu}{\mu}} \right)} \quad (4.47)$$

The values of μ and ν which satisfy these two equations cannot be found analytically and finding them numerically is an expensive operation. This forces us into using non-optimal methods to estimate the order parameter.

4.2.2 Non-optimal estimation of the order parameter of K-distributed intensity

Statistics from which the order parameter of K-distributed intensity may be determined include the second normalised intensity moment and the mean normalised density [47]. Non-optimal estimates of these statistics are described in this section and their performance in terms of estimating the order parameter compared.

4.2.2.1 Second normalised intensity moment

Let the m th moment of n statistically identical independent K-distributed intensity values with unknown order parameter ν be estimated by

$$\widehat{I}^m = \frac{1}{n} \sum_{k=0}^n I_k^m \quad (4.48)$$

This gives an unbiased estimate of the m th moment of the intensity. As the order parameter completely defines the second normalised moment of the intensity values given by (3.55)

$$I^{(2)} = 2 \left(1 + \frac{1}{\nu} \right) \quad (4.49)$$

the order parameter may be determined from an estimate of the second normalised moment given by

$$\widehat{I}^{(2)} = \widehat{I}^2 / \widehat{I}^2 \quad (4.50)$$

The bias and variance of a function of estimated parameters may be approximated by expanding the function around the expected values of the estimates [48] (7.20,7.21). Expanding $\widehat{I}^{(2)}$ around $I^{(2)}$ its bias may be approximated by [44]

$$\langle \widehat{I}^{(2)} \rangle - I^{(2)} \approx \frac{1}{n} I^{(2)} \left(-1 + 3I^{(2)} - 2 \frac{I^{(3)}}{I^{(2)}} \right) \quad (4.51)$$

and its variance may be approximated by

$$\left\langle \left(\widehat{I}^{(2)} - I^{(2)} \right)^2 \right\rangle \approx \frac{1}{n} \left(I^{(2)} \right)^2 \left(-1 + 4I^{(2)} - 4 \frac{I^{(3)}}{I^{(2)}} + \frac{I^{(4)}}{\left(I^{(2)} \right)^2} \right) \quad (4.52)$$

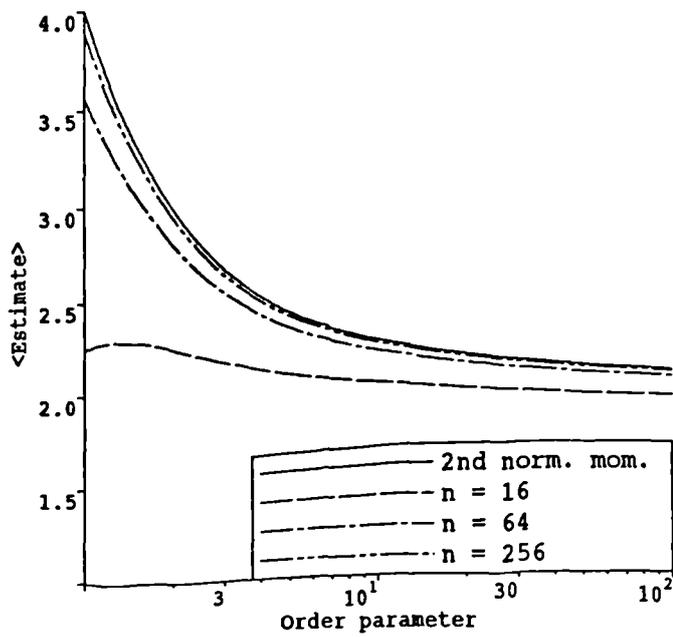


Figure 4.3: The second normalised moment $I^{(2)}$ and the approximated bias in its estimate $\widehat{I}^{(2)}$ for different sample sizes.

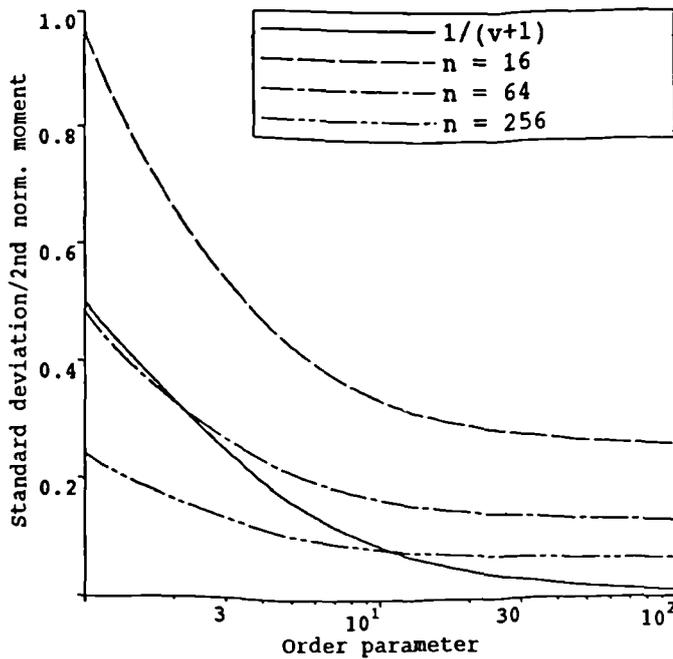


Figure 4.4: Approximated standard deviation in the estimate $\widehat{I}^{(2)}$ of the second normalised moment $I^{(2)}$ divided by $I^{(2)}$. Above the solid curve given by $1/(\nu + 1)$, the standard deviation is greater than the difference in the second normalised moment between when the order parameter equals ν and when it is infinite.

Figure 4.3 shows the second normalised intensity moment and the approximated bias in its estimate for different sample sizes and Figure 4.4 shows the approximated standard deviation of the estimate divided by the second normalised moment. The estimate is significantly biased when the sample size is small ($n = 16$). This bias causes there to be little difference between estimates of the second normalised moment of intensity values which are K-distributed with different order parameters. Except when the sample size is large and the order parameter is small, the standard deviation is greater than the difference in the second normalised moment between when the order parameter equals ν and when it is infinite, i.e., the estimate fluctuates over a range greater than the range of the second normalised moment. From these results it is clear that very large sample sizes are needed to determine the order parameter from such an estimate of the second normalised intensity moment.

Note, for the order parameter to be positive the estimate of the second normalised moment must be greater than two while in practice it only has to be greater than one. When the estimate is less than two the order parameter should be interpreted as being infinite because the second normalised moment is then smaller than its expected value when no texture is present.

4.2.2.2 Mean normalised density

When intensity is K-distributed the mean value of the normalised density D_1 (see Section 3.5) is determined solely by the order parameter ν of the intensity distribution (3.133)

$$\langle D_1 \rangle = \psi(\nu) + \psi(1) - \ln \nu \quad (4.53)$$

Thus the order parameter can also be determined from an estimate $\widehat{D}_1 = \widehat{\ln I} - \ln \widehat{I}$ of the normalised density. In fact this estimate gives the ML estimate of the unspckled gamma distributed ISCS [47]. Thus it may be expected to give a better estimate of the order parameter of the speckled K-distributed intensity than the second normalised moment. Expanding the estimate \widehat{D}_1 around $\langle D_1 \rangle$ its bias may be approximated by

$$\langle \widehat{D}_1 \rangle - \langle D_1 \rangle \approx \frac{1}{2n} (I^{(2)} - 1) = \frac{1}{n} \left(\frac{1}{2} + \frac{1}{\nu} \right) \quad (4.54)$$

and its variance may be approximated by

$$\left\langle (\widehat{D}_1 - \langle D_1 \rangle)^2 \right\rangle \approx \frac{1}{n} \left(\langle (\ln I)^2 \rangle - \langle \ln I \rangle^2 - 1 \right) = \frac{1}{n} \left(\frac{\pi^2}{6} + \psi'(\nu) - 1 \right) \quad (4.55)$$

Here we have used the relation

$$I^{(2)} = 2 \left(\frac{\langle I \ln I \rangle}{\langle I \rangle} - \langle \ln I \rangle \right) \quad (4.56)$$

Figure 4.5 shows the mean normalised density and the approximated bias in its estimate for different sample sizes and Figure 4.6 shows the approximated standard deviation of the estimate divided by the mean normalised density.

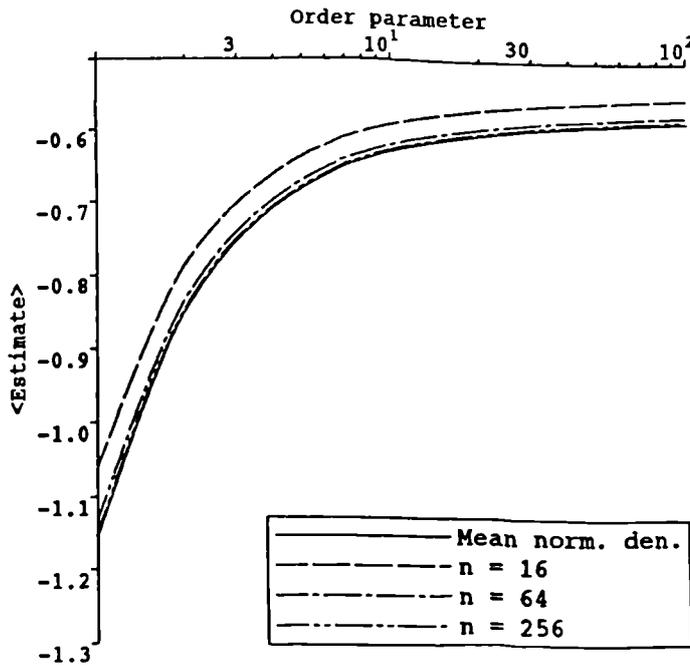


Figure 4.5: The mean normalised density (D_1) and the approximated bias in its estimate \widehat{D}_1 for different sample sizes.

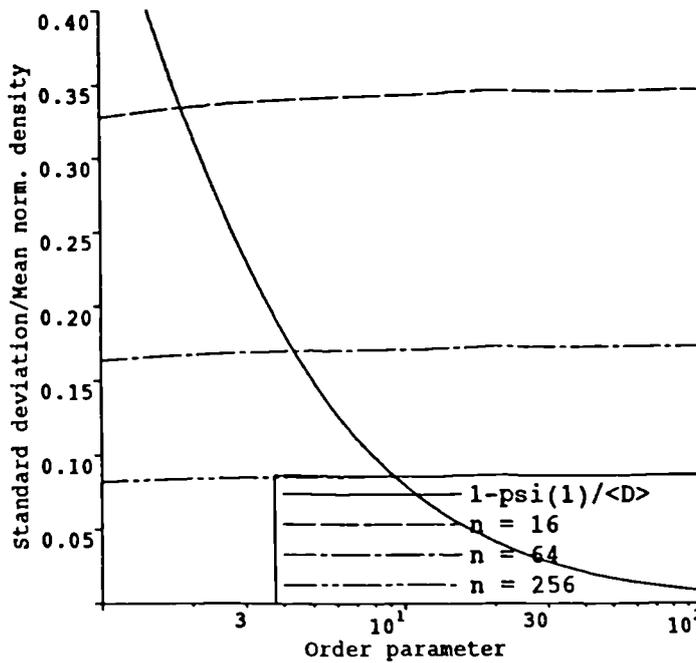


Figure 4.6: Approximated standard deviation in the estimate \widehat{D}_1 of the mean normalised density $\langle D_1 \rangle$ divided by $\langle D_1 \rangle$. Above the solid curve given by $1 - \psi(1) / \langle D \rangle$, the standard deviation is greater than the difference in the mean normalised density between when the order parameter equals ν and when it is infinite.

The bias in the estimate varies between $3/2n$ when $\nu = 1$ and $1/2n$ when $\nu \rightarrow \infty$, and is much less significant than the bias in the estimate of the second normalised intensity moment. The standard deviation in the estimate is also less significant particularly when the order parameter is small. These results suggest that the estimate of the mean normalised density is the more reliable method of estimating the order parameter. To check this the estimates were applied to two homogeneous regions, (a) the region whose intensity histogram was shown to fit an exponential distribution in Figure 3.3, and (b) the region whose intensity histogram was shown to fit a K-distribution of order parameter 10 in Figure 3.4. Both estimates were used to determine the order parameter in nine discrete 16×16 pixel sub-regions. The results are plotted in Figure 4.7. Note, the algorithm for estimating the order parameter using the mean normalised density generates a maximum value of 1000.

In both region (a) and (b) using either estimate there is a large variability in the order parameter estimated in different sub-regions. However, in each region the range of variability is larger estimating the moment from the second normalised intensity moment than estimating it from the mean normalised density. In region (a) whose histogram fits an exponential distribution, the estimate was greater than 100 only 4 out of 9 times using the second normalised intensity moment whilst it was greater than 100 6 out of 9 times using the mean normalised density. In region (b) whose histogram fits a K-distribution with order parameter 10 the estimate fell between 5 and 20 only 3 out of 9 times using the second normalised intensity moment whilst it fell between 5 and 20 6 out of 9 times using the mean normalised density. There appears to be a small degree of correlation between the two estimates. These results appear to confirm that the mean normalised density is the more reliable method of estimating the order parameter.

Note that setting $m = n = 2$ in the expression (3.87) for the joint moments of the real and imaginary components of the complex field gives

$$\frac{\langle Z_{\mathfrak{R}}^2 Z_{\mathfrak{I}}^2 \rangle}{\langle Z_{\mathfrak{R}}^2 \rangle \langle Z_{\mathfrak{I}}^2 \rangle} = 1 + \frac{1}{\nu} \quad (4.57)$$

Thus the order parameter may also be estimated from the real and imaginary components of the field; $\langle Z_{\mathfrak{R}}^2 Z_{\mathfrak{I}}^2 \rangle$, $\langle Z_{\mathfrak{R}}^2 \rangle$ and $\langle Z_{\mathfrak{I}}^2 \rangle$ may be estimated from $\widehat{Z}_{\mathfrak{R}}^2$, $\widehat{Z}_{\mathfrak{I}}^2$ and $\widehat{Z}_{\mathfrak{R}}^2 \widehat{Z}_{\mathfrak{I}}^2$ respectively. This estimate is obviously related to the second normalised intensity moment but its performance has yet to be assessed.

4.3 Summary

In this chapter:

- Methods of estimating the mean intensity and texture in a homogeneous region of a SAR image were discussed.
- The effect of arithmetic averaging on intensity and amplitude statistics was analysed. The arithmetic average of a set of statistically identical

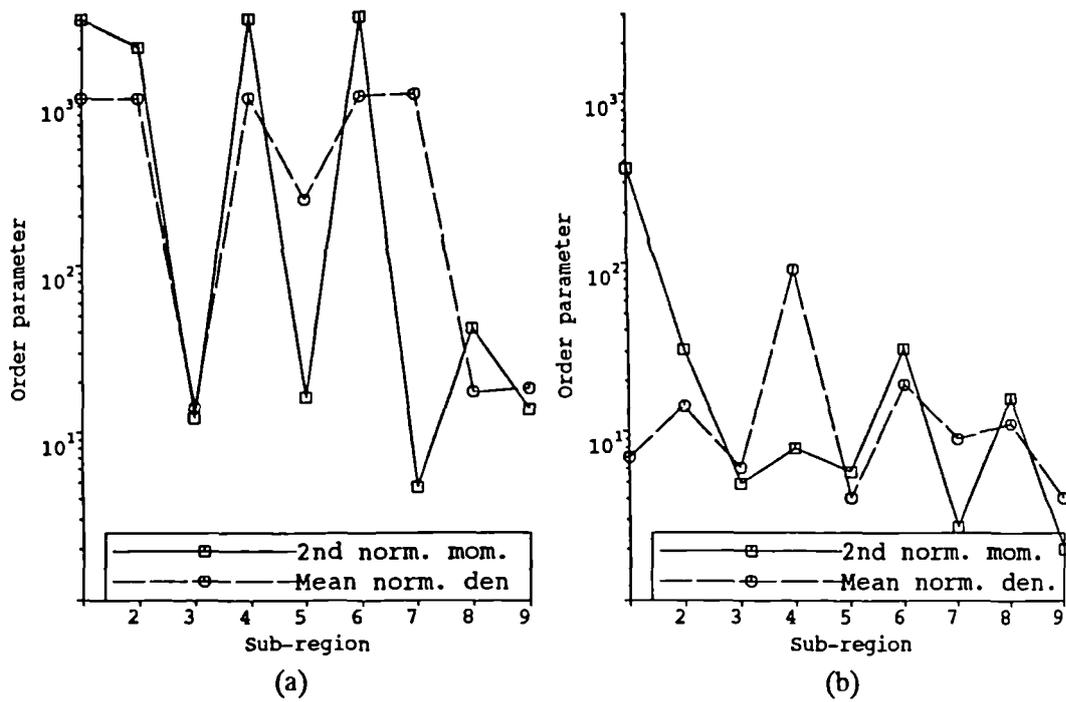


Figure 4.7: Result of estimating the order parameter in two homogeneous regions using the second normalised intensity moment and the mean normalised density. In each region the order parameter is estimated in nine discrete sub-regions each 16×16 pixels in size.

random variables gives an unbiased and consistent estimate of their mean. The effects of averaging also need to be known to determine the statistics of multi-look SAR data.

- The effects of correlation, texture and heterogeneous regions on the statistics of the average intensity were determined. All of these effects increased the coefficient of variation of the average.
- It was shown that when the imaged SCS is constant the average intensity is gamma distributed and that it gives the ML estimate of the mean intensity. The ML estimate of the mean amplitude is derived from the average intensity using (3.94).
- The ML estimates of the mean intensity and order parameter of K-distributed intensity were also derived but were found to be too complex to use in practice.
- Other methods of estimating the order parameter of K-distributed intensity, i.e., texture, were analysed. An estimate based on the mean normalised density was shown to be less biased and to be more stable than an estimate based on the second normalised intensity moment.

Chapter 5

Template matching

Template matching is an established way to detect a known feature in an image [14]. In this chapter the feasibility of matching templates representing map features to features in SAR images is investigated (see Appendix E for a discussion of the selection of digital map data suitable for matching features in SAR images). Normally no detailed radiometric information can be deduced from a given map feature about a feature in a SAR image matching it, only shape information. This precludes adopting a full **matched filter** approach. We can, however, often predict that certain types of linear target will be brighter (hedgcs, leading edges of woodland) or darker (roads and rivers) than their surroundings. By a linear target we mean a feature which appears in an image as a very thin segment bounded by parallel edges spaced on the order of a resolution cell apart. Features made up of intersections of such linear targets are suitable candidates for matching [27]. The shape of such a feature can be completely represented by a binary (i.e., 1-0) template. When the template is correlated with the image the position of the feature is established where the maximum correlation measure occurs. Ideally the probability of the maximum occurring at the 'true' match should be greater than the probability of it occurring at any partial matches. Feature shape determines what partial matches can occur, these in turn affect the shape of peaks in the template matched image. The exact manner in which partial matches affect feature detectability depends on the correlation measure being used.

Two correlation measures for matching templates representing map features to SAR image features are introduced in Section 5.1; standard cross-correlation which has no knowledge of speckle, and a correlation measure based on the Kolmogorov-Smirnov test of fit, which does take speckle into consideration. The behaviour of the two correlation measures at a match and as they move away from a match are analysed in Section 5.2. The analysis predicts that the Kolmogorov-Smirnov measure will perform better than cross-correlation at matching bright features but that the reverse is true when dark features are matched. In Section 5.3 the performance of the measures at matching bright features in simulated and real SAR data are compared. The reasons for restricting this performance analysis to *bright* features only are, that they are often much more distinctive than dark features in a SAR image and are thus the

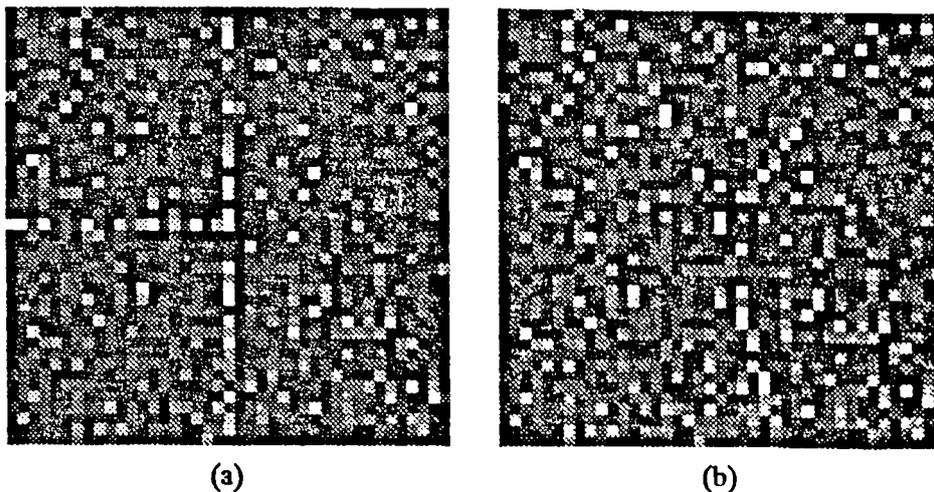


Figure 5.1: Simulated single-look SAR intensity images containing features suitable for matching.

more obvious candidates for matching, and that we particularly wished to check the predicted improvement in performance afforded by the Kolmogorov-Smirnov measure. This prediction is confirmed when both measures are used to match features in simulated images. However, neither measure performs well on real data. Reasons for this are considered and alternative methods for matching map features to SAR images are suggested in Section 5.4.

Only the matching of templates to single-look SAR intensity images will be considered. Initially it will be assumed that pixel values are spatially independent and that the imaged SCS is constant over extended targets, i.e., all clutter is purely due to speckle and pixel intensity is exponentially distributed. This shall be referred to as the simple image model. Examples are shown in Figure 5.1 of simulated SAR images generated according to this model; *feature-a* which consists of a 'T junction' intersect of *two* linear targets, and *feature-b* which consists of a 'Y junction' intersect of *three* linear targets. In each simulated image, the values of pixels making up features and the values of pixels making up their background are sampled from exponentially distributed populations with means μ_f and μ_b respectively where $\mu_f : \mu_b = 4 : 1$. Figure 5.2 shows templates to be used for matching features in these simulated images.

5.1 Correlation measures

5.1.1 Cross-correlation

The cross-correlation $X(\bar{y})$ between a template $t(\bar{y})$ and an image $I(\bar{y})$ is given by

$$X(y_a, y_r) = \frac{1}{n} \sum_{w_a=0}^{n_a-1} \sum_{w_r=0}^{n_r-1} t(w_a, w_r) I(w_a + y_a, y_r + w_r) \quad (5.1)$$



Figure 5.2: Templates for matching features shown in Figure 5.1.



Figure 5.3: Images produced when the templates in Figure 5.2 are matched to the images in Figures 5.1 using cross-correlation.

where the template is defined over a window $n_a \times n_r$ pixels in size. We have chosen to divide by n , the number of non-zero pixels in the binary template which is given by

$$n = \sum_{w_a=0}^{n_a-1} \sum_{w_r=0}^{n_r-1} t(w_a, w_r) \quad (5.2)$$

If the template selects pixels from a single population the value output by cross-correlation is an unbiased and consistent estimate of the mean of that population, and if the population is exponentially or gamma distributed the output gives the ML estimate of the mean of the population (see Chapter 4).

Figure 5.3 shows the images produced by matching the templates in Figure 5.2 to the images in Figure 5.1 using cross-correlation. If a feature is brighter than the background ($\mu_f > \mu_b$), on average its position will be indicated by the global maximum in the template matched image. Likewise the position of the feature will be indicated by the global minimum when it is darker than the background. In both Figures 5.3 (a) and 5.3 (b) the maximum correctly identifies the position of the feature at the centre of the image. However, the peak is more clearly defined in (b).

We are interested in the probability that the maximum cross-correlation value occurs at the match. Let $p_{\bar{y}}(X)$ be the probability distribution of the pixel value at position \bar{y} in the template matched image. Ignoring the correlation between values in the template matched image, the probability that the

maximum occurs at position \bar{y}_0 is given by [42]

$$\text{prob}(\bar{y}_0 = \text{max}) = \int_0^\infty p_{\bar{y}_0}(X) \left(\prod_{i \neq 0} \int_0^X p_{\bar{y}_i}(W) dW \right) dX \quad (5.3)$$

where \bar{y}_i represents all points in the image other than \bar{y}_0 . In practice there will be a large degree of correlation between values in the template matched image and this should be taken into account when interpreting results.

When the template selects pixels wholly from the feature or its background, the pixels contributing to the correlation measure may be considered to be sampled from a single population, whilst at partial matches the pixels may be considered to be sampled from two populations. Assuming the simple image model applies, when n pixels are sampled from a single population with mean μ the value output by cross-correlation will be gamma distributed (3.59) with order parameter n and mean μ (see Section 4.1.1). Whilst when p pixels with mean μ_f are selected from the feature and $n - p$ pixels with mean μ_b are selected from the background, the value output by cross-correlation will have the same distribution as the sum of two gamma distributed random variables with respective mean values $\frac{p}{n}\mu_f$ and $(1 - \frac{p}{n})\mu_b$, and order parameters p and $n - p$. The pdf of such a distribution is given by (C.1).

Figure 5.4 shows the results of applying this analysis to the features in Figure 5.1. For both features the probability that the maximum occurs at the match increases as the ratio of the mean intensity of the feature to its background increases. However, the probability that the maximum occurs at the match is about twice as great for *feature-b* as it is for *feature-a*. For *feature-a* there is never more than a 50% probability that the maximum occurs at the match. For *feature-b*, where the ratio of the mean intensity of the feature to its background is large, the probability that the maximum occurs at the match is greater than 90%.

This difference in detectability is due to the partial matches that can occur. With *feature-a* as the template moves away from the match, up to two-thirds of the pixels can still be selected from the feature. With *feature-b* only a third of pixels can still be selected from the feature. This can be seen in Figure 5.5 which shows the ACF's of the templates used for matching *feature-a* and *feature-b*. At each point the height of the ACF represents the number of matching pixels, the central peak represents the match and the ridges partial matches. The partial matches are much stronger in the case of *feature-a* as compared to *feature-b*. Using cross-correlation the probability that the maximum will occur at any partial match in the template matched image increases as the number of pixels selected from the feature at that partial match increases. Thus there is a greater probability that wrong detections will occur when *feature-a* as compared to *feature-b*, is matched. These results agree with the peak in Figure 5.3 (b) being more clearly defined than that in Figure 5.3 (a). It is clear that the partial matches that occur between a template and the feature it is used to match should be kept to a minimum, i.e., ideally the ACF of the template should resemble a delta function.

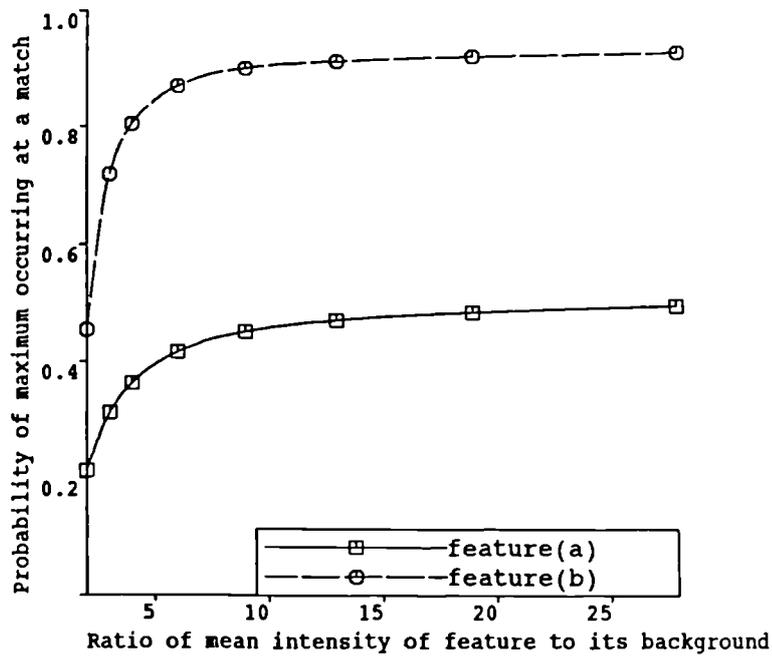


Figure 5.4: Probability of the maximum occurring at the match versus feature to background ratio.

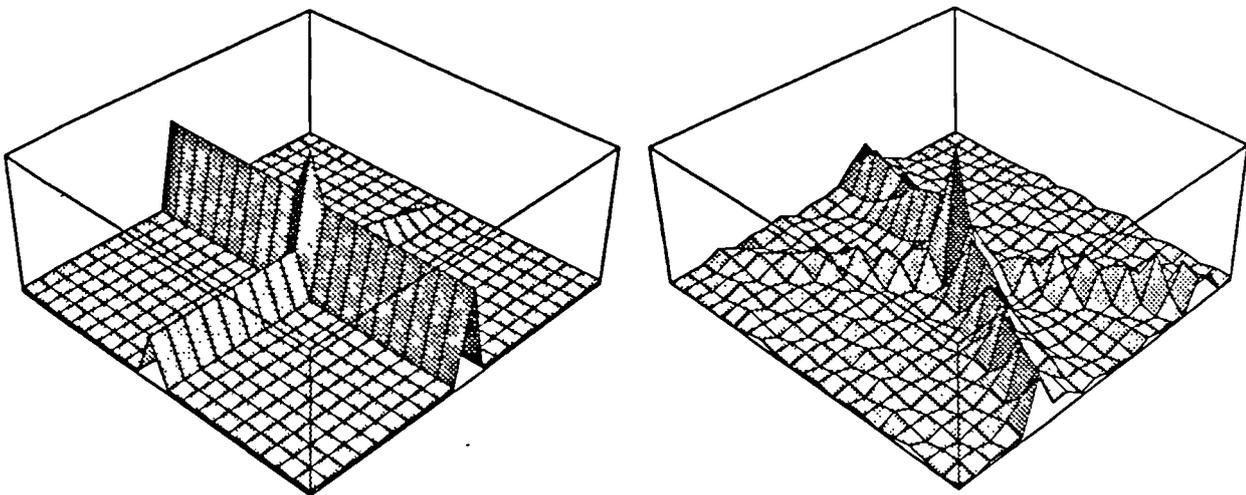


Figure 5.5: ACF's of templates used for matching *feature-a* and *feature-b*.



Figure 5.6: Images produced when the templates in Figure 5.2 are matched to the images in Figures 5.1 using the Kolmogorov-Smirnov measure.

Where the simple image model does not apply, i.e., where pixel intensity is K-distributed or correlated, the output of cross-correlation will have greater variance. This will increase the probability of the maximum occurring at partial matches.

5.1.2 The Kolmogorov-Smirnov correlation measure

Cross-correlation has no knowledge of the underlying populations in a single-look SAR intensity image. The **Kolmogorov-Smirnov test of fit** can be used to develop a correlation measure similar to cross-correlation in that it generates an estimate of the local mean over homogeneous areas, but which takes distributional information into account.

Assuming the simple image model applies a population of pixel intensity values is characterised by a single parameter; the mean value of an exponential distribution. At each position in the image the Kolmogorov-Smirnov test is used to find the mean value of the exponential distribution which best fits the pixel values selected by the template at that position. The Kolmogorov-Smirnov test of fit is preferred to the χ^2 test of fit [33, 49, 56] as the latter is not reliable with small sample sizes.

The position of bright (dark) features matching the template should again be indicated by the global maximum (minimum) in the template matched image. Figure 5.6 shows the template matched images produced by matching the templates in Figure 5.2 to the images in Figure 5.1 using the ‘Kolmogorov-Smirnov’ correlation measure; in both cases the match is correctly indicated by a central peak. Note that these are sharper than those produced using cross-correlation and that once again *feature-b* generates a more clearly defined peak than *feature-a* (see Figure 5.3).

5.2 Comparison of correlation measures

The Cross-correlation (CC) and Kolmogorov-Smirnov (KS) measures were compared by examining how the peaks in the template matched image produced by each measure differed. This was achieved by examining the behaviour of the

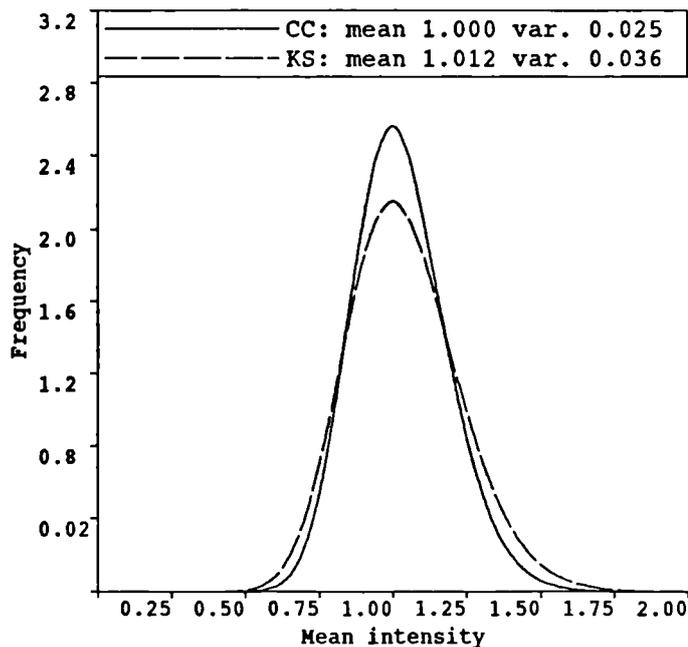


Figure 5.7: Probability distribution of the estimate of the mean intensity of a homogeneous feature with unit mean, produced by the CC and KS correlation measures, when a template, $n = 40$ in size, matches the feature.

measures at a match and as a template moves away from a match.

5.2.1 Behaviour of the measures at a match

At a match the value produced by both measures is an estimate of the mean intensity of the feature they are matching. For a template of size $n = 40$ matching a feature with unit mean intensity, the probability distributions of the estimates of the mean intensity produced by each measure were generated from 1,000,000 simulations with different realisations of exponential speckle. Figure 5.7 shows the distribution of the estimated values. The results for the CC measure were in agreement with theory, which predicts that it should be gamma distributed with degree 40 and unit mean. The probability distribution for the KS measure had a similar shape but a slightly larger variance and a small positive bias.

5.2.2 Behaviour of the correlation measures as a template moves away from a match

A test was carried out to determine the expected values output by the two measures over a range of partial matches. Suppose a template selects p pixels with mean μ_x and $n - p$ pixels with mean μ_y . Assuming the simple image model applies the samples will be distributed according to

$$f_I(I) = \frac{p}{n} \frac{1}{\mu_x} e^{-\frac{I}{\mu_x}} + \left(1 - \frac{p}{n}\right) \frac{1}{\mu_y} e^{-\frac{I}{\mu_y}} \quad (5.4)$$

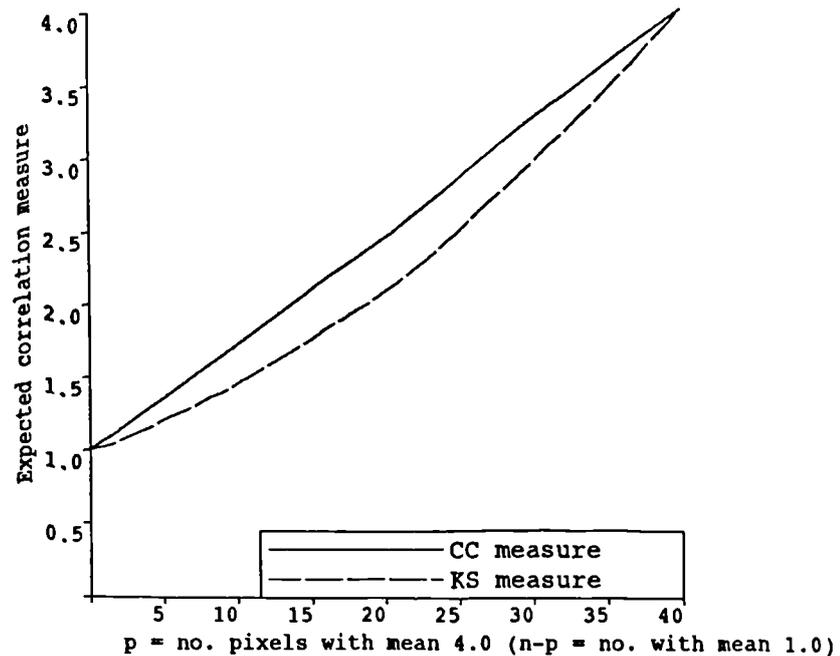


Figure 5.8: Expected cross-correlation and Kolmogorov-Smirnov measures generated by a template matching p pixels with mean $\mu_x = 4.0$ and $n - p$ pixels with mean $\mu_y = 1.0$ ($n = 40$).

The expected CC measure is given by

$$\mu_p = \frac{p}{n}\mu_x + \left(1 - \frac{p}{n}\right)\mu_y \quad (5.5)$$

The expected KS measure was calculated numerically by finding the mean of the exponential distribution which best fitted the distribution defined by (5.4).

Figure 5.8 shows a plot of the expected values for $p = 0, \dots, n$, when $\mu_x = 4.0$, $\mu_y = 1.0$ and $n = 40$. The expected value of the CC measure varies linearly from μ_y to μ_x as p increases. The expected value of the KS measure varies non-linearly; for small values of p the KS measure increases at a slower rate than the CC measure but as p gets larger the KS measure increases more quickly, eventually approaching μ_x at a faster rate than the CC measure. Thus when the feature being matched is brighter than its background ($\mu_f > \mu_b$) the expected KS measure will decrease more rapidly than the expected CC measure as we move away from the match. The reverse is the case whenever the feature being matched is darker than its background ($\mu_b > \mu_f$). As the ratio of the mean intensity between a feature and its background $\mu_f : \mu_b$ increases, the difference between the two measures increases.

The results indicate that when a template is matched to a bright feature in a dark background the KS measure will produce a sharper but slightly more variable maximum than the CC measure and that when a template is matched to a dark feature in a bright background the CC measure will produce a sharper

but slightly less variable minimum than the KS measure. Thus when the simple image model applies we would expect the KS measure to be better than the CC measure at matching bright features but to be worse at matching dark features.

5.3 Performance tests

To check the predicted improvement in performance afforded by the Kolmogorov-Smirnov measure the performance of both measures at detecting a variety of *bright* features in simulated and real data was compared.

5.3.1 Simulated data

Both measures were used to match templates to simulated images containing different types of bright features in a range of speckle realisations. Four types of image feature were considered:

1. Thin (single pixel wide) features with uniform background.
2. Thick (three pixel wide) features with a uniform background.
3. Thin features with a non-uniform background.
4. Thick features with a non-uniform background.

Examples of type 1 are shown in Figure 5.1. Each feature split its background into three areas; for uniform backgrounds the ratio of the mean intensity of the feature to the mean intensity of the three areas was 4:1:1:1, for non-uniform backgrounds the ratio of the mean intensity of the feature to that of the three areas was 8:3:2:1. This ratio was chosen so that the overall ratio of the energy in the background to that in the feature would be approximately the same as with a uniform background. For each type of feature, 1000 images containing such a feature but with different speckle realisations were generated. To see how the measures performed over homogeneous areas a similar series of images of pure speckle were also generated.

Measurements were carried out on each template matched image to determine how sharp the maximum peak was and how much it deviated from its expected position at the match. The measurements comprised of

1. Comparing the height of the maximum peak to the next highest peak and the highest col surrounding it to determine how distinct the maximum peak was.
2. Measuring the minimum gradient falling away from the maximum peak and the gradient to the highest surrounding col to determine how sharp the maximum peak was.
3. Calculating the standard deviation of the maximum from the match and counting the total number of times deviations occurred to determine how reliable the correlation measures were.

A summary of the results for the features shown in Figure 5.1 is given in Table 5.1.

The KS measure performs consistently better than the CC measure; peaks were sharper (steeper gradients), as predicted in Section 5.2.2, and better separated (greater ratios), and the maximum deviated less from its expected position at the match. Features in uniform backgrounds were only slightly better matched than features in non-uniform backgrounds. With thick features there exists more than one position where the template matches the feature. Thus we would expect thick features to be matched more often, but less accurately, than thin features. This is the case with *feature-a* but not with *feature-b*. With *feature-b* the proportion of pixels that can be selected from the feature at a partial match is a third for thin features but two-thirds for thick features; this increases the wrong detection rate. With *feature-a* this proportion only increases slightly from thin features to thick features.

In agreement with the results given in Section 5.1.1, *feature-b* was matched better than *feature-a*, though *feature-a* was matched correctly more often than predicted (79% as compared to 36% in Figure 5.4). Correlation between pixels in the template matched image reduces the effect partial matches have on the detectability. Global maxima in homogeneous images were not as distinct nor as sharp as global maxima in images containing matching features, they were also uniformly distributed over the image. The results of matching other feature shapes were similar.

5.3.2 Real Data

Three sub-images (a) to (c), each containing a feature that was thought to be detectable using template matching, were extracted from a RSRE X-band SAR image (see Figure 5.9). Each feature is a 'T junction' intersect of two linear targets. Even though such features were shown in Section 5.2 to be more susceptible to partial matches than 'Y junction' intersects of three linear targets, the former were chosen for matching as they were the most prevalent type of intersect found in the image. Templates for matching these features were extracted from 1:10,000 scale map data, raster scanned and transformed to the same frame of reference as the SAR data (see Figure 5.10), and by visual inspection, from the image itself (see Figure 5.11). This was to ensure there existed a template which perfectly matched the image feature.

The results of matching using both correlation measures and both sets of templates are shown in Figures 5.12 and 5.13. Using the templates extracted from the map, only feature (b) is matched correctly using the KS method and no features are matched correctly using the CC method. Using templates extracted from the image feature (b) is correctly matched using both methods, feature (a) is only correctly matched using the KS method whilst neither method manages to match feature (c) correctly. The poorer performance of the map templates may be due to them being thicker than those extracted from the image. A morphological thinning algorithm [32] was applied to the map templates to produce single pixel wide templates. The only improvement after thinning was

MATCHING USING CROSS-CORRELATION							
Width	Bkgd	Ratio1	Ratio2	Grad1	Grad2	SDev	WDet
<i>feature-a</i>							
1	Unif	0.877	0.761	0.808	0.723	3.738	20.8
1	NUnif	0.884	0.770	0.765	0.698	3.763	25.0
3	Unif	0.902	0.798	0.367	0.631	3.027	12.3
3	NUnif	0.905	0.800	0.373	0.625	3.219	14.3
Homogeneous		0.959	0.822	0.158	0.205	-	-
<i>feature-b</i>							
1	Unif	0.747	0.613	1.094	1.084	1.568	3.1
1	NUnif	0.770	0.633	1.045	1.034	1.628	3.8
3	Unif	0.868	0.762	0.373	0.585	2.162	18.2
3	NUnif	0.875	0.766	0.373	0.581	2.155	19.2
Homogeneous		0.956	0.799	0.166	0.203	-	-

MATCHING USING THE KOLMOGOROV-SMIRNOV MEASURE							
Width	Bkgd	Ratio1	Ratio2	Grad1	Grad2	SDev	WDet
<i>feature-a</i>							
1	Unif	0.829	0.647	1.222	1.092	3.078	15.5
1	NUnif	0.850	0.670	1.082	0.986	3.267	24.1
3	Unif	0.857	0.722	0.505	0.926	1.872	3.6
3	NUnif	0.870	0.729	0.525	0.910	2.388	8.6
Homogeneous		0.949	0.763	0.235	0.308	-	-
<i>feature-b</i>							
1	Unif	0.631	0.492	1.548	1.511	0.812	1.2
1	NUnif	0.697	0.536	1.473	1.370	1.202	3.4
3	Unif	0.792	0.676	0.488	0.800	1.336	8.3
3	NUnif	0.815	0.691	0.508	0.777	1.600	12.3
Homogeneous		0.946	0.746	0.248	0.287	-	-

Table 5.1: Averaged results of performance tests on 1000 simulated images. Width - width of matching feature in the image; 1 or 3 pixels. Bkgd - type of background; uniform or non-uniform. Ratio1 - value of 2nd highest peak : value of maximum peak. Ratio2 - value of highest col surrounding the maximum peak : value of maximum peak. Grad1 - least gradient dropping away from maximum peak. Grad2 - gradient from the maximum peak to its highest surrounding col. SDev - standard deviation of the maximum peak from the match (pixels). WDet - percentage of wrong detections.

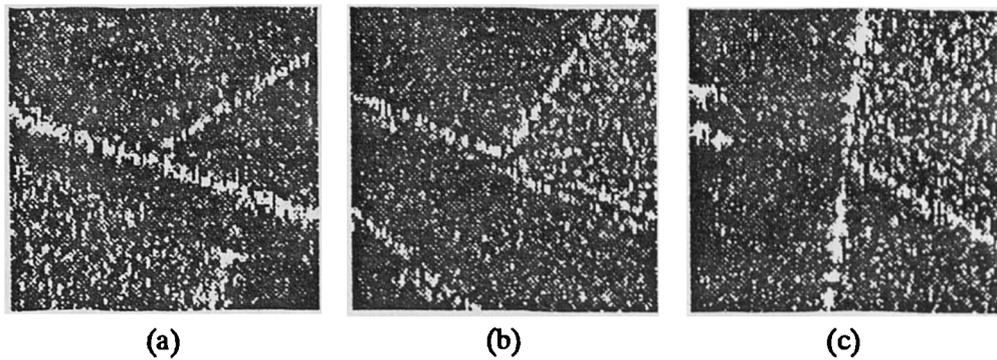


Figure 5.9: Features extracted from a RSRE X-band SAR image (100×100 pixel subimages), which were thought to be detectable using template matching.

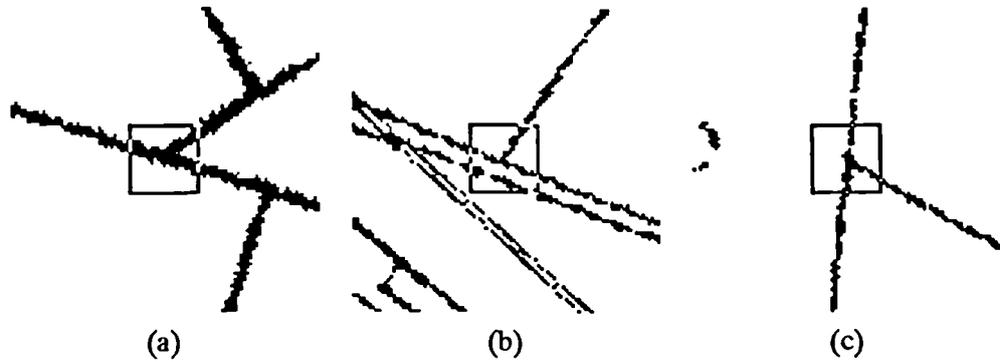


Figure 5.10: Templates (shown in square boxes) derived from map data for matching the features shown in Figure 5.9.

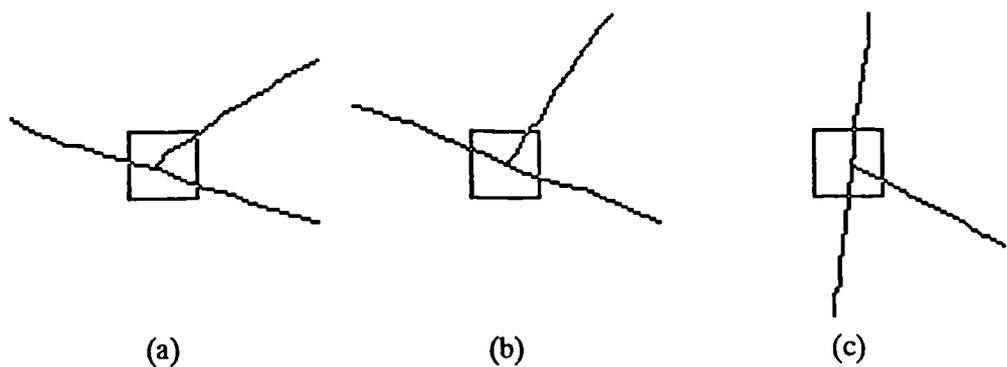


Figure 5.11: Templates (shown in square boxes) derived from the image itself for matching the features shown in Figure 5.9.

that feature (b) was correctly matched using the CC method as well as the KS method.

From this small number of tests it appears that neither measure performs as well on real data as on simulated data. Even though the SAR features being looked at are much brighter relative to their background ($\mu_f : \mu_b = 10 : 1$) than the features in the simulated images were ($\mu_f : \mu_b = 4 : 1$), a much larger number of wrong detections appear to be occurring.

To investigate why this is the case the statistics of pixel values sampled along the above features were analysed. As each feature consists of a 'T junction' intersect of two linear targets, separate measurements were carried out on the linear targets making up the long and short arms of the feature. Figure 5.14 shows profiles of how sampled values, a 21 pixel moving average of these values (an estimate of mean intensity) and a Kolmogorov-Smirnov fit of the distribution of these 21 pixels to an exponential distribution, varies along the linear targets. For each feature the point at which the two linear targets making it up intersect is in the middle of the longer profile shown on the left, and is at the left hand end of the shorter profile shown on the right.

The average intensity can be seen to vary along all three profiles. The variability in the exponential fit shows that this variance cannot be explained by the simple image model alone. There is a correlation between where the profiles are brightest and where detections occurred. With features (a) and (c) partial matches with bright parts of the features led to wrong detections. Feature (b) was only matched correctly because it is at its brightest close to the intersection. These observations show that the simple image model and the CC and KS measures are insufficient for template matching features in SAR images.

5.4 Discussion

The fact that map and SAR image data are mismatched data types places constraints on the use of simple matching techniques. While shape is common to both data types, without a priori knowledge, radiometric information is only available in the sense that certain types of feature are likely to be brighter or darker than their surroundings. When the simple image model applies, distributional information can be used to improve methods of locating the correct match. However, an important outcome of this chapter is that types of linear targets useful for matching tend not to fit the simple image model. They are made up of heterogeneous populations and/or a fluctuating imaged SCS. This has implications for any image analysis method based on using simple image models.

Because a constant imaged SCS cannot be assumed along linear targets and because cross-correlation is not a reliable means of detecting matches between map and image data, more powerful methods are required. The simplest such extension is to use an image model which allows for fluctuations in the imaged SCS along linear targets. The evidence in Figure 5.14 suggests that such an extension is needed to model the intensity along the types of linear targets (hedges and lines of trees?) picked out from the images shown in Figure 5.9. This is in-

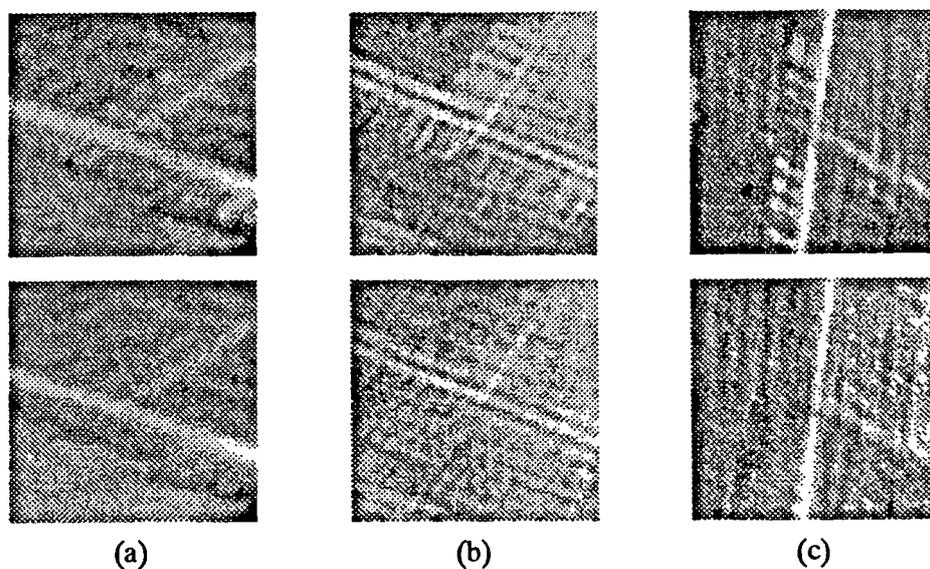


Figure 5.12: Result of matching the templates derived from the map data and shown in Figure 5.10, to the sub-images shown in Figure 5.9 using the CC method (top) and the KS method (bottom).

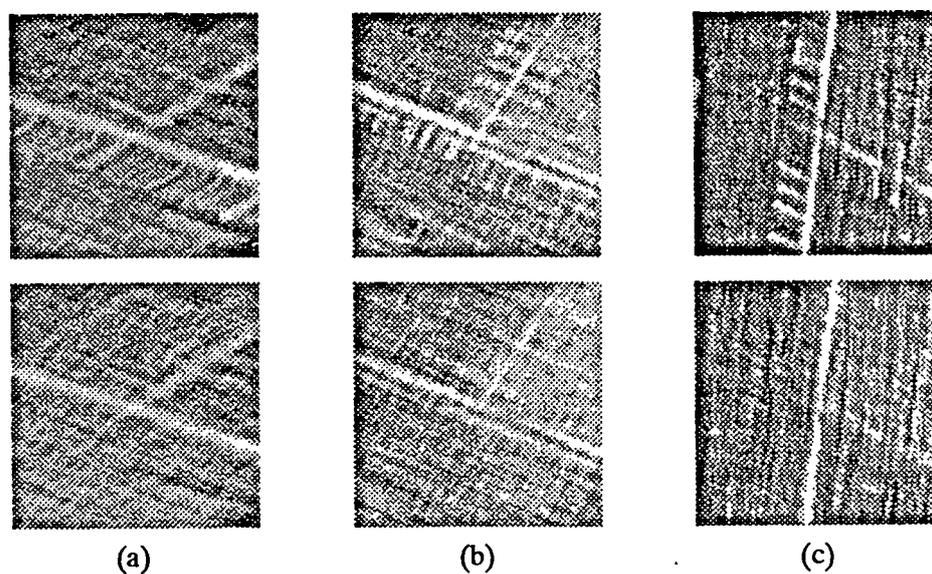


Figure 5.13: Result of matching the templates derived from the image itself and shown in Figure 5.11, to the sub-images shown in Figure 5.9 using the CC method (top) and the KS method (bottom).

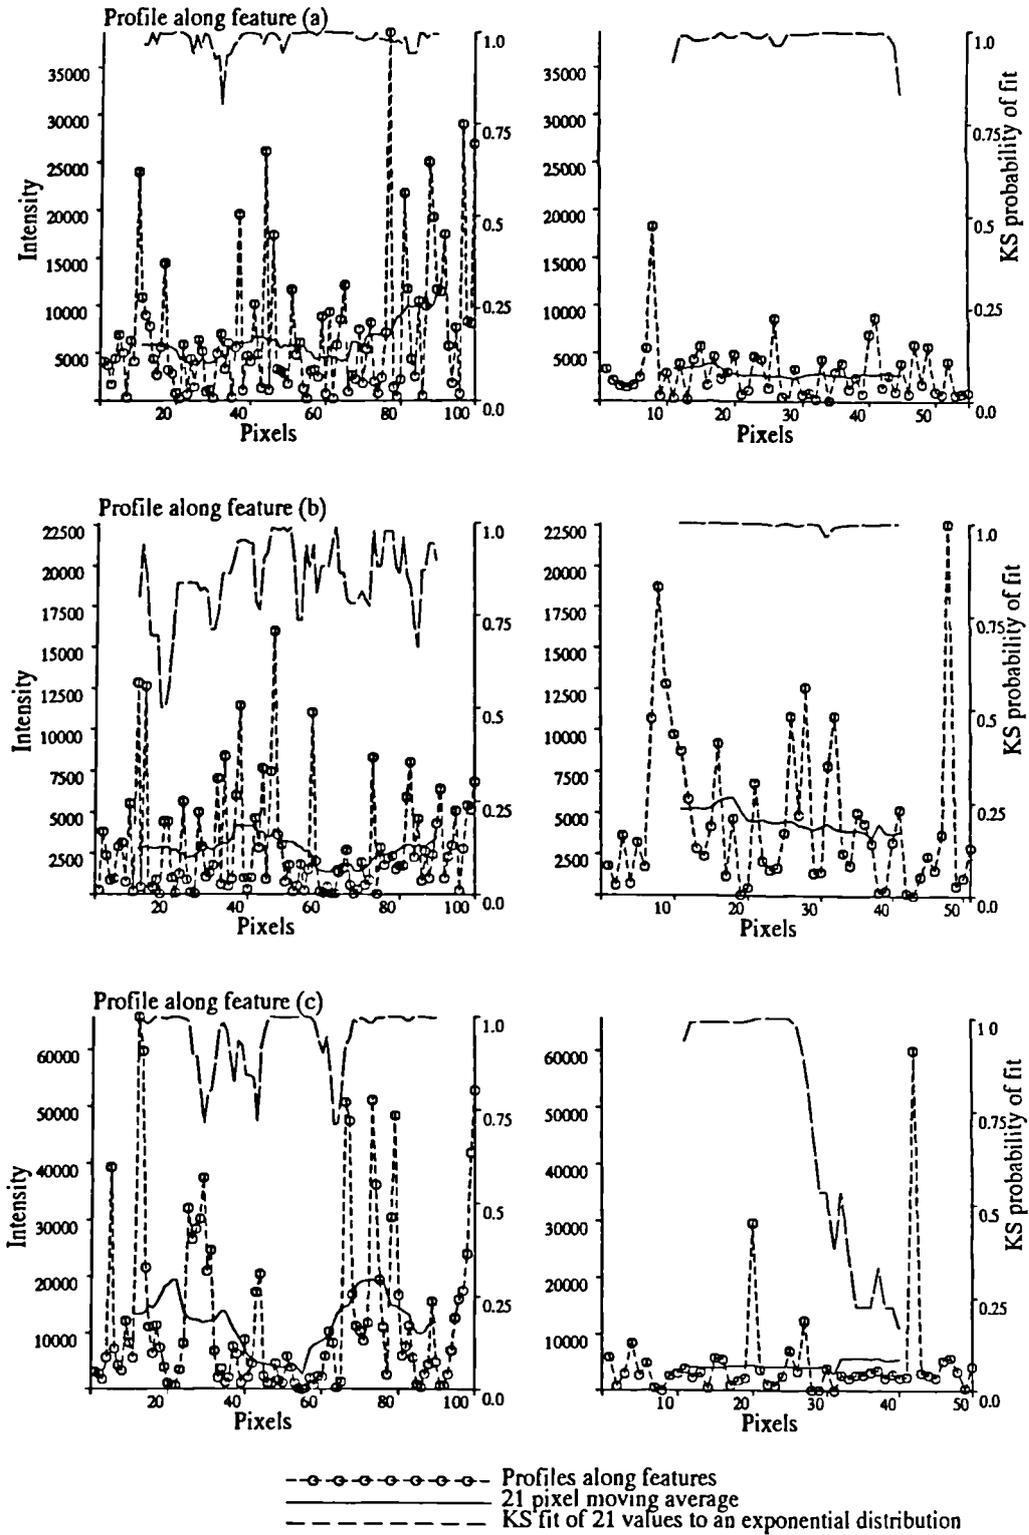


Figure 5.14: Profiles taken along linear targets.

investigated in the next chapter (Chapter 6) where the K-distributed model for the intensity over extended targets discussed in Chapter 3, is extended to describe the intensity along edges and linear targets. The fit of pixels intensity along edges and linear targets to both exponential and K-distributions is also more fully tested. Even if it is assumed that intensity is K-distributed along linear targets, using the Kolmogorov-Smirnov test, fitting a two parameter K-distribution is much slower than fitting a single parameter exponential distribution, which in itself is several orders of magnitude slower than cross-correlation. If a correlated K-distribution is required four parameters would be necessary [44]. The ensuing high processing overheads do not make this an attractive option.

The problem of mismatched data types can be removed by structuring the image by applying edge and line detectors prior to matching. The correlation process is then better defined, but new questions are raised about the representation of the SAR image. These questions concern the reliability of the detectors, their ability to cope with speckle, edge thickening and positional accuracy, and the generation of excessive edge/line detections in textured regions. Nonetheless, this seems to offer real hope because of the development of effective edge detectors and segmentation algorithms. These methods are investigated in Chapters 7 to 9.

A third approach which relies less critically on the geometrical accuracy needed for correlation methods is to use topological concepts. This still requires the SAR image to be represented in binary form. Graph theory, fuzzy logic or matching using regions then provide possible methods for matching the global structure of the map and data representations. Such methods have been used in the optical bands [41, 12] and for SAR [53], but must currently be regarded as tentative rather than proven.

5.5 Summary

In this chapter:

- Correlation measures for matching map templates to features in single-look SAR intensity images, in particular features consisting of intersections of linear targets which are brighter than their surroundings, have been investigated
- It was shown that feature shape determines the kind of partial matches that can occur and that in turn this affected feature detectability.
- When a simple image model was assumed, analysis predicted that a correlation measure based on the Kolmogorov-Smirnov test of fit would perform better than cross-correlation at matching bright features in simulated SAR images, but that the reverse would be true for dark features.
- The first part of this prediction was confirmed when the performance of the correlation measures at matching bright features in simulated SAR images generated according to the simple image model was compared.

- However, neither measure performed well at detecting bright features in real SAR data. Reasons for this were considered.
- The simple image model was found to be inadequate for modelling intensity along linear targets. Although more general image models are available (K-distributions), due to the extra parameters involved in defining them, a modified version of the Kolmogorov-Smirnov measure would be appreciably slower.
- Alternative methods for matching map features to SAR images using edge detection, segmentation and graph theory were suggested.

Chapter 6

Statistics along edges

In Chapter 3 it was generally assumed that the surfaces being imaged were stationary, i.e., the elementary scatterers within a resolution cell were from a single population. This assumption is valid for modelling the intensity statistics over extended targets. However, these models do not describe the intensity statistics over the whole of the image. For instance, although they describe the statistics to either side of an edge they do not describe the intensity statistics along the actual edge, whether it is abrupt or gradual.

By an edge we mean a *boundary in an image* across which a transition in mean intensity or texture occurs as opposed to a *physical transition* in the underlying scene. Certainly there is always a correspondence between edges in an image and physical transitions. However, this is not always a one to one mapping. A single physical transition may be characterised by several edges in an image, e.g., a region of radar shadow is characterised by two edges. For purposes of image analysis edges initially need to be characterised as image features rather than as physical features as this is how they are first encountered, and understood, by any low-level image interpretation scheme. It is only when higher order knowledge is applied that edges may be interpreted in terms of the physical transitions they represent.

Whilst a model for intensity statistics over extended targets is sufficient for optimally detecting a large number of features, e.g., detecting edges bounding extended targets (see Chapter 7), as seen in the last chapter a model for the intensity statistics along edges is needed to optimally detect linear targets. Edge detectors perform poorly at linear target detection, especially when the linear targets separate regions of the same intensity and texture. In addition as linear targets can confuse edge detection and segmentation algorithms they need to be detected and suitably flagged prior to carrying out these operations (see Chapter 8). To detect linear targets, measurements comparing pixel values *along the feature itself* to pixel values on either side of it are needed. Methods that have been investigated other than template matching include the Hough transform [26] and dynamic programming [55]. Whichever method is used, models for the statistics along linear features, as well as over extended targets, are required to set decision thresholds.

In this chapter the K-distribution model described in Chapter 3 for the in-

tensity over extended targets in SAR images is further developed to cover the intensity along edges. By analysing the mechanisms which give rise to edges in SAR images, models for the detected field along an edge are derived in Section 6.1. It is shown that the most general model is given by a sum of scattering components from M different populations plus components describing specular and secondary scattering, and system noise. This may be written

$$Z = \sum_{i=1}^M \sum_{j=1}^{n_i} z_{ij} + z_{sp} + z_{sc} + z_{sn} \quad (6.1)$$

where there are n_i independent scattering contributions z_{ij} from the i 'th population, z_{sp} and z_{sc} represent the specular and secondary scattering components and z_{sn} the system noise term. The statistics of the intensity along edges described by (6.1) are derived in Section 6.2. In Section 6.3 measurements of intensity statistics along edges and bright linear targets are compared to predictions made by the model.

6.1 Detected field along an edge

In this section we examine how to represent the detected field along an edge arising from a change in surface type or cover. We then go on to consider how to represent the detected field along edges generated by changes in terrain. For simplicity the system response $h(\mathbf{x}, \mathbf{x}_i)$ in (3.1) will be approximated by a two dimensional rectangular function which may be taken as representing the main lobe of a more realistic *sinc* shaped system response (3.141).

Edges are the result of changes in the underlying scattering population altering the statistics of the detected field. Thus by definition, along an edge the detected field will be the result of scattering from two or more populations. For the case of M such populations the detected field may be expressed

$$Z = \sum_{i=1}^M \sum_{j=1}^{n_i} z_{ij} + z_{sn} \quad (6.2)$$

where each n_i is a random variable representing the number of independent and statistically identical complex scatterer contributions $z_{i_1}, \dots, z_{i_{n_i}}$ from the i 'th population within the resolution cell. The mean proportion of scatterers from the i 'th population is given by $p_i = \mu_{n_i} / \mu_n$, where n_i has mean μ_{n_i} and $\mu_n = \mu_{n_1} + \dots + \mu_{n_M}$; obviously $p_1 + \dots + p_M = 1$. We have chosen to add an additional term z_{sn} to take account of system noise.

The most obvious form of edge arises from a spatial change in the surface type or cover (e.g., a transition from bare soil to grass). Assuming pixel values over the homogeneous regions to either side of an edge may be modelled as the result of scattering from single populations, pixel values along the edge may be modelled as the result of scattering from just two populations. This may be written

$$Z = \sum_{j=1}^{n_1} z_{1j} + \sum_{j=1}^{n_2} z_{2j} + z_{sn} \quad (6.3)$$

As natural edges are never perfectly straight the mixture of the two populations $p_1 : p_2$, will vary from pixel to pixel along the edge. In fact even if the edge is perfectly straight the mixture of the two populations will vary from pixel to pixel unless the edge lies parallel to or at 45 degrees diagonal to the sampling grid on which the image is defined.

The detected field from a resolution cell containing a linear target (e.g., a river or road) of width less than the resolution will be the result of scattering from three populations; the linear target and the extended targets to either side. This can be expressed

$$Z = \sum_{j=1}^{n_1} z_{1j} + \sum_{j=1}^{n_2} z_{2j} + \sum_{j=1}^{n_3} z_{3j} + z_{sn} \quad (6.4)$$

When the extended targets to either side have the same scattering properties the detected field is just given by (6.3). Of course if the resolution is much greater than the width of the linear target it will have no significant effect. The linear target will then appear as a simple edge (6.3) or will not be apparent at all if the extended targets to either side have the same scattering properties.

6.1.1 Terrain effects

Changes in surface slope affect the incidence angle of the radar beam with the surface and thus both the imaging geometry and the behaviour of scatterers. Within our definition of changes in surface slope we include not just topographic effects but also changes in vegetation height, land/water boundaries and man made features. Here we will consider changes in surface slope in the range direction only. This is not because changes in surface slope in the azimuth direction do not generate edge features, but rather because the most marked effects are due to changes in surface slope in the range direction (i.e., layover and shadowing).

If ϑ_0 is the incidence angle of the radar beam with a horizontal surface, the incidence angle of the beam with a surface inclined at an angle ϑ_s in the range direction is given by $\vartheta = \vartheta_0 - \vartheta_s$; here ϑ_s is positive when the surface is inclined towards the SAR platform and is negative when it is inclined away from it. If the SAR has slant range resolution d_{sr} the ground range resolution (2.29) is given by $d_{gr} = d_{sr} / \sin \vartheta$ and the illuminated length normal to the radar beam is given by $d_{il} = d_{sr} / \tan \vartheta = d_{gr} \cos \vartheta$. Both of these quantities decrease as the incidence angle increases. The decrease in the ground resolution reduces the number of scatterers per resolution cell, while the decrease in the illuminated length normal to the radar beam reduces the power incident on each resolution cell. Thus even if isotropic scattering is assumed, the total power received from a resolution cell will decrease as the incidence angle increases. In the majority of cases this effect is magnified by the fact that backscatter per unit area from most rough surfaces decreases with incidence angle [59].

At small incidence angles additional effects come into play. When the incidence angle is greater than 30 degrees most backscatter is due to surface roughness. However, at smaller incidence angles specular scattering will occur [15, 59].

As the incidence angle approaches zero, specular scattering will produce very large backscatter values. Added to the fact that both d_{gr} and d_{il} are very large in such cases, this will often cause saturation of the SAR receiver.

6.1.1.1 Changes in surface slope

The decrease in received power with incidence angle causes rapid changes in surface slope to produce rapid changes in the detected field, i.e., edges. Along the transition between regions with incidence angles ϑ_1 and ϑ_2 ($\vartheta_1 < \vartheta_2$) the ground resolution will lie in the interval

$$\frac{d_{sr}}{\sin \vartheta_1} > d_{gr} > \frac{d_{sr}}{\sin \vartheta_2} \quad (6.5)$$

and the illuminated length normal to the beam will lie in the interval

$$\frac{d_{sr}}{\tan \vartheta_1} > d_{il} > \frac{d_{sr}}{\tan \vartheta_2} \quad (6.6)$$

Thus we would expect the received power along the edge to lie between the values of the received power in the regions to either side. With the change in surface slope across the edge the effective density of scatterers will change, along with the behaviour of single scatterers if scattering is non-isotropic. Thus the detected field along an edge produced by a change in surface slope may be modelled as the sum of scattering from two populations with different densities (6.3) plus an additional contribution z_{sp} to describe any specular scattering from either surface.

$$Z = \sum_{j=1}^{n_1} z_{1j} + \sum_{j=1}^{n_2} z_{2j} + z_{sp} + z_{sn} \quad (6.7)$$

The above discussion assumes the incidence angles ϑ_1 and ϑ_2 are both positive and are less than $\pi/2$. When the incidence angle goes below zero, layover results, whilst when it is greater than $\pi/2$, radar shadowing occurs. The amount of layover and shadow in an image depends on the look angle. At small look angles there will be more layover and less shadow while at large look angles there will be more shadow and less layover.

6.1.1.2 Layover

Figure 6.1 shows how different regions of the surface contribute to range gates in a region of layover [10, 37]. Within the region of layover, range gates 2-5 in Figure 6.1, the total backscatter will be the result of scattering from the slope causing the layover (Surface 2) and the slopes in front of and behind it (Surfaces 1 and 3 respectively). This results in d_{il} , the illuminated length normal to the beam per resolution cell, being greater over the region of layover than to either side of it.

Let $d_{il_1} = d_{sr} / \tan \vartheta_1$, $d_{il_2} = d_{sr} / \tan \vartheta_2$ and $d_{il_3} = d_{sr} / \tan \vartheta_3$. Within the region of layover (range gates 3 and 4) $d_{il} = d_{il_1} + d_{il_2} + d_{il_3}$, whilst in front of and behind it (range gates 1 and 6) $d_{il} = d_{il_1}$ and d_{il_3} respectively, and in the

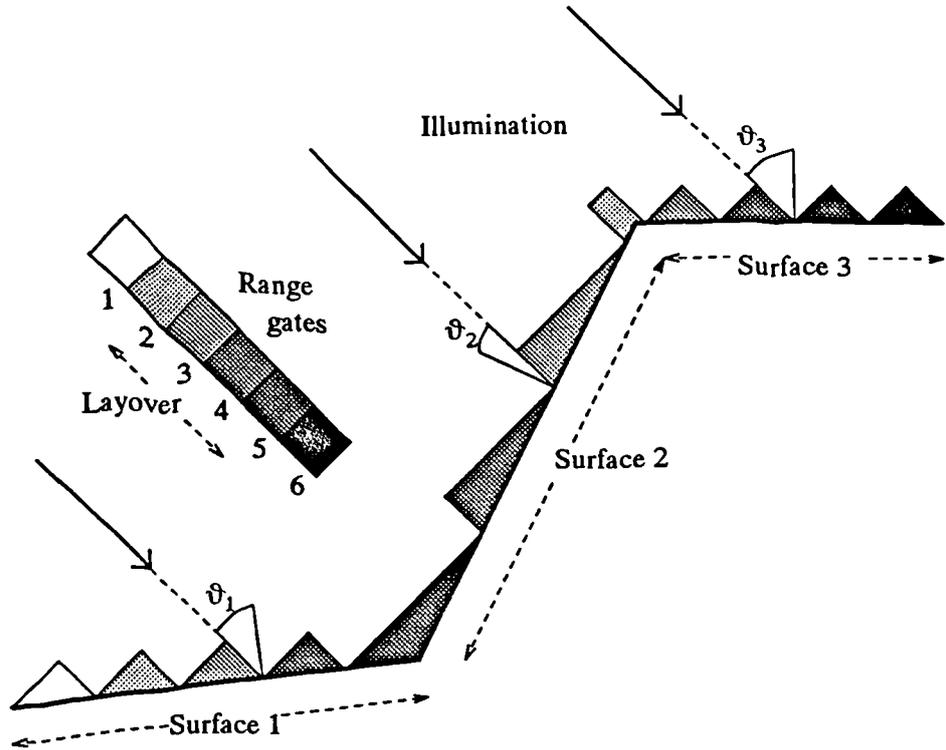


Figure 6.1: Example of imaging geometry resulting in layover. The regions on the surface contributing to each range gate are indicated by shading.

transition regions entering and leaving the region of layover (range gates 2 and 5) d_{il} will lie in the ranges $d_{il_1} < d_{il} < d_{il_1} + d_{il_2} + d_{il_3}$ and $d_{il_3} < d_{il} < d_{il_1} + d_{il_2} + d_{il_3}$ respectively. This increase in illuminated length, coupled with the fact that specular scattering will occur if ϑ_2 is small, causes the received power from regions of layover to be much greater than that in adjacent areas in an image. Thus layover is indicated by large changes in intensity. As layover is the result of scattering from three regions, pixel values along edges of regions of layover need to be modelled as the result of scattering from three populations.

The large changes in gradient associated with layover effects may also allow strong secondary scattering (double bounce) to occur; these effects will be strongest when the surface facets producing the secondary scattering lie at right angles to each other. However, due to the additional range delay secondary scattering induces, such effects will only appear after the region of layover (range gates 5 and 6). Thus these effects only need to be included in our model for pixel values along the trailing edge of the region of layover.

Combining all the above effects the detected field along the leading edge of a region of layover is given by

$$Z = \sum_{j=1}^{n_1} z_{1j} + \sum_{j=1}^{n_2} z_{2j} + \sum_{j=1}^{n_3} z_{3j} + z_{sp} + z_{sn} \quad (6.8)$$

where z_{sp} represents specular scattering. Likewise the detected field along the

trailing edge of a region of layover is given by

$$Z = \sum_{j=1}^{n_1} z_{1j} + \sum_{j=1}^{n_2} z_{2j} + \sum_{j=1}^{n_3} z_{3j} + z_{sp} + z_{sc} + z_{sn} \quad (6.9)$$

where z_{sc} represents secondary scattering.

6.1.1.3 Radar shadow

Where the incidence angle is greater than $\pi/2$ the radar beam does not intercept the surface. This continues to be the case until the surface once again comes into line of sight of the beam. Over this period several range gates may have been passed without any backscatter return; this produces a region of radar shadow where all that is detected is the system noise. Thus along an edge bounding a region of radar shadow the detected field is given by

$$Z = \sum_{j=1}^{n_1} z_i + z_{sn} \quad (6.10)$$

Often the leading edge of a region of radar shadow will correspond with the trailing edge of a region of layover. The detected field along such an edge may be described by (6.9) except with the sum in z_3 , removed

$$Z = \sum_{j=1}^{n_1} z_{1j} + \sum_{j=1}^{n_2} z_{2j} + z_{sp} + z_{sc} + z_{sn} \quad (6.11)$$

In this case secondary scattering may appear in the region of radar shadow. Where the leading edge of a region of layover corresponds with the trailing edge of a region of radar shadow the sum in z_1 , in (6.8) disappears and the detected field is given by

$$Z = \sum_{j=1}^{n_2} z_{2j} + \sum_{j=1}^{n_3} z_{3j} + z_{sp} + z_{sn} \quad (6.12)$$

When the surface in front of the region causing layover is in shadow no secondary scattering will occur. Over mountainous terrain an image may consist entirely of regions of layover and radar shadow. The detected field along edges is then described by

$$Z = \sum_{j=1}^{n_2} z_{2j} + z_{sp} + z_{sn} \quad (6.13)$$

All of the above expressions for the detected field along an edge can be seen to be special cases of the general expression given at the beginning of the chapter for the detected field along an edge (6.1), with M at most equalling three. The statistics of the intensity arising from (6.1) are analysed in the next section.

6.2 Intensity statistics along edges

6.2.1 Characteristic function of the intensity

When speckle is fully developed in each population contributing to the detected field along an edge, each independent scattering contribution can be treated as a circularly symmetric complex random variable. Assuming that the specular and secondary scattering contributions and the system noise term are also circularly symmetric random variables with respective amplitudes a_{sc} , a_{sp} and a_{sn} , using the **Hankel transform** the detected field expressed by (6.1) will then have characteristic function

$$\Phi_Z(\omega) = \prod_{i=1}^M \langle J_0(\sqrt{\rho_i}\omega) \rangle^{n_i} \langle J_0(a_{sp}\omega) \rangle \langle J_0(a_{sc}\omega) \rangle \langle J_0(a_{sn}\omega) \rangle \quad (6.14)$$

where $\rho_i = |z_i|^2$ represents the SCS of scatterers in the i 'th population and $J_0(x)$ is the zeroth-order Bessel function of the first kind.

The birth-death-immigration process adopted in Chapter 3 and [29] to model fluctuations in the number of independent contributions from a single population, can also be used to model the number of independent contributions n_i from each of M populations. It will be assumed that each n_i in (6.14) is controlled by an independent birth-death-immigration process which results in it having a negative binomial distribution (3.70) with mean μ_{n_i} and order parameter ν_i . Averaging over the n_i 's and normalising the SCS contributions ρ_i by μ_n , from (3.75) the expected value of the characteristic function is given by

$$\Phi_Z(\omega) = \prod_{i=1}^M \left[1 + \frac{p_i \mu_n}{\nu_i} \left(1 - \left\langle J_0 \left(\sqrt{\frac{\rho_i}{\mu_n}} \omega \right) \right\rangle \right) \right]^{-\nu_i} \langle J_0(a_{sp}\omega) \rangle \langle J_0(a_{sc}\omega) \rangle \langle J_0(a_{sn}\omega) \rangle \quad (6.15)$$

Substituting the series expansion for $J_0(x)$ defined by (F.1) into (6.15) and assuming the total mean number of scatterers is large, i.e., $\mu_n \rightarrow \infty$, the characteristic function of the detected field tends to

$$\Phi_Z(\omega) \rightarrow \prod_{i=1}^M \left[1 + \frac{p_i \langle \rho_i \rangle \omega^2}{4\nu_i} \right]^{-\nu_i} \langle J_0(a_{sp}\omega) \rangle \langle J_0(a_{sc}\omega) \rangle \langle J_0(a_{sn}\omega) \rangle \quad (6.16)$$

The characteristic function of a circular symmetric complex random variable z whose intensity $|z|^2$ is K-distributed with mean μ and order parameter ν is given by [29]

$$\Phi_z(\omega) = \left[1 + \frac{\mu \omega^2}{4\nu} \right]^{-\nu} \quad (6.17)$$

Therefore from (6.16) it can be seen that the contribution from the i 'th population to the detected field is K-distributed in intensity with mean $p_i \langle \rho_i \rangle$ and order parameter ν_i .

6.2.2 Intensity moments

The pdf of the intensity of the detected field may be derived from the characteristic function (6.16) using the inverse Hankel transform (F.11)

$$f_I(I) = \frac{1}{2} \int_0^\infty \Phi_Z(\omega) J_0(\omega\sqrt{I}) \omega d\omega \quad (6.18)$$

Solutions to this integral for the case $M = 1$ and $a_{sp} = a_{sc}$ equalling a constant are discussed in [17, 29]; i.e., specular and/or secondary scattering of the same constant amplitude with a homogeneous K-distributed background. This equates to the situation along an edge between a region of layover and a region of radar shadow (6.13).

Although it has not been possible to derive an analytic solution to (6.18) in the general case the intensity will have mean

$$\mu_I = \mu_{I_0} + \langle a_{sp}^2 \rangle + \langle a_{sc}^2 \rangle + \langle a_{sn}^2 \rangle \quad (6.19)$$

where

$$\mu_{I_0} = \sum_{i=1}^M p_i \langle \rho_i \rangle \quad (6.20)$$

is the mean intensity of the field Z_0 due to scattering from the M populations but excluding the specular and secondary scattering contributions and the system noise. In addition, the intensity moments can be generated using the results in Appendix D, or by repeatedly differentiating the characteristic function of the detected field (6.16) and using (F.12).

By substituting the intensity moments of each of the M populations given by (3.54) into (D.1) the intensity I_0 of the field due to scattering from the M populations will have moments

$$\begin{aligned} \langle I_0^m \rangle = & m! \sum_{k_1=0}^m \sum_{k_2=0}^{k_1} \dots \sum_{k_{M-1}=0}^{k_{M-2}} \frac{m!}{(m-k_1)!(k_1-k_2)! \dots k_{M-1}!} \times \\ & \frac{\Gamma(m-k_1+\nu_1) \beta_1^{\nu_1} \Gamma(k_1-k_2+\nu_2) \beta_2^{\nu_2} \dots \Gamma(k_{M-1}+\nu_{M-1}) \beta_M^{\nu_M}}{\Gamma(\nu_1) \Gamma(\nu_2) \dots \Gamma(\nu_M)} \end{aligned} \quad (6.21)$$

where $\beta_i = p_i \langle \rho_i \rangle / \nu_i$. These are the moments of the product of a unit mean exponentially distributed random variable and the weighted sum of M independent gamma distributed random variables where the i 'th random variable has mean $\langle \rho_i \rangle$ order parameter ν_i and weighting p_i . Thus, in accordance with the multiplicative model (3.37), the intensity I_0 is given by the product of a unit mean exponential speckle process and the imaged SCS due to scattering from the M populations given by

$$\Upsilon_0 = \sum_{i=0}^M p_i \rho_i \quad (6.22)$$

and which has characteristic function

$$\Phi_{\Upsilon_0}(\omega) = \prod_{i=1}^M \left[1 - \frac{p_i \langle \rho_i \rangle \omega}{\nu_i} \right]^{-\nu_i} \quad (6.23)$$

The multiplicative model for the intensity still applies because speckle is fully developed in each population and the total number of scattering contributions μ_n is large. The fact that scatterers are from heterogeneous populations has no effect on this result.

Assuming that only the phases of the specular and secondary scattering contributions are random, from (D.1) the field given by Z_0 plus the specular and secondary scattering contributions has normalised intensity moments given by

$$I^{(m)} = \frac{m!}{(1 + r_{sp} + r_{sc})^m} \sum_{k_1=0}^m \sum_{k_2=0}^{k_1} \frac{m!}{(m - k_1)!(k_1 - k_2)!k_2!} \frac{r_{sp}^{m-k_1}}{(m - k_1)!} \frac{r_{sc}^{k_1-k_2}}{(k_1 - k_2)!} \Upsilon_0^{(k_2)} \quad (6.24)$$

where $r_{sp} = a_{sp}^2/\mu_{I_0}$ and $r_{sc} = a_{sc}^2/\mu_{I_0}$. When μ_{I_0} is much greater than the intensity of both the specular and secondary scattering contributions, they will both have negligible effect on the total intensity I whose distribution will equal that of I_0 . When μ_{I_0} is much less than the intensity of either the specular or the secondary scattering contribution but is much greater than the other, the normalised intensity moments tend to unity and the intensity to a constant equalling a_{sp}^2 or a_{sc}^2 depending on which is dominant. When μ_{I_0} is much less than the intensity of both the specular and secondary scattering contributions the normalised intensity moments will tend to

$$I^{(m)} = \frac{m!}{(1 + r_{sp} + r_{sc})^m} \sum_{k_1=0}^m \binom{m}{k_1} \frac{r_{sp}^{m-k_1}}{(m - k_1)!} \frac{r_{sc}^{k_1}}{k_1!} \quad (6.25)$$

which equals

$$I^{(m)} = \frac{1}{2^m} \binom{2m}{m} = \frac{(2m - 1)(2m - 3) \dots 3 \cdot 1}{m!} \quad (6.26)$$

when the specular and secondary scattering contributions have equal amplitude.

6.2.3 The effect of system noise

The effect of system noise on the detected intensity has not yet been considered. It is normally assumed that the real and imaginary components of the system noise are independent and Gaussian distributed with the same variance σ_{sn}^2 and zero mean. The system noise is then exponentially distributed in intensity and is most easily represented as a constant noise term $\langle a_{sn}^2 \rangle = 2\sigma_{sn}^2$ added to the imaged SCS. This means that system noise is indistinguishable from a contribution from a scatterer population which has a constant imaged SCS of the same magnitude, and is the reason why an image appears speckled with constant imaged SCS when no signal is detected in regions of radar shadow. For an airborne SAR, but not for a spaceborne SAR, we would expect the mean backscatter to system noise ratio to be large and the effect of the latter can be ignored except over regions of low backscatter and radar shadow.

6.2.4 Statistics of the imaged SCS

In the remainder of this section the statistics of the imaged SCS due to scattering from M populations will be analysed to determine how much it differs from the imaged SCS due to scattering from a single population. The characteristic function of the imaged SCS given by (6.23) has the same form as the characteristic function of the average of intensity values sampled from exponential populations with different mean values discussed in Section 4.1.1. When $M = 2$ and ν_1 and ν_2 are integers, the pdf of the imaged SCS is given by (C.1). More generally, the normalised moments of the imaged SCS given by differentiating the characteristic function are

$$\Upsilon_0^{(2)} = 1 + q_2 \quad (6.27)$$

$$\Upsilon_0^{(3)} = 1 + 3q_2 + 2q_3 \quad (6.28)$$

$$\Upsilon_0^{(4)} = 1 + 6q_2 + 3q_2^2 + 8q_3 + 6q_4 \quad (6.29)$$

etc, where

$$q_k = \sum_{i=1}^M \frac{p_i^k \langle \rho_i \rangle^k}{\mu_{I_0}^k \nu_i^{k-1}} \quad (6.30)$$

The normalised moments lie between those of gamma distributed random variables with order parameters $M\hat{\nu}$ and ν_{min} , where $\hat{\nu} = (\nu_1 + \nu_M)/M$ and $\nu_{min} = \min(\nu_i)$. This is because

$$\frac{1}{(M\hat{\nu})^{k-1}} \leq q_k \leq \frac{1}{\nu_{min}^{k-1}} \quad (6.31)$$

The normalised moments tend to one, i.e., constant imaged SCS, when there are no underlying fluctuations ($\nu_i \rightarrow \infty$), as we would expect.

The normalised moments will tend to those of a gamma distributed random variable with order parameter $M\hat{\nu}$ when the contributions to q_k from different population are equal. This will occur under a number of conditions

1. When there is an equal mean proportion of scatterers from each population and the mean SCS of scatterers from a population given by $\langle \rho_i \rangle$ is inversely proportional to the variance in their density given by $1/\nu_i$, i.e., when p_i and $\langle \rho_i \rangle / \nu_i$ are constant. $\hat{\nu} / \mu_{I_0}$ will then equal $\nu_i / \langle \rho_i \rangle$.
2. Along an edge where there is a change in texture but none in intensity and the mean proportion of scatterers from a population is inversely proportional to the variance in their density, i.e., when $\mu_{I_0} = \langle \rho_i \rangle$ and p_i / ν_i is constant. $M\hat{\nu}$ will then equal ν_i / p_i .
3. Along an edge where there is a change in intensity but none in texture and the mean proportion of scatterers from a population is inversely proportional to its mean SCS, i.e., when $\hat{\nu} = \nu_i$ and $p_i \langle \rho_i \rangle$ is constant.

When the contribution to the q_k 's from the i 'th population is much greater than that from all the others, i.e., when

$$\frac{p_i^k \langle \rho_i \rangle^k}{\nu_i^{k-1}} \gg \frac{p_j^k \langle \rho_j \rangle^k}{\nu_j^{k-1}} \quad (6.32)$$

for all $j \neq i$, the normalised moments will tend to those of a gamma distributed random variable with order parameter ν_i . This will occur from a combination of

1. The majority of scattering coming from the i 'th population ($p_i \gg p_j$), i.e., there being no edge present.
2. The mean SCS of scatterers from the i 'th population being much greater than the mean SCS of scatterers from other populations ($\langle \rho_i \rangle \gg \langle \rho_j \rangle$).
3. The variance in the density of scatterers from the i 'th population being much greater than the variance in the density of scatterers from other populations ($\nu_i \ll \nu_j$).

6.3 Measurement of intensity statistics along edges

The above analysis indicates that along many types of edge in SAR images the imaged SCS will have moments similar to if not exactly the same as that of a gamma distribution. However, due to speckle this can only be checked by indirect measurement of the statistics of the intensity which we would expect to be K-distributed if the imaged SCS is gamma distributed. Measuring intensity statistics along edges involves locating their exact position, and taking samples along them. Locating the exact position of edges in speckle is not always simple. In addition, due to the one dimensional nature of edges, only small samples can be taken and sampling errors will be large. Thus rather than using sample moments to check statistics, the Kolmogorov-Smirnov test was used to test the degree of fit of edges to K and exponential distributions. Measurements were carried out along twelve 100 pixel long edges selected from a RSRE X-band airborne SAR image (single-look intensity). All of the edges consist of a simple transition in intensity and are shown in Figure 6.2. They are all due to changes in vegetation cover except the last two, which show the boundary between a runway and vegetation (grass).

For each edge its exact position was located as best as possible and pixel values were sampled along it. The Kolmogorov-Smirnov test was then used to find the K-distribution which best fitted the samples. Table 6.1 lists the mean and order parameter of the K distribution which best fitted the samples, the mean of the exponential distribution which best fitted the samples and the degrees of fit. For comparison with the parameters of the best fit K-distribution, the mean and order parameter were also estimated directly from the data (see Chapter 4). Only K-distributions up to order 99 were fitted; for high order parameters the K-distribution is indistinguishable from the exponential distribution.

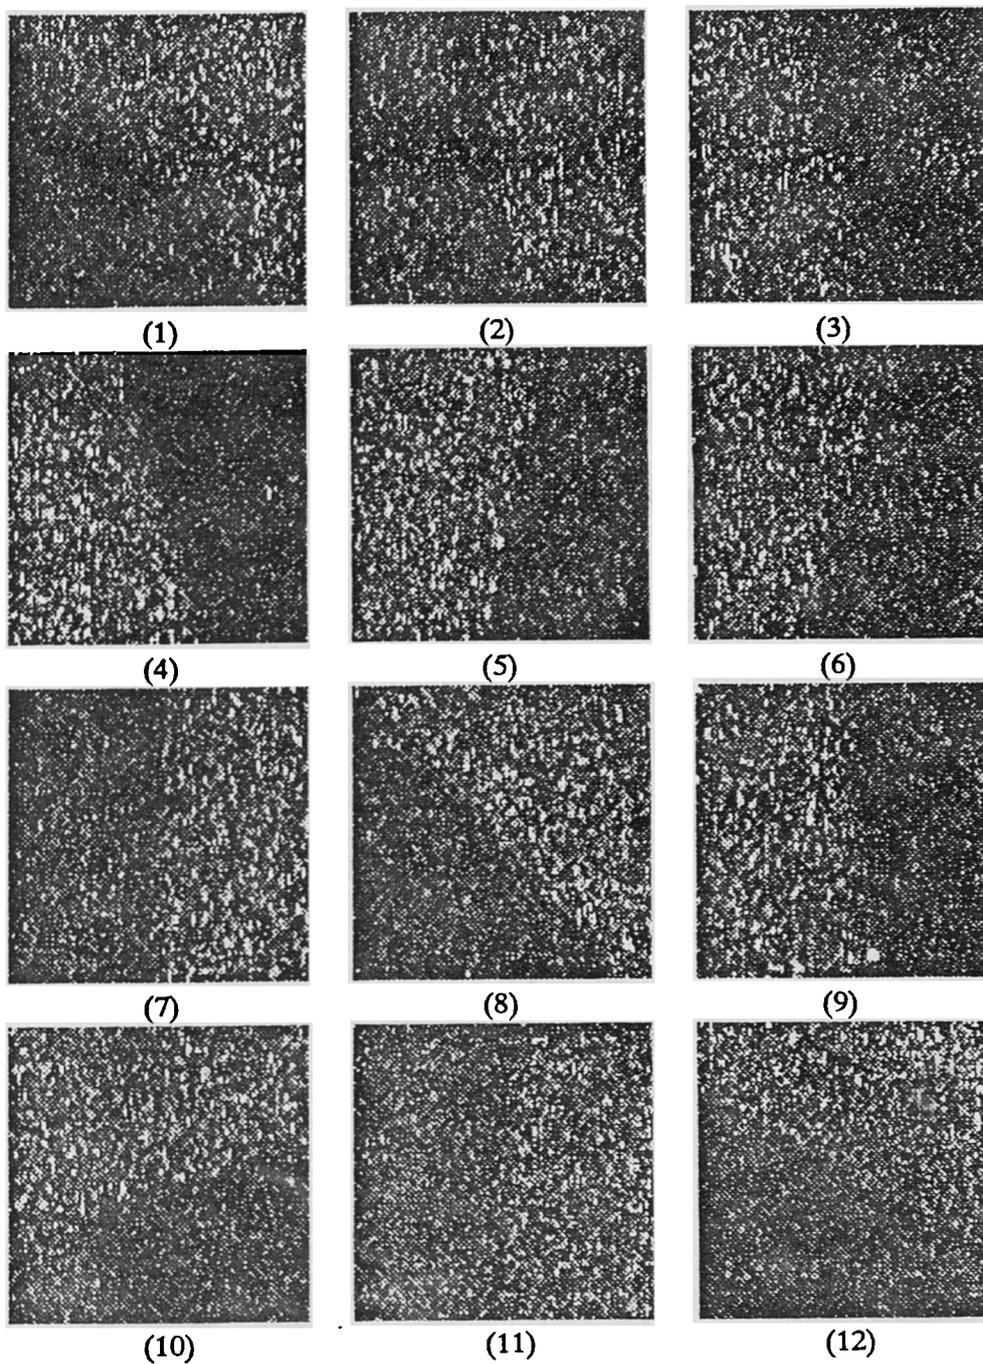


Figure 6.2: Edge features from a RSRE X-band SAR image (100×100 pixel subimages).

Profile	K-distribution			Exp. dist.		Estimates	
	μ	ν	fit	μ	fit	μ	ν
1	556.4	18	0.996	536.8	0.993	597.5	9
2	1020	95	0.990	1012	0.990	1036	1000
3	775.5	2	1.000	656.2	0.587	727.1	3
4	626.8	8	0.682	572.6	0.608	558.5	38
5	984.5	4	0.985	899.6	0.735	975.4	4
6	1310	2	0.999	981.6	0.766	1142	4
7	775.1	93	0.995	770.8	0.996	740.7	35
8	822.2	99	0.515	818.8	0.530	746.0	1000
9	986.5	9	0.957	905.4	0.935	1435	2
10	909.1	4	0.981	810.7	0.720	889.2	8
11	305.0	16	0.964	299.7	0.904	315.7	10
12	498.4	99	0.524	495.6	0.534	435.9	1000

Table 6.1: Best fit of samples from profiles along edges, to K and exponential distributions using the Kolmogorov-Smirnov test. For each profile the table lists: the mean and order parameter of the K-distribution which best fitted the samples and the degree of fit, the mean of the exponential distribution which best fitted the samples and the degree of fit, and the mean and order parameter estimated from the average intensity and mean normalised density respectively.

All profiles gave a good degree of fit (> 0.5); for all but three (4,8,12) the fit was very good (> 0.95). Four profiles (2,7,8,12) fitted exponential distributions better than K-distributions. As expected, profiles which best fitted K-distributions with high order parameters (> 10) gave a comparable degree of fit to an exponential distribution, while profiles which best fitted a K-distribution with low order parameter did not give as good a fit to an exponential distribution.

Except for profile 9 the mean of the best fit distribution never varied from the mean calculated from the sample moments by more than 16%. For profile 9 there was a large difference due to there being a point target 15dB brighter than the mean intensity lying along the edge. However as this only involved a few pixel values it had little effect on the overall fit of the profile; using the Kolmogorov-Smirnov test with a sample size of 100 a maximum deviation of 0.05 (5 pixels) from the fitted cumulative distribution function represents a degree of fit of 0.95, while a maximum deviation of 0.15 (15 pixels) represents a degree of fit of 0.05. For most profiles the estimated mean was smaller than the mean of the best fit distribution.

6.3.1 Measurements along linear targets

Knowledge of intensity statistics along edges is primarily required for detecting linear targets. To determine whether such features can be described by K-distributions the above measurements carried out along simple edges were also carried out along the six linear targets making up the three features used for template matching in the last chapter (see Figure 5.9).

For each of the features A to C the longer linear target making up the 'T junction' intersect will be referred to using the subscript one, i.e., A.1, B.1 and C.1, while the shorter one will be referred to using the subscript two, i.e., A.2, B.2 and C.2. Each of the linear targets is much brighter than the extended targets to either side; radar shadow is also apparent along the trailing edge of targets A.1 and C.2 (near range is towards the top of the images shown). All three features represent some form of boundary between fields. In the region where the image is taken (Wiltshire, England), fields boundaries are most often hedges or lines of trees; though they could be walls or fences in this case. As the RSRE X-band SAR has a very large look angle it is likely that specular scattering, but not layover, is significantly adding to the brightness of vertical features [4]. As these bright features are more distinctive than simple edges they are more easily located and we can be more certain that the profiles are correctly sampled along them rather than over the extended targets to either side. As already noted in the last section, to a good approximation we expect the imaged SCS to be gamma distributed, and thus the intensity to be K-distributed, when scattering results from several populations one of which (the hedge ?) is radiometrically much brighter than the others. Here we are ignoring any specular scattering component.

The results of measurements are listed in Table 6.2. Apart from profile C.2 a very good fit was found (> 0.9). One profile best fitted an exponential

Profile	K-distribution			Exp. dist.		Estimates	
	μ	ν	fit	μ	fit	μ	ν
A.1	6648	79	0.998	6656	0.997	6619	1000
A.2	3409	95	0.912	3394	0.918	3031	1000
B.1	3081	3	0.993	2650	0.658	2858	4
B.2	4887	35	0.998	4828	0.997	4643	6
C.1	12390	3	0.991	10910	0.543	12200	3
C.2	4571	96	0.171	4540	0.173	4884	4

Table 6.2: Best fit of samples from profiles along thin linear targets, to K and exponential distributions. The parameters listed for each profile are the same as those in Table 6.1.

Profile	K-distribution			Exp. dist.		Estimates	
	μ	ν	fit	μ	fit	μ	ν
a.1	405.5	99	0.959	404.9	0.964	359.8	1000
a.2	665.7	93	0.936	659.6	0.937	630.1	1000
b.1	501.2	7	0.981	457.1	0.963	441.7	6
b.2	1058	99	0.783	1051	0.788	918.2	1000
c.1	525.3	4	0.958	437.1	0.885	443.4	7
c.2	485.3	95	0.711	481.4	0.713	441.4	1000

Table 6.3: Best fit of samples from profiles parallel to linear targets over extended targets, to K and exponential distributions. The parameters listed for each profile are the same as those in Table 6.1.

distribution (A.2), whilst the others best fitted K-distributions. The mean of the best fit distribution never varied from the mean calculated from the sample moments by more than 12.5%. In general the estimated mean was smaller than the mean of the best fit distribution. Profile C.2 which did not fit any distribution was found to contain two point targets. Unlike profile 9 in Table 6.1 the number of pixels involved was sufficient to affect the fit.

6.3.2 Comparison to measurements over extended targets

As a control test on the method used to sample pixel values along edges and linear targets, the same methods were used to sample pixel values from the extended targets to either side of the linear targets shown in Figure 5.9, along profiles parallel to but displaced about ten pixels from the targets. The fit of the profiles to K and exponential distributions was then tested. Table 6.3 gives the results.

All the profiles gave a good fit (> 0.7) to either an exponential distribution or a K-distribution. However, no near perfect fits (> 0.99) were found while half the profiles along edges or linear features fell into this category. This is surprising and indicates that the regions to either side of the linear targets are not as homogeneous as we thought. Though the analysis in Section 6.2 indicated that the statistics along edges will often be like those of a K-distribution, we did not expect them to be closer to the statistics of a K-distribution than statistics measured over extended targets.

6.4 Summary

In this chapter:

- The model for intensity statistics over extended targets described in Chapter 3 has been extended to the case of intensity statistics along edges and thin linear targets.
- The physical mechanisms generating edge features were analysed in terms of how they effected the detected field. The mechanisms include changes in surface type and cover, and terrain effects occurring in the range direction; azimuth effects were not considered.
- Assuming a rectangular system response, the detected field along an edge was shown to be expressible in terms of sums of random contributions from different scatterer populations plus contributions representing specular and secondary scattering, and system noise (6.1).
- The statistics of the intensity resulting from (6.1) were derived. Assuming the number of scatterers from each population is controlled by a birth-death-immigration process, the imaged SCS along an edge is given by a weighted sum of gamma distributed random variables, each one of which represents scattering from a single population.
- The moments of the imaged SCS resulting from it being a weighted sum of gamma distributed random variables were generated. Under certain conditions the moments were similar to those of a single gamma distribution.
- Measurements carried out on intensity values sampled along edges and linear targets, and over extended targets in a RSRE X-band SAR image gave empirical confirmation of *some* of the results. Measurements along edges bounding regions of layover and radar shadow have yet to be undertaken.
- Using the Kolmogorov-Smirnov test, edges and linear targets were found to fit K-distributions with a range of order parameters. In some cases the fit was better than that over extended targets. This has not been fully explained.

Chapter 7

Edge detection

Edge detection is a key process in the analysis of SAR images. Prior detection of the edge structure of an image is crucial for a number of image analysis tasks, e.g., segmentation, classification and image map matching [51, 62]. In this chapter the detection of edges in SAR images using the local operator approach is investigated. As in Chapter 6 we define an edge as a boundary in an image across which a transition in mean intensity or texture occurs. The detection of intensity edges will be concentrated upon. We shall once again initially limit ourselves to the simple image model already introduced in Chapter 5 before considering the effects of correlated and K-distributed data in latter sections of the chapter.

In Section 7.1 the local operator approach to edge detection is described and an outline is given of the general procedure for identifying the most suitable operator for detecting edges in a image with specific noise properties (e.g., speckle). An overview of available operators for detecting edges in SAR images is then given in light of this in Section 7.2. Intensity edges may be detected by measuring the ratio or the difference of the mean intensities in adjacent regions of an image. The ML estimates of these measures under the simple image model are derived in Section 7.2.1.1, and are found to be respectively given by the ratio and difference of the average intensities in the two regions; both of these estimates correspond to operators previously used for edge detection in SAR images [57]. The distributions of the values output by these operators under the simple image model are analysed in Sections 7.2.1.2 and 7.2.1.3. Only the ratio operator acts as a constant false alarm rate (CFAR) edge detector; the false alarm rate of the difference operator increases with image brightness. This dependence can be removed by normalising the difference in average intensity between two adjacent regions by the sum of the mean intensities over the two regions. This normalisation factor must be estimated. The ML estimate of the sum of the mean intensities is given by the sum of the average intensities previously used to estimate the difference in mean intensity. However, it turns out that the operator given by the difference in two average intensities divided by their sum is equivalent to the ratio operator. Thus further analysis only needs to concentrate on one of these operators. The ratio operator is the obvious choice as its output distribution under the simple image model is already

known. Before looking at the ratio operator in more detail a variety of operators that have previously been proposed for edge detection in SAR and which are based on measuring image heterogeneity are described in Section 7.2.2, along with reasons why they are considered to be unsuitable.

By extending the analysis given in [57] the performance and limitations of the ratio operator are more fully determined in Section 7.3 including (a) the general conditions under which the operator performs as a CFAR edge detector; (b) the detectability of edges as a function of scale under the simple image model; and (c) the performance of the operator when the simple image model does not apply. In Section 7.4 various problems encountered in interpreting edges detected at different scales are discussed. The analysis is confirmed by measurements carried out on real and simulated data.

7.1 The local operator approach to edge detection

The local operator approach to edge detection involves applying an operator over a region in an image to estimate whether a transition in mean intensity or texture occurs within it. An edge image is generated by moving a window defining a region of interest over an image; at each position the operator is applied over the window. The output of the operator is assigned to the pixel in the edge image which spatially corresponds to the pixel in the original image the window is centred over. Finally, in most cases, the edge image is thresholded to produce an edge map; this is not the case with the Marr Hildreth operator discussed in Section 7.2.2.2 where the edge map is defined in terms of zero crossings. Pixels whose values pass the threshold are referred to as edge pixels. When interpreting the edge map it must be borne in mind that the representation it gives does not necessarily imply that an edge actually passes through any given edge pixel, only that an edge has been detected within a window centred on it. Because the edge may lie anywhere in the window, edges in the edge map can be as thick as the window used to detect them.

Identifying the most suitable operator for detecting edges in an image with specific noise properties involves several stages; (a) determining how an edge is manifested in terms of the parameters describing the underlying image, (b) determining what measures involving these parameters may be used to detect the edge, (c) determining the optimal operators for estimating these measures from image data, and (d) analysing and comparing the performance of the resulting operators to determine which is the most suitable edge detector. This last stage involves checking that the operator meets certain general edge detection criteria, namely:

1. There needs to be a well defined test (e.g., threshold) for checking whether the value output by the operator indicates an edge.
2. This test should lead to a high probability of detecting true edges and a low probability of false alarm (PFA).

3. To allow the same degree of credibility to be attached to each edge detected the PFA should be constant over the image (the CFAR criterion).
4. There must be a means for combining edges detected at different scales into a single well defined edge map, i.e., one with thin connected edges.

The first three criteria concern how edges are detected, whilst the last concerns how detected edges are interpreted.

Operators for detecting edges within a window fall into two distinct classes

1. Those which involve splitting the window into two equal halves along an axis and testing for the presence of an edge oriented along the axis by measuring the disparity between the two halves in terms of the underlying image parameters.
2. Those which involve measuring the heterogeneity of the whole region within the window. By their very nature edges correspond to heterogeneous as opposed to homogeneous areas.

Operators from the latter class which we shall term 'heterogeneity' measures, are independent of edge orientation, whilst those from the former class which we shall term 'disparity' measures, are dependent on edge orientation. Unless edge orientation is known a priori, to detect all possible edges within a region operators from the former class have to be applied over several windows each split at a different orientation. The values output at each orientation are then combined to give an overall measure of the presence of an edge. The range of orientations needed to do this depends on the shape of the window. Using a square window horizontal, vertical and two diagonal orientations are normally needed [57], whilst a thin rectangular window split along its longer side may only need to be oriented horizontally and vertically to detect all edges [61].

Changes in mean intensity and texture may take place over a variety of scales. A strong (high contrast/abrupt) edge may be detectable over a distance of a few pixels, whilst a weak (low contrast/gradual) edge may only be detectable over tens or even hundreds of pixels. Thus to detect all the edges in an image operators need to be applied over a range of scales (i.e., different sized windows) as well a range of orientations [18].

7.2 Operators for detecting edges in SAR images

We shall first concentrate on edge detection operators which are based on measuring the disparity between two half windows.

7.2.1 Disparity measures

The estimate of disparity output by an operator should be optimal, i.e., the maximum likelihood estimate, when the axis along which the window is split lies directly over an abrupt edge separating two homogeneous regions. The pixel values in each half window will then be sampled from a single region. When this

is not the case pixel values in at least one of the halves will be sampled from both regions.

Assuming intensity is K-distributed adjacent regions of an image may differ in terms of their mean intensity or their order parameter which is a measure of texture. The disparity between two mean intensities can be measured by either taking their ratio or their difference. Due to the infinite range of values the order parameter can take, the disparity between two order parameters is best measured by taking their ratio or the difference in their logs.

As already seen in Chapter 4, ML estimates involving K-distributed random variables are not easily derived. Therefore we shall once again limit ourselves to the simple image model and determine the ML estimates for the ratio and difference of the mean intensity of adjacent regions which confirm to the model. Such an image contains no texture and we will not pursue texture based edge detection any further in this chapter. However, the effect of texture on the detection of intensity edges will be considered in Section 7.3.3.

7.2.1.1 ML estimates of the ratio and difference of mean values

Let x_1, \dots, x_{n_x} and y_1, \dots, y_{n_y} be two sets of pixel intensity values each sampled from a homogeneous area with unknown mean in a single-look image. These values may be taken to represent the pixels sampled from the image by two halves of a window. Even though the number of pixels sampled by each half window will be equal, for generality we shall initially assume that $n_x \neq n_y$. When the simple image model applies the likelihood that the x_i 's have mean μ_x and that the y_i 's have mean μ_y is given by

$$F(\mu_x, \mu_y) = \prod_{i=1}^{n_x} \frac{e^{-x_i/\mu_x}}{\mu_x} \prod_{i=1}^{n_y} \frac{e^{-y_i/\mu_y}}{\mu_y} \quad (7.1)$$

For any function g of μ_x and μ_y the estimate \hat{g} of g which maximizes the likelihood, occurs where the first derivative of the log-likelihood $L = \ln(F)$ with respect to g given by

$$\frac{dL}{dg} = \frac{\partial L}{\partial \mu_x} \frac{\partial \mu_x}{\partial g} + \frac{\partial L}{\partial \mu_y} \frac{\partial \mu_y}{\partial g} \quad (7.2)$$

equals zero. The partial derivative of L with respect to μ_x is given by

$$\frac{\partial L}{\partial \mu_x} = \frac{n_x}{\mu_x} \left(\frac{\hat{x}}{\mu_x} - 1 \right) \quad (7.3)$$

where \hat{x} is the average intensity of the x_i 's given by (4.3). As the simple image model is already assumed this average intensity gives the ML estimate of the mean intensity μ_x of the x_i 's (see Chapter 4). The partial derivative of L with respect to μ_y , and \hat{y} are similarly defined. The first derivative of the log-likelihood with respect to g will then equal zero when

$$\frac{n_x}{\mu_x} \left(\frac{\hat{x}}{\mu_x} - 1 \right) \frac{\partial \mu_x}{\partial g} + \frac{n_y}{\mu_y} \left(\frac{\hat{y}}{\mu_y} - 1 \right) \frac{\partial \mu_y}{\partial g} = 0 \quad (7.4)$$

When g is defined as the ratio of the mean intensities μ_x/μ_y , its partial derivatives with respect to μ_x and μ_y are given by

$$\frac{\partial g}{\partial \mu_x} = \frac{1}{\mu_y} \quad (7.5)$$

and

$$\frac{\partial g}{\partial \mu_y} = \frac{-\mu_x}{\mu_y^2} \quad (7.6)$$

Substituting (7.5) and (7.6) into (7.4) gives

$$n_x \left(\frac{\hat{x}}{\mu_x} - 1 \right) = n_y \left(\frac{\hat{y}}{\mu_y} - 1 \right) \quad (7.7)$$

Setting the ML estimate of the ratio of the mean intensities \hat{g} to equal $g = \mu_x/\mu_y$, and rearranging terms gives

$$\hat{g} \left(\frac{\mu_y}{\hat{y}} + \frac{n_y}{n_x} \left(1 - \frac{\mu_y}{\hat{y}} \right) \right) = \frac{\hat{x}}{\hat{y}} \quad (7.8)$$

Assuming μ_y and its ML estimate \hat{y} are approximately equal, the ML estimate of the ratio of the mean intensities is approximately given by the ratio of the ML estimates of the mean intensities

$$\hat{g} \approx \frac{\hat{x}}{\hat{y}} \quad (7.9)$$

This becomes an *exact equality* whenever the number of independent samples in the two sets are equal ($n_x = n_y$). Hence under the simple image model the optimal operator for measuring the ratio of the mean intensities in two homogeneous regions is given by the ratio of the average intensities in the two regions.

When g is defined as the difference of the mean intensities $\mu_x - \mu_y$, its partial derivatives with respect to μ_x and μ_y are given by

$$\frac{\partial g}{\partial \mu_x} = 1 \quad (7.10)$$

and

$$\frac{\partial g}{\partial \mu_y} = -1 \quad (7.11)$$

Substituting (7.10) and (7.11) into (7.4) gives

$$\frac{n_x}{\mu_x} \left(\frac{\hat{x}}{\mu_x} - 1 \right) = \frac{n_y}{\mu_y} \left(\frac{\hat{y}}{\mu_y} - 1 \right) \quad (7.12)$$

Which can be rearranged to give

$$n_y \mu_x - n_x \mu_y = n_y \frac{\mu_x}{\mu_y} \hat{y} - n_x \frac{\mu_y}{\mu_x} \hat{x} \quad (7.13)$$

Assuming the ratio of the mean intensities μ_x/μ_y and its estimate \hat{x}/\hat{y} are approximately equal gives

$$n_y\mu_x - n_x\mu_y \approx n_y\hat{x} - n_x\hat{y} \quad (7.14)$$

Thus when the number of independent samples in the two sets are equal, the ML estimate of the difference of the mean intensities is approximately given by the difference of the ML estimates of the mean intensities

$$\hat{g} = \mu_x - \mu_y \approx \hat{x} - \hat{y} \quad (7.15)$$

Thus under the simple image model the optimal operator for measuring the difference of the mean intensities in two homogeneous regions is given by the difference of the average intensities in the two regions.

As it is the size of the ratio of the mean intensities that is primarily of interest for edge detection, rather than whether the ratio is greater than or less than one, the ML estimate of the ratio is usually normalised to be less than one, i.e., the normalised ratio r is defined

$$r = \min\left(\frac{\hat{x}}{\hat{y}}, \frac{\hat{y}}{\hat{x}}\right) \quad (7.16)$$

Likewise, the absolute value of the ML estimate of the difference in the mean intensities is usually taken as the absolute difference is primarily of interest, rather than whether the difference is positive or negative, i.e., the absolute difference d is defined

$$d = |\hat{x} - \hat{y}| \quad (7.17)$$

The operators resulting from these estimates correspond to operators previously used for detecting edges in SAR images [6, 18, 21, 38, 57].

To analyse and compare the performance of the normalised ratio and absolute difference operators defined by (7.16) and (7.17) we need to look at the distributions of the values output by them; in particular the relation between edge detection and false alarm rates. When \hat{x} is given by the average of n_x independent exponentially distributed intensity values with mean μ_x it will be gamma distributed with mean μ_x and order parameter n_x . Likewise when \hat{y} is given by the average of n_y independent exponentially distributed intensity values with mean μ_y it will be gamma distributed with mean μ_y and order parameter n_y . Thus when the simple image model applies, the distribution of the normalised ratio and absolute difference operators are given by the normalised ratio and absolute difference of two independent gamma distributed average intensities.

7.2.1.2 Distribution of the normalised ratio

Following from (A.7) the normalised ratio r of two independent gamma distributed average intensities with unknown means μ_x and μ_y and order parameters n_x and n_y , has pdf

$$f_r(r|\mu_x, \mu_y, n_x, n_y) =$$

$$\begin{aligned}
& \int_0^\infty \left[\frac{e^{-\frac{rx}{\beta_x}} (rx)^{n_x-1}}{\Gamma(n_x) \beta_x^{n_x}} \frac{e^{-\frac{x}{\beta_y}} x^{n_y-1}}{\Gamma(n_y) \beta_y^{n_y}} + \frac{e^{-\frac{x}{\beta_x}} x^{n_x-1}}{\Gamma(n_x) \beta_x^{n_x}} \frac{e^{-\frac{rx}{\beta_y}} (rx)^{n_y-1}}{\Gamma(n_y) \beta_y^{n_y}} \right] x dx \\
&= \frac{\Gamma(n_x + n_y)}{\Gamma(n_x) \Gamma(n_y)} \left[\frac{\left(\frac{\beta_y}{\beta_x}\right)^{n_x} r^{n_x-1}}{\left(1 + \frac{\beta_y}{\beta_x} r\right)^{n_x+n_y}} + \frac{\left(\frac{\beta_x}{\beta_y}\right)^{n_y} r^{n_y-1}}{\left(1 + \frac{\beta_x}{\beta_y} r\right)^{n_x+n_y}} \right] \quad (7.18)
\end{aligned}$$

where $\beta_x = \mu_x/n_x$ and $\beta_y = \mu_y/n_y$. This distribution is related to the variance ratio or F-distribution [1, 49]. Letting the ratio of the unknown mean intensities $\mu_x/\mu_y = R$ (the ratio we are trying to estimate) the pdf of the normalised ratio may be defined solely in terms of R and the number of independent samples n_x and n_y

$$f_r(r|R, n_x, n_y) = \frac{\Gamma(n_x + n_y)}{\Gamma(n_x) \Gamma(n_y)} \left[\frac{\left(\frac{r}{R} \frac{n_x}{n_y}\right)^{n_x}}{\left(1 + \frac{r}{R} \frac{n_x}{n_y}\right)^{n_x+n_y}} + \frac{\left(Rr \frac{n_y}{n_x}\right)^{n_y}}{\left(1 + Rr \frac{n_y}{n_x}\right)^{n_x+n_y}} \right] \frac{1}{r} \quad (7.19)$$

In terms of setting edge detection and false alarm thresholds the cdf of the normalised ratio is of more importance, it is given by

$$\begin{aligned}
p_r(r|R, n_x, n_y) &= \\
& \frac{\Gamma(n_x + n_y)}{\Gamma(n_x) \Gamma(n_y)} \int_0^r \left[\frac{\left(\frac{r}{R} \frac{n_x}{n_y}\right)^{n_x}}{\left(1 + \frac{r}{R} \frac{n_x}{n_y}\right)^{n_x+n_y}} + \frac{\left(Rr \frac{n_y}{n_x}\right)^{n_y}}{\left(1 + Rr \frac{n_y}{n_x}\right)^{n_x+n_y}} \right] \frac{1}{r} dr \\
&= \frac{\Gamma(n_x + n_y)}{\Gamma(n_x) \Gamma(n_y)} \left[\sum_{k=0}^{n_x-1} \binom{n_x-1}{k} (-1)^{n_y-k-1} \frac{1 - \left(1 + \frac{r}{R} \frac{n_x}{n_y}\right)^{-(n_x+n_y-k-1)}}{n_x + n_y - k - 1} \right. \\
& \quad \left. + \sum_{k=0}^{n_y-1} \binom{n_y-1}{k} (-1)^{n_x-k-1} \frac{1 - \left(1 + Rr \frac{n_y}{n_x}\right)^{-(n_x+n_y-k-1)}}{n_x + n_y - k - 1} \right] \quad (7.20)
\end{aligned}$$

When the number of independent samples making up the two averages are equal, as will normally be the case in edge detection, the pdf of the normalised ratio is given by [57]

$$f_r(r|R, n) = \frac{\Gamma(2n)}{(\Gamma(n))^2} \left[\frac{\left(\frac{r}{R}\right)^n}{\left(1 + \frac{r}{R}\right)^{2n}} + \frac{(Rr)^n}{(1 + Rr)^{2n}} \right] \frac{1}{r} \quad (7.21)$$

and its cdf is given by

$$p_r(r|R, n) = \frac{\Gamma(2n)}{(\Gamma(n))^2} \sum_{k=0}^{n-1} \binom{n-1}{k} (-1)^{n-k-1} \frac{2 - \left(1 + \frac{r}{R}\right)^{-(2n-k-1)} - (1 + Rr)^{-(2n-k-1)}}{2n - k - 1} \quad (7.22)$$

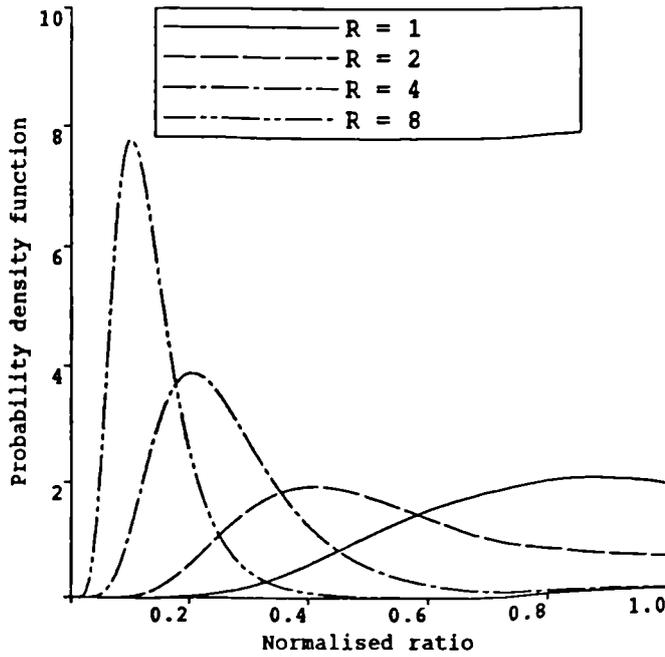


Figure 7.1: Pdf of the normalised ratio of two gamma distributed average intensities with means μ_x and μ_y and order parameter $n = 10$, where the edge ratio $R = \mu_x/\mu_y = 1, 2, 4, 8$.

where $n_x = n_y = n$. In addition, when no edge is present, i.e., $\mu_x = \mu_y$ and $R = 1$, the normalised ratio has pdf

$$f_r(r|R=1, n) = \frac{\Gamma(2n)}{(\Gamma(n))^2} \frac{2r^{n-1}}{(1+r)^{2n}} \quad (7.23)$$

and cdf

$$p_r(r|R=1, n) = \frac{2\Gamma(2n)}{(\Gamma(n))^2} \sum_{k=0}^{n-1} \binom{n-1}{k} (-1)^{n-k-1} \frac{1 - (1+r)^{-(2n-k-1)}}{2n-k-1} \quad (7.24)$$

This defines the probability of false alarm for the normalised ratio under the simple image model, i.e., the probability that the output of the operator will be less than r when no edge is present. This PFA depends only on the size of the window. Thus assuming the simple image model applies, the false alarm rate will be constant over an edge map generated using a single window size. Figure 7.1 shows the pdf of the normalised ratio when $R = 1, 2, 4, 8$ and $n = 10$ (e.g., a 5×5 window split into two equal halves by a single pixel wide line). As the edge ratio R increases the distribution of the normalised ratio becomes more separable from the distribution when no edge is present ($R = 1$).

Figure 7.2 shows the result of applying the normalised ratio operator to the SAR image shown in Figure 3.1 using a 9×9 window.

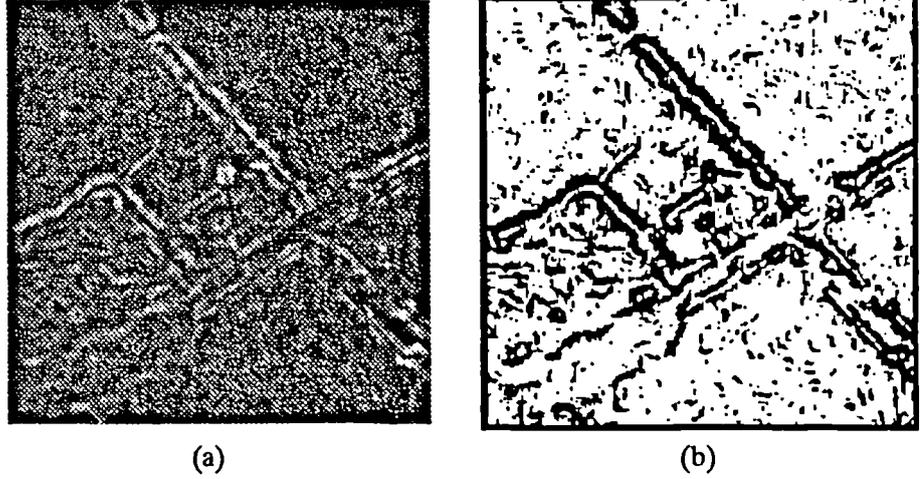


Figure 7.2: The result of applying the normalised ratio operator to the SAR image shown in Figure 3.1 using a 9×9 window, (a) the minimum ratio over horizontal, vertical and two diagonal orientations shown on a $-\log$ scale, (b) after thresholding at $\min(r) = 0.5$.

7.2.1.3 Distribution of the absolute difference

Following from (A.4) the absolute difference d of two gamma distributed average intensities with unknown means μ_x and μ_y and order parameters n_x and n_y has pdf

$$\begin{aligned}
 f_d(d|\mu_x, \mu_y, n_x, n_y) &= \\
 & \int_0^\infty \left\{ \frac{e^{-\frac{(d+x)}{\beta_x}} (d+x)^{n_x-1} e^{-\frac{x}{\beta_y}} x^{n_y-1}}{\Gamma(n_x) \beta_x^{n_x} \Gamma(n_y) \beta_y^{n_y}} + \frac{e^{-\frac{x}{\beta_x}} x^{n_x-1} e^{-\frac{(d+x)}{\beta_y}} (d+x)^{n_y-1}}{\Gamma(n_x) \beta_x^{n_x} \Gamma(n_y) \beta_y^{n_y}} \right\} dx \\
 &= \frac{\alpha^{n_x+n_y-1}}{\Gamma(n_x) \Gamma(n_y) \beta_x^{n_x} \beta_y^{n_y}} \left\{ e^{-\frac{d}{\beta_x}} \sum_{k=0}^{n_x-1} \binom{n_x-1}{k} \Gamma(n_x+n_y-k-1) \left(\frac{d}{\alpha}\right)^k \right. \\
 & \quad \left. + e^{-\frac{d}{\beta_y}} \sum_{k=0}^{n_y-1} \binom{n_y-1}{k} \Gamma(n_x+n_y-k-1) \left(\frac{d}{\alpha}\right)^k \right\} \quad (7.25)
 \end{aligned}$$

where $\beta_x = \mu_x/n_x$, $\beta_y = \mu_y/n_y$ and $1/\alpha = 1/\beta_x + 1/\beta_y$. The cdf of the absolute difference is given by

$$\begin{aligned}
 p_d(t|\mu_x, \mu_y, n_x, n_y) &= \\
 & \int_0^d \frac{\alpha^{n_x+n_y-1}}{\Gamma(n_x) \Gamma(n_y) \beta_x^{n_x} \beta_y^{n_y}} \left\{ e^{-\frac{r}{\beta_x}} \sum_{k=0}^{n_x-1} \binom{n_x-1}{k} \Gamma(n_x+n_y-k-1) \left(\frac{d}{\alpha}\right)^k \right. \\
 & \quad \left. + e^{-\frac{r}{\beta_y}} \sum_{k=0}^{n_y-1} \binom{n_y-1}{k} \Gamma(n_x+n_y-k-1) \left(\frac{d}{\alpha}\right)^k \right\} dr
 \end{aligned}$$

$$\begin{aligned}
&= \sum_{k=0}^{n_x-1} \binom{n_y+k-1}{k} \frac{\alpha^{n_y+k}}{\beta_x^k \beta_y^{n_y}} \left(1 - e^{-\frac{d}{\beta_x}} \sum_{j=0}^{n_x-k-1} \frac{\left(\frac{d}{\beta_x}\right)^j}{j!} \right) \\
&\quad + \sum_{k=0}^{n_y-1} \binom{n_x+k-1}{k} \frac{\alpha^{n_x+k}}{\beta_x^{n_x} \beta_y^k} \left(1 - e^{-\frac{d}{\beta_y}} \sum_{j=0}^{n_y-k-1} \frac{\left(\frac{d}{\beta_y}\right)^j}{j!} \right) \quad (7.26)
\end{aligned}$$

When the number of independent samples making up each average is equal the pdf of the absolute difference is given by [57]

$$f_d(d|\mu_x, \mu_y, n) = \frac{\alpha^{2n-1}}{\Gamma(n)^2 \beta_x^n \beta_y^n} \left(e^{-\frac{d}{\beta_x}} + e^{-\frac{d}{\beta_y}} \right) \sum_{k=0}^{n-1} \binom{n-1}{k} \Gamma(2n-k-1) \left(\frac{d}{\alpha}\right)^k \quad (7.27)$$

and its cdf is given by

$$\begin{aligned}
p_d(t|\mu_x, \mu_y, n) &= \sum_{k=0}^{n-1} \binom{n+k-1}{k} \alpha^{n+k} \left\{ \frac{1}{\beta_x^k \beta_y^n} \left(1 - e^{-\frac{d}{\beta_x}} \sum_{j=0}^{n-k-1} \frac{\left(\frac{d}{\beta_x}\right)^j}{j!} \right) \right. \\
&\quad \left. + \frac{1}{\beta_x^{n_x} \beta_y^k} \left(1 - e^{-\frac{d}{\beta_y}} \sum_{j=0}^{n-k-1} \frac{\left(\frac{d}{\beta_y}\right)^j}{j!} \right) \right\} \quad (7.28)
\end{aligned}$$

where $n_x = n_y = n$. In addition when no edge is present, i.e., $\mu_x = \mu_y = \mu$, the absolute difference has pdf

$$f_d(d|\mu, n) = \frac{e^{-\frac{d}{\beta}}}{4^{n-1} \Gamma(n)^2 \beta} \sum_{k=0}^{n-1} \binom{n-1}{k} \Gamma(2n-k-1) \left(\frac{2d}{\beta}\right)^k \quad (7.29)$$

and cdf

$$p_d(t|\mu, n) = \sum_{k=0}^{n-1} \binom{n+k-1}{k} \frac{1}{2^{n+k-1}} \left(1 - e^{-\frac{d}{\beta}} \sum_{j=0}^{n-k-1} \frac{\left(\frac{d}{\beta}\right)^j}{j!} \right) \quad (7.30)$$

where $\beta = \mu/n$.

The PFA of the absolute difference is defined by $1 - p_d(t|\mu, n)$, i.e., the probability that the absolute difference is greater than t when no edge is present. This probability depends on the mean intensity as well as the window size. Thus, for any given window size the absolute difference operator does not act as a constant false alarm rate edge detector. Figure 7.3 shows how the distribution of the absolute difference over a homogeneous region is broadened as mean intensity increases

When $n_x = n_y = n$ the moments of the absolute difference may be expressed

$$\begin{aligned}
\langle d^m \rangle &= \frac{\alpha_x^{n-1} \alpha_y^{n-1} (\mu_x + \mu_y)^m}{\Gamma(n) n^m} \times \\
&\quad \sum_{k=0}^{n-1} \frac{\Gamma(2n-k-1) \Gamma(m+k+1)}{\Gamma(n-k) \Gamma(k+1)} \left(\frac{\alpha_x^{m+1}}{\alpha_y^k} + \frac{\alpha_y^{m+1}}{\alpha_x^k} \right) \quad (7.31)
\end{aligned}$$

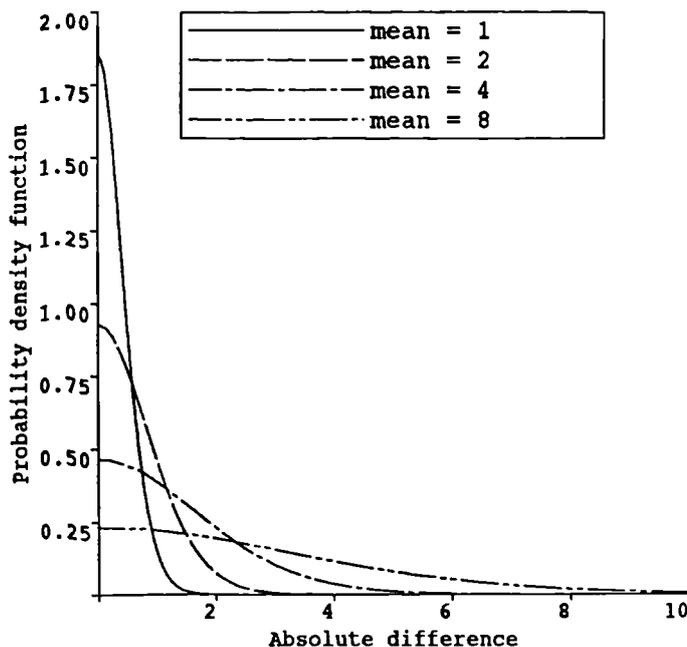


Figure 7.3: Pdf of the absolute difference of two gamma distributed average intensities with the same mean μ and order parameter $n = 10$ when $\mu = 1, 2, 4, 8$.

where $\alpha_x = \alpha/\beta_x = (\beta_x/\beta_y + 1)^{-1}$ and $\alpha_y = \alpha/\beta_y = (\beta_y/\beta_x + 1)^{-1}$. Because $\beta_x/\beta_y = R$, apart from the $\mu_x + \mu_y$ term the moments depend only on the edge ratio R and the number of independent samples n . The effect of the $\mu_x + \mu_y$ term is to make the m 'th moment of the absolute difference proportional to the mean intensity raised to the power m . The standard deviation of the false alarm distribution will then be proportional to the mean intensity. Thus after an edge image generated using the absolute difference operator is thresholded there will be more false alarms in parts of the edge map which correspond to brighter parts of the original SAR image than in parts which correspond to darker parts of the original SAR image. This is the reason why all edge detectors which are based on using difference operators to estimate image gradient perform poorly on SAR images, e.g., the standard Sobel, Prewitt and Roberts operators [6, 21, 28, 57]. The rationale behind all of these operators assumes that image noise is additive and is independent of the mean intensity.

Figure 7.4 shows the result of applying the difference operator to the SAR image shown in Figure 3.1 using a 9×9 window. There are more false alarms in the more intense woodland (bottom left) than in the less intense agricultural fields (top right).

The dependence of the false alarm rate on the mean intensity can be removed by normalising the absolute difference by the sum of the mean intensities.

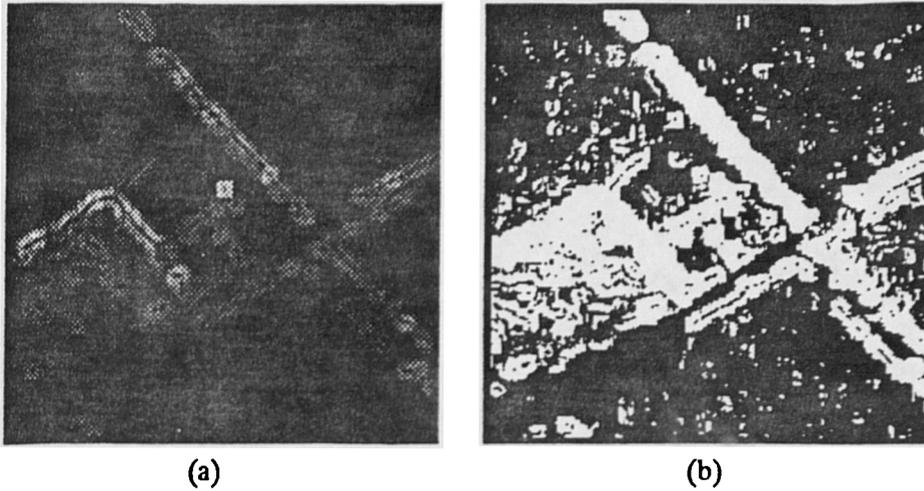


Figure 7.4: The result of applying the absolute difference operator to the SAR image shown in Figure 3.1 using a 9×9 window, (a) the maximum absolute difference over horizontal, vertical and two diagonal orientations, (b) after thresholding at $\max(d) = \hat{I}/2$ where \hat{I} is the average intensity over the whole of the image.

From (7.27) the pdf of the absolute difference divided by $\mu_x + \mu_y$ is given by

$$f_{d/(\mu_x + \mu_y)}(d|\mu_x, \mu_y, n) = \frac{\alpha_x^{n-1} \alpha_y^{n-1}}{\Gamma(n)^2} \left(e^{-\frac{d}{\alpha_x}} + e^{-\frac{d}{\alpha_y}} \right) \times \sum_{k=0}^{n-1} \binom{n-1}{k} \Gamma(2n-k-1) \left(\frac{nd}{\alpha_x \alpha_y} \right)^k \quad (7.32)$$

Figure 7.5 shows plots of this distribution when $n = 10$ and $R = 1, 2, 4, 8$.

In practice as the normalising factor given by the sum of the mean intensities is unknown, it has to be estimated as well as the absolute difference of the mean intensities. The ML estimate of the sum of the mean intensities $\mu_x + \mu_y$ is given by the sum of the average intensities. Thus the normalised absolute difference may be estimated using

$$d_0 = \frac{|\hat{x} - \hat{y}|}{\hat{x} + \hat{y}} \quad (7.33)$$

Dividing the numerator and denominator of the RHS of (7.33) by whichever of \hat{x} and \hat{y} is the greater, the normalised absolute difference d_0 may be reformulated in terms of the normalised ratio r

$$d_0 = \frac{1-r}{1+r} \quad (7.34)$$

This defines a one-to-one mapping between the normalised absolute difference and the normalised ratio. Thus in terms of edge detection and false alarm rates the two operators will have identical performance. In fact if the function g

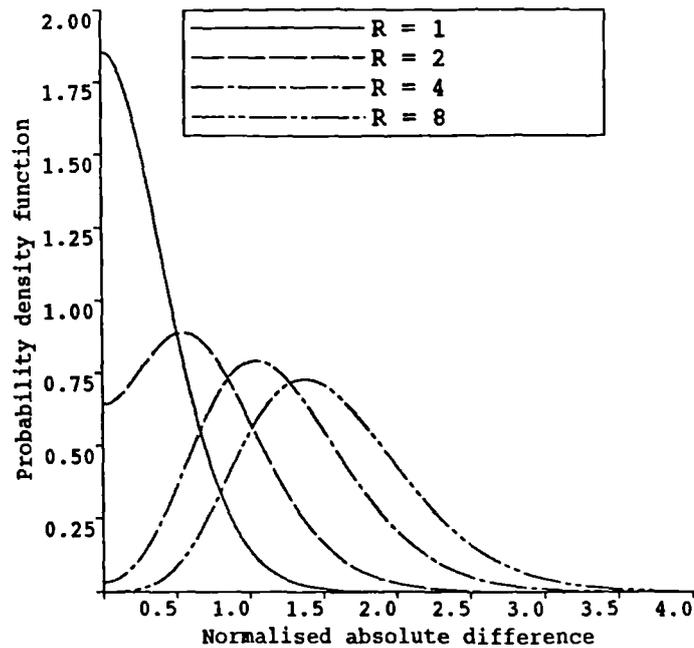


Figure 7.5: Pdf of the normalised absolute difference of two gamma distributed average intensities with order parameter $n = 10$ and edge ratio $R = 1, 2, 4, 8$.

in (7.4) is set to equal $|\mu_x - \mu_y| / (\mu_x + \mu_y)$ its ML estimate turns out to be equal to the ML estimate of μ_x / μ_y given by r . Thus further analysis only needs to concentrate on one of these operators. The normalised ratio operator is the obvious choice as its output distribution under the simple image model is already known (7.21).

7.2.2 Heterogeneity measures

Before looking at the ratio operator in more detail we shall briefly examine operators that have previously been proposed for edge detection in SAR which belong to the other class of edge detection operator referred to at the start of this section, i.e., those which are based on measuring image heterogeneity.

7.2.2.1 Coefficient of variation based operators

Under the simple image model the coefficient of variation of the intensity over a homogeneous area is given by the inverse of the number of looks and increases in heterogeneous areas (i.e. close to edges). Thus edges may be detected using an estimate of the coefficient of variation given by [57, 60]

$$\widehat{v}_x^2 = \frac{\widehat{I^2}}{\widehat{I}^2} - 1 \quad (7.35)$$

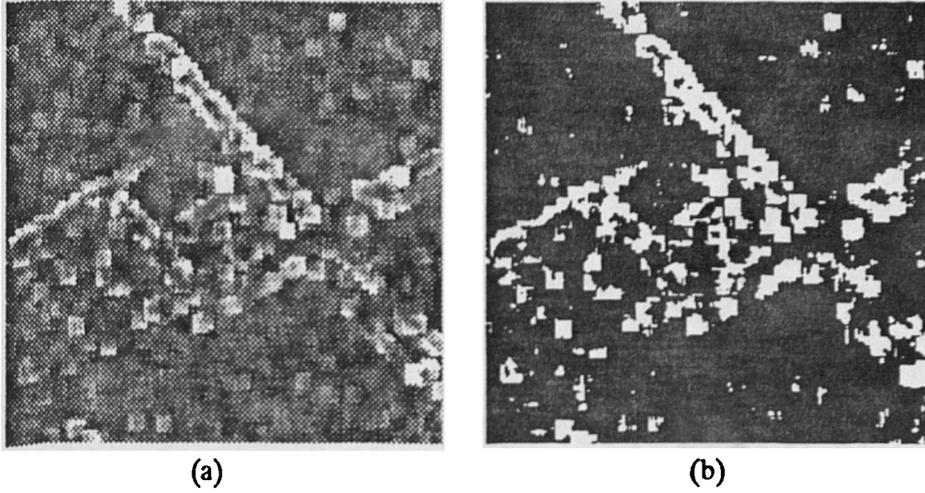


Figure 7.6: The result of applying the coefficient of variation operator to the SAR image shown in Figure 3.1 using a 9×9 window, (a) the output image, (b) after thresholding at 1.5.

This operator is effectively the same as that used to estimate texture from the second normalised intensity moment in Chapter 4. Thus it will act as a CFAR operator when the simple image model applies but the false alarm rate will be particularly sensitive to texture. Likewise, when used to estimate texture it will be strongly affected by the presence of edges. Even when the simple image model is assumed, the theoretical distribution of values output by the estimate of the coefficient of variation operator given by (7.35) is unknown. Thus, it is not straightforward to attach a probability of an edge being present to a given threshold. Figure 7.6 shows the result of using the coefficient of variation operator to detect edges in the SAR image shown in Figure 3.1. The resulting image is very 'blocky', this is because a single bright pixel will generate a strong signature wherever it lies within the processing window.

A related but slightly more sophisticated method is to use a likelihood ratio to test whether intensity values in a region are sampled from one gamma distributed population with mean intensity μ_0 , or two gamma distributed populations with mean intensities μ_1 and μ_2 [18], i.e., whether the region is homogeneous or heterogeneous. The likelihood ratio is given by

$$F = \prod_{i=1}^n \frac{e^{-I_i/\beta_0} I_i^{m-1}}{\Gamma(m) \beta_0^m} \bigg/ \prod_{i=1}^n \frac{1}{2} \left(\frac{e^{-I_i/\beta_1} I_i^{m-1}}{\Gamma(m) \beta_1^m} + \frac{e^{-I_i/\beta_2} I_i^{m-1}}{\Gamma(m) \beta_2^m} \right) \quad (7.36)$$

where m is the number of looks and $\beta_j = \mu_j/m$. In the heterogeneous case it is assumed that the two populations are in equal proportion. The mean intensities μ_0 , μ_1 and μ_2 are derived from the first and second moments of the data using, $\mu_0 = \hat{I}$,

$$\mu_1 = \hat{I} + \sqrt{\hat{I}^2 \frac{m}{m+1} - \hat{I}^2} \quad (7.37)$$

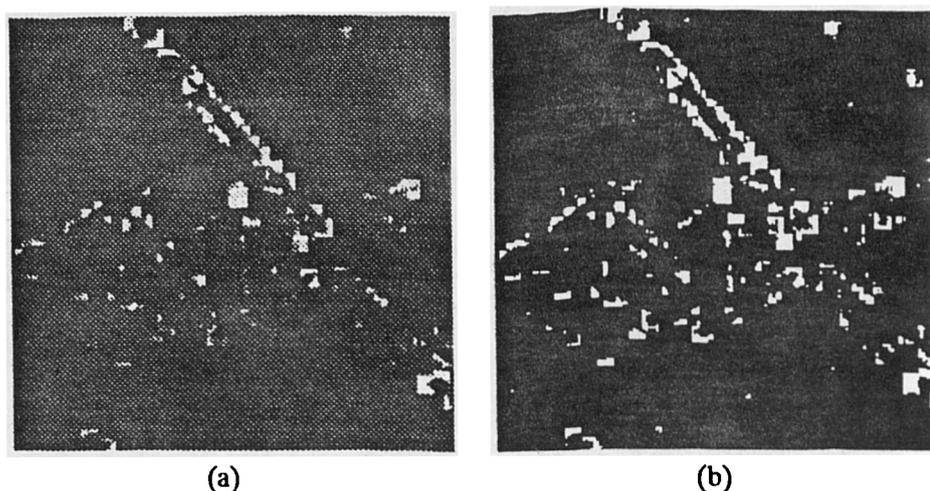


Figure 7.7: The result of applying the likelihood ratio operator to the SAR image shown in Figure 3.1 using a 9×9 window, (a) the likelihood ratio image shown on a log scale, (b) after thresholding at 2.

and

$$\mu_2 = \hat{I} - \sqrt{\hat{I}^2 \frac{m}{m+1} - \hat{I}^2} \quad (7.38)$$

The operator performs reasonably on multi-look data but various problems exist [38, 57]. The method for deriving μ_1 and μ_2 assumes that $\widehat{I}^2/\widehat{I}^2$ lies in the range $1 + 1/m$ to $2 + 2/m$, but this does not always occur and the likelihood ratio cannot be determined in such cases. When $\widehat{I}^2/\widehat{I}^2$ is less than $1 + 1/m$ it is assumed the region is homogeneous whilst when it is greater than $2 + 2/m$ it is assumed an edge is present. The theoretical distribution of values output by the likelihood ratio is unknown and hence, as with the coefficient of variation, thresholds are not easily set. Finally, this method is much slower than other edge detection operators over comparable window sizes because of the large number of operations involved in calculating the likelihood ratio. Figure 7.7 shows the result of applying the likelihood ratio operator to the SAR image shown in Figure 3.1. The resulting edge map is sparser and less 'blocky', than that produced using the coefficient of variation and shown in Figure 7.6 (b).

7.2.2.2 The Marr Hildreth operator

The absolute difference operator discussed in Section 7.2.1 basically detected edges by treating the intensity image I as a surface and then estimating the gradient of the surface given by its first derivative ∇I

$$\nabla I = \left(\frac{\partial I}{\partial x}, \frac{\partial I}{\partial y} \right) \quad (7.39)$$

where x and y define an orthonormal basis. In terms of edge detection we are primarily interested in the magnitude of the gradient given by

$$|\nabla I|^2 = \left(\frac{\partial I}{\partial x}\right)^2 + \left(\frac{\partial I}{\partial y}\right)^2 \quad (7.40)$$

To estimate this magnitude the absolute difference operator has to be applied in two orthogonal orientations. Alternatively, the magnitude may be estimated by applying the absolute difference operator over a number of orientations and taking the maximum. Edges are indicated by ridges in the resulting image.

Edges may also be detected by looking at the second derivative of an image in particular the Laplacian defined as

$$\nabla \cdot \nabla I = \nabla^2 I = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2} \quad (7.41)$$

Whereas edges are indicated by maxima in the first derivative they are indicated by zero-crossing points in the second derivative, i.e., where the rate of change of the gradient goes from being positive to being negative or vice versa. The Laplacian is independent of edge orientation and is estimated over the whole of a window. Thus it may loosely be termed a heterogeneity measure. Using a 5×5 window the Laplacian may be estimated by convolving the image with the mask

$$\begin{array}{ccccc} 1 & 1 & 0 & 1 & 1 \\ 1 & 0 & -2 & 0 & 1 \\ 0 & -2 & -4 & -2 & 0 \\ 1 & 0 & -2 & 0 & 1 \\ 1 & 1 & 0 & 1 & 1 \end{array}$$

To reduce the effect of image noise on estimation the Laplacian operator is normally combined with a smoothing filter. When this smoothing filter is Gaussian the well known Marr-Hildreth operator results [40].

The performance of the Marr-Hildreth operator on SAR images has been investigated in [6, 20] and will only be discussed briefly here. The zero-crossing method of identifying edges in the output edge image guarantees that the edge map has the desirable property of consisting of single pixel wide closed boundaries. However, even with smoothing a large number of false alarms are generated by the Laplacian operator and further processing is needed to distinguish between edges and false alarms. It has been demonstrated in [6] that for multi-look SAR data (intensity is assumed to be correlated and Gaussian distributed) the Marr-Hildreth operator is a CFAR edge detector, and that a 'dead' zone exists to either side of true edges, isolating them from false alarms. It has also been shown in [20], that when the same edge is detected at different scales, i.e., with smoothing filters of different width, the edges detected do not lie over each other; as the width of the filter increases zero crossings are shifted towards the darker side of the edge. Despite the large false alarm rate the Marr-Hildreth edge detector may have some use for edge thinning and linking when combined with

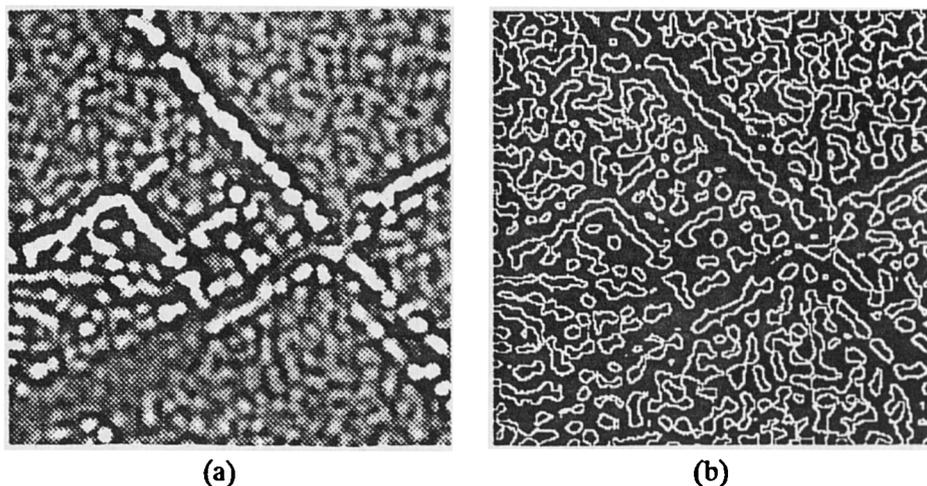


Figure 7.8: The result of applying the Marr-Hildreth operator using a Gaussian smoothing filter of width three pixels to the SAR image shown in Figure 3.1, (a) output image, (b) zero crossings.

other edge detectors such as the normalised ratio discussed above. Figure 7.8 shows the result of applying the Marr-Hildreth edge detector to the SAR image shown in Figure 3.1. The large number of false alarms and the ‘dead’ zone to either side of true edges are clearly apparent.

Even though the heterogeneity measures discussed in this Section benefit from being independent of edge orientation this advantage is outweighed by the problems caused by threshold selection and texture in the case of the coefficient of variation and likelihood ratio operators, and by the excessive false alarm rate in the case of the Marr-Hildreth operator. Further analysis will concentrate solely on the normalised ratio operator.

7.3 Performance analysis of the normalised ratio operator

The most promising edge detection operator identified in the previous section was the normalised ratio operator. This operator has been proposed by several authors [6, 38, 57] as a means of detecting edges in SAR. The fullest explication of the operator is found in [57]. In this section the performance and limitations of the normalised ratio operator are more fully determined including, (a) the general conditions under which the operator performs as a CFAR edge detector; (b) the detectability of edges as a function of scale under the simple image model; and (c) the performance of the operator when the simple image model does not apply.

7.3.1 Conditions for a constant false alarm rate

The third criterion listed in Section 7.1 requires the false alarm rate to be constant over the image (i.e., the normalised ratio must have the same distribution over all homogeneous regions). Assuming the multiplicative image model (3.38) applies this will be the case if and only if both speckle and the imaged surface texture are strict-sense stationary process over the whole image. The average of n statistically identical intensity values with mean μ is given by

$$\hat{I} = \mu \widehat{\xi\chi} \quad (7.42)$$

where ξ and χ represent unit mean speckle and imaged texture respectively, and $\widehat{\xi\chi}$ the average of n samples of their product, i.e., clutter. Given two such averages from adjacent areas in the same homogeneous region their normalised ratio will be independent of μ and will only depend on ξ and χ . Thus if ξ and χ are strict-sense stationary over the whole image the normalised ratio estimated using a given window size will have the same distribution over any homogeneous area.

Obviously, the conditions for a CFAR are satisfied when the simple image model applies because speckle is then assumed to be spatially independent and the imaged SCS is constant (no texture). The conditions for a CFAR are still satisfied when speckle is correlated but only so long as the correlation structure is the same over the whole image. Because the real and imaginary components of speckle are Gaussian its correlation structure is completely described by its ACF, which is determined solely by the system response to a point target. The correlation width in range, in terms of pixels, is determined by the pixel sampling rate in range relative to the slant range resolution, i.e., the width of the range component of the ACF. Likewise, the correlation width in azimuth is determined by the pixel sampling rate in azimuth relative to the azimuth resolution. The false alarm rate will no longer be constant if either the shape of the system response, or the pixel sampling rate in range or azimuth, change over the image. For example, the correlation width in range will vary across an image which has been corrected to ground range. The false alarm rate will also no longer be constant when an image contains *differently* textured regions. Where the ACF of speckle is known to vary across an image due to system effects it may be possible to design an adaptive threshold to produce a CFAR. Compensating for different textures would be much more difficult.

Because the normalised absolute difference given by (7.33) is equivalent to the normalised ratio operator it will operate as a CFAR edge detector under the same condition as the normalised ratio operator. The difference of two average intensities \hat{x} and \hat{y} has variance $\sigma_x^2 + \sigma_y^2$ which equals $(\sigma_x^2 + \sigma_y^2)/n$ where n is the number of independent samples making up each average. Thus the standard deviation of the false alarm rate will be constant over differently textured regions when the absolute difference is normalised by a measure based on the local standard deviation given by $\sqrt{\sigma_x^2 + \sigma_y^2}$ as opposed to the local mean. Edge detection using such an operator is discussed in Chapter 9.

7.3.2 Threshold selection and edge detectability

The cumulative false alarm distribution of the normalised ratio given by (7.24) can be used to select a threshold for a desired PFA with a given window size, (7.22) can then be used to determine the detectability of an edge with given edge ratio R using that threshold.

The alarm distribution given by (7.24) only refers to a normalised ratio estimated at a single orientation. Given normalised ratios estimated at horizontal, vertical and two diagonal orientations the minimum corresponds to the most likely value for the edge ratio R . Thus the PFA for a single orientation $p_r(r|R=1, n)$ needs to be related to the PFA of the minimum of the normalised ratios estimated at all four orientations $p_{\min(r)}(r|R=1, n)$. If the normalised ratios estimated at the four orientations were independent of each other, the PFA for the minimum ratio would be given by

$$p_{\min(r)}(r|R=1, n) = 1 - (1 - p_r(r|R=1, n))^4 \quad (7.43)$$

In practice as the normalised ratios are estimated over the same area they are correlated. In [57] it is suggested that to take this correlation into account the PFA of the minimum should be derived using

$$p_{\min(r)}(r|R=1, n) = 1 - (1 - p_r(r|R=1, n))^3 \quad (7.44)$$

Simulated data was used to check this relationship.

The normalised ratio operator was applied at horizontal, vertical and two diagonal orientations to 5×5 simulated homogeneous images each generated according to the simple image model. The distributions of the four normalised ratios output and the minimum ratio were built up from one million such simulations each with different speckle realisations. The cumulative distributions of the normalised ratios and their theoretical cdf given by (7.24) are plotted in Figure 7.9. As we would expect, for each separate orientation there is a near perfect fit between the cumulative distribution of the normalised ratio and the theoretical distribution. The cumulative distributions only diverge when the number of samples is small at low probabilities. The cumulative distribution of the minimum ratio and the theoretical cdf's given by (7.43) and (7.44) are plotted in Figure 7.10. The cumulative histogram of the minimum ratio closely fitted both theoretical cdf's. For large PFA's ($> 10^{-1}$) the cumulative histogram was closer to (7.44), but for small PFA's ($< 10^{-2}$) it was closer to (7.43). This implies that (7.43) rather than (7.44) should be used to determine thresholds for the minimum ratio when low PFA's are required. Note that for low PFA's $p_{\min(r)} \simeq 4p_r$. This is because $1 - (1 - p_r)^q \simeq p_r q$ when $p_r \ll 1$.

If the proportion of edge pixels in an image is p_e , a PFA of p_e would yield one edge pixel for every false alarm. To produce a low noise edge map a PFA of $p_e \times 10^{-2}$ or even $p_e \times 10^{-3}$ is needed. The value of p_e will vary from image to image. Assuming it to lie in the range 10^{-1} to 10^{-3} , the PFA needs to be somewhere in the range 10^{-3} to 10^{-6} .

The cumulative distribution (7.22) was used to compare the false alarm rate against the detectability of different strength edges under the simple image

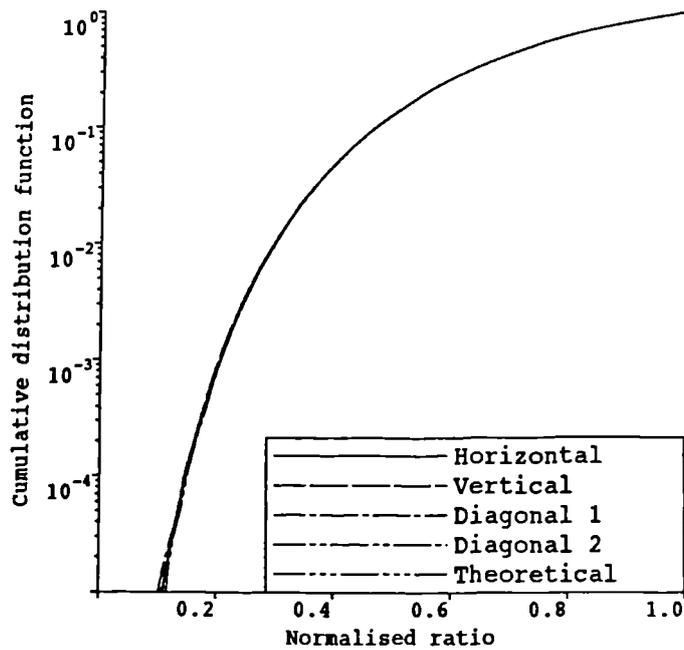


Figure 7.9: Cumulative distributions of the normalised ratios calculated at horizontal, vertical and two diagonal orientations over one million 5×5 simulated homogeneous images with different exponential speckle realisations, and their theoretical cdf given by (7.24).

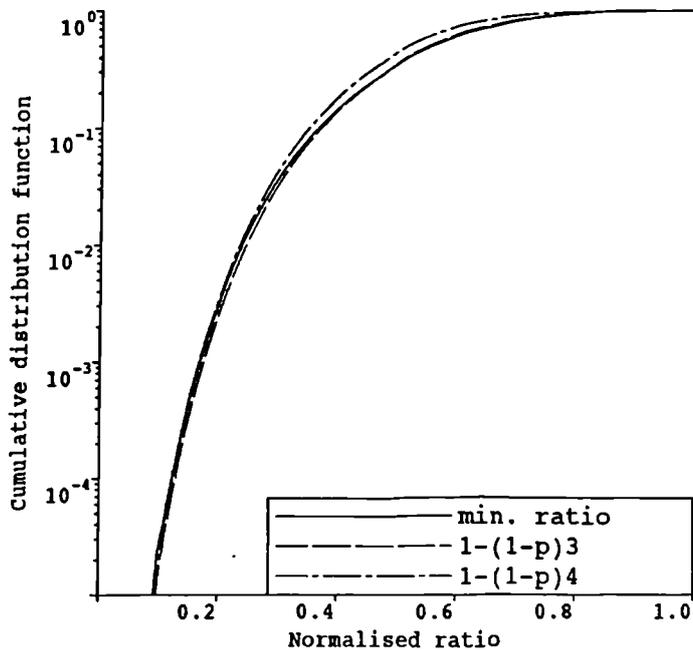


Figure 7.10: Cumulative distribution of the minimum normalised ratio calculated at horizontal, vertical and two diagonal orientations over one million 5×5 simulated homogeneous image with different exponential speckle realisations, and theoretical cdf's given by (7.43) and (7.44).

model, for a variety of window sizes. Figure 7.11 shows a plot of edge detectability versus false alarm rate when $R = 2, 4, 8, 16$ (see key for window sizes used). Only the larger windows ($n \geq 11$) have a high probability (> 0.5) of detecting weak edges ($R = 2$) whilst retaining a low PFA ($< 10^{-3}$). As edge strength increases, windows as small as 5×5 will detect edges with a high probability whilst retaining the same low PFA. 3×3 windows appear to be unsuitable for detecting edges with a high probability whilst retaining a low PFA. This implies that 5×5 is the minimum window size that should be used for detecting edges in a single-look image with the normalised ratio operator.

7.3.3 Correlated and K-distributed intensity

When intensity is correlated and/or K-distributed it is not possible to determine analytically the distribution of the output of the normalised ratio operator analogous to (7.21). However, in (7.21) the only parameter affecting the false alarm rate is n , the number of independent samples making up the average intensities \hat{x} and \hat{y} . The coefficient of variation of the averages is given by $1/n$. This value controls how broad the distribution of the normalised ratio is and thus how many false alarms fall below any set threshold. This suggests that by looking at how correlated and K-distributed intensity influence the coefficient of variation of the averages, some idea of how they alter the false alarm rate may be gained.

For some value $a > 1$ the average of n correlated and K-distributed intensity values will have the same first and second moments as the average of n/a uncorrelated and gamma distributed intensity values, i.e., their coefficient of variation is given by a/n . By reducing n in (7.21) by a the equivalent simple image model gives an indication of how correlated and K-distributed intensity alters the false alarm rate. As the higher order moments of the averages may differ, care is needed in using this method to select thresholds. When determining the PFA it is the tails of distributions that we are interested in, thus the central limit theorem cannot be employed to argue that distributions are the same when large numbers of pixels are averaged.

When the pixel intensity values being averaged are K-distributed with mean μ , order parameter ν and covariance $C_I(\bar{z})$ defined by (4.33), from (3.57) and (4.35) the coefficient of variation of the average of a $n_a \times n_r$ rectangular region of statistically identical intensity values is given by

$$v_I^2 = \frac{1 + \frac{2}{\nu}}{n_a n_r} \left\{ 1 + 2 \sum_{w_a=1}^{n_a-1} c_a(w_a) \left(1 - \frac{w_a}{n_a} \right) \right\} \times \left\{ 1 + 2 \sum_{w_r=1}^{n_r-1} c_r(w_r) \left(1 - \frac{w_r}{n_r} \right) \right\} \quad (7.45)$$

Figure 7.12 shows the relative increase a in the coefficient of variation of an average due to speckle being correlated (it is assumed that correlation is isotropic and that it does not include a texture component, i.e., $c_a(x) = c_r(x) = \text{sinc}^2(x/d)$ where d is the resolution in pixels). Figure 7.13 shows the corresponding increase

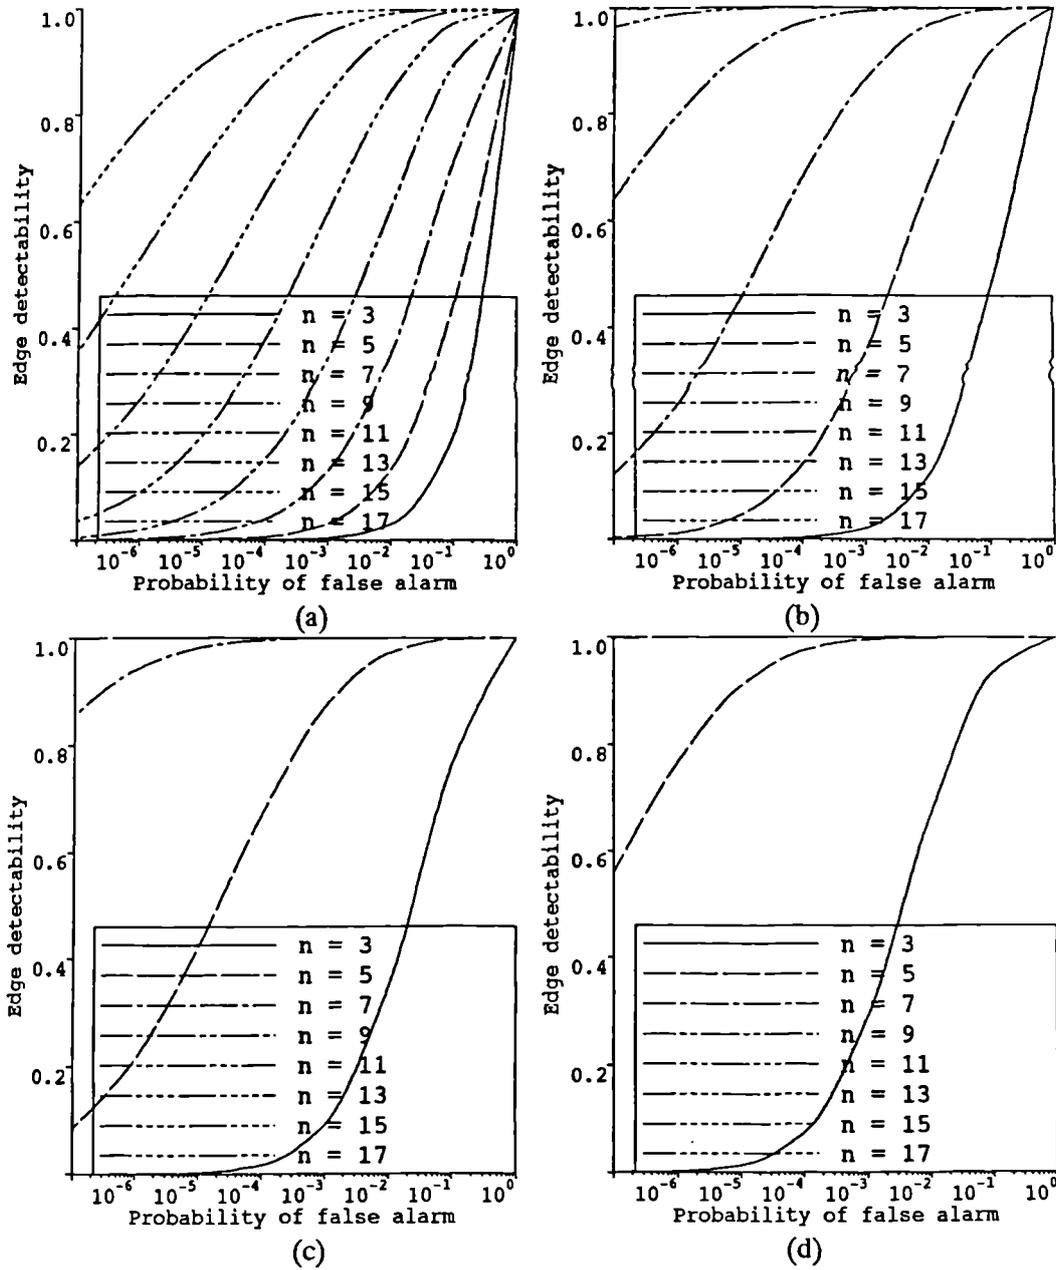


Figure 7.11: Comparison of edge detectability and probability of false alarm under the simple image model, when the normalised ratio operator is applied over a range of $n \times n$ windows (see key) to edges of different strength (a) $R = 2$, (b) $R = 4$, (c) $R = 8$ and (d) $R = 16$.

in the false alarm rate under the equivalent simple image model with thresholds set for a given PFA assuming no correlation. Even small amounts of correlation lead to a significant increase in the coefficient of variation of an average. This increase causes the false alarm rate to be several orders of magnitude greater than would be the case under the simple image model. The effect is greater when larger windows and lower PFA's are employed.

When intensity is K-distributed with order parameter ν the coefficient of variation of the average intensity is increased by a factor $1 + \frac{2}{\nu}$ as compared to when the simple image model applies. Figure 7.14 shows the corresponding increase in the false alarm rate. The increase in the coefficient of variation and thus the false alarm rate is only significant when ν is small (< 10); although the effect is not as strong as that due to correlation it can still increase the false alarm rate by several orders of magnitude. Once again this effect is greater when lower PFA's are employed; window size has little effect.

7.3.4 Measured False Alarm Rates

The effect of correlation on the false alarm rate (FAR) was confirmed by measurement. The normalised ratio operator was applied at different scales to 50×50 pixel homogeneous regions extracted from ERS-1, RSRE and AIRSAR images, and to a homogeneous 50×50 pixel simulated image generated according to the simple image model.

The RSRE and AIRSAR data is single-look whilst the ERS-1 image is a 3 look fast delivery product (FDP). Figure 7.15 shows histograms of the intensity values within the homogeneous regions extracted from the ERS-1, RSRE and AIRSAR images. As expected the intensity values extracted from the ERS-1 image fit a gamma distribution with order parameter 3, while the intensity values extracted from the RSRE and AIRSAR images fit exponential distributions. The intensity ACF's within the extracted regions were also measured and are shown in Figure 7.16. These revealed that the ERS-1 image was equally correlated in both range and azimuth, the RSRE image was strongly correlated in range and but was uncorrelated in azimuth, and the AIRSAR image was slightly less correlated in azimuth than the ERS-1 image but was more correlated in range.

The resulting edge images were thresholded at values corresponding to a range of PFA's under the simple image model (10^{-3} to 10^{-6}). In each edge map, the actual FAR was determined by measuring the proportion of pixels set as being edges. The results when $n = 5, 9, 11, 17$ are shown in Table 7.1. Over the 50×50 pixel areas tested the smallest measurable FAR is 4×10^{-4} . Thus, when the simple image model applies we would expect no false alarms to occur below an expected PFA of 10^{-4} .

In the ERS-1, RSRE and AIRSAR data, the measured FAR's were several orders of magnitude greater than the expected PFA assuming the simple image model applied; this discrepancy tends to increase as the window size increases and as the PFA decreases, in accordance with the results shown in Figure 7.13. With simulated data the measured FAR's and expected PFA's agreed.

The normalised ratio operator was re-applied to the ERS-1, RSRE and AIR-

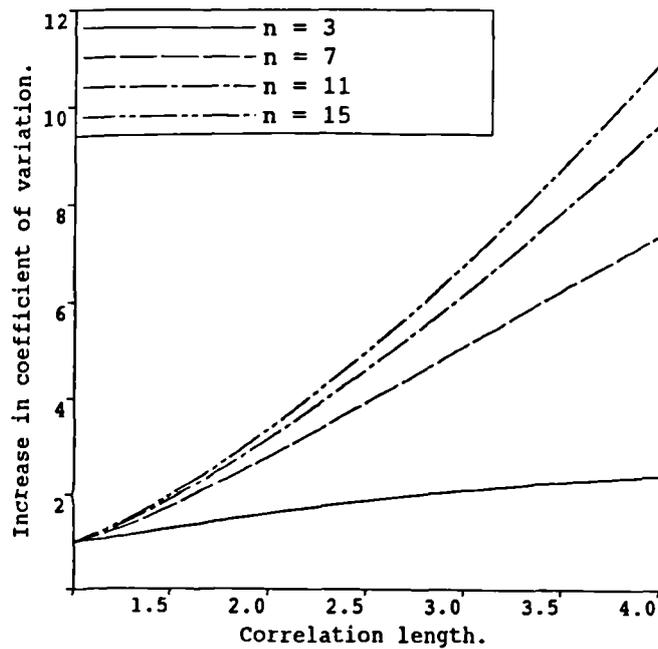


Figure 7.12: Relative increase in the coefficient of variation of an average over an $m \times n$ homogeneous area due to speckle being correlated ($m = \frac{n-1}{2}$).

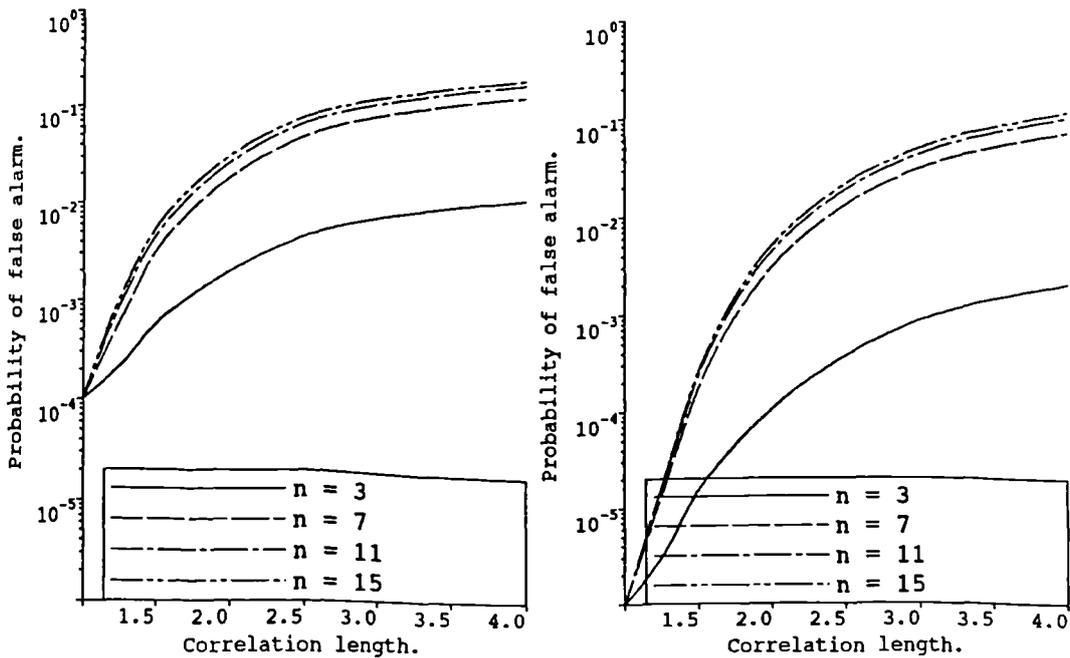


Figure 7.13: Increase in the PFA associated with the effects shown in Figure 7.12 when the normalised ratio operator is applied over a range of $n \times n$ windows with thresholds set for an expected PFA under the simple image model of (a) 10^{-4} and (b) 10^{-6} .

ERS-1				
Expected PFA	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$ for different window sizes			
	5 × 5	9 × 9	13 × 13	17 × 17
10 ⁻³	14.8	27.2	38.4	42.4
10 ⁻⁴	36	92	164	184
10 ⁻⁵	80	240	720	960
10 ⁻⁶	0.0	800	2400	800

RSRE				
Expected PFA	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$ for different window sizes			
	5 × 5	9 × 9	13 × 13	17 × 17
10 ⁻³	6.8	12.8	42.8	58.4
10 ⁻⁴	4	20	156	268
10 ⁻⁵	0.0	40	440	960
10 ⁻⁶	0.0	0.0	400	1600

AIRSAR				
Expected PFA	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$ for different window sizes			
	5 × 5	9 × 9	13 × 13	17 × 17
10 ⁻³	27.6	74.0	92.0	100.0
10 ⁻⁴	116	392	448	348
10 ⁻⁵	240	2120	2320	840
10 ⁻⁶	800	9200	11600	800

Simulated				
Expected PFA	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$ for different window sizes			
	5 × 5	9 × 9	13 × 13	17 × 17
10 ⁻³	0.0	1.70	0.0	0.865
10 ⁻⁴	0.0	0.0	0.0	0.0
10 ⁻⁵	0.0	0.0	0.0	0.0
10 ⁻⁶	0.0	0.0	0.0	0.0

Table 7.1: Expected PFA's under the simple image model, and measured FAR's when the normalised ratio operator is applied to homogeneous regions in ERS-1, RSRE, AIRSAR and simulated images. Measurements were carried out over 50 × 50 pixel regions, thus the minimum measurable FAR is 4 × 10⁻⁴; smaller FAR's are indicated by zeros.

ERS-1					
Expected PFA	Measured FAR for different window sizes				
	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$	5 × 5	9 × 9	13 × 13	17 × 17
10^{-3}	1.19	0.0	0.0	0.0	0.0
10^{-4}	0.0	0.0	0.0	0.0	0.0

RSRE					
Expected PFA	Measured FAR for different window sizes				
	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$	5 × 5	9 × 9	13 × 13	17 × 17
10^{-3}	5.82	0.0	0.0	0.0	0.0
10^{-4}	8.96	0.0	0.0	0.0	0.0

AIRSAR					
Expected PFA	Measured FAR for different window sizes				
	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$	5 × 5	9 × 9	13 × 13	17 × 17
10^{-3}	11.9	9.6	29.3	55.4	
10^{-4}	35.7	0.0	0.0	138	

Table 7.2: Expected PFA's under the simple image model, and measured FAR's when the normalised ratio operator is applied to homogeneous regions in ERS-1, RSRE and AIRSAR images after they had been sub-sampled. In each case no false alarms were detected below an expected PFA of 10^{-4} .

ERS-1				
Expected PFA	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$ for different window sizes			
	5 × 5	9 × 9	13 × 13	17 × 17
10 ⁻³	0.4	0.8	2.4	1.2
10 ⁻⁴	0.0	0.0	0.0	0.0

RSRE				
Expected PFA	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$ for different window sizes			
	5 × 5	9 × 9	13 × 13	17 × 17
10 ⁻³	0.4	0.4	1.2	2.4
10 ⁻⁴	0.0	0.0	0.0	0.0

AIRSAR				
Expected PFA	$\frac{\text{Measured FAR}}{\text{Expected PFA}}$ for different window sizes			
	5 × 5	9 × 9	13 × 13	17 × 17
10 ⁻³	2.0	9.2	9.6	0.0
10 ⁻⁴	4	12	0.0	0.0

Table 7.3: Expected PFA's under the simple image model, and measured FAR's when the normalised ratio operator is applied to homogeneous regions in ERS-1, RSRE and AIRSAR images using thresholds corrected for correlation. In each case no false alarms were detected below an expected PFA of 10⁻⁴.

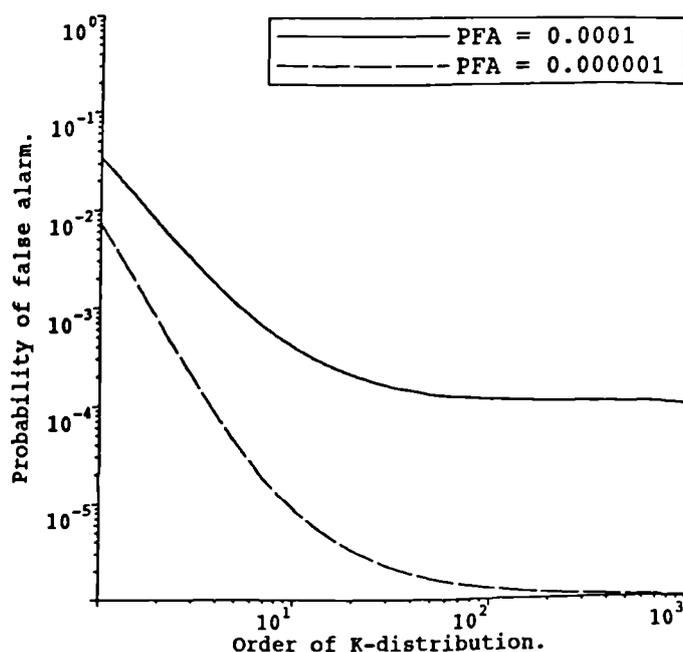
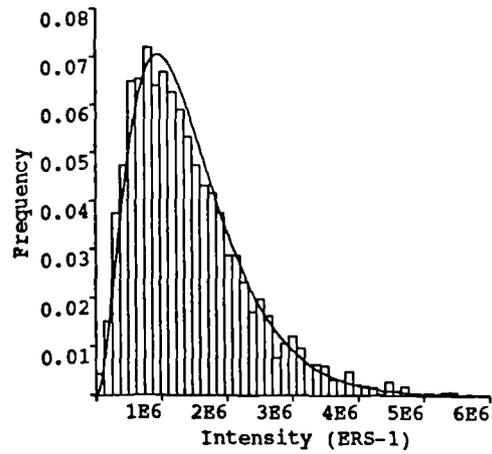


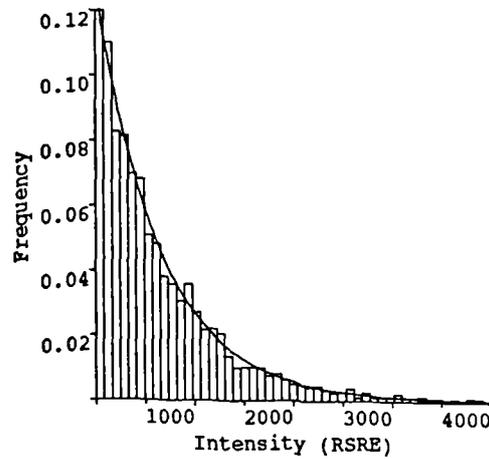
Figure 7.14: Increase in the PFA due to intensity being K-distributed when the normalised ratio operator is applied with thresholds set for different PFA's (see key) under the simple image model.

SAR images using two methods to compensate for the known system induced correlation, (a) sub-sampling, and (b) thresholds corrected according to the equivalent simple image model discussed in Section 7.3.3. The measured FAR's are shown in Tables 7.2 and 7.3 respectively. Both methods lead to large reductions in FAR's. In many cases no false alarms occurred, and none at all below an expected PFA of 10^{-4} . Where false alarms occur in the ERS-1 and RSRE images the measured FAR is of the same order of magnitude as the expected PFA. However, with the AIRSAR data the measured FAR's were still significantly higher than the expected PFA's, particularly when correlation was compensated for by sub-sampling. This implies that the extracted region is not as homogeneous as it appears. The measured FAR's were generally closer to the expected PFA's when correlation was compensated for by correcting thresholds as compared to sub-sampling.

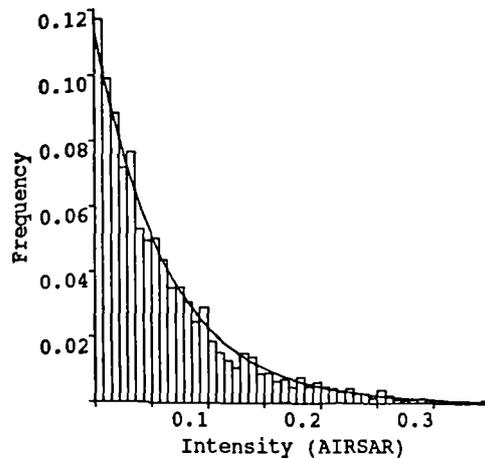
Whilst both of the above methods enable the desired false alarm rate to be met they both lead to a reduction in edge detectability. This is illustrated in Figure 7.17 which shows the result of applying the normalised ratio operator to the RSRE image shown in Figure 3.1 using 9×9 and 17×17 windows, with thresholds initially set for an expected PFA of 10^{-4} under the simple image model, and (a) no compensation for correlation, (b) sub-sampling to remove correlation, and (c) thresholds corrected for correlation. Sub-sampling also leads to an overall loss of information. If a $m \times n$ sub-sampling is required to remove correlation, the information loss can be minimised by sub-sampling to produce



(a)



(b)



(c)

Figure 7.15: Histograms of intensity values within a homogeneous regions in (a) an ERS-1 FDP image, (b) a RSRE image, and (c) an AIRSAR image, with overlaid expected distributions, (a) a gamma distribution with order parameter 3, (b) and (c) an exponential distribution.

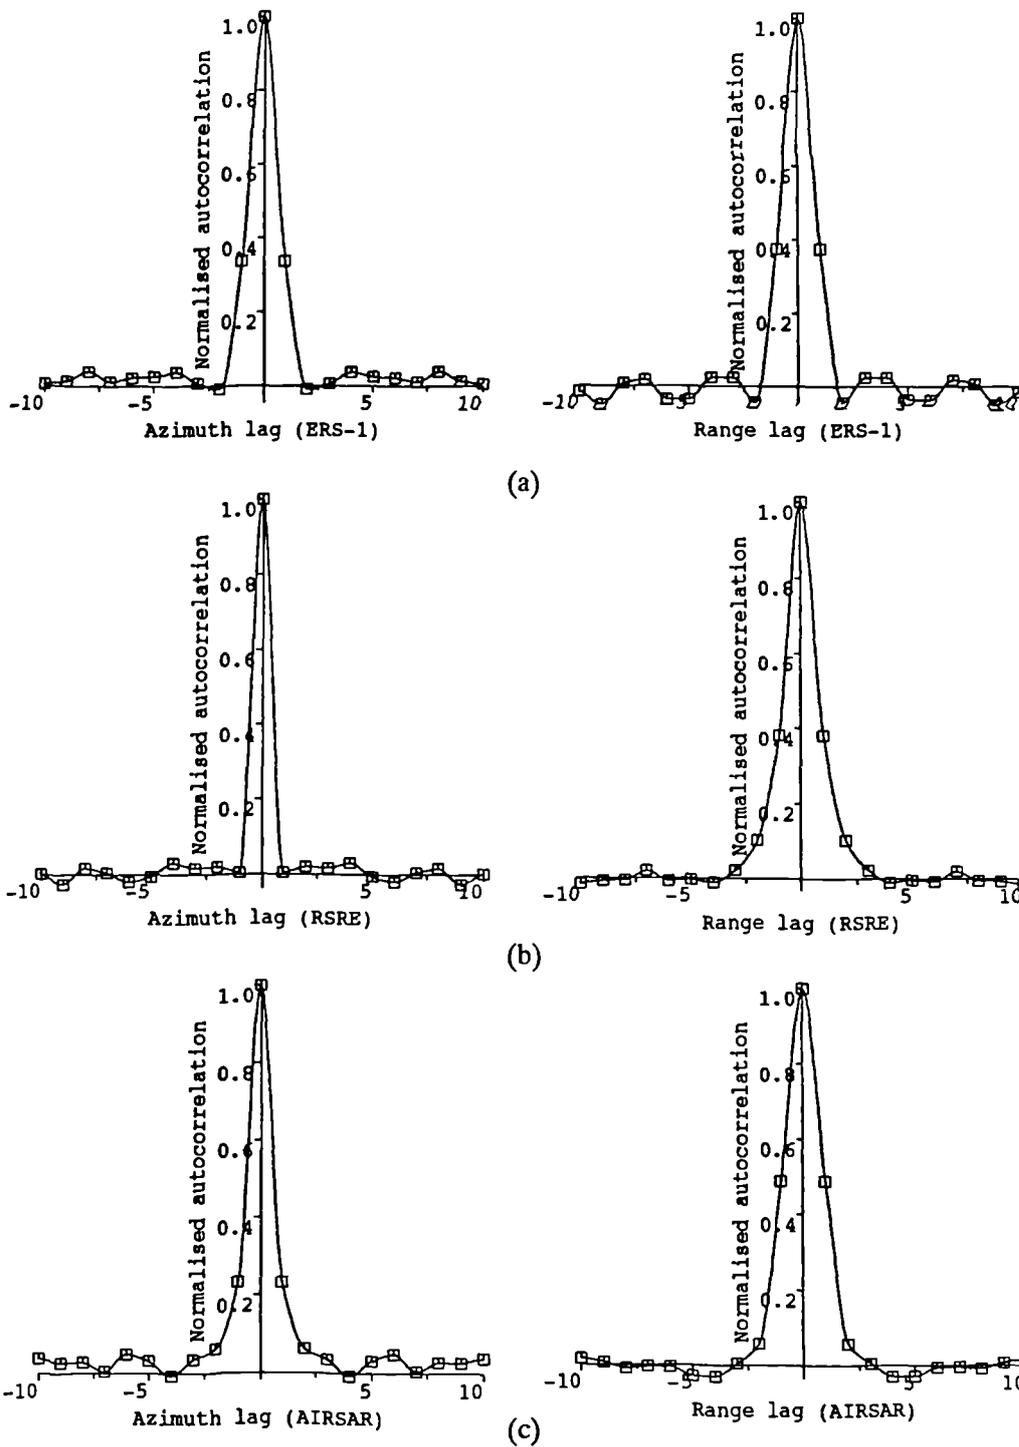


Figure 7.16: Intensity ACF's measured within homogeneous regions in (a) an ERS-1 FDP image, (b) a RSRE image, and (c) an AIRSAR image.

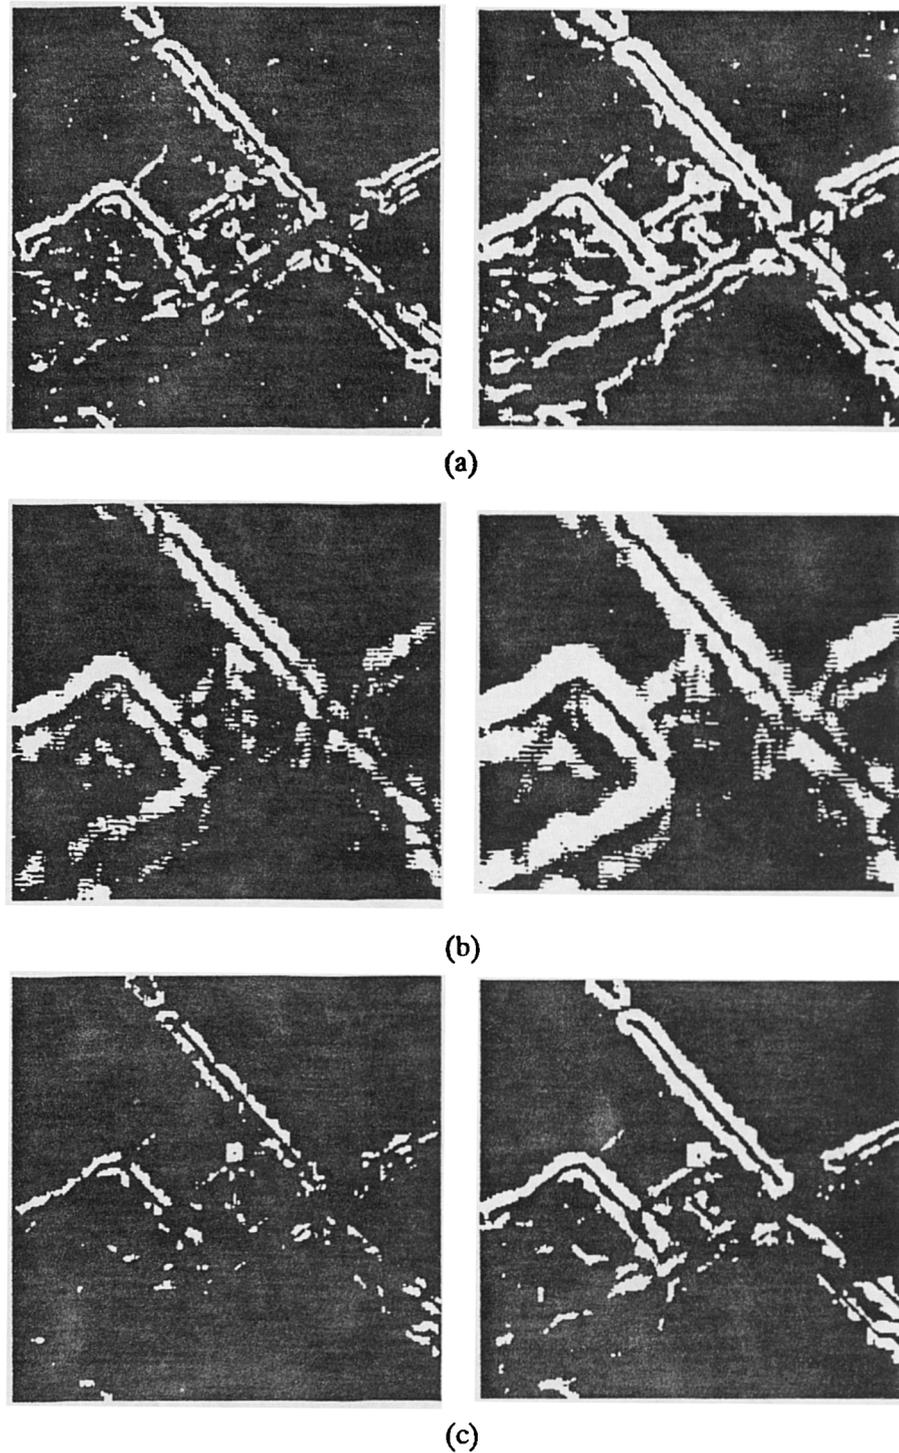


Figure 7.17: The result of applying the normalised ratio operator to the RSRE image shown in Figure 3.1 using 9×9 (left) and 17×17 windows (right), with thresholds initially set for an expected PFA of 10^{-4} under the simple image model, and (a) no compensation for correlation, (b) sub-sampling to remove correlation, and (c) thresholds corrected for correlation.

$m \times n$ different sub-sampled images, applying edge detection to each of them separately and then recombining them to recover as much information as possible (though not all). This has been carried out on the images shown in Figure 7.17 (b). The sub-sampling actually takes place within the windows used to calculate averages.

In a 256×256 edge map such as shown in Figure 7.17 we would expect to see seven false alarms when a PFA of 10^{-4} is set. There are significantly more false alarms than this in the edge maps shown in Figure 7.17 (a) but there appear to be very few, if any, false alarms in the edge maps shown in Figures 7.17 (b) and (c). Whilst both methods of compensating for correlation appear to be successful in this respect the remaining edges in the sub-sampled edge map suffer from edge thickening, whilst those in the threshold corrected edge map are thinned. Edges present in the uncorrected edge map are *missing from the sub-sampled and threshold corrected edge maps*, e.g., the edge defining the boundary of radar shadow in the bottom left of Figure 3.1 and the thin linear feature to the top right of the woodland on the left of the image.

The above results confirm that system induced correlation between pixels has a marked effect on the performance of the normalised ratio operator but can be compensated for using measurements of the intensity ACF. Similar results were produced when other homogeneous regions were tested. If the correlation is non-isotropic, as in the case of the RSRE data, edges at different orientations may not be equally detectable.

7.4 Edge scale size and detectability

The above analysis has shown that by going a long way towards satisfying the first three edge detection criteria listed in Section 7.1, the normalised ratio operator offers a powerful potential for detecting edges in SAR data. Performance is degraded by system induced correlation and texture. No consideration has yet been given to the fourth criterion concerning how edges detected at different scales are to be combined into a single well defined edge map. Problems which arise at this stage of edge detection using the normalised ratio operator are (a) edge thickening, and (b) multiple edges occurring within the window. These interrelated problems are particularly acute when large windows are employed to detect weak edges. To be able to take full advantage of the potential of the normalised ratio operator for edge detection these problems need to be resolved.

7.4.1 Edge thickening

This is a general problem with edge detection using the local operator approach and is apparent in the edge maps shown in Figure 7.17. As mentioned at the beginning of Section 7.1, detected edges in the edge map can be as thick as the window used to detect them. Edge thickening occurs when too large a window is used to detect a strong edge. This leads to edge structures on a scale smaller than the window size being obscured, and difficulties in relating edges detected at different scales as they have different thicknesses. Edges need to be thinned.

Ideally we would like to be able to thin edges down to the *most probable location of the real edge*. This does not always involve thinning the edge down to a single pixel wide line. In cases where an edge is detected over a gradual intensity gradient, the location of the real edge is in fact ill defined (this illustrates the general error of assuming that images consist of 'flat' homogeneous segments separated by abrupt edges). Also, as edge signatures are often asymmetrical, just thinning the edge down to a central line of pixels will lead to misleading results. Attempts to thin edges detected by the normalised ratio operator using the strength and orientation of edge pixels to determine the most probable location of the real edge have so far failed to produce satisfactory results. The primary problem is developing a method for thinning thick (strong) edges whilst retaining thin (weak) edges. A morphological method is suggested in [57] but it only produces a limited amount of thinning; edges are thinned down to the thickness of the smallest window the normalised ratio operator was applied over (minimum = 3). Also, the method leads to some loss of edge connectivity and holes appear within some edges.

7.4.2 Multiple edges occurring within the window

The model upon which the normalised ratio operator is based is no longer valid when more than one edge occurs within a window. The operator should only be applied at a given scale where no edges are detectable at smaller scales; this strategy will also minimise edge thickening. To do this edges must be detected in order of increasing scale. The segmentation algorithm discussed in Chapter 9 uses such a strategy when detecting edges [62].

The smallest scale edges in the image are those defining point like features. Such features have to be much brighter than surrounding speckle peaks, otherwise they are indistinguishable from speckle. When edges are detected for using a $n \times n$ window, a point like features consisting of a single very bright pixel will generate a distinctive $n \times n$ signature in the resulting edge map. Where the point like feature consists of several bright pixels more complex signatures are generated. Accurate detection of these features is crucial for subsequent detection of larger features, otherwise their signatures seriously confuse the interpretation of edge maps. This is particularly the case where point like features occur close together (e.g., in urban areas). This topic is dealt with in the next chapter.

Figure 7.18 shows the edges detected when the normalised ratio operator was applied over a range of scales ($n = 5, 7, 9$) to a 100×100 pixel region extracted from an ERS-1 FDP image of Flevoland (Netherlands). Thresholds have been corrected to take correlation into account so that the PFA equals 10^{-4} (i.e., there should be no false alarms in an image of this size). Whilst strong edges are well defined it is evident that small scale detail is obscured due to edge thickening and multiple edges, in particular point features.

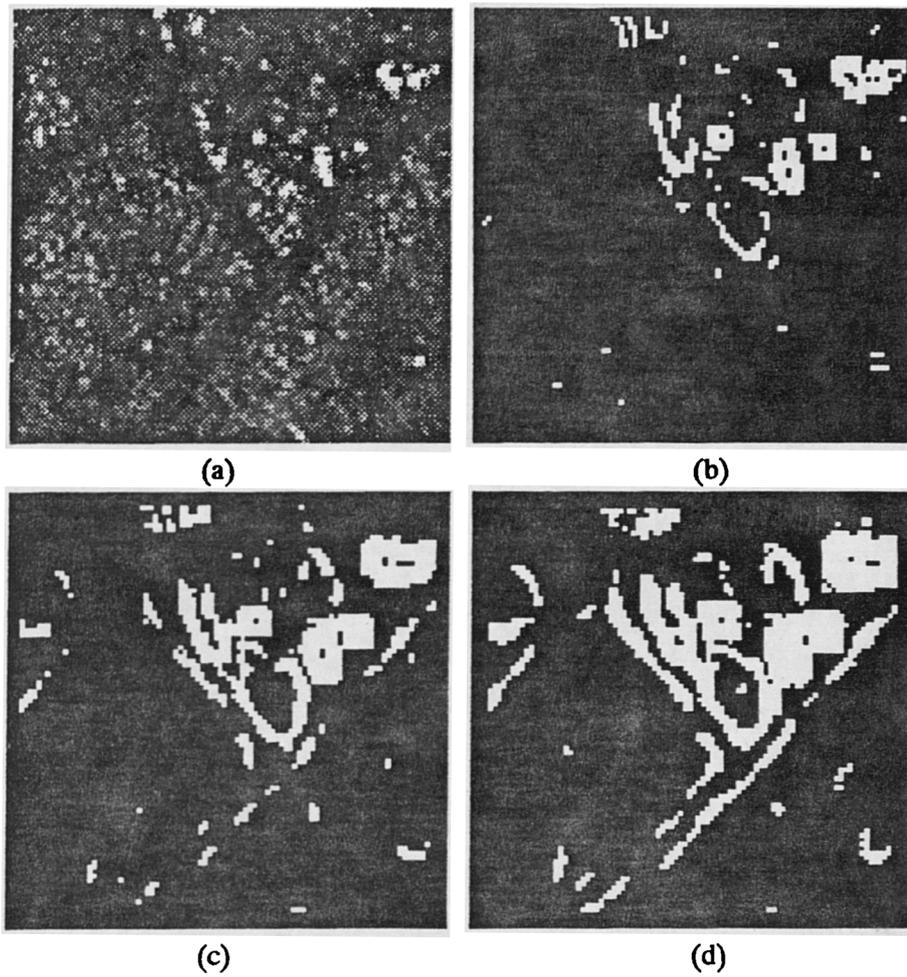


Figure 7.18: (a) ERS-1 FDP image and result of applying the normalised ratio operator using (b) 5×5 , (c) 7×7 and (d) 9×9 windows with thresholds set for a PFA of 10^{-4} .

7.5 Summary

In this chapter:

- Edge detection in SAR images using the local operator approach was investigated in terms of general edge detection criteria.
- Intensity edges may be detected by measuring the difference or ratio of intensities between adjacent regions. The ML estimates of these measures under the simple image model were determined and were found to be given by the normalised ratio and absolute difference operators.
- The false alarm distributions of both of these operators under the simple image model has previously been analysed [57]. Only the normalised ratio operator acts as a constant false alarm rate edge detector, the false alarm rate of the absolute difference operator increases with mean image intensity.
- The difference operator acts as a constant false alarm rate edge detector when it is normalised by the mean intensity. However, it is then equivalent to the ratio operator.
- The coefficient of variation, likelihood ratio and Marr-Hildreth edge detection operators were also investigated but were found to suffer from problems of threshold selection and excessive false alarms.
- The performance of the normalised ratio operator was more fully investigated.
- The general conditions under which the operator acts as a constant false alarm rate edge detector were determined. The operator acts as CFAR detector when clutter is a stationary process, i.e., when there is no change in correlation or texture over an image.
- The theoretical performance of the operator under the simple image model as a function of scale size and edge ratio was determined. The minimum window size for detecting strong edges in a single-look image whilst maintaining a low false alarm rate was found to be 5×5 . Significantly larger windows are needed to detect weak edges whilst maintaining a low false alarm rate.
- The effect of correlation and texture (in the form of K-distributed intensity) on the false alarm rate was determined. Both increase the false alarm rate by increasing the standard error of estimates of the local mean, correlation having the more significant effect. This was confirmed by measured false alarm rates.
- Two methods of compensating for correlation were investigated, subsampling and correcting threshold to take correlation into account. The latter was found to have the more beneficial effect both in terms of edge detection and false alarm rates.

- When applied to real images the performance of the operator depends on the scale size of the edges being detected. Edge thickening occurs when too large a window is used to detect strong edges this leads to small scale edges being obscured. In addition, point targets can confuse edge detection. It is suggested that these problems may be overcome by only detecting edges at larger scales when no edges have been detected at smaller scales; this would involve an initial stage of point target detection.

Chapter 8

Point target detection

Radiometrically bright point targets appear as peaks (local maxima) in a SAR image. However, a large number of peaks are also generated by speckle. Thus, methods are required to distinguish peaks representing point targets (information) from speckle peaks (noise). These methods are required even when we are not interested in the point targets themselves; as happens in many applications. This is because point target peaks which are much brighter than speckle peaks will confuse edge detection operators which have no knowledge of such features (see Chapter 7). Thus, point targets need to be identified and suitably flagged prior to edge detectors being applied.

So as to be able to predict the magnitude of speckle peaks and thus be able to determine how bright a point target peak has to be, to be distinguishable from speckle, the statistics of speckle peaks under the simple image model are analysed in Section 8.1. Different approaches to point feature detection are then considered in the light of this analysis in Section 8.2. In Section 8.3 the results of using these methods for detecting point-like targets in an ERS-1 FDP image are presented.

8.1 Statistics of speckle peaks

To distinguish peaks caused by bright point targets from speckle peaks we are interested in knowing

1. In a homogeneous region of a SAR image, how intense are the maximum and minimum values relative to the mean intensity and size of the region?
2. What fluctuations in intensity occur between a pixel and its immediate neighbours?

In the following analysis, which aims to answer these questions, we shall limit ourselves to the simple image model already introduced in previous chapters, i.e., we shall assume that image intensity within a homogeneous region of a n -look image is gamma distributed with order parameter n and that pixel values are spatially independent.

8.1.1 Distributions of maximum and minimum values

Before being able to analyse the statistics of the maximum and minimum intensity within a region we need to know the general distributions of the maximum and minimum values of an arbitrary set of independent random variables each with known distribution.

Let x_1, \dots, x_m be a set of independent, positive random variables where each x_i has pdf $f_{x_i}(x)$ and cdf $p_{x_i}(x)$. The probability that the maximum value of the variables is less than x is given by the probability that they are all less than x , and the probability that their minimum value is greater than x is given by the probability that they are all greater than x . Thus the maximum and minimum have cdf's

$$p_{\max}(x) = \prod_{i=1}^m p_{x_i}(x) \quad (8.1)$$

$$p_{\min}(x) = 1 - \prod_{i=1}^m (1 - p_{x_i}(x)) \quad (8.2)$$

and pdf's

$$f_{\max}(x) = \frac{d}{dx} p_{\max}(x) = \sum_{j=1}^m \left(f_{x_j}(x) \prod_{\substack{i=1 \\ i \neq j}}^m p_{x_i}(x) \right) \quad (8.3)$$

$$f_{\min}(x) = \frac{d}{dx} p_{\min}(x) = \sum_{j=1}^m \left(f_{x_j}(x) \prod_{\substack{i=1 \\ i \neq j}}^m (1 - p_{x_i}(x)) \right) \quad (8.4)$$

When the random variables are statistically identical with pdf $f_x(x)$ and cdf $p_x(x)$, the cdf's of the maximum and minimum are given by

$$p_{\max}(x) = p_x(x)^m \quad (8.5)$$

$$p_{\min}(x) = 1 - (1 - p_x(x))^m \quad (8.6)$$

and the pdf's are given by

$$f_{\max}(x) = m f_x(x) p_x(x)^{m-1} \quad (8.7)$$

$$f_{\min}(x) = m f_x(x) (1 - p_x(x))^{m-1} \quad (8.8)$$

8.1.2 Maximum and minimum intensity

Substituting the pdf (3.59) and cdf (3.60) of the n -look gamma distributed intensity into (8.7) and (8.8), the maximum and minimum intensity of m pixels sampled from a homogeneous region with mean μ will have respective pdf's

$$f_{\max}(x) = m \frac{e^{-\frac{x}{\beta}} x^{n-1}}{\Gamma(n) \beta^n} \left(1 - e^{-\frac{x}{\beta}} \sum_{k=0}^{n-1} \frac{\left(\frac{x}{\beta}\right)^k}{k!} \right)^{m-1} \quad (8.9)$$

$$f_{\min}(x) = m \frac{e^{-\frac{x}{\beta}} x^{n-1}}{\Gamma(n) \beta^n} \left(\sum_{k=0}^{n-1} \frac{\left(\frac{x}{\beta}\right)^k}{k!} \right)^{m-1} \quad (8.10)$$

where $\beta = \mu/n$.

In the single-look case

$$f_{\max}(x) = \frac{m}{\mu} e^{-\frac{x}{\mu}} \left(1 - e^{-\frac{x}{\mu}} \right)^{m-1} \quad (8.11)$$

$$f_{\min}(x) = \frac{m}{\mu} e^{-\frac{x}{\mu}} \quad (8.12)$$

and the expected values of the maximum and minimum intensity are given by

$$\langle \max \rangle = \mu (\psi(m+1) + \gamma_E) \quad (8.13)$$

$$\langle \min \rangle = \frac{\mu}{n} \quad (8.14)$$

As $\ln m$ gives a good approximation to $\psi(m+1)$ even for small values of m (≥ 10) [1], the maximum intensity of pixels in a homogeneous region can be considered to increase logarithmically with the regions size. The expected value of the minimum intensity is inversely proportional to the size of the region. The relation between the size of a homogeneous region and the expected maximum and minimum intensity values, are plotted in Figure 8.1.

A measure of the dynamic range of data is given by the ratio of the maximum and minimum values. The ratio of the expected maximum and minimum single-look intensity is given by

$$\frac{\langle \max \rangle}{\langle \min \rangle} = m (\psi(m+1) + \gamma_E) \approx m (\ln m + \gamma_E) \quad (8.15)$$

Because $m \gg \ln m$ for large m , relative to the mean intensity, most of the dynamic range arises from very small as opposed to very large intensities, i.e., relative to the mean intensity speckle troughs are much larger than speckle peaks. Given the shape of the exponential distribution this is hardly surprising.

Note: the process of determining the maximum and minimum of a block of pixels is equivalent to applying morphological dilation and erosion [39] respectively, to that block of pixels.

8.1.3 Distribution of maximum and minimum ratios

We now turn to look at the magnitude of the intensity fluctuations occurring between a pixel and its immediate neighbours. These can be analysed in terms of the maximum and minimum intensity ratio between the pixel and its immediate neighbours.

If the positive random variables x_i are independent and have cdf's $p_{x_i}(x)$, the probability that the ratio x_0/x_i ($i \neq 0$) is greater than r given x_0 is

$$\text{prob}\left(\frac{x_0}{x_i} > r\right) = \text{prob}\left(\frac{x_0}{r} > x_i\right) = p_{x_i}\left(\frac{x_0}{r}\right) \quad (8.16)$$

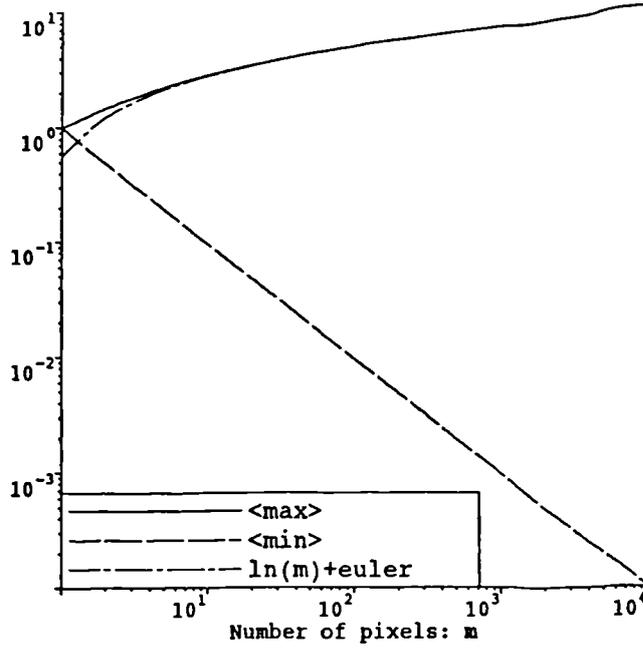


Figure 8.1: Relation between the number of pixels contained in a homogeneous region with unit mean intensity in a single-look image, and the expected values of the maximum and minimum intensity, $\langle \max \rangle$ and $\langle \min \rangle$.

Thus the probability $p_{\exists}(\tau)$, that amongst m random variables x_1, \dots, x_m there exists a x_i such that the ratio x_0/x_i is less than τ is given by

$$p_{\exists}(\tau) = 1 - \text{prob} \left(\frac{x_0}{x_i} > \tau : i = 1, \dots, m \right) = 1 - \int_0^{\infty} f_{x_0}(x) \prod_{i=1}^m p_{x_i} \left(\frac{x}{\tau} \right) dx \quad (8.17)$$

where $f_{x_0}(x)$ is the pdf of x_0 . Likewise, the probability $p_{\forall}(\tau)$, that for each x_i the ratio x_0/x_i is less than τ is given by

$$p_{\forall}(\tau) = \text{prob} \left(\frac{x_0}{x_i} < \tau : i = 1, \dots, m \right) = \int_0^{\infty} f_{x_0}(x) \prod_{i=1}^m \left(1 - p_{x_i} \left(\frac{x}{\tau} \right) \right) dx \quad (8.18)$$

Both $p_{\exists}(\tau)$ and $p_{\forall}(\tau)$ define cdf's (going from zero to infinity they monotonically increase from zero to one). Their corresponding pdf's $f_{\exists}(\tau)$ and $f_{\forall}(\tau)$, are given by differentiation with respect to τ .

The complement of $p_{\exists}(\tau)$, i.e., $1 - p_{\exists}(\tau)$, gives the probability that for each x_i the ratio x_0/x_i is greater than τ . Whilst the complement of $p_{\forall}(\tau)$, gives the probability that there exists a x_i such that the ratio x_0/x_i is greater than τ . Setting $\tau = 1$ in $1 - p_{\exists}(\tau)$ and $p_{\forall}(\tau)$ gives the respective probabilities that x_0 is the maximum value of the $m + 1$ variables, and that it is the minimum value. If the random variables are statistically identical both of these probabilities must equal $1/(m + 1)$. This is because there is then an equal likelihood that any one variable is the maximum/minimum.

Assuming that pixel values are spatially independent in an image, a value of $m = 4$ is used to determine the distribution of ratios between a pixel and its immediate horizontal and vertical neighbours (4-connected neighbourhood), and a value of $m = 8$ is used to determine the distribution of ratios between a pixel and its immediate horizontal, vertical and diagonal neighbours (8-connected neighbourhood).

8.1.4 Intensity fluctuations

Under the simple image model the probability that the intensity ratio between a pixel x_0 within a homogeneous region with mean intensity μ , and at least one of its m neighbours is less than r is given by substituting the pdf (3.59) and cdf (3.60) of the gamma distributed intensity into (8.17)

$$\begin{aligned} p_{\exists}(r) &= 1 - \int_0^{\infty} \frac{e^{-\frac{x}{\beta}} x^{n-1}}{\Gamma(n) \beta^n} \left(1 - e^{-\frac{x}{r\beta}} \sum_{k=0}^{n-1} \frac{\left(\frac{x}{r\beta}\right)^k}{k!} \right)^m dx \\ &= \frac{1}{\Gamma(n)} \sum_{\lambda=0}^m \binom{m}{\lambda} (-1)^\lambda \int_0^{\infty} e^{-x(1+\frac{\lambda}{r})} x^{n-1} \left(\sum_{k=0}^{n-1} \frac{\left(\frac{x}{r}\right)^k}{k!} \right)^\lambda dx \end{aligned} \quad (8.19)$$

where $\beta = \mu/n$. Likewise, substituting the pdf and cdf of the gamma distributed intensity into (8.18) gives the probability that the ratio between a pixel x_0 and each of its m neighbours is less than r

$$p_{\forall}(r) = \frac{1}{\Gamma(n)} \int_0^{\infty} e^{-x(1+\frac{m}{r})} x^{n-1} \left(\sum_{k=0}^{n-1} \frac{\left(\frac{x}{r}\right)^k}{k!} \right)^m dx \quad (8.20)$$

As the value of the ratios are independent of the mean intensity, so the above probabilities are also independent of the mean intensity. In the single-look case

$$p_{\exists}(r) = 1 - \sum_{\lambda=0}^m \binom{m}{\lambda} \frac{(-1)^\lambda}{1 + \frac{\lambda}{r}} \quad (8.21)$$

$$p_{\forall}(r) = \frac{1}{1 + \frac{m}{r}} \quad (8.22)$$

When $m = 8$ the complements of (8.21) and (8.22), (i.e., $1 - p_{\exists}(r)$ and $1 - p_{\forall}(r)$) represent the respective probabilities that in a homogeneous area in a single-look image the intensity ratio between a pixel and each of its 8-connected neighbours is greater than r , and that the ratio between a pixel and at least one of its 8-connected neighbours is greater than r . These probabilities are plotted in Figure 8.2 with respect to the ratio r .

As we would expect $1 - p_{\exists}(r)$ is always less than $1 - p_{\forall}(r)$. Both probabilities are one at $r = 0$. As r increases the probability that a pixel is r times more intense than at least one of its 8-connected neighbours decreases slowly. A much faster decrease occurs with the probability that a pixel is r times more intense

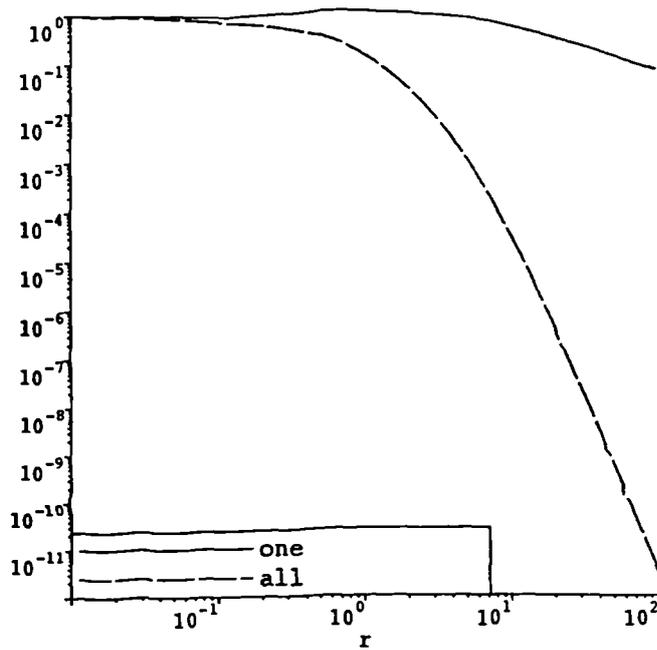


Figure 8.2: Probability that a pixel within a homogeneous region in a single-look image is r times more intense than (i) one and (ii) all of its 8-connected neighbours.

than all of its 8-connected neighbours. At $r = 1$, the probabilities $1 - p_{\exists}(r)$ and $1 - p_{\forall}(r)$ are equivalent to the probability that the pixel is a local maximum ($1/9$) and the probability that it is not a local minimum ($8/9$), respectively. When $r > 1$, $1 - p_{\exists}(r)$ gives a measure of the minimum fluctuation in intensity existing between a speckle peak and its 8-connected neighbours and $1 - p_{\forall}(r)$ gives a measure of the maximum fluctuation in intensity existing between a speckle peak and its 8-connected neighbours.

The above only describes the intensity fluctuations occurring within a homogeneous region. The statistics of intensity fluctuations occurring across an edge are also of interest. Along an edge separating regions of mean intensity μ_x and μ_y , the probability that the normalised ratio between a pixel x on one side of the edge and a neighbouring pixel y on the other side, is less than r is

$$\text{prob} \left(\min \left(\frac{x}{y}, \frac{y}{x} \right) < r \right) = p_r(r|R, n) \quad (8.23)$$

where $p_r(r|R, n)$ is the cdf of the normalised ratio of gamma distributed random variables (7.22). This cdf depends on $R = \mu_x/\mu_y$, the ratio of the mean intensities on either side of the edge, and n , the number of looks. The probability that among m such pairs of pixels (i.e., along an edge m pixels long), there exists a pair whose normalised ratio is less than r is given by $1 - (1 - p_r(r|R, n))^m$. Whilst the probability that the normalised ratio between each pair is less than r is given by $p_r(r|R, n)^m$.

Both probabilities are plotted in Figure 8.3 for the single-look case for different values of the ratio of the mean intensities on either side of the edge ($R = 2, 4, 8, 16$) with $m = 10$. As the normalised ratio is always less than or equal to one both probabilities are one at $r = 1$. Both probabilities show similar behaviour to those plotted in Figure 8.2. The probability that the normalised ratio of at least one pair of pixels is less than r decreases slowly with r , whilst the probability that the normalised ratio of each pair of pixels is less than r decreases quickly with r . As the ratio of the mean intensity on either side of the edge increases, both probabilities increase, the difference between them decreases and the probability that there exists a pair of pixels whose normalised ratio is less than r tends to certainty.

8.2 Approaches to point feature detection

In this Section we consider a variety of approaches to point feature detection based on the analysis given in the previous section and observations made in Chapter 7.

1. **Thresholding:** The very brightest features in an image may be detected by global thresholding. Global image statistics can be used to select thresholds automatically (e.g., ten standard deviations above the mean). To take into account variations in mean intensity over the image, an adaptive threshold can be employed. By estimating the mean intensity and using (8.9) a threshold can be set for a given PFA of a speckle peak being detected over the region assuming it is homogeneous and the simple image model applies. If a low false alarm rate is to be maintained thresholding will only detect the very brightest parts of point features rather than the full extent of the feature they represent.
2. **Edge detection:** As already mentioned in Chapter 7, point features generate distinctive signatures in edge maps. The signatures of point features can be separated out from the signatures of true edges using the fact that point features have no orientation, i.e., they consist of those pixels that are detected as being edges at all orientations. Using the normalised ratio operator this is equivalent to retaining the maximum as opposed to the minimum of the normalised ratios estimated over different orientations. Thresholds can be selected using similar methods to those used in edge detection. Determining the true extent of point features from the 'blocky' signatures which result is not easy.
3. **Coefficient of variation:** A simple measure of local heterogeneity in SAR images is the coefficient of variation (see Chapter 7). Point features generate extremely bright and easily identifiable signatures in the resulting coefficient of variation image (much brighter than those generated by edges). As with the maximum normalised ratio, these signatures are very 'blocky' and interpreting them is not easy. Once again, as in edge detection, threshold selection is not straightforward. Thresholds certainly need

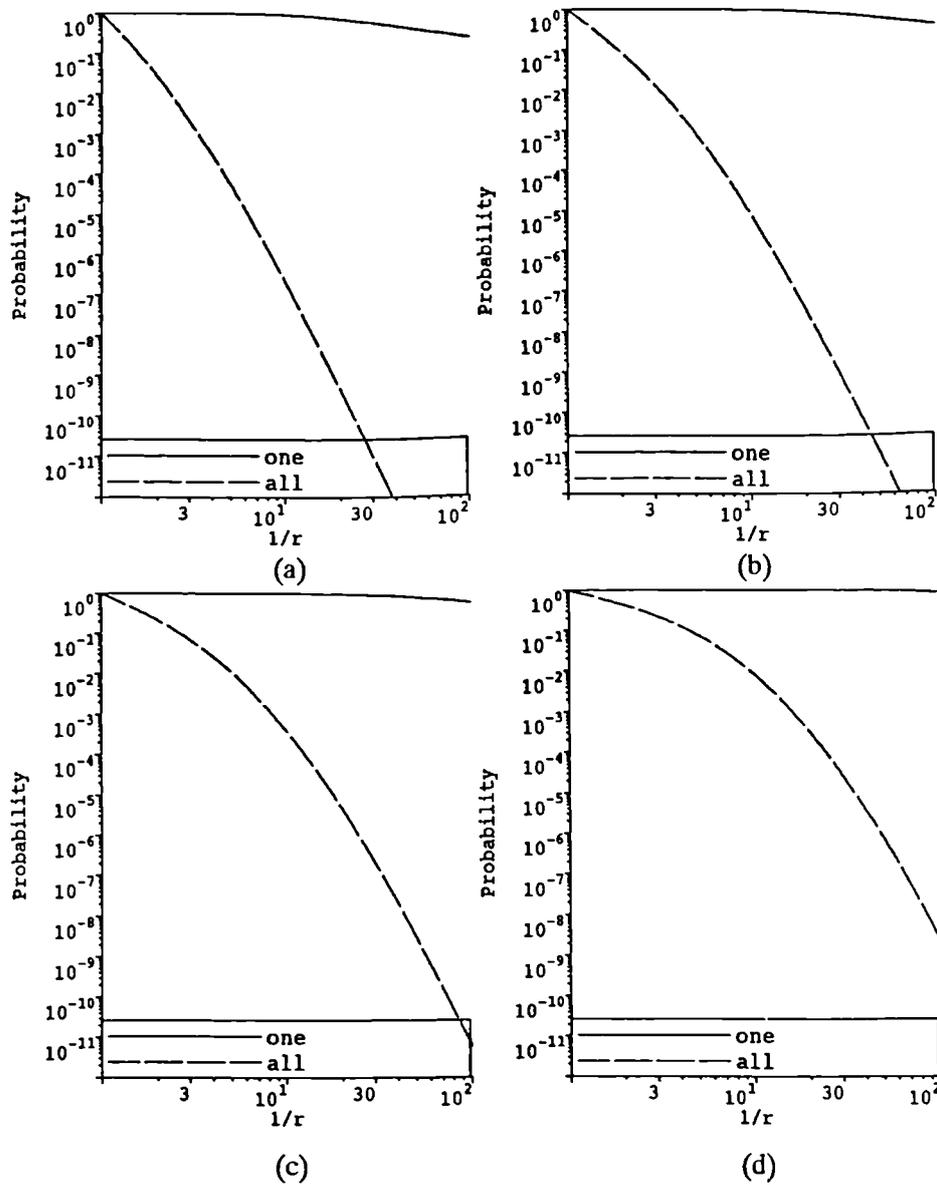


Figure 8.3: Probability that along an edge ten pixels long separating regions of relative mean intensity R in a single-look image, that the normalised ratio of (i) one and (ii) all pairs of pixels lying across the edge is less than τ ; (a) $R = 2$, (b) $R = 4$, (c) $R = 8$, and (d) $R = 16$.

to be higher than those used for edge detection.

4. **Morphological methods:** Morphological opening using a $n \times n$ window smooths all peaks in an image less than n pixels wide [39]. Therefore dividing an image by its morphological opening enhances peaks whilst normalising the rest of the image. However, it is unclear at what value to threshold the resulting image at to distinguish between point features and speckle peaks.
5. **Magnitude of intensity fluctuations:** To be apparent at all, point features must be of much greater magnitude than peaks in surrounding speckle. This implies that intensity fluctuations occurring around a point feature must be greater than those occurring in speckle alone. It has been investigated whether point features, unlike edges, can be detected just on the basis of the magnitude of the intensity fluctuations occurring from pixel to pixel.

8.2.1 Point feature detection based on the magnitude of intensity fluctuations

While Figure 8.2 and 8.3 show that large intensity fluctuations exist between neighbouring pairs of pixels in a single-look image, they also show that they do not connect to form boundaries across which there is a large fluctuation in intensity between all pairs of pixels lying across the boundary. This is the case both within extended targets and across edges. However, by their very nature we expect large intensity fluctuations to exist on all sides of point features.

An algorithm has been devised which uses this feature to detect small bright segments. The ratio between all pairs of 8-connected pixels in an image is calculated (this ratio is normalised to be greater than one). Neighbouring pixels are defined as belonging to the same segment if their normalised ratio is less than a given threshold. In the resulting segmentation the normalised ratio between any two neighbouring pixels on either side of a segment boundary must be greater than the threshold. The distribution given by (8.19) can be used to set the threshold τ such that the probability of a speckle peak being τ times brighter than all surrounding pixels is extremely low; as with the normalised ratio operator used for edge detection, this false alarm rate will be constant so long as speckle is a strict-sense stationary multiplicative noise process (see Chapter 7).

8.3 Results of point target detection

Figures 8.5 to 8.10 show the results of applying the above methods to the ERS-1 FDP image shown in Figure 8.4 (the same as that shown in Figure 7.18). This image is taken over an agricultural region in Flevoland, Netherlands, but contains a number of bright point-like targets which are most likely caused by scattering from man-made features such as buildings. It must be stressed that as visual inspection of the original image is the only available method for

determining whether detected features represent point targets or false alarms the comparison of these results is highly subjective.

In Figures 8.5 and 8.6 thresholds are set for a given number of standard deviations above the estimated mean intensity μ . This mean is estimated over the whole image for global thresholding (Figure 8.5) and over a 16×16 pixel window for adaptive thresholding (Figure 8.6). Since the ERS-1 FDP image is a 3-look intensity image, assuming the simple image model applies the intensity standard deviation over a homogeneous area is given by $\mu/\sqrt{3}$ (see Chapter 4). Thus the threshold for k standard deviation above the mean is given by $\mu(1 + k/\sqrt{3})$. The PFA's were derived using (8.9) where m is the size of the window over which the mean is estimated; we have assumed that pixel values are independent. The features detected by both global and adaptive thresholding may be directly related to point-like targets in the original image. Global thresholding appears to be more successful at detecting the brightest targets and less susceptible to false alarms than adaptive thresholding. This might not be the case in an image which contained larger variations in mean background intensity.

The blocky signatures resulting from using the maximum normalised ratio and coefficient of variation are clearly apparent in Figures 8.7 and 8.8 respectively. The gaps in the signatures in Figure 8.7 occur when the window lies directly over a point-like target. This is because the point-like target is then split equally between the two half windows over which the normalised ratio is estimated, whatever the orientation (see Chapter 7). The larger holes in the signatures in Figure 8.8 occur when the point-like feature takes up most of the window within which the coefficient of variation is estimated. While both of these methods give a good indication of the presence of point-like targets, it is difficult to directly relate the more complex signatures to the true location of the point-like targets they result from.

Figure 8.9 shows the result of dividing the original image by its morphological opening. More point-like features appear to be detected with a low false alarm rate in Figure 8.9 (b) than using straight thresholding (Figures 8.5 and 8.6). However, thresholds were selected using purely visual criteria. More objective methods are needed for doing this.

The results of point target detection based on the magnitude of intensity fluctuations are shown in Figure 8.10. Left: the segmentations generated by merging adjacent pixels whose intensity ratio lies below a given threshold; also any segment whose mean intensity lies below the mean intensity of the whole image is merged with surrounding segments. Right: the result of thresholding the mean intensity of segments just above the mean intensity of the whole image. Most of the image is represented as a single segment with smaller segments indicating brighter regions of the image. At a threshold of 2.0 and higher these mainly correspond to point-like targets in the original image.

The results presented in Figures 8.5 to 8.10 are somewhat limited as they only show the result of the different methods applied to a single 3-look image. However, they help to illustrate many of the points made in Section 8.2. Simple thresholding techniques are useful for detecting many point-like features. The difficulties involved in interpreting the signatures produced using the maximum

normalised ratio and coefficient of variation do not make these methods an attractive option. Point target detection based on morphological opening and the magnitude of intensity fluctuations both produce promising results which warrant further investigation of these methods.

8.4 Summary

In this chapter:

- The problem of distinguishing peaks representing bright point targets from speckle peaks has been investigated.
- The statistics of speckle peaks under the simple image model were analysed both in terms of the maximum and minimum intensity within an area of given size, and in terms of the intensity fluctuations occurring between neighbouring pixels.
- A number of approaches to point target detection were described including methods based on: simple thresholding, edge detection operators, morphological opening and the magnitude of intensity fluctuations between adjacent pixels.
- The results of applying these methods to an ERS-1 FDP image were presented. The edge detection operators (normalised ratio and coefficient of variation) produced the poorest results. The results using the morphological and intensity fluctuation approaches are as good as, if not marginally better than, those produced by simple thresholding. Further work is needed to determine whether this is generally the case when the methods are applied to a variety of SAR images containing point-like targets.

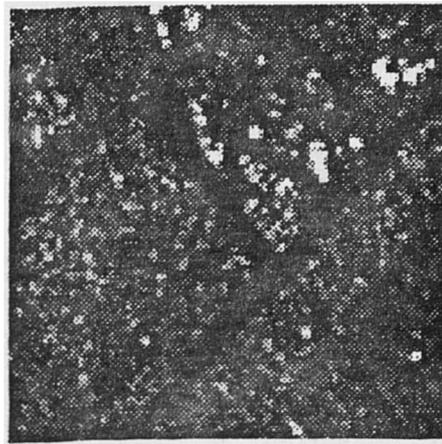
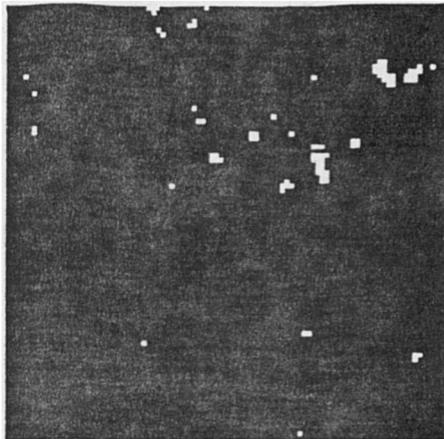
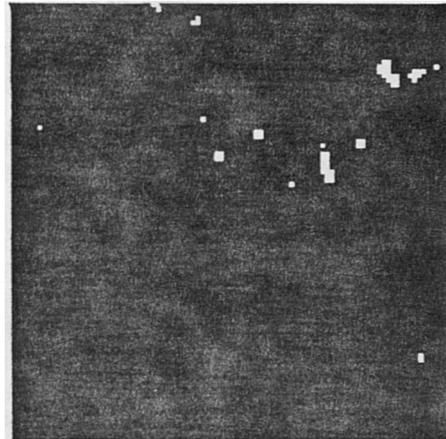


Figure 8.4: ERS-1 FDP image of Flevoland, Netherlands showing bright point-like features.

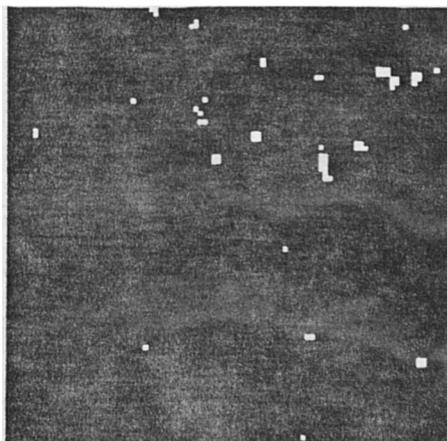


(a)

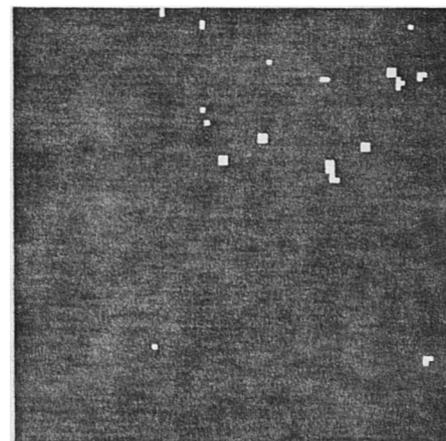


(b)

Figure 8.5: Point target detection using thresholding at (a) 7, and (b) 10, standard deviations above the image mean assuming the image is homogeneous. This corresponds to a PFA of (a) 0.297, and (b) 3.40×10^{-3} .



(a)



(b)

Figure 8.6: Point target detection using thresholding at (a) 7, and (b) 10, standard deviations above the local mean estimated over a 16×16 pixel area assuming the image is locally homogeneous. This corresponds to a PFA of (a) 0.898×10^{-2} , and (b) 0.872×10^{-4} .

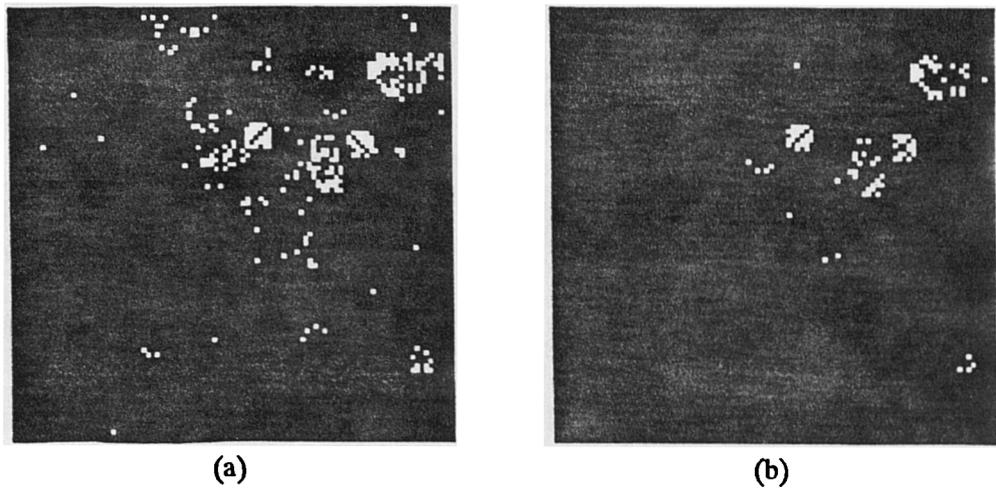


Figure 8.7: Point target detection using the maximum normalised ratio estimated over a 5×5 pixel window with thresholds set for a PFA of (a) 10^{-2} , and (b) 10^{-4} .

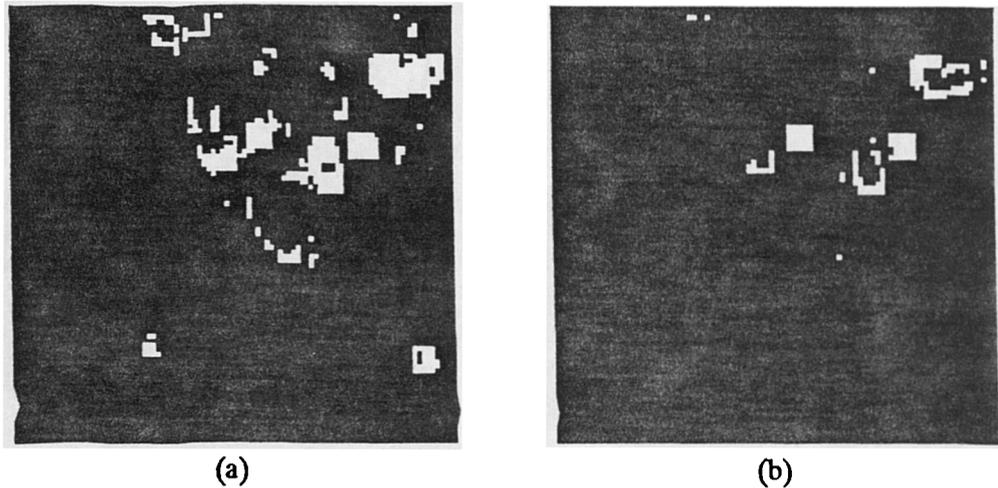


Figure 8.8: Point target detection using the coefficient of variation estimated over a 5×5 pixel window and thresholds at (a) 1.5, and (b) 3.

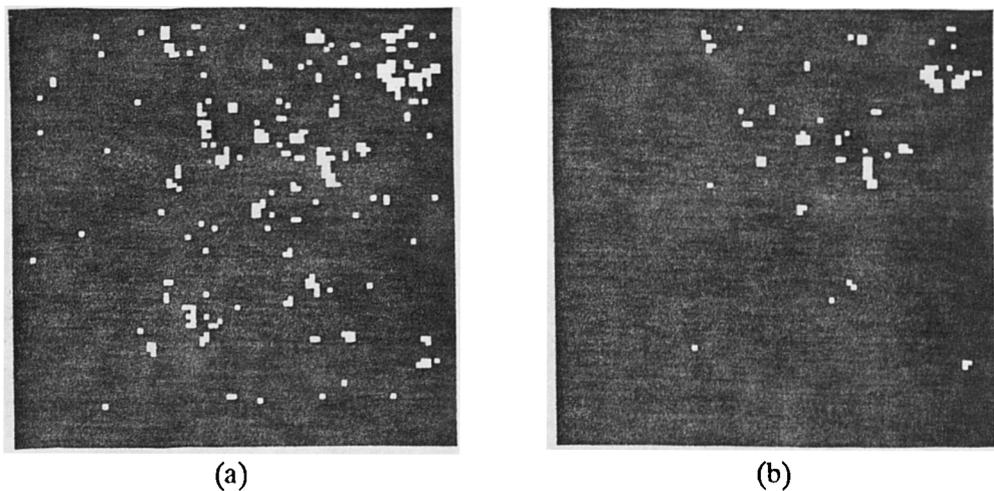


Figure 8.9: Point target detection using the original image divided by its morphological opening calculated over a 5×5 pixel window and then thresholded at (a) 10, and (b) 20.

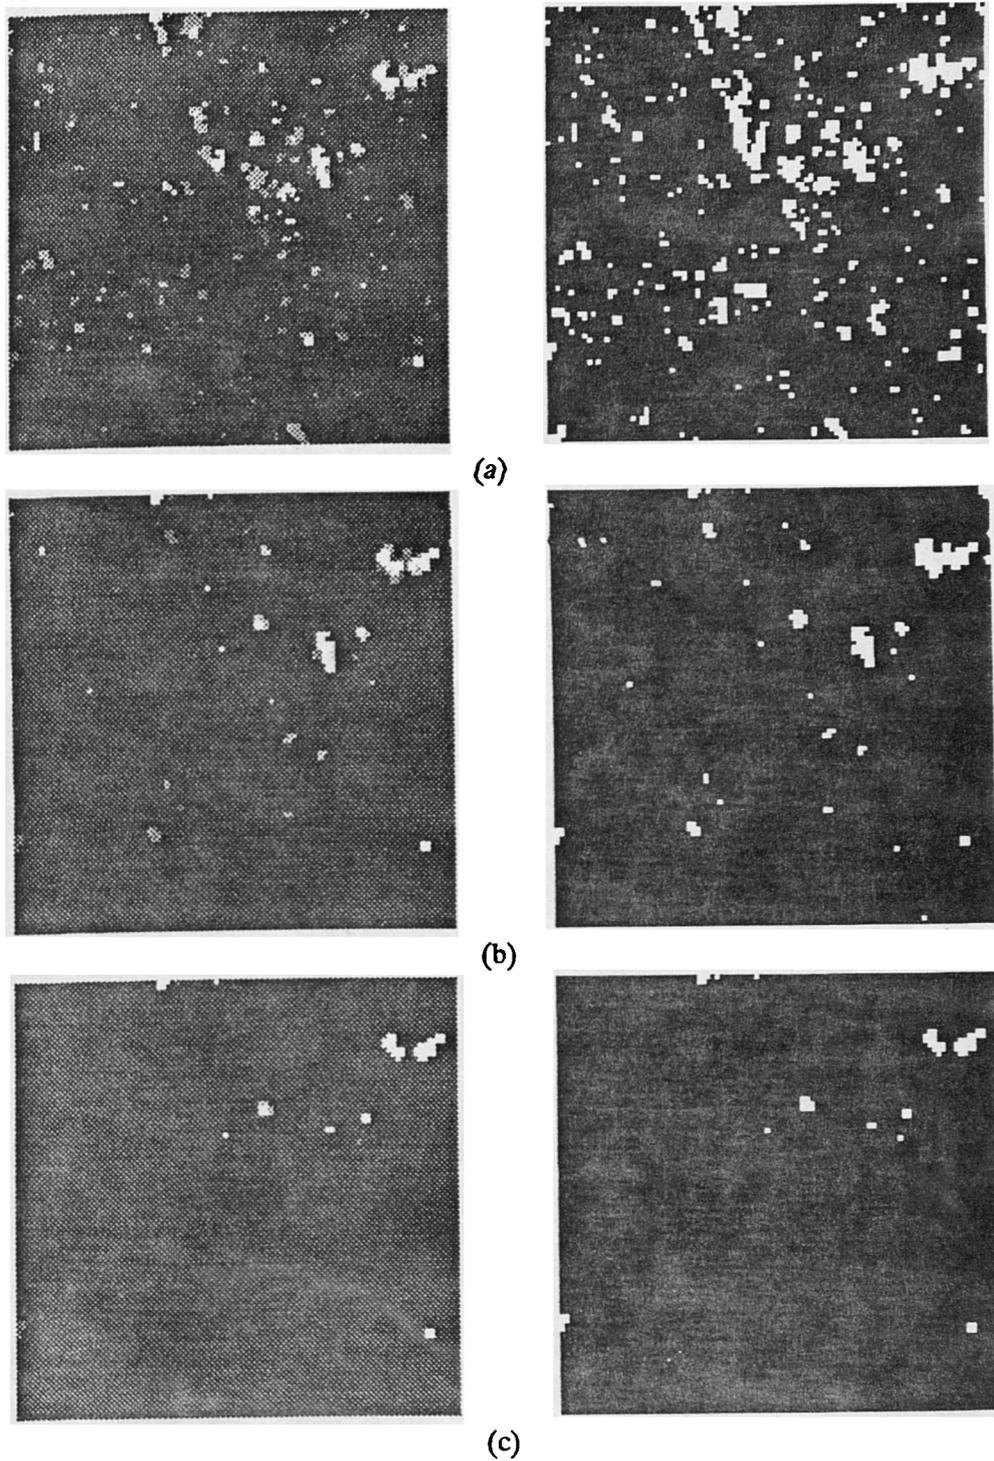


Figure 8.10: Point target detection based on the magnitude of intensity fluctuations. Left: segmentation resulting from merging adjacent pixels whose intensity ratio lies below (a) 1.5, (b) 2.0, and (c) 2.5; using (8.19) these thresholds correspond to PFA's of (a) 2.49×10^{-2} , (b) 5.82×10^{-3} , and (c) 1.46×10^{-3} . Right: the result of thresholding the images shown on the left just above the mean intensity of the original image.

Chapter 9

Segmentation

In order to fully exploit the near-global coverage and regular revisit time afforded by satellite-borne SAR, automatic interpretation methods are necessary. For many applications, segmentation of the image into statistically homogeneous regions is particularly important, since it provides a means of imposing structure on the image, which can then be utilised for matching and change detection purposes. Two kinds of change are of interest. Firstly, structural changes may occur between images, as will be normal, for example, in the dynamic environments near ice edges. Secondly, the internal backscattering properties of individual segments may vary, as will occur in agricultural regions because of crop development. Segmentation is fundamental in detecting the first form of change, but is equally important for the second form. This is because, after segmentation, area-based measures can be used to derive the properties of extended targets. Such measures are necessary to combat the effects of speckle, and by using known segments, the averaging involved can be carried out only over pixels from the same region.

Segmentation may help to remove the effects of speckle, but speckle provides a major obstacle to the provision of reliable and robust methods of segmentation. Nonetheless, powerful methods have been developed in recent years, driven mainly by the properties of high resolution airborne data, but based on principles which should be more generally applicable. In this chapter we investigate one such algorithm [61, 62]. The algorithm itself is described in Section 9.1 along with various modifications and known weaknesses.

ERS-1 data provides an excellent dataset on which to test whether the segmentation algorithm can provide information which can be used effectively in applications, and to motivate further developments. Section 9.2 shows the application of the algorithm to a set of ERS-1 images, gathered during 1992 over the Feltwell area of East Anglia. The particular problem of interest in this section is whether segmentation provides a useful tool for analysis of a sequence of ERS-1 images of agricultural areas. We can identify a range of subproblems within this central question:

- Does the segmentation find all the features located by eye?
- Does it locate features that the eye misses?

- Does it find ‘too many’ features?
- Does it provide a means of comparing/matching images?
- Does it permit the identification of structure not present in map data (e.g., splitting or merging of fields)?

9.1 Segmentation algorithm

The algorithm makes use of the fact that edges and segments define each other. It involves an iterative process of edge detection and segment growing. Detected edges are used to limit segment growing, then the resulting segmentation is used to generate an improved edge detection. The algorithm is robust in that there are no scene dependent parameters to set. Thus no operator intervention is required to obtain a full segmentation. Because the full segmentation may contain too much ‘information’ for a given application, segment merging is often required to degrade the segmentation down to the required level of information. This involves the setting of a single parameter which has precise meaning in terms of the probability of an edge being present.

In Section 9.1.1 we first describe how the algorithm uses a measure of segmentation accuracy to determine when to halt, before looking at the edge detection stage in Section 9.1.2, and the segment growing and merging stages in Section 9.1.3.

9.1.1 Segmentation accuracy and halting

After each segment growing stage the accuracy of the segmentation (i.e., segment homogeneity) is measured in terms of the average contrast \hat{v} within segments; this average is weighted by the size of segments and is defined as

$$\hat{v} = \sum_{i=1}^m \frac{n_i}{n} \frac{\sqrt{\hat{\sigma}_i^2}}{\hat{\mu}_i} \quad (9.1)$$

where m is the number of segments, n_i is the number of pixels making up the i 'th segment, $n = n_1 + \dots + n_m$ is the total number of pixels in the image, and $\hat{\mu}_i$ and $\hat{\sigma}_i^2$ are the estimated mean and variance of the pixel values $x_{i_1}, \dots, x_{i_{n_i}}$ making up the i 'th segment given by

$$\hat{\mu}_i = \hat{x}_i \quad (9.2)$$

and

$$\hat{\sigma}_i^2 = \frac{n_i}{n_i - 1} (\hat{x}_i^2 - \hat{x}_i^2) \quad (9.3)$$

Here we have assumed that pixel values are uncorrelated. This measure of segmentation accuracy provides a means of monitoring the segmentation process and deciding when to halt. As segments become more homogeneous the contrast

within each segment decreases tending to a constant which is dependent on the type of data. Thus the accuracy measure should decrease as the segmentation improves.

In the standard implementation of the algorithm, iteration halts when there is an increase in the accuracy measure from one iteration to the next. The segmentation generated by the previous iteration is then output, i.e., the segmentation corresponding to the first minimum in the accuracy measure. Tests have shown that when the algorithm is allowed to continue iterating until a second or third minimum in the accuracy measure is reached, there is only a small decrease in the value of the accuracy measure, if at all, and there is no overall improvement in the segmentation.

Figure 9.1 shows the result of halting the segmentation algorithm at the first, second and third minima in the accuracy measure when applied to a RSRE amplitude image of an agricultural scene, and Figure 9.2 plots how the accuracy measure varied over different iterations. Although the three segmentations differ, visually no single segmentation appears to be better than any of the others (see Section 9.1.2.5 below for a discussion of more objective methods of comparing segmentations). The accuracy drops steeply over the first two iterations but then flattens out with only minor fluctuations. The first, second and third minima occur at the 5th, 9th and 13th iterations with respective values 0.4925, 0.4899 and 0.4909. From (3.97) we would expect the contrast within a homogeneous region in a single-look amplitude image, and thus the accuracy measure of a correct segmentation, to equal $\sqrt{4/\pi - 1} \approx 0.5227$. However, the minimum values in Figure 9.2 are all slightly smaller. This is most likely due to correlation being ignored when estimating the standard deviation within segments. Correlation induces a negative bias in the estimate of the standard deviation whilst the estimate of the mean remains unbiased; this has the overall effect of inducing a negative bias in the contrast

9.1.2 Edge Detection

9.1.2.1 The normalised gradient operator

The algorithm detects edges using a normalised gradient operator. At each point in an image the gradient d across a rectangular window centred at that point is estimated by splitting the window into two equal halves along its longer side, estimating the mean values μ_a and μ_b within each half by averaging the pixel values within it, and then taking the difference of these two averages. This estimated gradient which equals

$$\hat{d} = \hat{x}_a - \hat{x}_b \quad (9.4)$$

has mean value $\mu_{\hat{d}} = \mu_a - \mu_b$ and standard deviation

$$\sigma_{\hat{d}} = \sqrt{\sigma_a^2/n_a + \sigma_b^2/n_b} \quad (9.5)$$

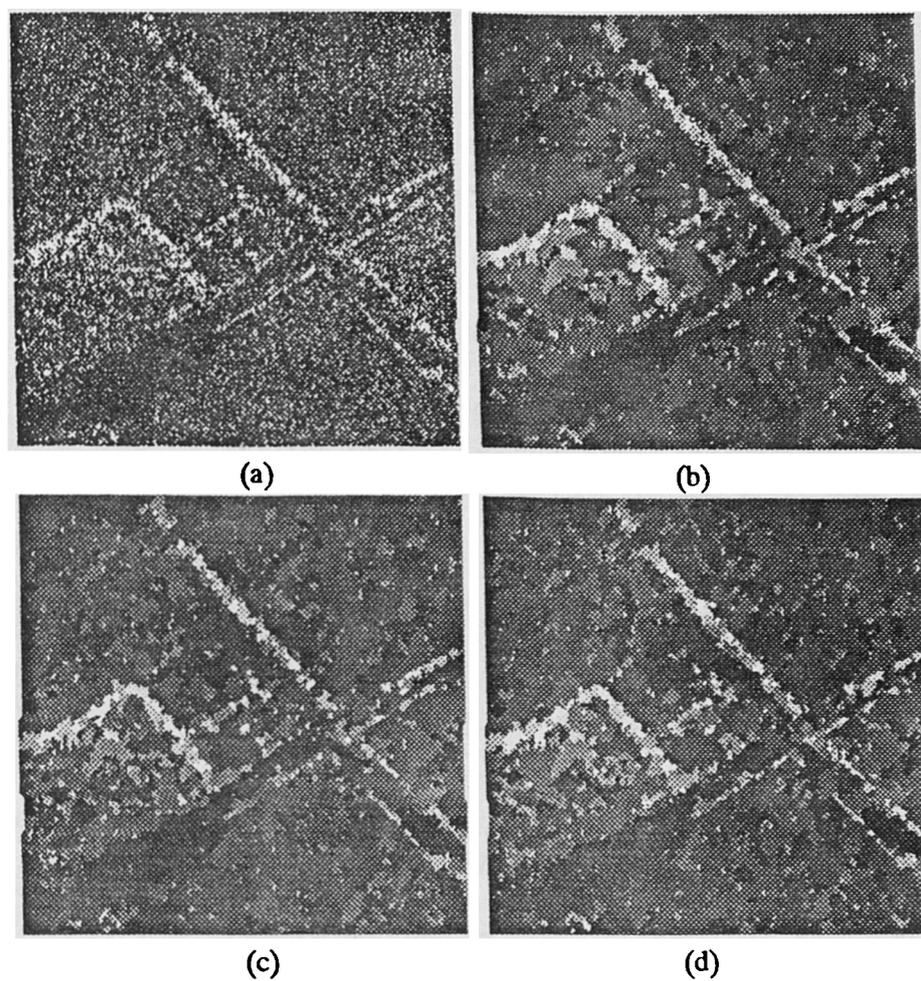


Figure 9.1: (a) A RSRE single-look amplitude X-band image of an agricultural scene (the same as that shown in Figure 3.9) and the result of segmentation when the algorithm is halted at the (b) first, (c) second and (d) third minima in the accuracy measure. Pixels within each segment are set to the average intensity within the segment, with ensuing suppression of speckle.

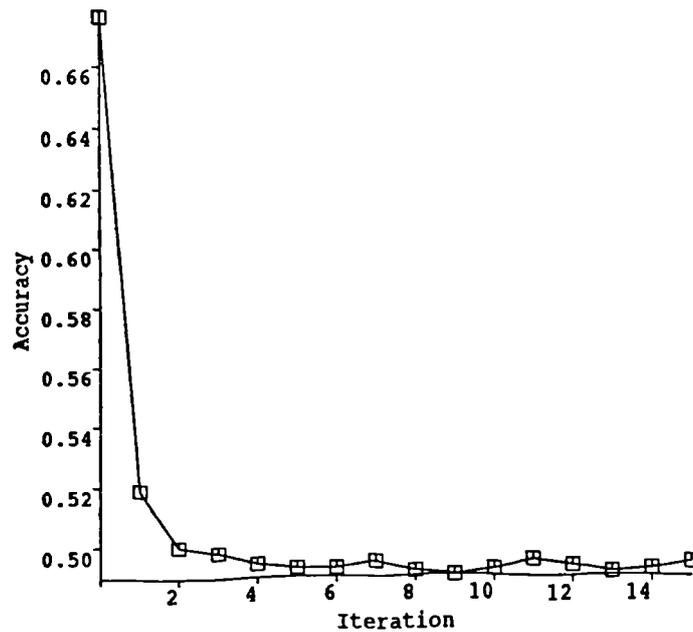


Figure 9.2: Variation in the accuracy measure over different iterations.

where σ_a^2 and σ_b^2 are the variances in the two half windows, and n_a and n_b are the number of independent samples in each half window; normally $n_a = n_b$ but this is not always the case (see later). Thus the false alarm distribution will be dependent on the local variance. To make the false alarm rate constant across the image the gradient needs to be normalised by dividing it by an estimate of its standard deviation $\sigma_{\hat{d}}$. Finally, as it is the magnitude of the gradient that is of interest rather than the direction of slope, the absolute value of the estimated normalised gradient is usually taken. The greater the absolute value output by this normalised gradient operator, the greater the likelihood that an edge is present within the window.

The normalisation factor $\sigma_{\hat{d}}$ is given by the square root of the sum over both half windows of the estimated variance within each half window divided by its size (9.5). Each of these estimated variances is given by the average of the estimated variances of the segments generated by the previous iteration over which the half window lies, weighted by the proportion of the half window each segment occupies (the estimated variance within each segment given by (9.3) will already be known from calculating the accuracy measure at the end of the previous iteration). In this manner the previous segmentation is fed into the next stage of edge detection. For the first iteration the whole image is treated as a single segment; this generates a large normalisation factor and only the strongest edges are detected. In the original algorithm standard deviations as opposed to variances are averaged [61]. This results in the standard deviation of the gradient $\sigma_{\hat{d}}$ being incorrectly given as

$$\sigma_{\hat{d}} = \sigma_a / \sqrt{n_a} + \sigma_b / \sqrt{n_b} \quad (9.6)$$

as opposed to (9.5). When it is assumed that pixel values in both segments are statistically identical this leads to a correcting factor of $\sqrt{2}$ having to be applied to false alarm thresholds.

To detect different strength edges the algorithm uses rectangular windows of dimensions 3×7 , 5×19 , 7×19 , 7×27 , 9×45 and 13×57 pixels. A 3×3 normalised Sobel operator [28] is used to detect the very smallest edges. In order to reduce edge thickening the windows are applied in order of increasing size. The larger windows are only used to detect low contrast edges where no high contrast edges have been detected by smaller windows. Where an edge has already been detected at a smaller scale the window is reduced in size so as to exclude the edge. Values output by edge detection are ignored when either n_a or n_b are reduced by more than 75%. To detect edges at different orientations each window is applied horizontally and vertically; using long thin windows precludes the need for applying windows at diagonal orientations.

9.1.2.2 Edge thresholding

The edge images produced are thresholded to produce an edge map. An edge pixel is set when the estimated gradient \hat{d} exceeds its estimated standard deviation by a factor k . A value of 2.3 was chosen for k by comparing the performance of the normalised gradient operator with that of visual identification of edges [61]. There are a number of practical and theoretical problems with this approach to edge detection.

Setting thresholds empirically rather than using assumed clutter distributions does not explicitly involve attaching false alarm probabilities to detected edges. However, for the purpose of analysing the accuracy of the edges input to the segment growing stage, the false alarm rates implicit in these thresholds need to be determined for the range of likely clutter distributions. Assuming the gradient estimated over homogeneous regions is approximately Gaussian distributed, 2.3 standard deviations equates to a PFA of 2×10^{-2} . This is quite high compared to the PFA's suggested for edge detection in Chapter 7. Thus, we can expect most edges to be detected along with a large number of false alarms. Segment growing will then be limited by false alarms as well as true edges; later the segment merging stage is used to remove edges which are due to false alarms. The effect on segmentation of using a lower PFA (i.e., a larger value of k), or of reducing the PFA over successive iterations, has yet to be investigated.

9.1.2.3 Point target detection

SAR images contain point-like features which the edge detection stage fails to detect. To take account of this, an additional step tests for point-like features after edge detection has taken place, using operators similar to the normalised gradient operator over windows with dimensions 1×4 and 3×4 pixels. Rather than representing distinct point features, many of the features detected by these simple methods may arise purely from speckle. Such features, which hinder interpretation, are clearly apparent in the segmentations shown in Figure 9.1. It was shown in Chapter 8 that alternative methods of detecting such features

based on morphological and ratio operations are available. However, the performance of these alternative point feature detectors needs to be more fully assessed before they are incorporated into the segmentation algorithm. As point features constitute the smallest scale edges in the image, to minimise edge thickening they should be detected prior to all other edges, and the results input to the edge detection stage.

9.1.2.4 Alternative edge detectors

If the gradient is normalised by the sum of the average values within the two half windows $\hat{x}_a + \hat{x}_b$, as opposed to its estimated standard deviation, the resulting operator is equivalent to the normalised ratio operator discussed in Chapter 7. When the simple image model is assumed this operator gives the ML estimate of the ratio of the mean intensities in two homogeneous regions. However, in textured regions it is less robust than the normalised gradient operator (e.g., when intensity is K-distributed), and it cannot benefit from information from the previous segmentation as the normalised gradient operator can. The robustness of the algorithm has been tested by using the ratio operator as the first edge detection step, since in pure speckle this should give an improved initial estimate of the true edges. It was found that this had little effect on the final segmentation.

Figure 9.3 compares using the normalised gradient and normalised ratio operators for the initial stage of edge detection when segmenting the RSRE image shown in Figure 9.1 (a). Although there are more false alarms in the initial edge detection using the normalised gradient operator as opposed to the normalised ratio operator, there is little difference in the quality of the final segmentations. In fact the normalised gradient operator generates a lower final accuracy measure, this implies that it is the better segmentation; this appears to be borne out by the segmentations themselves. The normalised ratio operator does not appear to offer an obvious advantage in this case.

The benefits of using the normalised ratio operator for the initial stage of edge detection are more obvious when the image contains extremely bright point targets. The total standard deviation of the image is then too high to allow most edges to be detected on the first iteration and the majority of the image is interpreted as a single segment.

Figure 9.4 shows a single-look SAR image taken with the AIRSAR airborne system over an arid region. Various point targets (corner reflectors) have been deployed over the scene for use in polarimetric calibration. Figure 9.5 compares using the normalised gradient and normalised ratio operators for the initial stage of edge detection when segmenting this image. The normalised gradient operator produces a poorer initial segmentation than the normalised ratio operator. However, surprisingly the final segmentation produced from the latter appears to be only slightly better than that produced from the former. The benefits of using the normalised ratio operator for the initial stage of edge detection do not seem to carry through to later iterations.

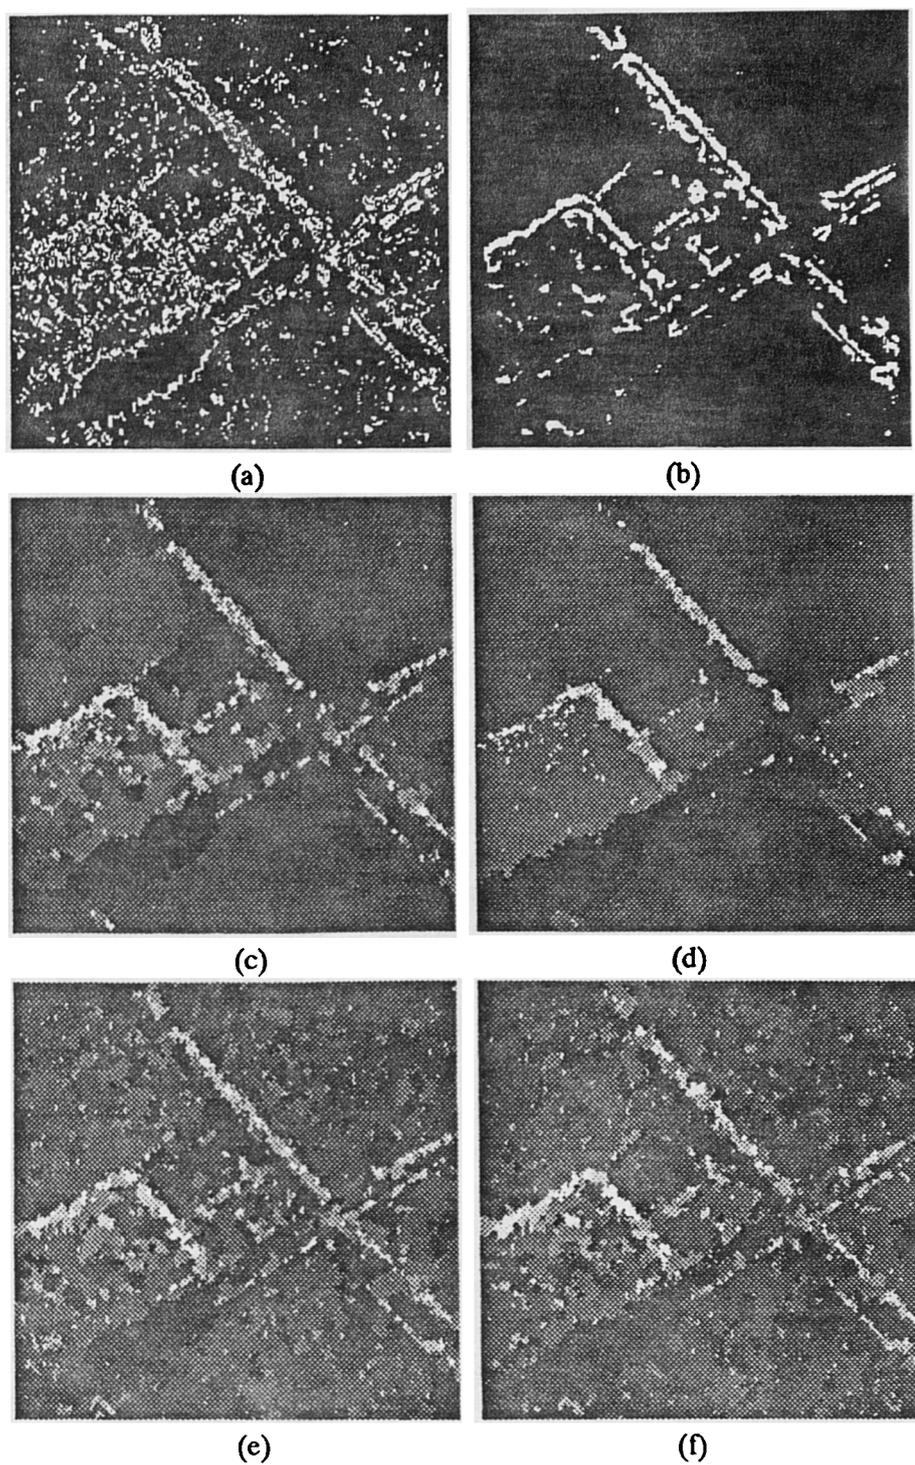


Figure 9.3: Comparison of using the normalised gradient and normalised ratio operators for the initial stage of edge detection; (a) and (b) the edge maps produced using the normalised gradient and normalised ratio operators, (c) and (d) initial segmentations, and (e) and (f) final segmentations with accuracy measures 0.4925 and 0.4955.

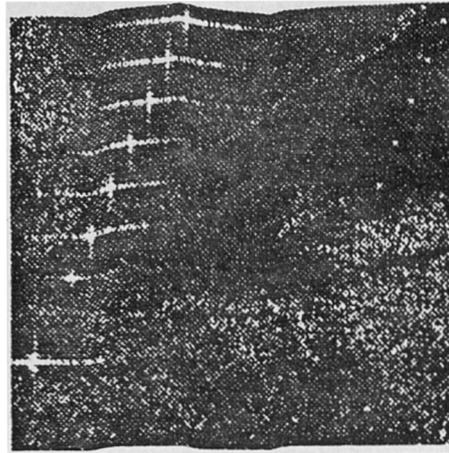


Figure 9.4: AIRSAR HH polarised C-band image containing bright point targets.

9.1.2.5 Comparison of segmentations

The visual comparison of different segmentations afforded by, for instance Figure 9.3, is purely subjective and more quantitative measures are needed. A global measure is of course given by the final accuracy measure. To test local accuracy the statistics within segments and between adjacent segments need to be analysed to check that segments are indeed homogeneous and that there are real differences between adjacent segments. The edge removal algorithm discussed in Section 9.1.3 is one possible method of merging segments which do not differ significantly. Segment homogeneity can be checked by dividing the original image by the segmented image, if the segmentation is correct there should be no structure in the resulting image.

Figure 9.6 shows the result of dividing the segmentations shown in Figures 9.3 (e) and (f) by the original amplitude image. Both images are dominated by speckle and have no significant structure. In fact, there is slightly more small scale structure in Figure 9.6 (b) than in Figure 9.6 (a), but this is lost in reproduction. It is possible to resegment these ratio images to extract any remaining structure, the result of doing this is shown in Figure 9.7. Whilst some features are detected there is little overall structure in the resegmentations which mainly consist of a single segment.

9.1.2.6 Intensity segmentation

Although the empirical approach to threshold selection does not impose any particular noise (i.e., clutter) model on the image, better results are produced when distributions are less asymmetric, i.e., using amplitude (Rayleigh speckle) or multi-look intensity (gamma distributed speckle) data as opposed to single-look intensity (exponential speckle) data. The segmentations shown so far have all used single-look amplitude as opposed to intensity data, Figure 9.8 shows the

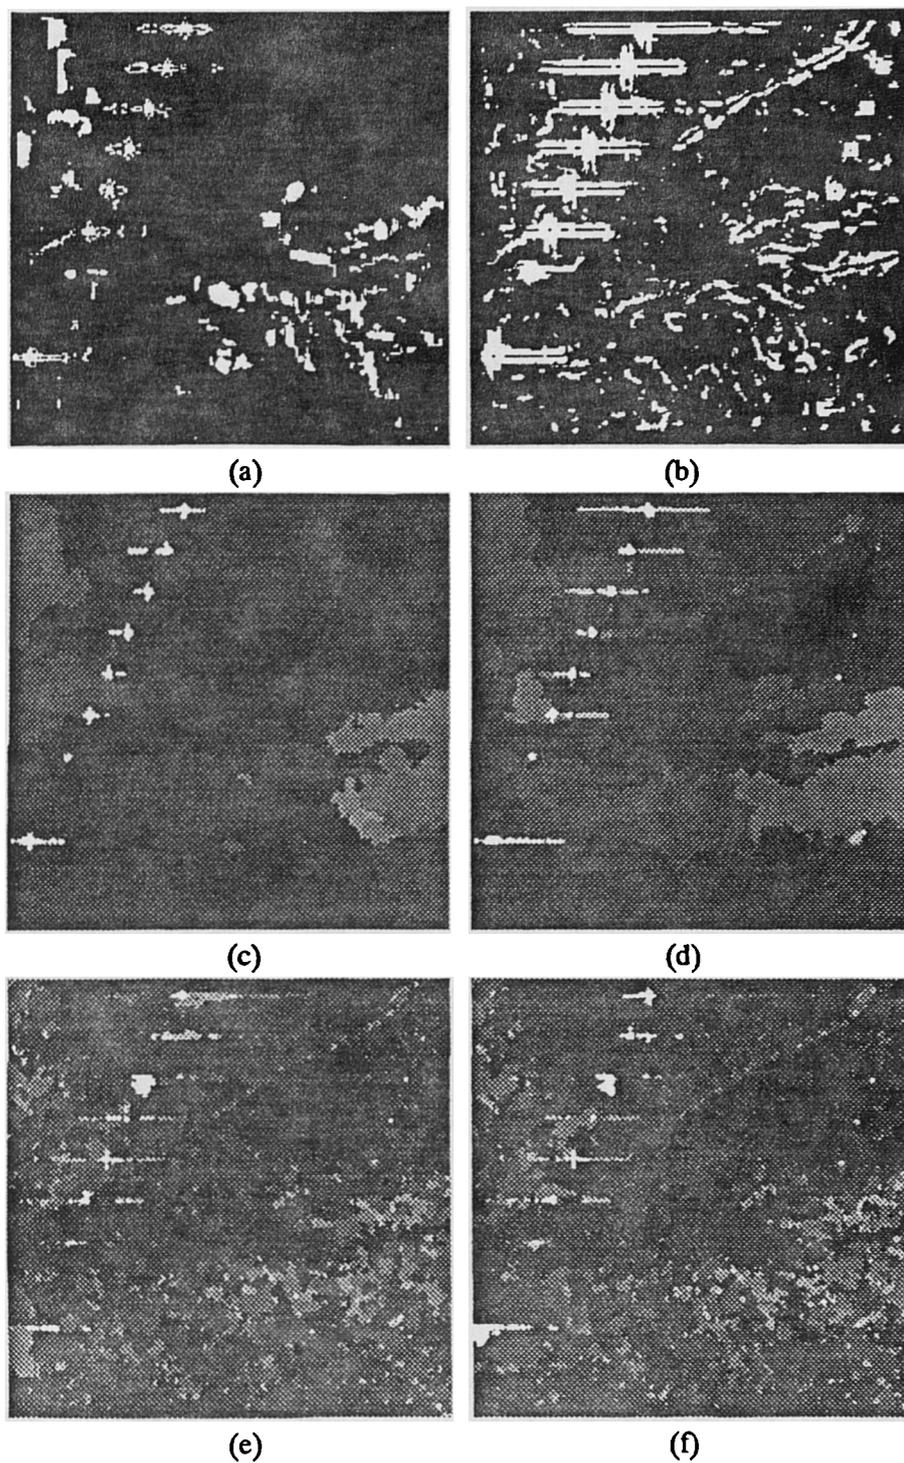


Figure 9.5: Comparison of using the normalised gradient and normalised ratio operators for the initial stage of edge detection when segmenting the image shown in Figure 9.4; (a) and (b) the edge maps produced using the normalised gradient and normalised ratio operators, (c) and (d) initial segmentations, and (e) and (f) final segmentations with accuracy measures 0.5765 and 0.5935.

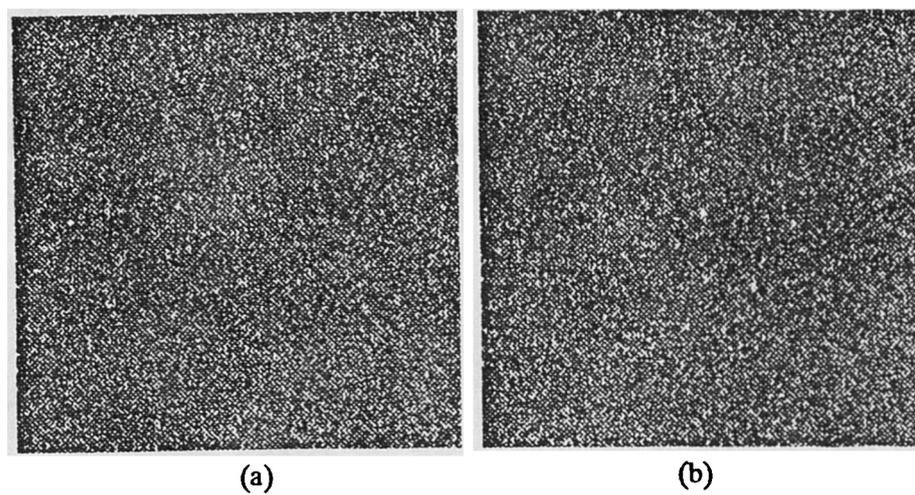


Figure 9.6: The result of dividing the segmentations shown in Figures 9.3 (e) and (f) by the original amplitude image.

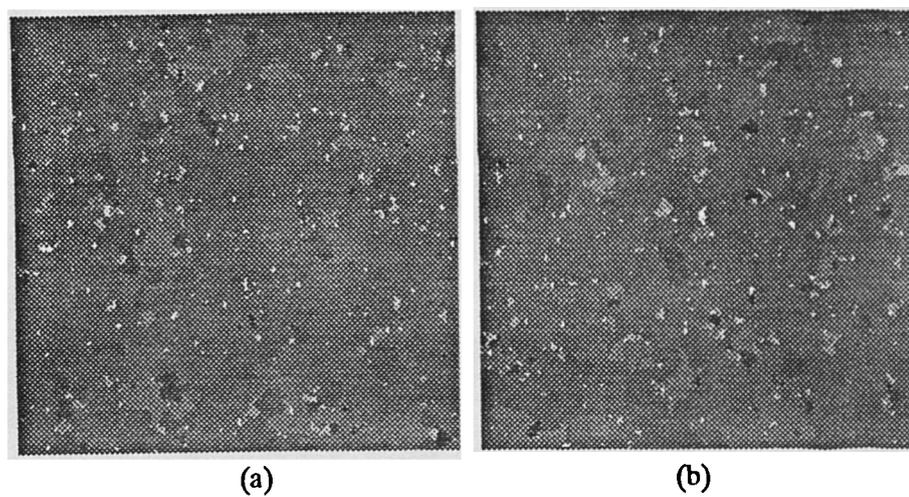


Figure 9.7: The result of resegmenting the ratio images shown in Figure 9.6.

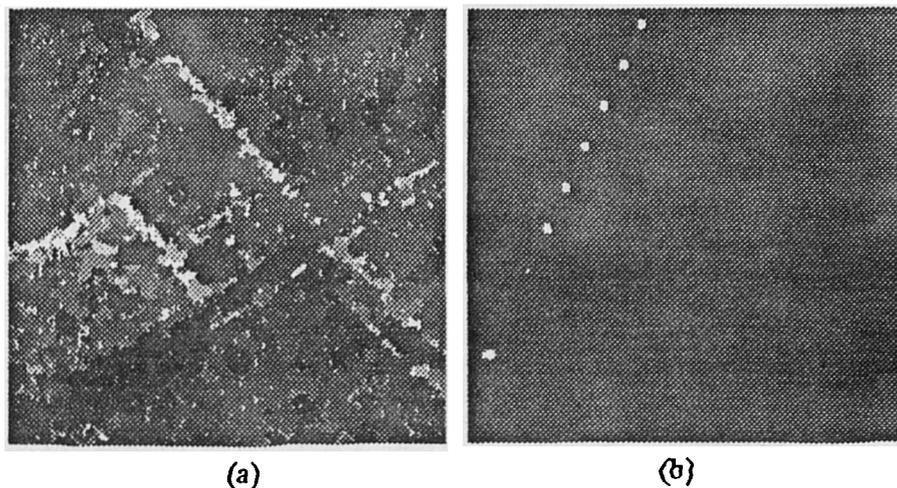


Figure 9.8: The result of intensity segmentation of the images shown in (a) Figure 9.1(a) and (b) Figure 9.4. The final accuracy measure of the former segmentation is 0.9579 whilst that of the latter is 5.451.

result of segmenting the images shown in Figures 9.1 (a) and 9.4 in intensity as opposed to amplitude. The results should be compared with those of amplitude segmentation shown in Figures 9.1 (b) and 9.5 (e) respectively. A passable segmentation is generated of the RSRE image, though linear features (hedges, rows of trees ?) are not as well detected as in the amplitude segmentation. With the AIRSAR image the segmentation algorithm only manages to detect the brightest point targets; there is no improvement when the algorithm is allowed to iterate beyond the first minimum.

9.1.3 Segment growing and merging

9.1.3.1 Segment growing

The edge map output by the edge detection stage is used to limit the growth of segments representing homogeneous regions. Segments are grown by fitting discs inside regions where no edges are present. Discs 64, 32, 16, 8, 4, 2 and 1 pixel in diameter are fitted in order of decreasing size. All the discs of a given size which overlap/abut are merged to form a single segment. Where a segment defined in terms of discs all of the same size overlaps/abuts segments defined in terms of larger discs, the segment is first reduced in size so that it is disjoint with all the surrounding segments, before being merged with the segment with the closest mean value. To aid this, look-up tables of the number and mean value of pixels within each segment are kept.

The algorithm segments extended targets into homogeneous regions. However in some cases over-segmentation appears to occur. Extended targets (e.g., fields) that we would like to be detected as single segments for certain applications, and that appear visually to be so, are split into several segments. To solve

this problem the algorithm has a final stage of segment merging. As the problem of over-segmentation is more evident in the edge map of the segmentation than in the mean image, segmentations displayed in this and following sections will show the edges separating segments as well as the mean value within them. The edges define the underlying structure of segmentations as compared to the purely visual appearance provided by the mean values which is potentially misleading. An alternative to carrying out merging after the last iteration is to apply it after each iteration step, this approach will not be considered any further.

9.1.3.2 Merging based on moments

The original algorithm merged segments whose estimated first, second and third moments differed by less than a given number k of standard errors. Where there was more than one candidate for merging a given segment, the segment with the closest mean value was chosen. A problem encountered with this approach was determining the optimum value of k to use. When different images were segmented with a range of values of k the best results were not obtained with the same values of k . The optimum value of k appeared to relate to the average number of independent samples per segment, which equals the average field size divided by the system resolution; larger segments require more merging, i.e., a larger value of k . It was only recently discovered that this problem was caused by the estimated standard errors of moments not being correctly normalised by the number of independent samples within a segment.

Figure 9.9 shows the result of using the corrected algorithm to segment one of the ERS-1 SLC (16/04/92) images to be discussed in Section 9.2, with merging being carried out after the last iteration only. The method is effective at merging adjacent segments with similar statistics. From these initial results it is clear that a full comparison with the methods described below is needed, along with an analysis of the effect of carrying out merging after each iteration.

It should be noted that merging based on moments has been used as the basis for an alternative SAR segmentation algorithm to the one described here [8]. The SAR image is broken up into a grid of small segments (4×4). Adjacent segments with similar moments are then merged. Merging continues until the moments of all adjacent segments differ by a given degree. The segmentation performance of this algorithm has yet to be compared with the algorithm currently under study.

9.1.3.3 Merging based on edge strength

Because of the problems encountered with segment merging based on moments most work on segment merging has been carried out using an alternative approach. An edge strength is attached to each segment boundary in the final segmentation; edges which fail to pass a given threshold are then deleted.

Given two adjacent segments i and j , a strength can be attached to the edges separating them by taking the normalised difference of their average values

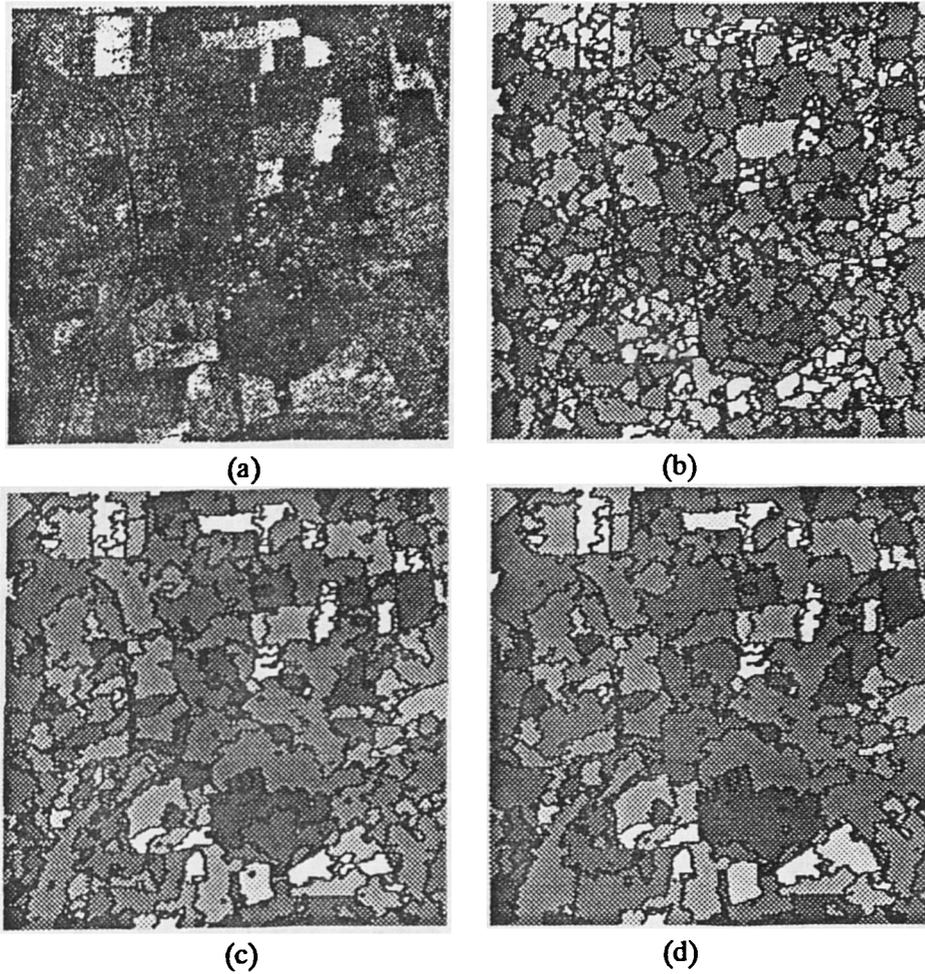


Figure 9.9: (a) ERS-1 image and its segmentation with (b) no merging, merging with (c) $k = 3.29$, and (d) $k = 4.42$. If it is assumed that the estimated moments are Gaussian distributed these thresholds correspond to false alarm rates of 10^{-3} and 10^{-5} respectively.

defined by

$$|\widehat{d}_0| = \frac{|\widehat{\mu}_i - \widehat{\mu}_j|}{\sqrt{\widehat{\sigma}_i^2 + \widehat{\sigma}_j^2}} \quad (9.7)$$

or the normalised ratio of their average values defined by

$$\widehat{r}_0 = \min \left(\frac{\widehat{\mu}_i}{\widehat{\mu}_j}, \frac{\widehat{\mu}_j}{\widehat{\mu}_i} \right) \quad (9.8)$$

However, such measures do not take into account the size of segments; large segments are more easily distinguished than small segments with the same intensity ratio or difference. Segment size can be accounted for by calculating the probability of the edge strength arising purely from clutter, given the size of the segments. Such a false alarm probability can be determined for the intensity ratio following the approach used to set false alarm rates for the normalised ratio operator in Chapter 7. It is assumed that each average intensity is gamma distributed with order parameter equalling the number of independent samples; assuming the simple image model applies this is given by the number of looks times the size of the segment. Correlation induced by the imaging system can be compensated for by reducing the order parameter.

Simply deleting all edges which lie below (above) a given threshold will lead to some segment boundaries losing closure; the mean value within a segment is then ill-defined. This problem is circumvented by carrying out thresholding dynamically. First the strength of every edge in the image is calculated. The weakest edge is then deleted if it lies below (above) the threshold, and the two segments it separates are merged. As the statistics of this merged segment will differ from those of the two segments it is comprised of, the strength of each of the remaining edges which bound it have to be recalculated. When this is complete the weakest edge is once again searched for and the whole process is repeated. This continues until there are no edges falling below (above) the threshold.

This dynamic thresholding process was implemented using two doubly linked lists to represent the topology of the segmentation, one of which lists all the segments and the other lists all the edges. Each element in the segment list has records containing the statistics needed to calculate the edge strength (e.g., its size and mean value), and a list of pointers to the edges in the edge list which bound it. Each element in the edge list contains a record of its strength and pointers to the two segments in the segment list it separates. When the weakest edge is found it is removed from the edge list. The records of one or other of the segments it points to in the segment list are then updated so as to represent the merged segment, and the other segment is deleted. Updating involves recalculating the statistics of the merged segment from those of the two segments it is comprised of and merging the edge lists of the two segments so that the merged segment only shares one edge with each neighbouring segment; any additional edges are deleted. In addition, to keep a tally of the new segmentation, pixels in the underlying segmentation representing the deleted segment need to be assigned to the merged segment. An extra record defining the bounding

box of each segment reduces the processing needed to achieve this. Finally, the strength of the edges bounding the merged segment are recalculated and the weakest remaining edge searched for.

Figure 9.10 shows the result of applying the different approaches to edge removal to the segmentation of the ERS-1 image shown in Figure 9.9 (b). The result of dynamic thresholding using the edge ratio and difference measures are very similar (a)-(d). Both methods remove weak edges irrespective of the size of the segments they separate. Strong edges separate all remaining segments, most of which are either very large or very small. A much better result is produced when the size of segments is taken into account by equating the edge ratio with a false alarm rate (e); the minimum strength of edges separating segments is then inversely related to their size and most point-like features have been deleted. It is clear from (f) that while global thresholding will identify the most distinctive edges it is not of use for segment merging. Because of the more even merging performance offered by equating the edge ratio with a PFA and the fact that this allows us to attach a quantitative measure of edge credibility to merged segments (assuming the simple image model applies), this measure will be adopted for dynamic thresholding of edges in the next Section.

A problem related to that of segment merging is merging edge pixels into segments. Presently the algorithm attaches edge pixels to the adjoining segment with the closest mean value. This can produce unsatisfactory results when detected edges are more than a single pixel wide. This problem would be removed if edges were represented by boundaries between pixels, i.e., on a grid displaced by half a pixel in both directions relative to the pixel grid.

Finally, it should be mentioned that the segmentation algorithm performs poorly at detecting linear features as single entities; in fact it has no knowledge of such features. Methods for detecting such features need to be incorporated into the algorithm.

9.2 Segmentation of ERS-1 data

In this section the results of using the segmentation algorithm for agricultural change detection in ERS-1 data is described. The test site is a 5km by 5km area surrounding the village of Feltwell in East Anglia, UK. This area is topographically flat.

9.2.1 The ERS-1 dataset

Both slant range single-look complex (SLC) images, and precision ground range (PRI) 3-look intensity images taken during the growing season of 1992 were used. Because the pixel sampling rate in the SLC images was approximately 4.5 times greater in azimuth (3.9 m) than in ground range (17 m) blocks of 5 pixels in azimuth were averaged in intensity to rectify the geometry of the image. The pixel sampling rate is 12.5 m in both azimuth and ground range in the PRI images. The sub-images covering the test site extracted from the SLC image have dimensions 1280 pixels in azimuth by 256 pixels in slant range (256 × 256

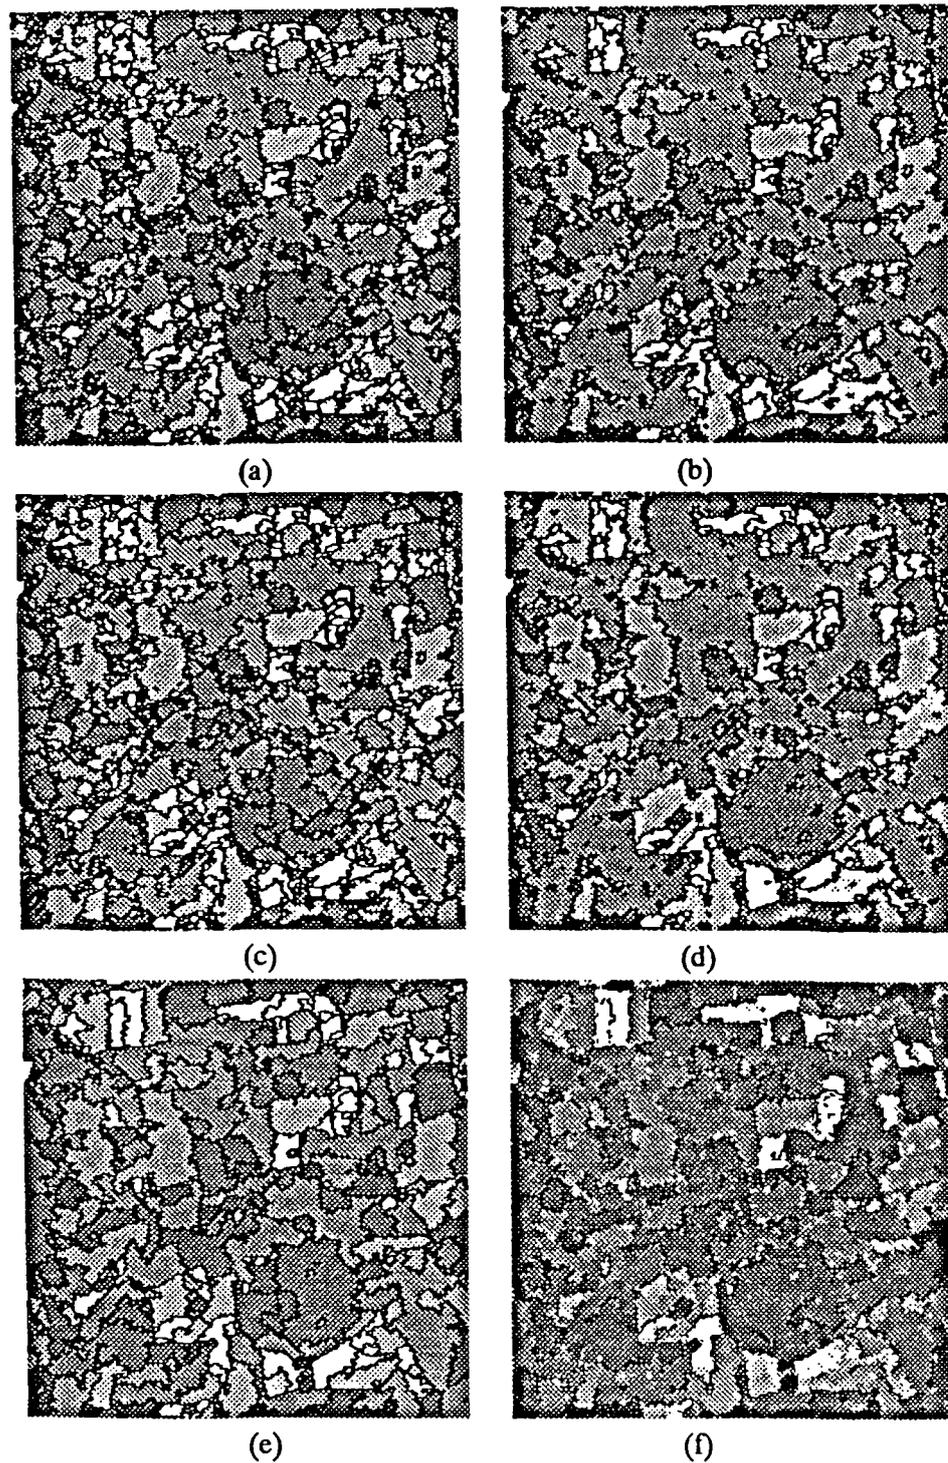


Figure 9.10: Different approaches to edge removal applied to the segmentation of the ERS-1 image shown in Figure 9.9 (b), (a) and (b) dynamically thresholding the normalised intensity ratio above 0.8 and 0.6, (c) and (d) dynamically thresholding the normalised intensity difference below 0.1 and 0.2, (e) dynamically thresholding the normalised intensity ratio above a false alarm rate of 10^{-5} , and (f) globally thresholding the normalised intensity ratio above a false alarm rate of 10^{-5} . In (f) the disconnected edges are overlaid on the original segmentation.

after averaging), whilst the PRI sub-images have dimension 400 pixels in both azimuth and ground range.

Figure 9.11 shows an example SLC image of the test site taken on 16/04/92, before and after averaging, and the corresponding PRI image. Figure 9.12 shows histograms of pixel intensity values sampled from a single homogeneous field in these three images. As we would expect, the intensity of the single-look SLC data fits an exponential distribution and the intensity of the 3-look PRI data fits a gamma distribution with order parameter 3. The intensity of the averaged SLC data also fits a gamma distribution with order parameter 3, this is because the 5 pixel average only covers 3 resolution cells in azimuth.

Figure 9.13 shows the intensity ACF's measured within the homogeneous field in the SLC, averaged SLC, and PRI images. Adjacent pixels in the SLC data are highly correlated in azimuth (correlation coefficient 0.5) but only slightly correlated in range (correlation coefficient 0.2); as averaging is carried out over discrete blocks of five pixels in azimuth, the azimuth correlation is reduced but the range correlation is unaffected. The measured resolution of SLC data is 9.8 m in slant range and 5.6 m in azimuth [35, 63]. Adjacent pixels in the PRI data are highly correlated in both azimuth and range (correlation coefficient > 0.5). This is to be expected given that the pixel sampling rate of 12.5 m in both range and azimuth is much less than the measured resolution of PRI data (25 m in ground range (mid-swath) and 22 m in azimuth [35]).

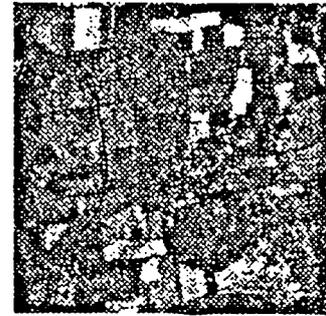
The series of six ERS-1 SLC and PRI images discussed in this Section are shown in Figures 9.14 and 9.15. The SLC images shown on the left and right hand sides of Figure 9.14 have different geometries, because the images shown on the left are all taken on descending passes, whilst those on the right are all taken on ascending passes. ERS-1 has a Sun-synchronous orbit with an inclination angle of 98.5 degrees and it crosses the equator at 10:30 am local time on a descending pass. Thus, on a descending pass over Feltwell which lies at 52.5 degrees north, ERS-1 will be heading 13 degrees west of south at around 10:45 am local time, while on an ascending pass it will be heading 13 degrees west of north at around 22:00 pm local time. It should also be noted that SLC sub-images extracted from different parts of the swath will have different pixel spacings in ground range. The images taken on descending passes are from far-range while those taken on ascending passes are from near-range.

Due to problems in obtaining data, only four of the PRI images (16/04/92, 12/05/92, 09/06/92 and 18/08/92) correspond to the same pass as one of the six SLC images. Two additional PRI images were also obtained, both taken on ascending passes on 26/04/92 and 22/09/92. The images were processed by the German, UK and central PAF's where different gains were applied; images processed at different PAF's have to be scaled differently to have the same mean value for display purposes.

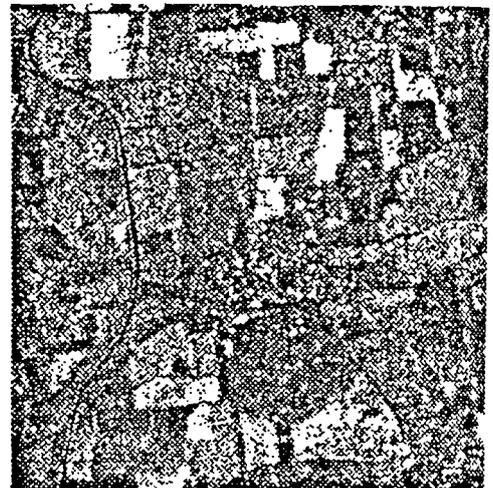
The area around Feltwell is intensively farmed and the regular pattern of agricultural fields is readily visible in the images, as is the man-made drainage channel (running from north to south on the left of the image) and the village of Feltwell itself (the collection of small bright points in the centre of the image). We know from workers involved in ground data collection that major



(a)

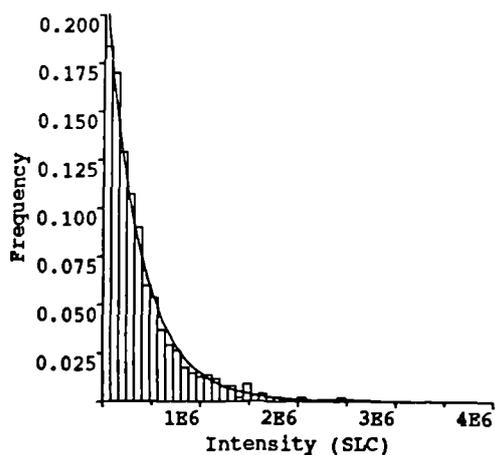


(b)

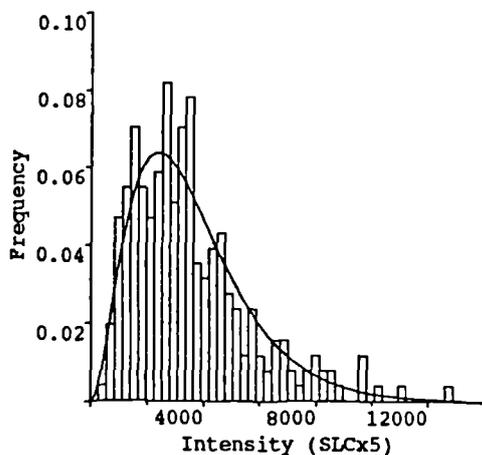


(c)

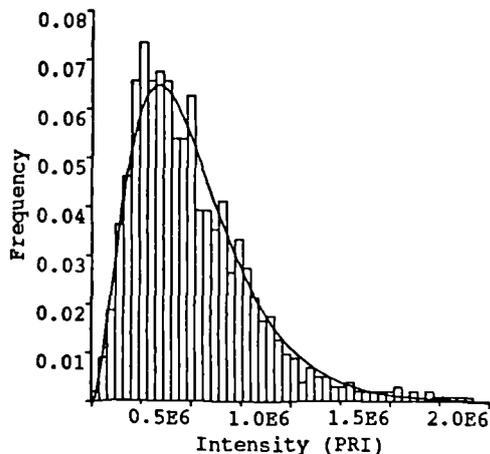
Figure 9.11: ERS-1 image of the Feltwell test site taken on 16/04/92, (a) SLC image, (b) SLC image averaged by 5 pixels in azimuth, and (c) PRI image. The images are shown in amplitude and are sized relative to their actual dimensions in pixels.



(a)



(b)



(c)

Figure 9.12: Histogram of intensity values within a homogeneous field in the (a) SLC image, (b) averaged SLC image, and (c) PRI image. The solid curves represent fitted theoretical distributions (a) exponential, and (b) and (c) gamma with order parameter three.

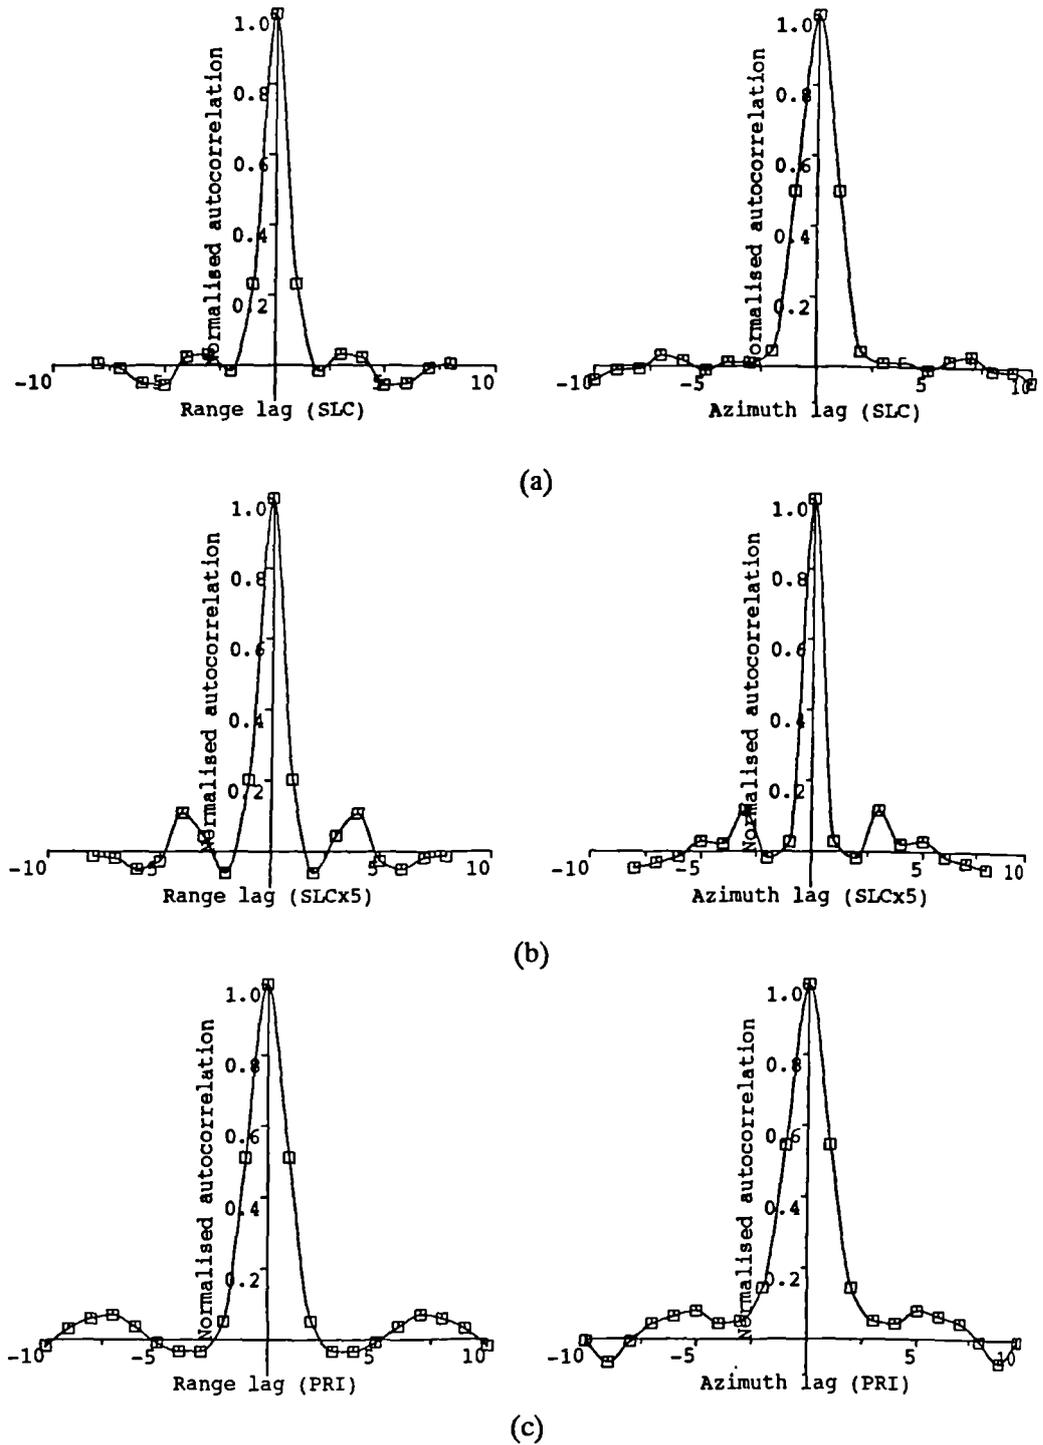


Figure 9.13: Intensity ACF measured within a homogeneous field in the (a) SLC image, (b) averaged SLC image, and (c) PRI image.

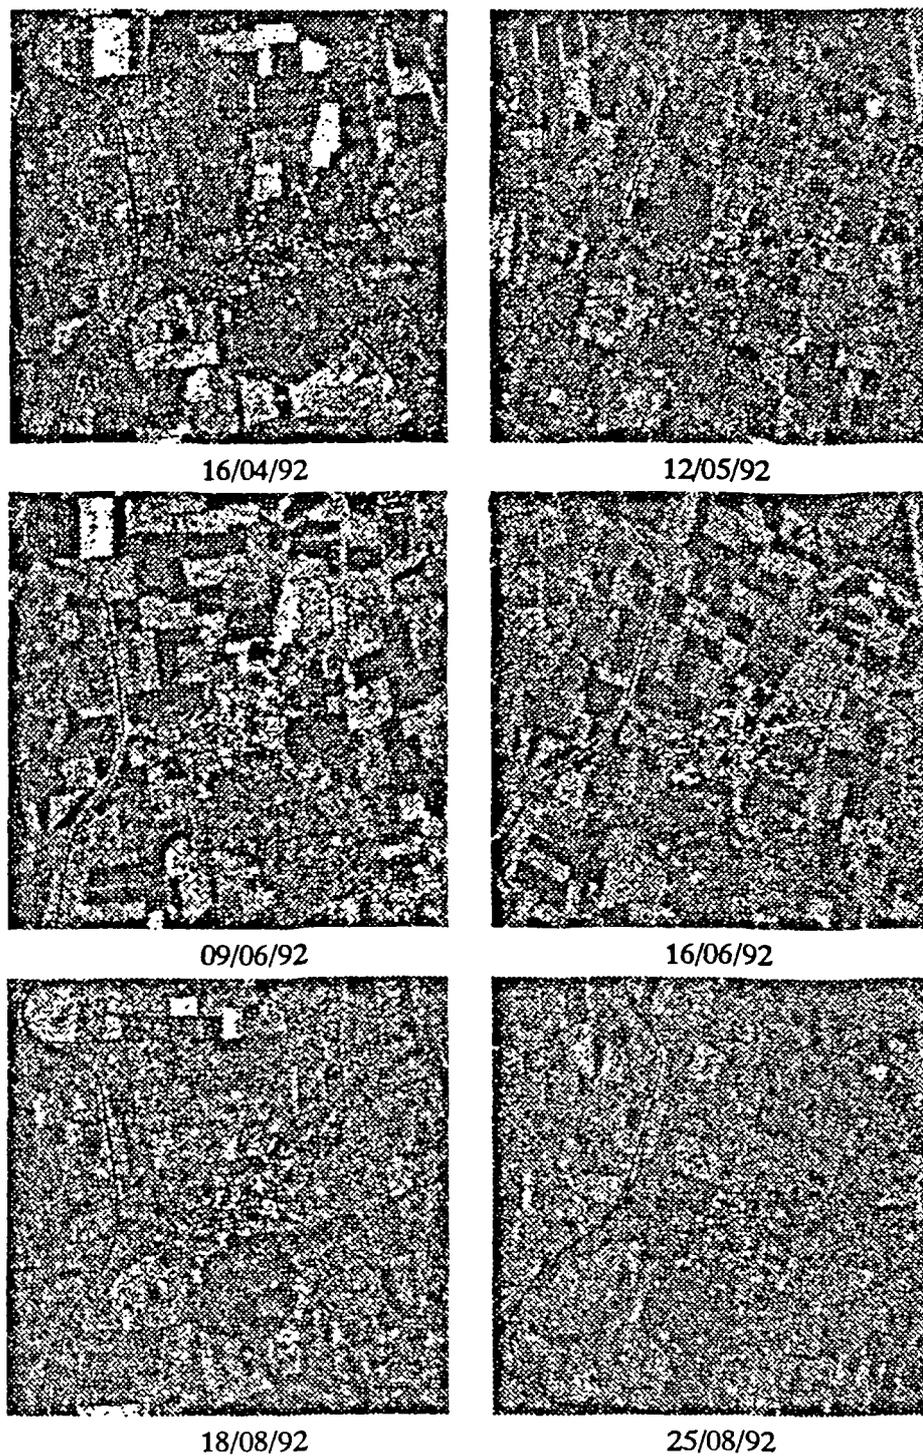


Figure 9.14: ERS-1 SLC images of Feltwell taken during the 1992 growing season. The left hand images are taken on a descending pass (flight ↓, illumination ←), and the right hand images are taken on an ascending pass (flight ↑, illumination →). The tops of the images are approximately aligned with the northerly direction.

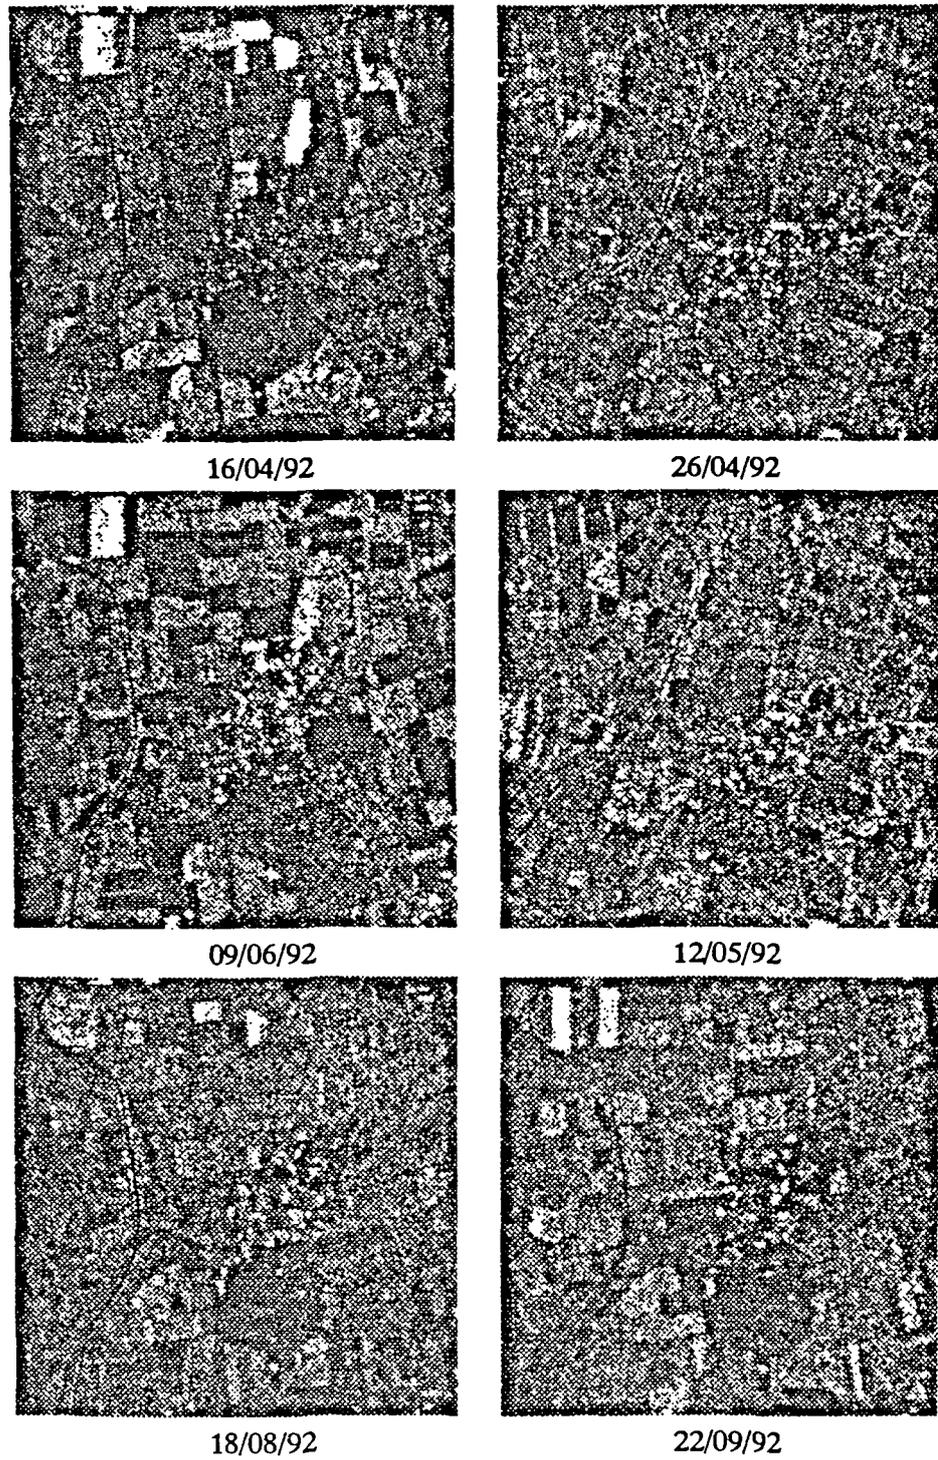


Figure 9.15: ERS-1 PRI images of Feltwell taken during the 1992 growing season.

structural change did not occur on the ground between the dates of the images (though some changes are to be expected; for example, a few fields may be split or merged as a result of farming practices). Nonetheless, *image* structure shows a remarkable amount of variation, as a result of the variations in the backscattering from land units. For example, the well defined field structures to the immediate north-west of the village seen in the images taken in June are almost completely absent in the earlier and later images, and the bright fields at the top left hand corner of the images taken on 16/04/92 and 09/06/92 are not as distinctive on other dates. The only static feature throughout the series of images is the drainage channel. If this were blocked out one could easily believe that the images are of different scenes. We are not in a position to fully discuss the physical causes of the differences observed in these images, as the ground data has not yet been fully collated and made available. Rather than make hypotheses, we will only discuss the segmentation results in what follows.

9.2.2 Segmentation of SLC data

Figure 9.16 shows the output from the segmentation process when applied to the SLC data with no segment merging. Visually, and in the context of detecting changes in field structures, the images are over-segmented. However, a large amount of this over-segmentation occurs in regions of the image which are too dark or too bright for the human eye to distinguish intensity differences and may represent real differences within fields. Close visual inspection has failed to reveal any segments which do not correspond to homogeneous regions in the original images.

Various weaknesses of the segmentation algorithm are apparent. It has failed to recognise the drainage channel as a single feature; the detection of long thin features is a generic problem in SAR segmentation, which requires attention. The expected regular shapes of the boundaries of agricultural fields are not accurately reproduced, since the segmentation faithfully follows the underlying data in which the field edges are, on the whole, by no means regular. In the context of an agricultural application, it may also be considered a weakness that the segmentation preserves many of the point-like features from the original images.

Some degree of segment merging is clearly required before field structures can be compared. Figures 9.17 and 9.18 show the result of dynamically removing all edges whose edge ratio corresponds to a PFA greater than 10^{-3} and 10^{-5} respectively. It has been assumed that pixels are independent when calculating the PFA corresponding to each edge ratio.

After merging with thresholds corresponding to a PFA of 10^{-5} the majority of fields are recognizable as single segments, and most segments are identifiable with distinct homogeneous regions in the original image. However, very bright and very dark regions of the image are still broken up into several segments even though the eye would tell us otherwise (intensity differences are in fact apparent when these regions are scaled). Most of the small segments have been removed, in particular those representing point-like features. Setting thresholds

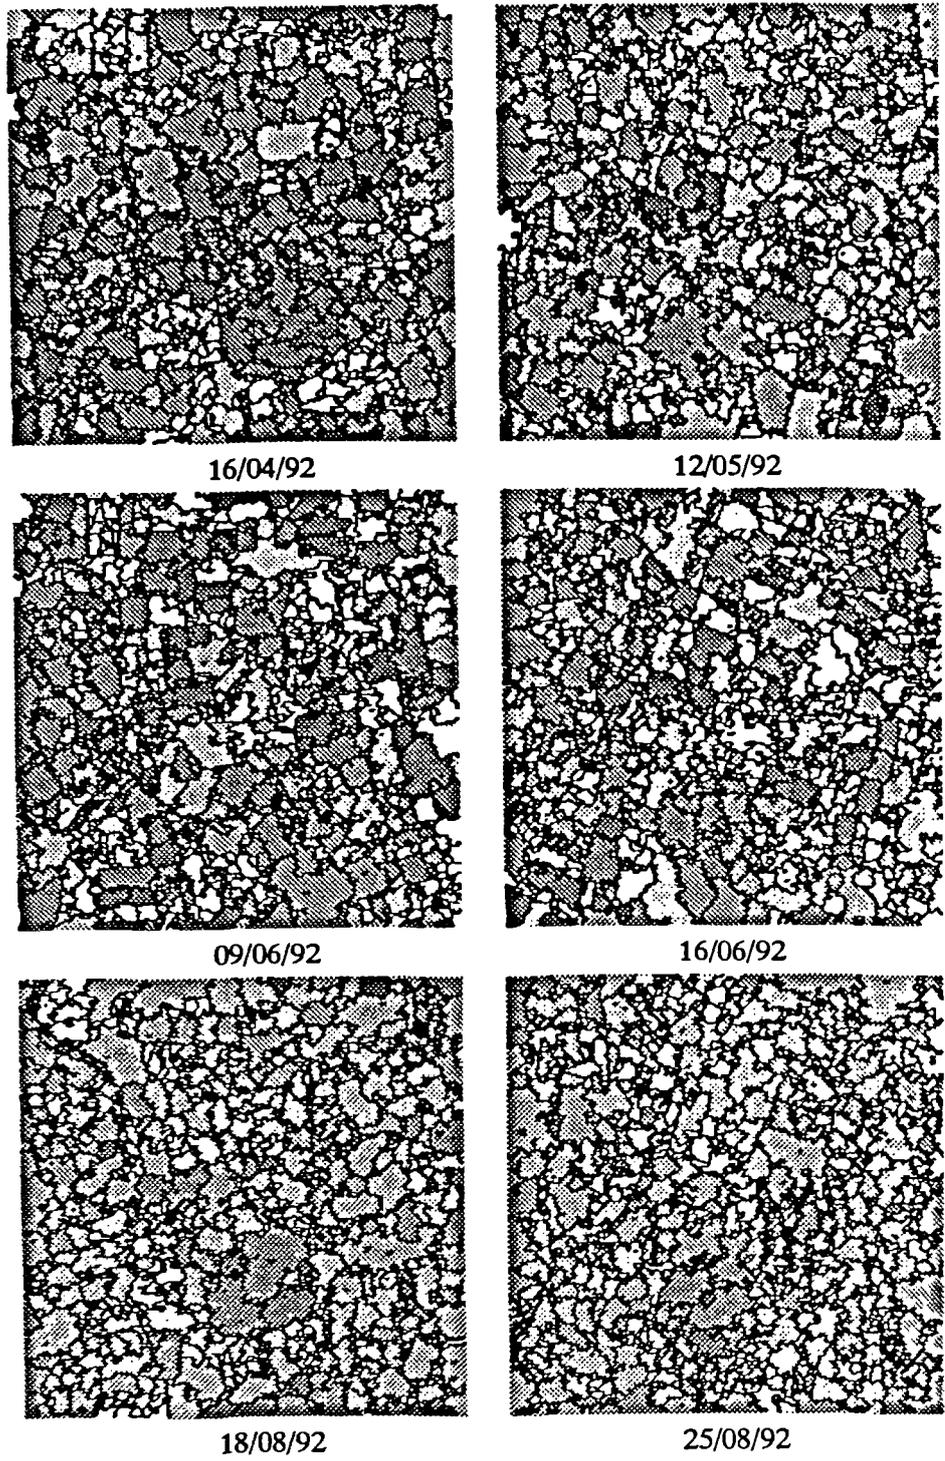


Figure 9.16: Segmentation of the SLC images shown in Figure 9.14.

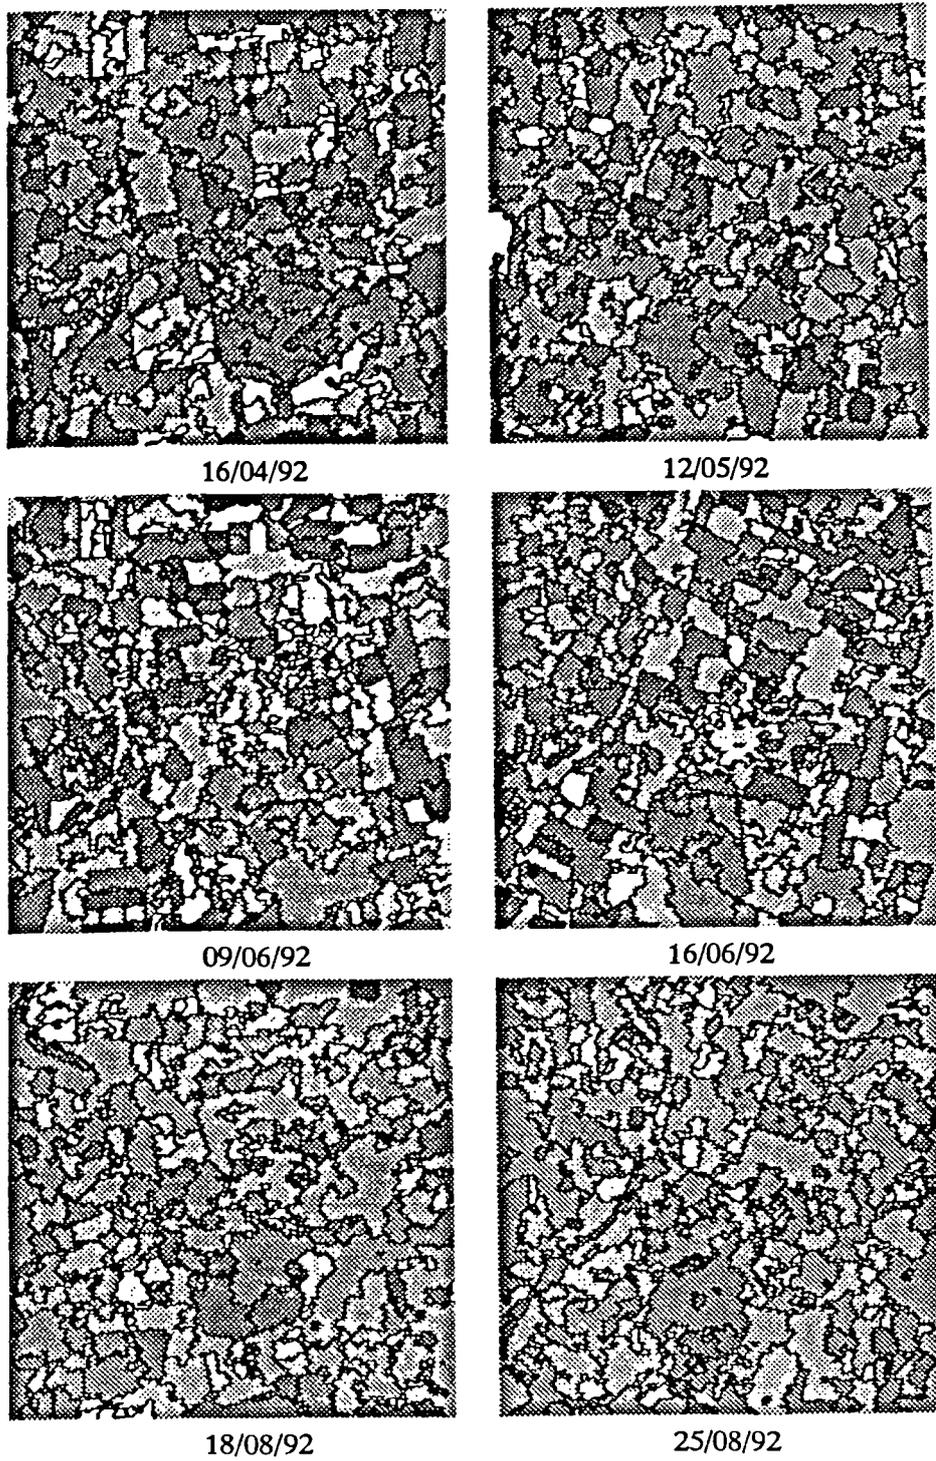


Figure 9.17: Result of edge removal using a PFA of 10^{-3} .

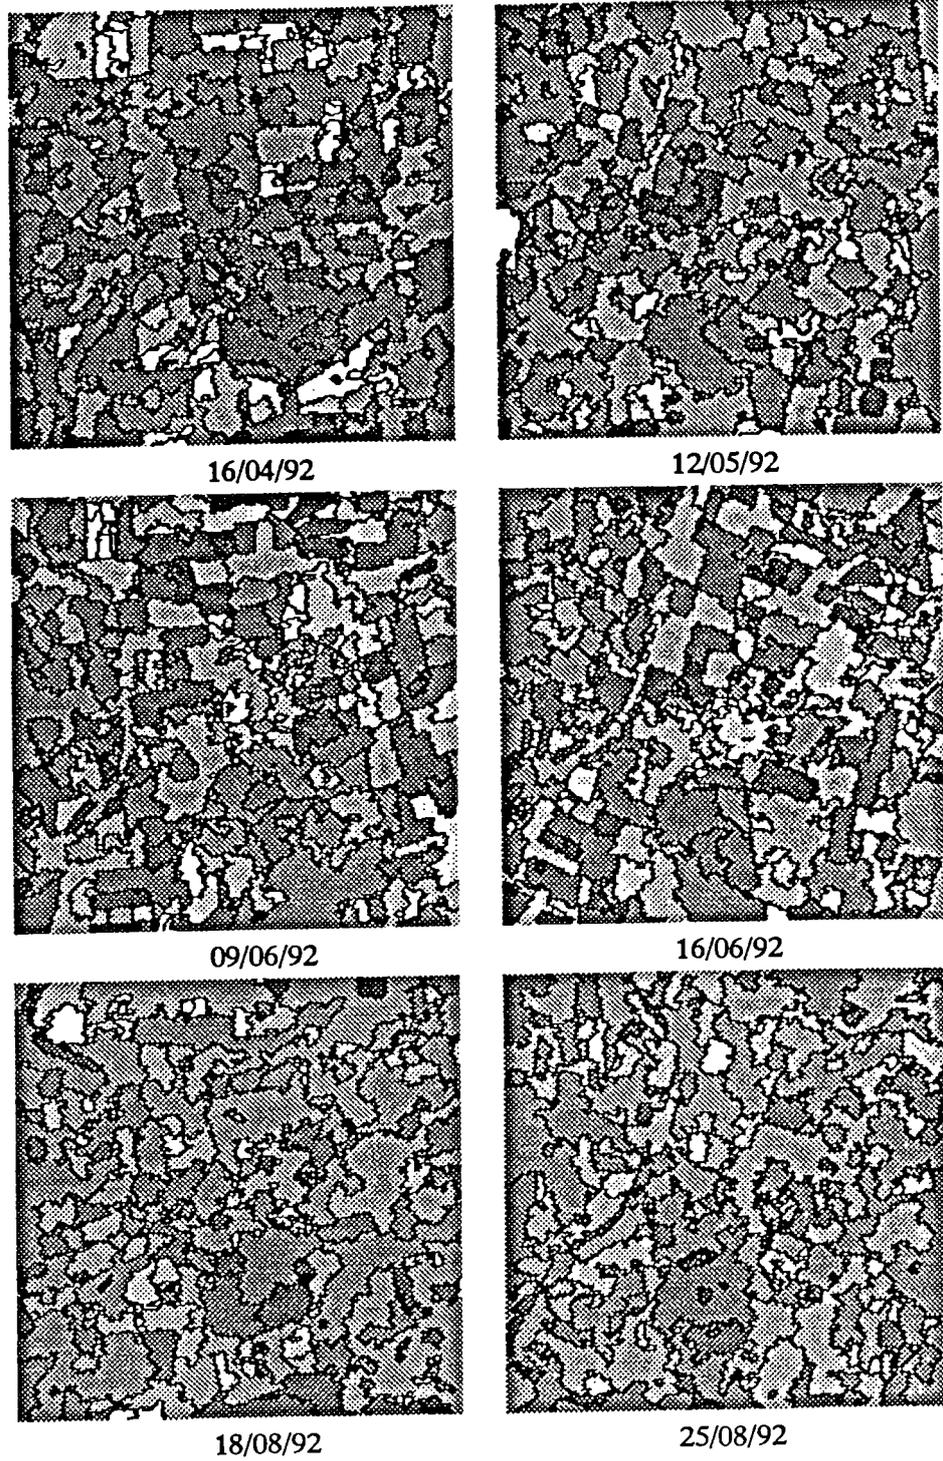


Figure 9.18: Result of edge removal using a PFA of 10^{-5} .

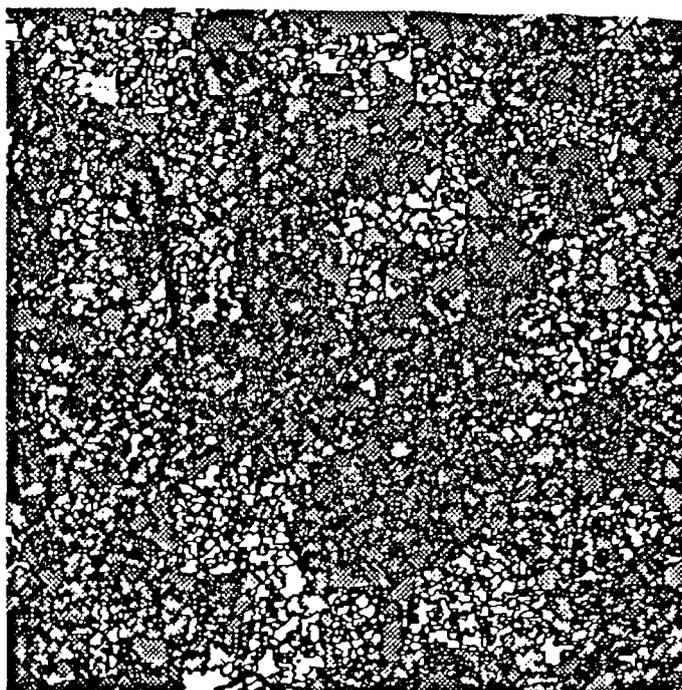


Figure 9.19: Segmentation of the PRI image of Feltwell taken on 16/04/92 and shown in Figure 9.15.

for a PFA of 10^{-3} still produces a significant amount of merging but a large number of extraneous edges and small segments still remain.

After segmentation and merging, the structural differences between the different SLC images are even more apparent than in the original unsegmented images. This is partially because the drainage channel, which gave a visual clue linking all the original images, is not preserved in their segmentations. The strongest similarities are seen between the segmentations of the images taken only a week apart on 09/06/92 and 16/06/92, where a number of field structures can be matched, e.g., the rotated 'J' shaped field slightly left of centre. Field structures in the 09/06/92 segmentation can also be matched to field structures in the 16/04/92 segmentation, e.g., the bright field top left. However, it is difficult to find any matches with the 12/05/92 segmentation which lies between the former dates. The segmentations of the 18/08/92 and 25/08/92 images reveal no clearly recognizable common structure even though the images were only taken a week apart, nor do they compare with any of the other segmentations.

9.2.3 Segmentation of PRI data

When the segmentation algorithm was applied to the PRI images they were segmented down to extremely small segments. Figure 9.19 shows the result using the 16/04/92 image. In addition, segment merging failed to remove a large number of edges. It was realised that this behaviour was caused by the much larger

degree of correlation between pixels in the PRI than in the SLC images (and for that matter the RSRE, AIRSAR and ERS-1 FDP images used previously in segmentation trials). Correlation reduces the number of independent samples available for estimating mean values and in turn this increases the variance of the estimated gradient. This results in edge detection thresholds being passed more often than expected and the generation of excessive false alarms. These false alarms limit the size, and increase the number, of segments that are produced in the segment growing stage. The images used previously were, of course, spatially correlated but not to a large enough extent to draw attention to this problem and the effect of correlation had generally been ignored. Correlation also affected the segment merging stage by reducing the number of independent samples within each segment. This broadened the false alarm distribution of the edge ratio. Thus, fewer edges failed to pass the false alarm threshold.

To overcome this problem the PRI images were sub-sampled prior to segmentation, the results of which are shown in Figure 9.20. The information loss due to sub-sampling causes the segmentation to look less detailed than those produced from the SLC data. As with the SLC data there are considerable differences in image structure between images taken on different dates. In fact, the six PRI images and their segmentations can be split into three distinct pairs in terms of similarity, (a) 16/04/92 and 09/06/92, (b) 26/04/92 and 12/05/92 and (c) 18/08/92 and 22/09/92. It is interesting to note that the images in each pair are either both taken in the morning on a descending pass or in the evening on an ascending pass. This may be due to diurnal changes in surface moisture content. It should also be noted that both images in pair (b) were taken between those in pair (a). A surprising result of the high degree of correlation in the PRI images is that the drainage channel which had not been preserved in the SLC segmentation, is preserved in the segmentation of the correlated PRI image but is no longer preserved after sub-sampling; this has yet to be fully explained.

The SLC and PRI images taken on the same pass are generated from the same raw data and thus we would expect similar structural features to be found in segmentations of both images. When the segmentations of the SLC and PRI images taken on 16/04/92, 12/05/92, 09/06/92 and 18/08/92 are compared, it is found that whilst the same general features are detected segment boundaries do not match. Of course, some degree of dissimilarity is to be expected given that different weighting functions have been used to process the raw data to produce the SLC and PRI images.

Merging was applied to the segmentations of the PRI images shown in Figure 9.20 to see whether the dissimilarities between SLC and PRI segmentations were reduced when weak edges were removed. Figures 9.21 and 9.22 show the result of merging with thresholds corresponding to a PFA of 10^{-3} and 10^{-5} respectively. The process of matching segments is made easier by the merging process. However, there still appear to be as many similarities as dissimilarities between segmentations. Whilst many similarities can be found between the segmentations of SLC and PRI images generated on 16/04/92 and 09/06/92, there is little similarity between the segmentations of the SLC and PRI images generated on 12/05/92 and 18/08/92. These observations have severe implications for

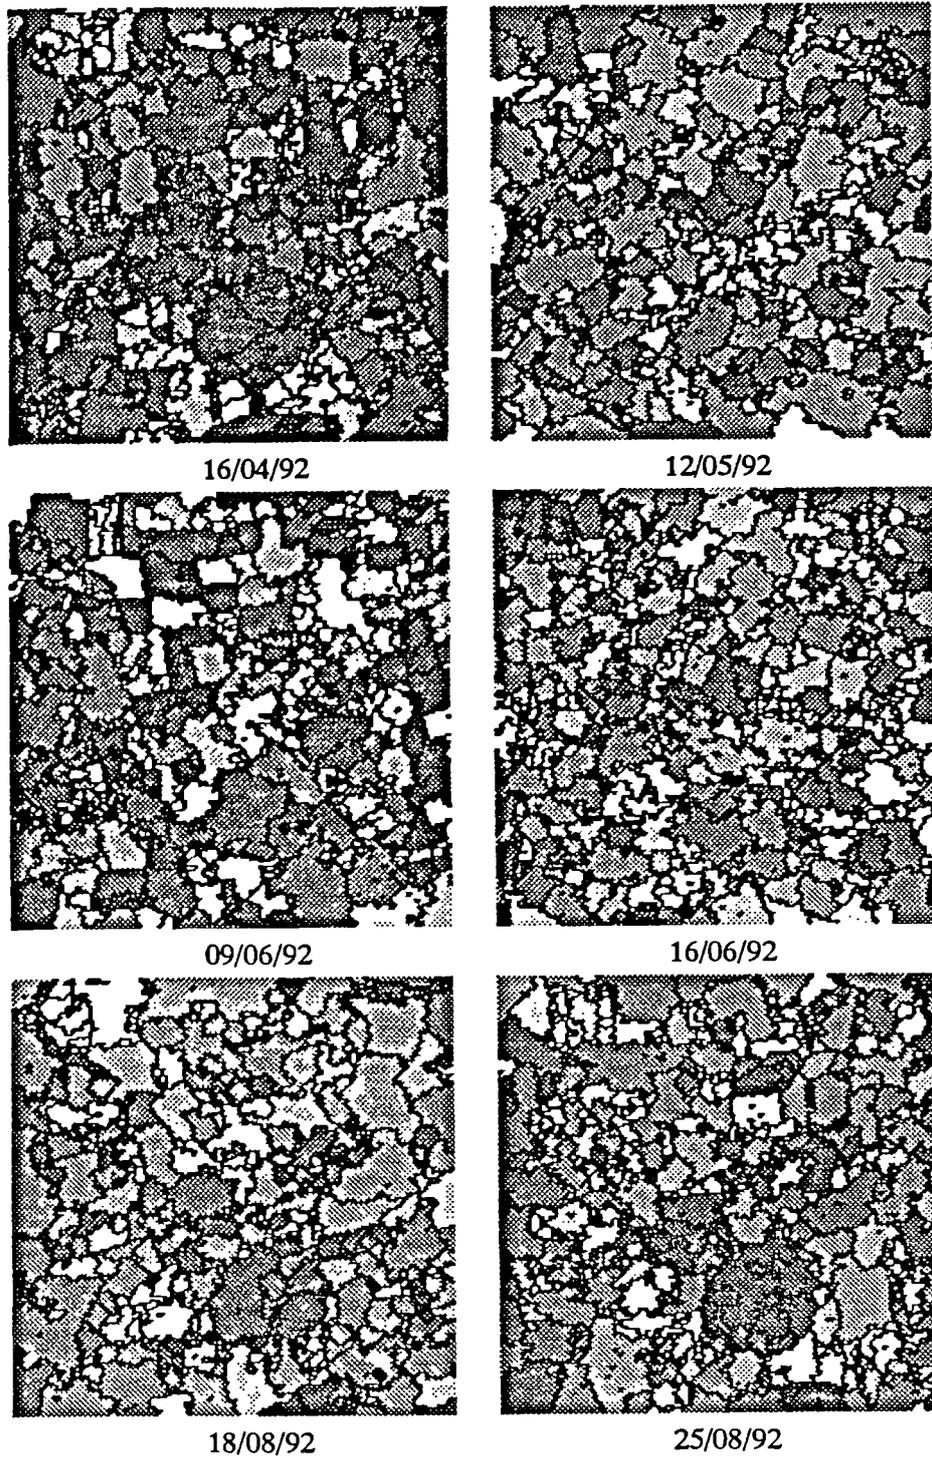


Figure 9.20: Segmentation of the SLC images shown in Figure 9.15 after sub-sampling over a 2×2 window.

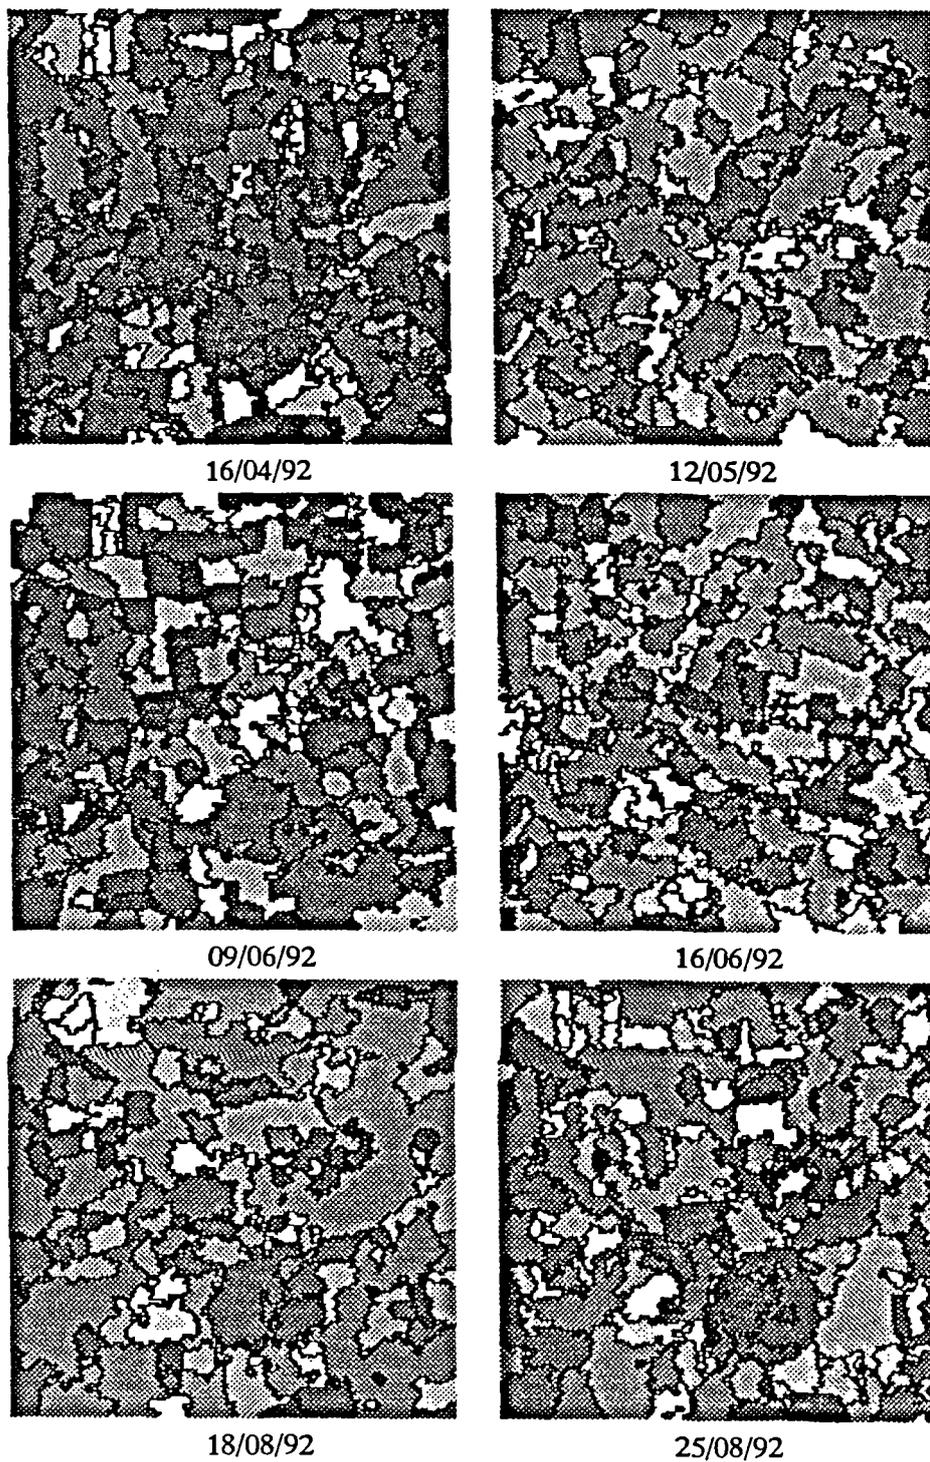


Figure 9.21: Result of edge removal using a PFA of 10^{-3} .

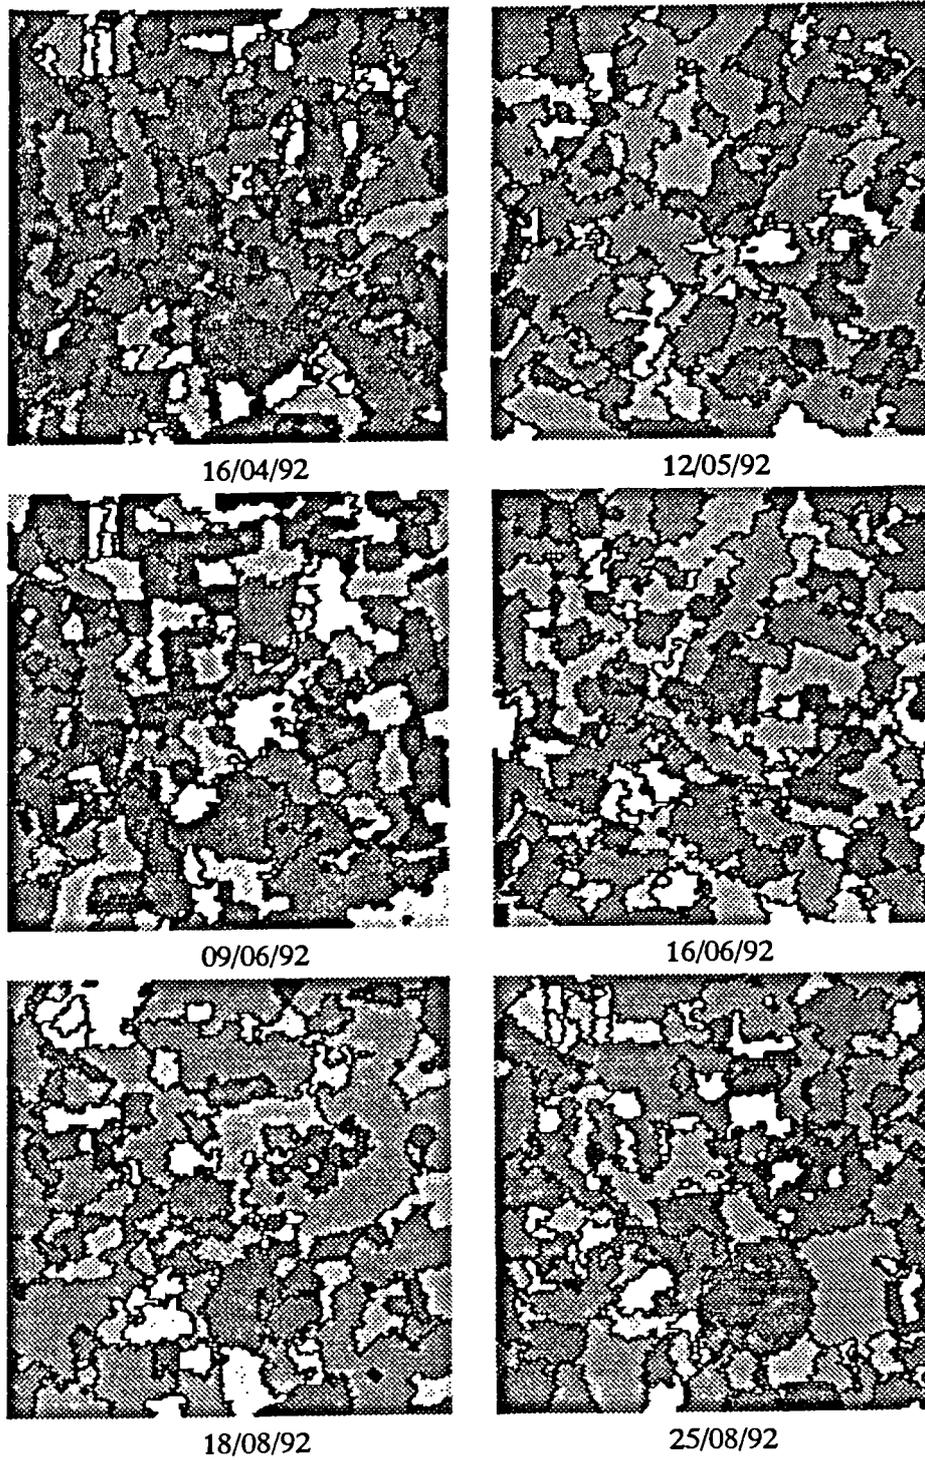


Figure 9.22: Result of edge removal using a PFA of 10^{-5} .

the stability of the segmentation algorithm under different image realisations. This was further borne out when different sub-samplings of a single PRI image were segmented.

Because sub-sampling of the PRI images was carried out over a 2×2 window four distinct sub-sampled images can be generated from a single image. Two distinct sub-samplings of the 16/04/92 PRI image and the result of segmentation with and without merging are shown in Figure 9.23. While there is little difference between the two sub-sampled images there are marked differences between their segmentations, not all of which are removed by merging. Similar differences result when the other two sub-samplings are segmented. These results show that the segmentation algorithm is sensitive to small differences between SAR images of the same scene.

9.3 Summary

In this chapter:

- The segmentation of SAR images using an algorithm originally developed for segmenting high resolution airborne imagery has been investigated. The algorithm involves an iterative process of edge detection and segment growing. Detected edges are used to limit segment growing, then the resulting segmentation is used to generate an improved edge detection.
- The edge detection and segment growing stages were described along with how a measure of the average contrast within segments is used to monitor the segmentation process and to decide when to halt. Weaknesses identified in the original algorithm and a number of modifications made to it during the course of writing this thesis, were also discussed.
 - Normally, iteration halts when the first minimum in the average contrast is reached. Instead the algorithm was allowed to iterate until the second and third minima; no overall improvement was noted in the resulting segmentations.
 - The normalised gradient operator used for edge detection relies on an estimate of the standard deviation within segments in the previous iteration. However, initially no previous segmentation is available and the image is treated as a single segment; this results in poor edge detection performance on the first iteration. The normalised ratio operator, which relies only on estimates of the local intensity, was substituted for the normalised gradient operator on the first iteration. Better initial segmentations were generated but surprisingly this did not carry through to the final segmentations. There was little difference between the final segmentations produced from the initial segmentations.
 - Segmentations produced by the algorithm contain many point-like features. For many applications we are interested in detecting such

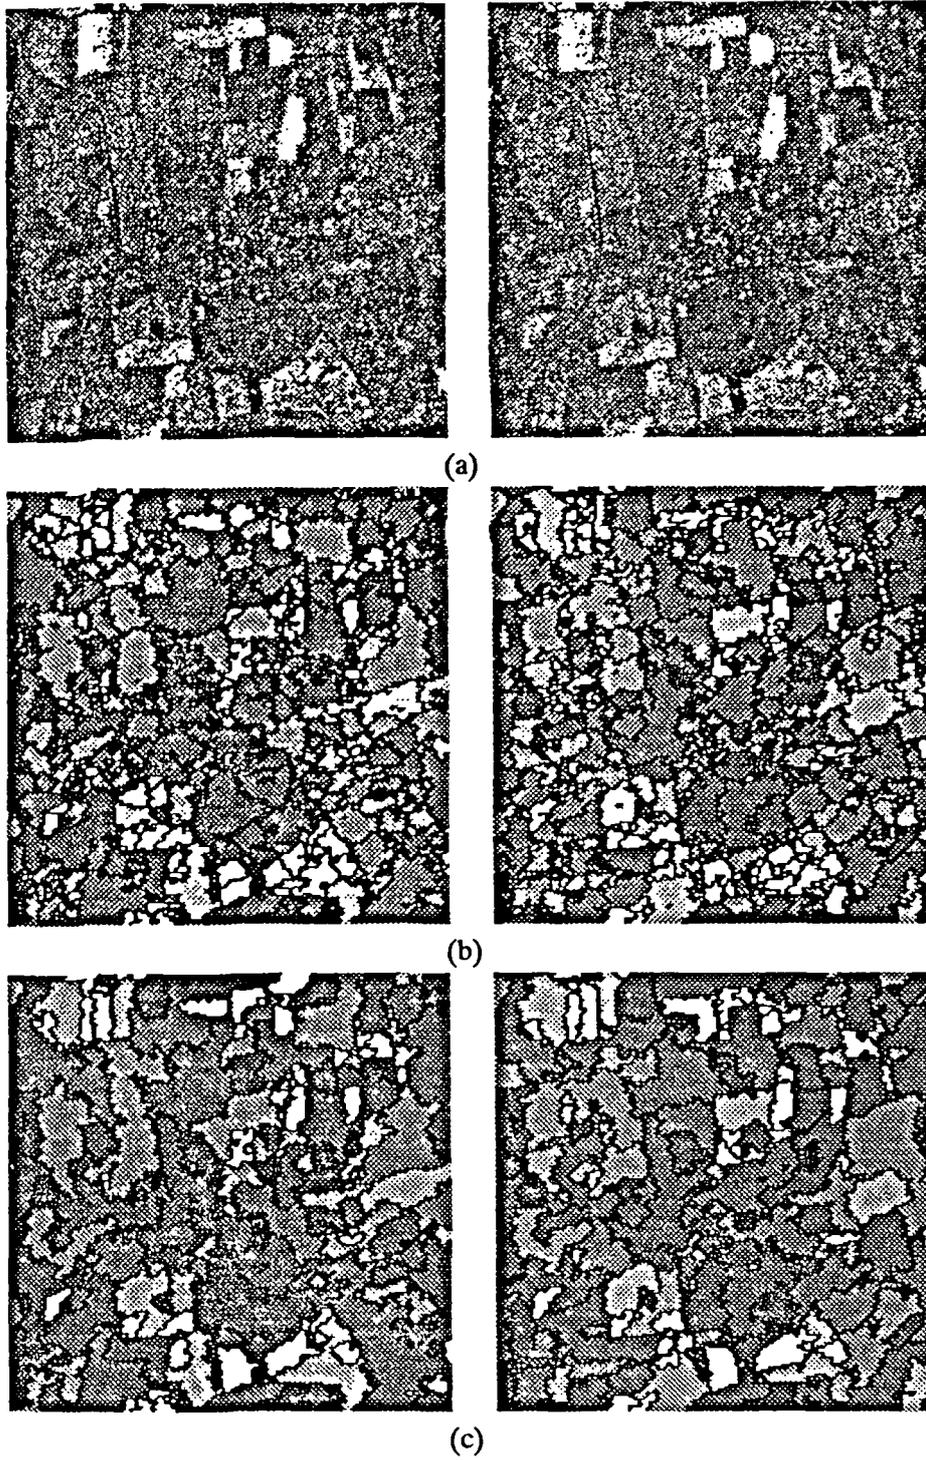


Figure 9.23: (a) two distinct sub-samplings of the 16/04/92 PRI image shown in Figure 9.15, (b) result of segmentation, and (c) result of subsequent segment merging with thresholds set for a PFA of 10^{-5} .

features only in so far it helps the detection of other types of feature (see Chapter 8). Whilst some point-like features represent distinct features it is suspected that many of them arise purely from speckle. More quantitative methods for distinguishing true point features from those that should be incorporated into segments need to be included in the algorithm, e.g., the methods described in Chapter 8. Ideally this should be the first stage of the edge detection process.

- Linear features such as rivers and lines of communication supply important visual clues when interpreting a SAR image, especially when it comes to matching. Currently the segmentation algorithm has no knowledge of such features. Thus, it is crucial that linear feature detectors be developed for incorporation into the algorithm.
- The thresholds presently used for edge detection correspond to a high false alarm rate. The effect of setting a lower false alarm rate or decreasing the false alarm rate over subsequent iterations has yet to be investigated.
- The need for objective methods of measuring segmentation accuracy was stressed. Homogeneity within segments can be checked by dividing the original image by an image showing the mean value within segments. The resulting ratio image should consist purely of clutter; this image can be resegmented to reveal any remaining structure. Measures of statistical difference may be used to check that adjacent segments differ significantly.
- The full segmentations produced by the algorithm are too complete for many applications, i.e., they suffer from over-segmentation. This is more evident when the edges of segments rather than the mean value within them is displayed. Since detailed visual inspection suggests that no real boundaries are omitted in the full segmentation, a well-founded way of merging segments is required which allows the segmentation to be degraded while preserving the desired level of ‘information’.
 - The original algorithm carries out segment merging by comparing their moments. This method was found to be highly dependent on segment size. Unfortunately it was only recently discovered that this was due to an error in how the standard error of moments was calculated. Thus most work on segment merging was carried out using an alternative approach and a proper analysis of the corrected algorithm still awaits.
 - A strength was attached to each edge; the weakest edges, up to a given threshold, were then removed and the segments they separated merged. This had to be carried out dynamically as the strength of the remaining edges bounding two merged segments had to be recalculated prior to the next weakest edge being removed, otherwise edges lost closure.
 - Dynamic thresholding was carried out using edge strengths based on both the ratio and difference of the average intensity within segments.

However, neither of these measures take into account the size of segments. We wished to retain both small segments separated by strong edges and large segments separated by weak edges.

- A more suitable measure is given by the false alarm probability corresponding to the edge strength given the size of the segments. Such a probability can be calculated for the edge ratio assuming the simple image model applies, and intrinsically weights both the strength of the edge and the size of the segments it separates. Correlation needs to be taken into account when calculating this measure.
- The segmentation algorithm was used to investigate agricultural change detection in ERS-1 images over Feltwell, East Anglia.
 - A set of six ERS-1 SLC and PRI images taken during the growing season of 1992 was used. Large structural differences between images taken at different dates were confirmed by segmentation.
 - Initial segmentations of the SLC data were over-segmented in the context of field structures. However, detailed visual inspection convinced us that the algorithm reliably found the segment boundaries in the image. It also found many boundaries that the eye did not recognise until it is told where to look. The eye is less able to detect edges in regions which are radiometrically very bright or very dark (relative to how the image is scaled), but the segmentation performance is unaffected by variations in mean image intensity.
 - Much of this over-segmentation was removed by merging. While the remaining segments faithfully reproduced the structure of the original images, there were large differences between segmentations taken from different dates. In addition, as the segmentation had failed to preserve a very prominent linear feature in the images, it was difficult to find any match at all between certain dates.
 - Previously unencountered problems arose when the segmentation algorithm was applied to the highly correlated PRI data. The high degree of correlation greatly increased the number of false alarms produced by edge detection and the resulting segmentation was heavily over-segmented. These problems were circumvented by sub-sampling the PRI images prior to segmentation. Ideally this problem should be overcome by taking correlation into account when setting edge detection and segment merging thresholds.
 - As with the SLC segmentations the PRI segmentations confirmed the large temporal differences in image structure.
 - The segmentations of SLC and PRI images generated from the same raw data were also compared. While the same general features were apparent, segment boundaries did not always match indicating that the algorithm is unstable. This was further borne out when segmentations of different sub-sampling of the same PRI image were com-

pared. These issues warrant further investigation if segmentation is to be used as a quantitative change detection tool.

Chapter 10

Conclusion and future work

In this concluding chapter we shall examine to what extent the results presented in previous chapters answer the four main questions set out in the introduction regarding the image map matching problem. We shall also highlight topics which have surfaced along the way and which warrant further investigation.

10.1 Background theory

Before addressing the image map matching problem the general background theory needed for carrying out any SAR image analysis was described in Chapters 2 to 4. Chapter 2 looked at the general theory of SAR imaging. It was shown how the system response to a point target of a fully focused SAR using a chirp pulse, could be described by a *sinc* function in both azimuth and range. The width between the main lobe and first null of this function defined the resolution.

The statistical models needed to describe the fluctuations in the detected intensity caused by speckle and imaged surface texture were discussed in Chapter 3. Starting from the assumption that the detected field is given by the sum of contributions from a large number of discrete elementary scatterers, whose phase components are uniformly distributed and independent (fully developed speckle), it was shown that the detected field could be modelled as the product of a Gaussian speckle process and the square root of a gamma distributed imaged surface cross section. The intensity is then K-distributed and the phase is uniformly distributed. Reasons for treating the imaged surface cross section as being gamma distributed were considered. It was shown that the order parameter of the distribution can be related to the degree of imaged surface texture. When no texture is detected (infinite order parameter) the imaged surface cross section is constant and the intensity is exponentially distributed. The second order intensity statistics were also derived. The intensity ACF is dependent on the ACF's of both the imaged surface cross section and the speckle process; the ACF of the latter is determined solely by the system response.

In Chapter 4 we went on to look at the estimation of the mean intensity and the degree of imaged surface texture, i.e., the order parameter of K-distributed intensity. It was shown that maximum likelihood estimates were only practical in the case of constant imaged surface cross section. The mean intensity within

a homogeneous region is then given by the average intensity and will be gamma distributed when the averaged intensities are independent. Correlation and fluctuations in the imaged surface cross section were shown to increase the variance of the average intensity. The result of averaging in both intensity and amplitude was also related to the statistics of multi-look SAR images. Non-optimal methods of estimating the degree of imaged surface texture were compared. An estimate based on the mean log-intensity was shown to be more reliable than an estimate based on the second normalised moment.

10.2 Direct matching

The first question in the introduction asked can map features be matched directly to SAR image features. The results presented on template matching in Chapter 5 gave a categoric no to this. A correlation measure which had knowledge of speckle was shown to perform better than straight cross-correlation under a simple image model (constant imaged surface cross section and independent speckle), both theoretically, and when used to match features in simulated images generated according to the simple image model. However, neither measure performed well when used to match features in real images. It was found that the simple image model did not adequately predict how a given feature would appear in a SAR image. Possible extensions to the model to allow for texture and correlation were considered. However, the processing overheads involved in fitting multi-parameter distributions to matched features did not make this an attractive option. This led on to the conclusion that the mis-match between the two data types had to be reduced prior to matching. Possible approaches include applying edge detection and segmentation algorithms to structure the image. This answered the second question set out in the introduction.

Before going on to look at these methods, further analysis was carried out in Chapter 6 of the statistics along edges including the types of linear features selected for matching. The discrete scatterer model was extended to describe the different scattering processes contributing to the detected field along an edge. The intensity statistics arising from this model were then derived. In the absence of strong specular or secondary scattering, the intensity is given by the product of an exponential speckle process and a weighted sum of gamma distributed imaged surface cross section contributions from the different scattering populations making up the edge. The weighting is determined by the mean number of scatterers from each population and the variance in their density. It was shown that this sum would itself be gamma distributed when either a single contribution dominated or when all the contributions have equal weighting; the intensity will then be K-distributed. These results implied that we can expect pixel intensity to be approximately K-distributed along many types of edge in single-look images. This was checked by measurement. The statistics of intensity values sampled along a number of edges and linear features were fitted to K-distributions. In virtually all cases a good fit was achieved; the best fit K-distributions had a range of order parameters. In fact, the intensity along edges and linear features was found to give a better fit to K-distributions than intensity

values sampled over extended homogeneous regions.

10.3 Edge detection

Edge detection was investigated in Chapter 7. Amongst a number of local operators suggested for detecting edges in SAR the normalised ratio operator was identified as the most promising in terms of general edge detection criteria. When the simple image model applies this operator acts as a constant false alarm rate edge detector and gives the maximum likelihood estimate of the intensity ratio between homogeneous regions, in addition, the false alarm rate corresponding to a given edge detection threshold can be quantified. This is a major advantage of using this operator. Other operators suffered from a combination of problems of threshold selection, false alarm rates being dependent on image intensity, and excessive false alarms.

When the performance of the normalised ratio operator was analysed in more detail a variety of problems were identified regarding the use of the operator. These problems are listed below along with proven/suggested solutions.

1. *When thresholds are selected so as to produce a desired false alarm rate under the simple image model, system induced correlation and texture (K-distributions) will cause the false alarm rate to be significantly greater than expected. In most cases the effect of correlation will be much greater than that due to texture. In a sense this is fortunate. System induced correlation is predictable, often only requiring measurement of the intensity ACF over a single homogeneous region, while texture which may vary over the image, is difficult to parameterise (see the discussion of texture estimation in Chapter 4). The measured intensity ACF may be used to determine the degree of sub-sampling needed, or how to adapt thresholds, to meet expected false alarm rates. Because of the information loss and edge thickening produced by sub-sampling threshold adaption is to be preferred.*
2. To detect edges at different scales, i.e., strong edges separating small segments as well as weak edges separating large segments, the normalised ratio operator has to be applied over a range of window sizes. Edge thickening results when edge detection is carried out within too large a window. This leads to small scale edges being obscured and problems in relating edges detected at different scales. To reduce thickening edge detection windows need to be applied in order of increasing size, with edges only being detected with a given window size where no edges have been detected using smaller windows [61].
3. Because the smallest scale edges in an image are those defining point-like targets, point target detection needs to be carried out prior to edge detection and the results fed into the edge detection process described in the previous item. A number of point target detection techniques were investigated in Chapter 8. Some of these produced promising results but

further work is need to check the consistency of the techniques over a range of images.

4. Edge thickening is the result of detected edges being represented as a property of specific pixels (the central pixel within the window over which the edge was detected). It may be possible to circumvent this problem using an alternative representation, e.g., representing edges as boundaries between pixels (such a representation is more in keeping with the aims of segmentation). The feasibility and benefits, if any, of doing this need to be investigated.
5. Detected edges are often disconnected. One possible method of carrying out linking is to identify the closed boundaries detected by the Marr-Hildreth operator which correspond to detected edges [6]; this also enables edges to be thinned to a single pixel wide line. The usefulness of this method has yet to be investigated.

10.4 Segmentation

Segmentation of SAR images using an algorithm developed at RSRE was investigated in Chapter 9. As a means of structuring an image this segmentation algorithm has various advantages over edge detection using the normalised ratio operator. Some of these are due to the normalised gradient operator used by the algorithm to detect edges, while others are due to the iterative edge detection - segment growing nature of the algorithm.

1. The segmentation algorithm detects edges using a gradient operator normalised by the local standard deviation as opposed to the local mean. Thus, the operator will approximately act as a constant false alarm rate edge detector in regions which are K-distributed with different order parameters.
2. Edge thickening and attendant problems are reduced by only detecting edges at any given scale where no edges have been detected at smaller scales (see the second item in the previous section).
3. The segment growing stage provides a means of linking disconnected edges (see the fifth item in the previous section).
4. The algorithm benefits from an iterative learning process by using previous segmentations to generate an improved edge detection.

The segmentation algorithm was used to investigate change detection in a series of ERS-1 image taken over the Feltwell agricultural test site. The results demonstrated the usefulness of the algorithm as a tool for extracting image structure and highlighting the structural differences between multi-temporal images. However, they also illustrated a number of its weaknesses.

1. The algorithm suffers from having no knowledge of linear features which are poorly detected by the edge detection stage. This is a serious problem as linear features may supply important visual clues when interpreting an image, particularly when matching a series of multi-temporal images. Thus, an important aim of future work is to develop suitable linear feature detectors for incorporation into the algorithm.
2. The algorithm detects an excessive number of small segments. While some of these certainly represent true point-like targets, it is suspected that many arise purely from speckle and these should be combined into surrounding segments. Therefore the methods used to detect point targets (small segments) need to be improved possibly using the methods described in Chapter 8. Also due to the problems caused by point targets for edge detection these methods need to be integrated more fully into the edge detection stage.
3. The segmentations produced by the algorithm are too complete for many applications, i.e., regions that we would expect to be detected as a single homogeneous segment (e.g., an agricultural field) and that appear to be so are split into several segments. In some cases this over-segmentation may be due to intensity differences between very dark or very bright regions which the eye cannot see unless the image is scaled. To degrade the segmentation down to the desired level of 'information' requires a final stage of segment merging. The original algorithm included a means of doing this based on comparing the moments of segments. However, due to an implementation error which was only very recently discovered, this method did not work as intended. This led to an alternative method of merging being developed. A degree of credibility is attached to each edge separating segments, the least credible edges up to a given threshold are then removed and the segments they separate are merged. This needs to be carried out dynamically to retain edge closure. This alternative method of merging which works successfully, still needs to be compared with the corrected version of the original method based on comparing moments. The effect of carrying out merging after each iteration rather than at the end also needs to be investigated.
4. It was found that segmentation and merging performance were strongly affected by correlation which had previously not been taken into account when selecting thresholds. This was discovered by chance when highly correlated ERS-1 PRI images were segmented. The resulting segmentations were excessively over-segmented. Known system induced correlation can be compensated for by sub-sampling prior to segmentation. However, unless data is 100% correlated this will lead to some information loss. Therefore it is desirable to develop means of predicting how to adapt thresholds to compensate for known correlation. The methods already developed for doing this when detecting edges using the normalised ratio operator may be of help. It is important that correlation is correctly taken

into account not just to reduce over-segmentation, but also to enable the differences between adjacent segments to be properly quantified and an overall accuracy to be attached to the segmentation.

5. The comparison of segmentations of ERS-1 SLC and PRI images of the same scene as well as different sub-samplings of the same image, indicate that the algorithm is unstable. This has serious implications for the use of segmentation for structural change detection. If insignificant differences between images can lead to significantly different segmentations there is little chance of quantifying larger differences. The source of this instability needs to be identified and if possible methods of rectifying it developed. It is suspected the instability may arise from false alarm edges influencing the segment growing stage, thus algorithm stability may be improved by using thresholds corresponding to a lower PFA. This needs investigation. Another possibility is to reduce/increase edge detection thresholds over subsequent iterations.
6. Most of the segmentations presented in Chapter 9 were compared using purely subjective visual criteria. This is insufficient as a means of objectively comparing different segmentation methods and quantifying differences between multi-temporal segmentations. Possible approaches to checking the homogeneity within regions and the differences between adjacent segments were outlined. These methods need to be more fully developed into clearly defined routines for measuring segmentation accuracy in an objective manner. A difficulty with defining what is meant by a 'good segmentation' is that criteria will vary from application to application. One possible solution is to simulate SAR images containing known features corresponding to the objects of interest to a particular application. The segmentation algorithm can then be run on these images to objectively measure its ability to detect such features. An implementation of such a test suite is described in [13].

10.5 Image map matching

Segmentation using the RSRE algorithm clearly offers a greater degree of image structuring than edge detection based on the normalised ratio operator. The former produces an edge map consisting of single pixel wide closed boundaries separating segments which supposedly represent distinct homogeneous regions, while the latter only produces a disconnected edge map with edges of varying thickness; the thickness of an edge being related to its strength. Therefore the segmentation algorithm would appear to be the more obvious choice of method for structuring a SAR image prior to carrying out image map matching. However, for some applications there is a danger in the segmentation algorithm imposing too much structure on an image as is seen in the over-segmented images shown in Chapter 9.

The segmentation algorithm is based on an image model which does not allow for gradual changes in image intensity; the output segmentation con-

sists of homogeneous regions (segments) separated by abrupt closed boundaries (edges). Over structuring may occur when disconnected edges are linked into closed boundaries by the segment growing stage. Over many types of terrain such an image model is highly suitable, e.g., when segmenting an image showing agricultural fields, but may not be so suitable in areas containing no hard boundaries, e.g., regions covered by gradually varying natural vegetation. The fact that segmentation imposes such an image model should always be remembered when interpreting the results of the segmentation process. Although edge detection is based on a similar image model it is less important to take it into consideration when interpreting results; so long as no edge linking is applied. Because apart from false alarms due to speckle, which may be controlled by thresholding, no edges are detected where they do not exist.

The type of linear feature we were previously interested in matching in Chapter 5 is presently poorly detected by the segmentation algorithm. Thus, image map matching of this type can only proceed in line with improvements in linear feature detection. An alternative which is more suited to the output of the present algorithm is region based matching. Segmentation combined with a suitable degree of merging may be used to produce a representation of the image which corresponds to a map showing land cover types. Matching may then proceed on the basis of a number of possible criteria, namely the shape, size, and topology of regions. Such an approach has been suggested previously but does not appear to have been properly investigated [27]. This approach does involve various assumptions concerning the land units shown in the map data, that they have not significantly changed shape (merged/split) by the time the image was taken, and that they are radiometrically distinguishable (we automatically assume that the map has been transformed to the same geometry as the image).

The results presented in Chapter 9 on the segmentation of ERS-1 images over the Feltwell agricultural test site imply that this last assumption can not always be taken for granted. While it is reasonable to assume that field boundaries will not have significantly changed over the five months during which the images were taken, radiometric changes within fields radically alter the perceived image structure both by eye and machine. Because of this, given an accurate field map of this test site, matching could not always be carried out using region based methods alone. In terms of matching this series of images to map data a more sensible feature for matching is the drainage channel, which as we have already said first requires improved methods for detecting such features. Water features such as this along with lines of communication are less prone to radiometric change than agricultural and forested areas and are thus temporally more stable features for matching (note, the appearance of such features can alter radically when imaged from different directions [25]). Matching using such linear features is particularly needed when the detection of radiometric change occurring within land units is of primary interest, e.g., monitoring crops or changes in surface moisture.

10.6 Future work

While the RSRE segmentation appears to supply a suitable means of structuring an image prior to matching, the above observations all the more emphasize the need to be able to reliably detect prominent linear features and to be able to match such features to map data. Only then can map data be used as an independent reference for change detection. Future work should concentrate on these two aspects of the image map problem, the first in particular as matching can only proceed once reliable linear feature detectors are available. The development of linear feature detectors will also significantly improve the usefulness of segmentation as an image analysis tool. Secondly it is important that the problems of point target detection, algorithm instability, checking of segmentation accuracy and correlation should all be addressed. Also, despite the problems discussed above the feasibility of carrying out region based matching should be investigated.

The advantages offered by segmentation over edge detection call into question the need to further develop edge detection using the normalised ratio operator as an image analysis tool in its own right. The suggested improvements in this algorithm are either shared in common with, or are already implemented in the segmentation algorithm. It may possibly be beneficial to implement these improvements to make available an 'optimal' non-iterative edge detector. It would also be interesting to investigate the problems involved in and benefits of representing detected edges as boundaries between pixels rather than as 'edge pixels'.

In this thesis all segmentation has been carried out on single images. However, additional radiometric and structural information is contained in multi-frequency and multi-polarimetric SAR datasets (e.g., AIRSAR quad-polarised three frequency data). To make full use of this information requires the development of multi-dimensional segmentation techniques. This would first involve identifying the information content of the multi-dimensional data so that suitable segmentation parameters could be defined. Operators for carrying out multi-dimensional edge detection based on these parameters would then need to be developed. After edge detection segment growing could proceed as in the single-dimensional case. The segmentation accuracy measure needed to control iteration would need to take account of the multi-dimensionality of the dataset.

Hopefully the increased structural information contained in multi-dimensional datasets, extracted via multi-dimensional segmentation, should make image map matching more feasible. Though possibly also more complex. However, the problems inherent in the single-dimensional segmentation algorithm, especially those of stability, should be addressed prior to such advanced methods being developed.

Appendix A

Relation between the input and output distributions of basic binary operators

Given two independent positive random variables x_1 and x_2 with pdf's $f_{x_1}(x)$ and $f_{x_2}(x)$ respectively, the pdf's of their sum, difference, absolute difference, product, ratio and normalised ratio are:

1. sum $y = x_1 + x_2$

$$f_y(y) = \int_0^y f_{x_1}(y-x) f_{x_2}(x) dx \quad (\text{A.1})$$

which by the convolution theorem equals

$$= \frac{1}{2\pi} \int_0^\infty \Phi_{x_1}(\omega) \Phi_{x_2}(\omega) e^{-j\omega x} d\omega \quad (\text{A.2})$$

where $\Phi_{x_1}(\omega)$ and $\Phi_{x_2}(\omega)$ are the characteristic functions of x_1 and x_2 .

2. difference $y = x_1 - x_2$

$$f_y(y) = \int_{\max(-y,0)}^\infty f_{x_1}(y+x) f_{x_2}(x) dx \quad (\text{A.3})$$

3. absolute difference $y = |x_1 - x_2|$

$$f_y(y) = \int_0^\infty [f_{x_1}(y+x) f_{x_2}(x) + f_{x_1}(x) f_{x_2}(y+x)] dx \quad (\text{A.4})$$

4. product $y = x_1 x_2$

$$f_y(y) = \int_0^\infty f_{x_1}\left(\frac{y}{x}\right) f_{x_2}(x) \frac{1}{x} dx \quad (\text{A.5})$$

5. ratio $y = \frac{x_1}{x_2}$

$$f_y(y) = \int_0^\infty f_{x_1}(yx) f_{x_2}(x) x dx \quad (\text{A.6})$$

6. normalised ratio $y = \min\left(\frac{x_1}{x_2}, \frac{x_2}{x_1}\right)$

$$f_y(y) = \int_0^\infty [f_{x_1}(yx) f_{x_2}(x) + f_{x_1}(x) f_{x_2}(yx)] x dx \quad (\text{A.7})$$

Appendix B

Characteristic function of a Rayleigh distribution

The characteristic function of a Rayleigh distributed random variable y is given by

$$\Phi_y(\omega) = 1 + j \frac{\omega \sqrt{\pi \mu}}{2} e^{-\frac{\omega^2 \mu}{4}} - \frac{\omega \mu}{2} e^{-\frac{\omega^2 \mu}{4}} \int_0^\omega e^{\frac{b^2 \mu}{4}} db \quad (\text{B.1})$$

where $\mu = \langle y^2 \rangle$. The proof is as follows.

The characteristic function of y is defined from its pdf (3.90) by

$$\Phi_y(\omega) = \int_0^\infty \frac{2ye^{-\frac{y^2}{\mu}}}{\mu} e^{j\omega y} dy$$

This may be expressed in terms of the characteristic function of $z = \sqrt{2/\mu}y$

$$\Phi_y(\omega) = \int_0^\infty ze^{-\frac{z^2}{2}} e^{j\omega \sqrt{\frac{\mu}{2}}z} dz = \Phi_z\left(\omega \sqrt{\frac{\mu}{2}}\right) \quad (\text{B.2})$$

where

$$\Phi_z(\omega) = \int_0^\infty ze^{-\frac{z^2}{2}} e^{j\omega z} dz$$

The characteristic function of y will be determined from that of z which may be rewritten as

$$\begin{aligned} \Phi_z(\omega) &= \int_0^\infty ze^{-\frac{z^2}{2}} e^{j\omega z} dz = e^{-\frac{\omega^2}{2}} \int_0^\infty ze^{-\frac{(z-j\omega)^2}{2}} dz \\ &= e^{-\frac{\omega^2}{2}} \int_{-j\omega}^{\infty-j\omega} (a+j\omega) e^{-\frac{a^2}{2}} da \end{aligned}$$

where $a = z - j\omega$.

Defining

$$\begin{aligned} \phi_1 &= \int_{-j\omega}^{\infty-j\omega} (a+j\omega) e^{-\frac{a^2}{2}} da \\ \phi_2 &= \int_0^{-j\omega} (a+j\omega) e^{-\frac{a^2}{2}} da \end{aligned}$$

$$\begin{aligned}\phi_3 &= \int_X^0 (a + j\omega) e^{-\frac{a^2}{2}} da \\ \phi_4 &= \int_{X-j\omega}^X (a + j\omega) e^{-\frac{a^2}{2}} da\end{aligned}$$

ϕ_1 , ϕ_2 , ϕ_3 and ϕ_4 constitute an integral round a closed path. By Cauchy's theorem their sum equals zero and

$$\Phi_z(\omega) = e^{-\frac{\omega^2}{2}} \lim_{X \rightarrow \infty} \phi_1 = -e^{-\frac{\omega^2}{2}} \lim_{X \rightarrow \infty} (\phi_2 + \phi_3 + \phi_4)$$

Thus we can determine $\Phi_z(\omega)$ by finding the limit of ϕ_2 , ϕ_3 and ϕ_4 as $X \rightarrow \infty$.

Because ϕ_2 is independent of X , letting $a = -jb$ its limit is given by

$$\lim_{X \rightarrow \infty} \phi_2 = \phi_2 = \int_0^\omega (\omega - b) e^{\frac{b^2}{2}} db = 1 - e^{\frac{\omega^2}{2}} + \omega \int_0^\omega e^{\frac{b^2}{2}} db \quad (\text{B.3})$$

The limit of ϕ_3 equals

$$\lim_{X \rightarrow \infty} \phi_3 = -1 - j\sqrt{\frac{\pi}{2}}\omega \quad (\text{B.4})$$

Letting $a = j\omega(b-1) + X$, ϕ_4 is given by

$$\phi_4 = j\omega e^{-\frac{X^2}{2}} \int_0^1 (j\omega b + X) e^{\frac{\omega^2(b-1)^2}{2}} e^{-j\omega(b-1)X} db$$

thus

$$|\phi_4| \leq |\omega| e^{-\frac{X^2}{2}} \int_0^1 |j\omega b + X| e^{\frac{\omega^2(b-1)^2}{2}} db \leq |\omega| e^{-\frac{X^2}{2}} e^{\frac{\omega^2}{2}} \int_0^1 |j\omega b + X| db$$

as the RHS $\rightarrow 0$ as $X \rightarrow \infty$, the limit of ϕ_4 equals zero

$$\lim_{X \rightarrow \infty} \phi_4 = 0 \quad (\text{B.5})$$

Combining (B.3,B.4,B.5) z has characteristic function

$$\begin{aligned}\Phi_z(\omega) &= -e^{-\frac{\omega^2}{2}} \lim_{X \rightarrow \infty} (\phi_2 + \phi_3 + \phi_4) \\ &= -e^{-\frac{\omega^2}{2}} \left(1 - e^{\frac{\omega^2}{2}} + \omega \int_0^\omega e^{\frac{b^2}{2}} db - 1 - j\sqrt{\frac{\pi}{2}}\omega \right) \\ &= 1 + j\sqrt{\frac{\pi}{2}}\omega e^{-\frac{\omega^2}{2}} - \omega e^{-\frac{\omega^2}{2}} \int_0^\omega e^{\frac{b^2}{2}} db\end{aligned}$$

Therefore from (B.2) y has characteristic function

$$\Phi_y(\omega) = 1 + j\frac{\omega\sqrt{\pi\mu}}{2} e^{-\frac{\omega^2\mu}{4}} - \frac{\omega\mu}{2} e^{-\frac{\omega^2\mu}{4}} \int_0^\omega e^{\frac{b^2\mu}{4}} db$$

This completes the proof of (B.1).

Appendix C

Pdf of the sum of two gamma distributed random variables

The pdf of the sum y of two gamma distributed random variables x_1 and x_2 with means μ_1 and μ_2 , and integer order parameters ν_1 and ν_2 is given by

$$f_y(y) = \frac{1}{\beta_1^{\nu_1} \beta_2^{\nu_2}} \left\{ e^{\frac{-y}{\beta_1}} \alpha^{\nu_2} \sum_{k=0}^{\nu_1-1} \binom{\nu_2-1+k}{k} \frac{y^{\nu_1-1-k} (-\alpha)^k}{\Gamma(\nu_1-k)} + e^{\frac{-y}{\beta_2}} (-\alpha)^{\nu_1} \sum_{k=0}^{\nu_2-1} \binom{\nu_1-1+k}{k} \frac{y^{\nu_2-1-k} \alpha^k}{\Gamma(\nu_2-k)} \right\} \quad (C.1)$$

where $\beta_1 = \mu_1/\nu_1$, $\beta_2 = \mu_2/\nu_2$ and $1/\alpha = 1/\beta_2 - 1/\beta_1$. This will be proven by induction on ν_1 .

When $\nu_1 = 1$, substituting the pdf's of x_1 and x_2 given by (3.59) into (A.1), their sum has pdf

$$\begin{aligned} f_y(y) &= \int_0^y \frac{e^{\frac{-(y-x)}{\beta_1}}}{\beta_1} \frac{e^{\frac{-x}{\beta_2}} x^{\nu_2-1}}{\beta_2^{\nu_2} \Gamma(\nu_2)} dx \\ &= \frac{1}{\beta_1 \beta_2^{\nu_2}} \left\{ e^{\frac{-y}{\beta_1}} \alpha^{\nu_2} + e^{\frac{-y}{\beta_2}} (-\alpha) \sum_{k=0}^{\nu_2-1} \frac{y^{\nu_2-1-k} \alpha^k}{\Gamma(\nu_2-k)} \right\} \end{aligned}$$

This proves (C.1) for $\nu_1 = 1$.

To complete the proof we will assume (C.1) to be true for $\nu_1 = \nu$ and prove it to be true for $\nu_1 = \nu + 1$. As the order parameters are integers each x_i is the sum of ν_i exponential random variables with mean β_i . Thus the sum of two gamma distributed random variables with means $(\nu + 1)\beta_1$ and $\nu_2\beta_2$, and order parameters $\nu + 1$ and ν_2 , is equal to the sum of two gamma distributed random variables with means $\nu\beta_1$ and $\nu_2\beta_2$ and order parameters ν and ν_2 , plus an exponential random variable with mean β_1 . Assuming (C.1) is true for $\nu_1 = \nu$ and once again using (A.1), the pdf of this sum is given by

$$f_y(y) = \int_0^y \frac{e^{\frac{-(y-x)}{\beta_1}}}{\beta_1} \frac{1}{\beta_1^{\nu} \beta_2^{\nu_2}} \left\{ e^{\frac{-x}{\beta_1}} \alpha^{\nu_2} \sum_{k=0}^{\nu-1} \binom{\nu_2-1+k}{k} \frac{x^{\nu-1-k} (-\alpha)^k}{\Gamma(\nu-k)} \right.$$

$$+ e^{\frac{-y}{\beta_2}} (-\alpha)^\nu \sum_{k=0}^{\nu_2-1} \left(\binom{\nu-1+k}{k} \frac{x^{\nu_2-1-k} \alpha^k}{\Gamma(\nu_2-k)} \right) dx$$

re-arranging terms

$$f_y(y) = \frac{1}{\beta_1^{\nu+1} \beta_2^{\nu_2}} e^{\frac{-y}{\beta_1}} \left\{ \alpha^{\nu_2} \sum_{k=0}^{\nu_2-1} \left(\binom{\nu_2-1+k}{k} \frac{(-\alpha)^k}{\Gamma(\nu-k)} \int_0^y x^{\nu-1-k} dx \right. \right. \\ \left. \left. + (-\alpha)^\nu \sum_{k=0}^{\nu_2-1} \left(\binom{\nu-1+k}{k} \frac{\alpha^k}{\Gamma(\nu_2-k)} \int_0^y e^{\frac{-x}{\alpha}} x^{\nu_2-1-k} dx \right) \right\}$$

and integrating, the pdf is given by

$$f_y(y) = \frac{1}{\beta_1^{\nu+1} \beta_2^{\nu_2}} e^{\frac{-y}{\beta_1}} \left\{ \alpha^{\nu_2} \sum_{k=0}^{\nu_2-1} \left(\binom{\nu_2-1+k}{k} \frac{y^{\nu-k} (-\alpha)^k}{\Gamma(\nu+1-k)} \right. \right. \\ \left. \left. + (-\alpha)^\nu \alpha^{\nu_2} \sum_{k=0}^{\nu_2-1} \left(\binom{\nu-1+k}{k} \right. \right. \right. \\ \left. \left. \left. - e^{\frac{-y}{\alpha}} (-\alpha)^\nu \alpha^{\nu_2} \sum_{k=0}^{\nu_2-1} \left(\binom{\nu-1+k}{k} \right) \sum_{j=0}^{\nu_2-1-k} \frac{1}{j!} \left(\frac{y}{\alpha} \right)^j \right) \right\} \quad (C.2)$$

Substituting the identities

$$\sum_{k=0}^{\nu_2-1} \binom{\nu-1+k}{k} = \binom{\nu_2-1+\nu}{\nu}$$

and

$$\sum_{k=0}^{\nu_2-1} \binom{\nu-1+k}{k} \sum_{j=0}^{\nu_2-1-k} \frac{1}{j!} \left(\frac{y}{\alpha} \right)^j = \sum_{k=0}^{\nu_2-1} \binom{\nu+k}{k} \frac{1}{\Gamma(\nu_2-k)} \left(\frac{y}{\alpha} \right)^{\nu_2-1-k}$$

into (C.2) gives

$$f_y(y) = \frac{1}{\beta_1^{\nu+1} \beta_2^{\nu_2}} \left\{ e^{\frac{-y}{\beta_1}} \alpha^{\nu_2} \sum_{k=0}^{\nu_2-1} \left(\binom{\nu_2-1+k}{k} \frac{y^{\nu-k} (-\alpha)^k}{\Gamma(\nu+1-k)} \right. \right. \\ \left. \left. + e^{\frac{-y}{\beta_1}} \alpha^{\nu_2} \binom{\nu_2-1+\nu}{\nu} (-\alpha)^\nu \right. \right. \\ \left. \left. - e^{\frac{-y}{\beta_2}} (-\alpha)^\nu \alpha^{\nu_2} \sum_{k=0}^{\nu_2-1} \left(\binom{\nu+k}{k} \frac{1}{\Gamma(\nu_2-k)} \left(\frac{y}{\alpha} \right)^{\nu_2-1-k} \right) \right\} \\ = \frac{1}{\beta_1^{\nu+1} \beta_2^{\nu_2}} \left\{ e^{\frac{-y}{\beta_1}} \alpha^{\nu_2} \sum_{k=0}^{\nu_2-1} \left(\binom{\nu_2-1+k}{k} \frac{y^{\nu-k} (-\alpha)^k}{\Gamma(\nu+1-k)} \right. \right. \\ \left. \left. + e^{\frac{-y}{\beta_2}} (-\alpha)^{\nu+1} \sum_{k=0}^{\nu_2-1} \left(\binom{\nu+k}{k} \frac{y^{\nu_2-1-k} \alpha^k}{\Gamma(\nu_2-k)} \right) \right\}$$

When $\nu_1 = \nu + 1$ this has the same form as (C.1). This completes the proof.

Appendix D

Intensity moments of sums of circularly symmetric complex random variables

Given n independent circularly symmetric complex random variables z_1, \dots, z_n we wish to prove that their sum $z = z_1 + \dots + z_n$ has intensity moments

$$\langle I^m \rangle = \sum_{k_1=0}^m \sum_{k_2=0}^{k_1} \dots \sum_{k_{n-1}=0}^{k_{n-2}} \left(\frac{m!}{(m-k_1)!(k_1-k_2)! \dots (k_{n-2}-k_{n-1})!k_{n-1}!} \right)^2 \times \langle I_1^{m-k_1} \rangle \langle I_2^{k_1-k_2} \rangle \dots \langle I_{n-1}^{k_{n-2}-k_{n-1}} \rangle \langle I_n^{k_{n-1}} \rangle \quad (\text{D.1})$$

where $I = |z|^2$ and $I_k = |z_k|^2$. The proof will be given by induction on n .

First we will prove (D.1) for $n = 2$. The intensity I of the sum $z = z_1 + z_2$ is given by

$$I = zz^* = I_1 + \alpha + \alpha^* + I_2$$

where $\alpha = z_1 z_2^*$. The m 'th moment of the intensity may then be expressed

$$\begin{aligned} \langle I^m \rangle &= \langle \langle (I_1 + I_2) + (\alpha + \alpha^*) \rangle^m \rangle = \sum_{k=0}^m \binom{m}{k} \langle (I_1 + I_2)^{m-k} (\alpha + \alpha^*)^k \rangle \\ &= \sum_{k=0}^m \sum_{j=0}^{m-k} \sum_{i=0}^k \binom{m}{k} \binom{m-k}{j} \binom{k}{i} \langle I_1^{m-k-j} I_2^j \alpha^{k-i} (\alpha^*)^i \rangle \end{aligned}$$

The ensemble average term in this summation is only non-zero when the phase of α^{k-i} and $(\alpha^*)^i$ cancel. This will only occur when $k = 2i$. Noting that $\alpha\alpha^* = I_1 I_2$ and rearranging factorial terms the m 'th moment becomes

$$\langle I^m \rangle = \sum_{i=0}^{m/2} \sum_{j=0}^{m-2i} \binom{m}{i+j} \binom{m-i-j}{i} \binom{i+j}{i} \langle I_1^{m-i-j} \rangle \langle I_2^{j+i} \rangle$$

Letting $j = k_1 - i$ and reversing the order of summation gives

$$\langle I^m \rangle = \sum_{k_1=0}^m \sum_{i=0}^{\min(k_1, m-k_1)} \binom{m}{k_1} \binom{m-k_1}{i} \binom{k_1}{i} \langle I_1^{m-k_1} \rangle \langle I_2^{k_1} \rangle$$

which using the identity [1] (24.1.1)

$$\binom{m}{k_1} = \sum_{i=0}^{\min(k_1, m-k_1)} \binom{m-k_1}{i} \binom{k_1}{i}$$

equals

$$\langle I^m \rangle = \sum_{k_1=0}^m \binom{m}{k_1}^2 \langle I_1^{m-k_1} \rangle \langle I_2^{k_1} \rangle \quad (\text{D.2})$$

This proves (D.1) for $n = 2$.

Assuming (D.1) is true for $n \leq p$ we will prove it is true for $n = p + 1$. Let $z_a = z_1$ and $z_b = z_2 + \dots + z_{p+1}$ where z_1, \dots, z_{p+1} are independent circularly symmetric complex random variables. From (D.2) $z = z_a + z_b = z_1 + \dots + z_{p+1}$ has intensity moments

$$\langle I^m \rangle = \sum_{k_1=0}^m \binom{m}{k_1}^2 \langle I_a^{m-k_1} \rangle \langle I_b^{k_1} \rangle$$

Using (D.1) to express $\langle I_b^{k_1} \rangle$ in terms of the moments of I_2, \dots, I_{p+1} gives

$$\langle I^m \rangle = \sum_{k_1=0}^m \binom{m}{k_1}^2 \langle I_1^{m-k_1} \rangle \sum_{k_2=0}^{k_1} \dots \sum_{k_n=0}^{k_{n-1}} \left(\frac{k_1!}{(k_1 - k_2)! \dots (k_{p-1} - k_p)! k_p!} \right)^2 \times \\ \langle I_2^{k_1 - k_2} \rangle \dots \langle I_p^{k_{p-1} - k_p} \rangle \langle I_{p+1}^{k_p} \rangle$$

Cancelling factorial terms gives the desired expression for the moments

$$\langle I^m \rangle = \sum_{k_1=0}^m \sum_{k_2=0}^{k_1} \dots \sum_{k_n=0}^{k_{n-1}} \left(\frac{m!}{(m - k_1)! (k_1 - k_2)! \dots (k_{p-1} - k_p)! k_p!} \right)^2 \times \\ \langle I_1^{m-k_1} \rangle \langle I_2^{k_1 - k_2} \rangle \dots \langle I_p^{k_{p-1} - k_p} \rangle \langle I_{p+1}^{k_p} \rangle$$

This complete the inductive proof of (D.1).

Appendix E

Selection of map data

In this appendix we discuss the selection of digital map data suitable for matching to features in SAR images. Binary templates representing features to be matched can be generated from digital map data in either vector or raster format. Generally it is easier to produce digital map data in raster as opposed to vector format. Given that SAR data is intrinsically in raster format, matching and collectively manipulating map and SAR data would be simpler if the map data was in raster as opposed to vector format. However, this does not preclude the use of vector data, as it can easily be scan-converted into raster data at any resolution when required (the opposite process of raster to vector conversion is not simple). The advantages and disadvantages of using vector data need to be compared with those of using raster data.

Generally map data takes up less storage when it is in vector as opposed to raster format; this depends on the scanning resolution. A certain degree of 'intelligence' can be added to vector data by labelling each vector with information on what type of feature it represents. Vector data of this type can be selectively scan-converted with only features of interest being converted, that is map features which represent the actual shape of objects; lines of communication, water features and vegetation boundaries. Abstract and symbolic features (text, symbols, grid lines and contours), may be useful at a later stage for interpretation but they would only confuse the registration process; it is best if such features are not present in the map data to be used for registration. In addition, it is possible to control how different types of features are represented in the scan-converted raster map, for example, a road can be represented by a single line or two parallel lines of any desired thickness. Thus, using 'intelligent' vector data it is possible to produce raster data with smooth lines of desired thickness and with only features of interest present. This is in contrast to directly scanning printed maps: all features on the printed map large enough to be detected by scanning, including imperfections such as small artefacts and noisy linework, will appear on the raster map. When using the latter the presence of text, symbols, grids, contours, dashed lines and imperfections in the data needs to be remembered.

To normalise and to clean up the linework on directly scanned raster data all

lines may be thinned to being a single pixel wide¹. Small artefacts, that is small noise features present in the original map or caused by dust when scanning, may be removed by deleting all strings of connected pixels made up of less than a set number of pixels. A similar method could be used to remove text and symbols. However, it is likely that interesting features represented by dashed lines would also be removed (a technique for converting dashed lines into solid lines would be useful). Given that the position of grid lines is well defined it should be possible to remove them. As contours cannot be distinguished from lines representing features of interest it is best to scan map data without contours.

Prior to matching, information about the projection and scale of the map data and known distortions in the SAR data arising from the imaging geometry, is used to transform the map data. Once transformed in this manner, map templates can be matched to image features without having to take scale and rotation into consideration. In addition this information can be used to limit the area of the image over which a match is searched for. Known distortions in spaceborne and airborne SAR are described in [9] and [45] respectively. The registration process is used to correct any unknown distortions in the SAR data. Distortions can arise from atmospheric and topographic effects and sensor motion. By transforming the map data and not the image as in geocoding, difficulties arising from interpolation and sampling in SAR data [50] are avoided. This rough initial registration provides a framework from which to start the main task of registration.

It is desirable that the map data to be used for interpretation and change detection be of large enough scale that the smallest detectable image features could be accurately represented in the map. Within Britain, 1:10,000 Ordnance Survey mapping scanned at 200dpi (ground resolution 1.25m per pixel), is more than adequate. The digital map data referred to in Chapter 5 is of this type.

¹A thinning algorithm which ensures that connectedness is conserved and that the thinned line lies along the middle of the original line should be used; a morphological algorithm meeting these criteria is described in [32].

Appendix F

Glossary

This glossary briefly defines various mathematical functions and relationships referred to in the main body of the thesis. For further details on topics listed see; [1] for mathematical functions, [7] for relationships involving the Fourier transform, and [11, 48] for the statistics of random variables and processes. For ease of reference topics listed in the glossary are printed in **bold** in the main body of the thesis wherever the reader may require further explication of the topic (e.g., where it is first introduced), and when they are mentioned in the glossary under topic headings other than their own.

Autocorrelation and Autocovariance

The autocorrelation function (ACF) $R(t_1, t_2)$ of a random process $x(t)$ at t_1 and t_2 is defined by

$$R(t_1, t_2) = \langle x(t_1) x(t_2) \rangle$$

and the autocovariance $C(t_1, t_2)$ of $x(t)$ at t_1 and t_2 is defined by

$$C(t_1, t_2) = \langle x(t_1) x(t_2) \rangle - \langle x(t_1) \rangle \langle x(t_2) \rangle$$

The normalised autocorrelation function $r(t_1, t_2)$ at t_1 and t_2 is derived by dividing the autocorrelation function at t_1 and t_2 by the expected values of $x(t)$ at t_1 and t_2

$$r(t_1, t_2) = \frac{R(t_1, t_2)}{\langle x(t_1) \rangle \langle x(t_2) \rangle} = \frac{\langle x(t_1) x(t_2) \rangle}{\langle x(t_1) \rangle \langle x(t_2) \rangle}$$

The normalised autocovariance $c(t)$ is defined likewise

$$c(t_1, t_2) = \frac{C(t_1, t_2)}{\langle x(t_1) \rangle \langle x(t_2) \rangle} = \frac{\langle x(t_1) x(t_2) \rangle}{\langle x(t_1) \rangle \langle x(t_2) \rangle} - 1$$

The correlation coefficient is given by dividing the autocovariance by the variance.

When $t_1 = t_2 = t_0$ the above functions generate the following first order statistics of $x(t)$ at t_0 . The autocorrelation function gives its second moment

$$R(t_0, t_0) = \langle x(t_0)^2 \rangle$$

the autocovariance gives its variance

$$C(t_0, t_0) = \langle x(t_0)^2 \rangle - \langle x(t_0) \rangle^2 = \sigma_{x(t_0)}^2$$

the normalised autocorrelation function gives its second normalised moment

$$r(t_0, t_0) = \frac{\langle x(t_0)^2 \rangle}{\langle x(t_0) \rangle^2} = x(t_0)^{(2)}$$

and the normalised autocovariance gives its coefficient of variation

$$c(t_0, t_0) = \frac{\langle x(t_0)^2 \rangle - \langle x(t_0) \rangle^2}{\langle x(t_0) \rangle^2} = v_{x(t_0)}^2$$

When the random process $x(t)$ is stationary its autocorrelation function only depends on $\tau = t_1 - t_2$, i.e.

$$R(\tau) = R(t_1, t_2) = R(t_3, t_4)$$

where $\tau = t_1 - t_2 = t_3 - t_4$.

Bessel functions

Bessel functions are solutions to the differential equations [1]

$$z^2 \frac{d^2 w}{dz^2} + z \frac{dw}{dz} + (az^2 - v^2) w = 0$$

where $a = \pm 1$. When $a = 1$ solutions include the Bessel function of the first kind $J_\nu(z)$, and when $a = -1$ solutions include the modified Bessel function of the second kind $K_\nu(z)$. Here we list the main properties of Bessel functions used in preceding chapters.

When $\nu = 0$ the Bessel function of the first kind has an integral representation [1]-(9.1.21)

$$J_0(z) = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{jz \cos \theta} d\theta$$

and a series expansion [1]-(9.1.10)

$$J_0(z) = \sum_{k=0}^{\infty} \frac{(-1)^k z^{2k}}{4^k (k!)^2} \quad (\text{F.1})$$

The modified Bessel function of the second kind has an integral representation [1]-(9.6.24)

$$K_\nu(z) = \int_0^{\infty} e^{-z \cosh t} \cosh(\nu t) dt \quad (\text{F.2})$$

where $|\angle z| < \pi/2$. The first derivative of $K_\nu(z)$ with respect to z is given by [1]-(9.6.26)

$$K'_{\nu-1}(z) = -\left(\frac{K_\nu(z) + K_{\nu-2}(z)}{2}\right) \quad (\text{F.3})$$

Useful integrals involving $K_\nu(z)$ include [1]-(11.3.27)

$$\int_0^z t^\nu K_{\nu-1}(t) dt = -z^\nu K_\nu(z) + 2^{\nu-1} \Gamma(\nu) \quad (\text{F.4})$$

and [1]-(11.4.22)

$$\int_0^\infty t^m K_{\nu-1}(t) dt = 2^{m-1} \Gamma\left(\frac{m+\nu}{2}\right) \Gamma\left(\frac{m-\nu+2}{2}\right) \quad (\text{F.5})$$

When $\nu = n - 1/2$, where n is an integer, $K_\nu(z)$ has a series expansion given by [1]-(10.2.15)

$$K_{n-\frac{1}{2}}(z) = \sqrt{\frac{\pi}{2z}} e^{-z} \sum_{k=0}^{n-1} \frac{\Gamma(n+k)}{\Gamma(n-k)\Gamma(k+1)} (2z)^{-k} \quad (\text{F.6})$$

Central limit theorem

The central limit theorem states that under certain conditions the sum $y = x_1 + \dots + x_n$ of n independent continuous random variables tends to being Gaussian distributed with mean $\mu_y = \mu_1 + \dots + \mu_n$ and variance $\sigma_y^2 = \sigma_1^2 + \dots + \sigma_n^2$ as n increases.

Sufficient conditions for the central limit theorem to apply are [48]:

1. The variance of the sum must tend to infinity as the number of variables tends to infinity.
2. For some number $m > 2$ the m th moment of all the variables must be finite.

Characteristic function

The characteristic function $\Phi_x(\omega)$ of a random variable x is given by the Fourier transform of its probability density function $f_x(x)$

$$\Phi_x(\omega) = \int_{-\infty}^{\infty} f_x(x) e^{j\omega x} dx = \langle e^{j\omega x} \rangle$$

and the inverse Fourier transform of the characteristic function gives the pdf

$$f_x(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \Phi_x(\omega) e^{-j\omega x} d\omega$$

Substituting for $e^{j\omega x}$ by its series expansion given by [1]-(4.2.1)

$$e^{j\omega x} = \sum_{k=0}^{\infty} \frac{(j\omega x)^k}{k!} \quad (\text{F.7})$$

the characteristic function of x is related to its moments by

$$\Phi_x(\omega) = \sum_{k=0}^{\infty} \frac{(j\omega)^k}{k!} \langle x^k \rangle$$

Thus the m th moment of x can be generated from its characteristic function by taking the m th derivative of its characteristic function

$$\frac{d^m \Phi_x(\omega)}{d\omega^m} = \sum_{k=0}^{\infty} \frac{j^k \omega^k}{k!} \langle x^{k+m} \rangle$$

and setting $\omega = 0$. The m th moment is then given by

$$\langle x^m \rangle = j^{-m} \frac{d^m \Phi_x(0)}{d\omega^m}$$

Circular symmetry

A bivariate function is termed to be circularly symmetric if its value only depends on the radial distance from the origin [48], i.e. the function $f(x, y)$ is circularly symmetric if

$$f(x_1, y_1) = f_x(x_2, y_2)$$

whenever

$$x_1^2 + y_1^2 = x_2^2 + y_2^2$$

A complex random variable is said to be circularly symmetric if its probability density function is circularly symmetric; this will be the case if and only if the phase and amplitude are independent. The phase will then be uniformly distributed over the interval $(-\pi, \pi]$. The real and imaginary components of a circularly symmetric random variable are uncorrelated but they can only be independent if and only if they are Gaussian, i.e. if the amplitude is Rayleigh distributed.

Coefficient of variation

The coefficient of variation of a random variable x with mean μ_x and variance σ_x^2 , is a measure of the width of its distribution relative to its mean value. It is given by

$$v_x^2 = \frac{\sigma_x^2}{\mu_x^2}$$

The square root of the coefficient of variation, i.e., standard deviation/mean, is normally used to measure contrast in an image degraded by multiplicative noise.

Convolution theorem

Given two functions $f_1(x)$ and $f_2(x)$ with Fourier transforms $F_1(\omega)$ and $F_2(\omega)$ respectively, the convolution theorem [48] states that their convolution $f(x) = f_1(x) * f_2(x)$ defined by

$$f(x) = \int_{-\infty}^{\infty} f_1(y) f_2(x-y) dy$$

has a Fourier transform $F(\omega)$ equal to the product of the Fourier transforms of the two functions

$$F(\omega) = F_1(\omega) F_2(\omega)$$

It follows from the convolution theorem that the characteristic function of the sum of two independent random variables will be given by the product of their characteristic functions. This is because the probability density function of the sum of two independent random variables is given by the convolution of their probability density functions (A.1), and because the characteristic function of a random variable is just the Fourier transform of its probability density function.

This result can be extended to the sum y of n independent random variables x_i , with the characteristic function of their sum $\Phi_y(\omega)$ being given by the product of the characteristic functions $\Phi_{x_i}(\omega)$ of the variables

$$\Phi_y(\omega) = \langle e^{j\omega y} \rangle = \langle e^{j\omega(x_1 + \dots + x_n)} \rangle = \langle e^{j\omega x_1} \rangle \dots \langle e^{j\omega x_n} \rangle = \prod_{i=1}^n \Phi_{x_i}(\omega)$$

Cumulative distribution function

The probability that a random variable x is less than y is given by its cumulative distribution function (cdf)

$$\text{prob}(x < y) = p_x(y)$$

The cdf of x is derived from its probability density function by integrating the latter from $-\infty$ to y

$$p_x(y) = \int_{-\infty}^y f_x(x) dx$$

Digamma function

The Digamma function [1] is defined

$$\psi(z) = \frac{d}{dz} \ln \Gamma(z) = \frac{\Gamma'(z)}{\Gamma(z)}$$

where $\Gamma(z)$ is the gamma function. For positive integer values it is given by

$$\psi(n) = -\gamma_E + \sum_{k=1}^{n-1} \frac{1}{k}$$

when $n \geq 2$, and

$$\psi(1) = -\gamma_E$$

where $\gamma_E = 0.57722\dots$ is Euler's constant. As $n \rightarrow \infty$

$$\psi(n) \rightarrow \ln n$$

The first derivative of the Digamma function $\psi'(z)$ is known as the Trigamma function. For positive integer values it is given by

$$\psi'(n) = \frac{\pi^2}{6} - \sum_{k=1}^{n-1} \frac{1}{k^2}$$

when $n \geq 2$, and

$$\psi'(1) = \frac{\pi^2}{6}$$

As $n \rightarrow \infty$

$$\psi'(n) \rightarrow 0$$

Ensemble average

The ensemble average (or expected value) of a function $u(x)$ of a continuous random variable x , is indicated by the operator $\langle \cdot \rangle$, and is defined

$$\langle u(x) \rangle = \int_{-\infty}^{\infty} u(x) f_x(x) dx$$

where $f_x(x)$ is the probability density function of x .

In the case of a discrete random variable n the ensemble average

$$\langle u(n) \rangle = \sum_{n=-\infty}^{\infty} u(n) f_n$$

where f_n is the probability density function of n .

Estimation

Given a set of random variables x_1, \dots, x_n described by some unknown parameter α , an estimate $\hat{\alpha}$ of α based on the observations x_1, \dots, x_n is said to be unbiased if

$$\langle \hat{\alpha} \rangle = \alpha$$

and the estimate is said to be consistent if

$$\hat{\alpha} \rightarrow \alpha$$

as $n \rightarrow \infty$.

Fourier transform

The version of the Fourier transform $F(\omega)$ of a function $f(x)$ used in this thesis is defined

$$F(\omega) = \int_{-\infty}^{\infty} f(x) e^{-j\omega x} dx$$

and the inverse Fourier transform is defined

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} F(\omega) e^{j\omega x} d\omega$$

Gamma function

The gamma function $\Gamma(z)$ is defined

$$\Gamma(z) = \int_0^{\infty} e^{-x} x^{z-1} dx$$

and has the property

$$\Gamma(z+1) = z\Gamma(z)$$

For positive integer values n

$$\Gamma(n+1) = n! = n \cdot (n-1) \dots 3 \cdot 2 \cdot 1$$

and for $z = \frac{1}{2}$

$$\Gamma\left(\frac{1}{2}\right) = \sqrt{\pi}$$

Gaussian distribution

A Gaussian (or normal) distributed random variable x with mean μ and variance σ^2 has probability density function

$$f_x(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (\text{F.8})$$

and central moments (i.e., with mean normalised to zero)

$$\langle (x - \mu)^m \rangle = \begin{cases} \Gamma\left(\frac{m+1}{2}\right) / \Gamma\left(\frac{1}{2}\right) (\sqrt{2}\sigma)^m & m \text{ even} \\ 0 & m \text{ odd} \end{cases} \quad (\text{F.9})$$

Hankel transform

Given a circularly symmetric complex random variable $z = x + jy$ with joint probability density function $f_z(x, y)$ its characteristic function given by

$$\Phi_z(u, v) = \langle e^{j\omega \cdot z} \rangle = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_z(x, y) e^{j(ux+vy)} dx dy$$

where $w = u + jv$, is also circularly symmetric, i.e., the characteristic function only depends on the amplitudes $a = |z|$ and $b = |\omega|$. Such a characteristic function can be expressed in terms of the Hankel transform [7]

$$\Phi_z(b) = \int_0^{\infty} f_a(a) J_0(ab) da = \langle J_0(ab) \rangle \quad (\text{F.10})$$

where $J_0(ab)$ is the zeroth-order Bessel function of the first kind and $f_a(a)$ is the probability density function of the amplitude a and is related to the joint probability density function of z by

$$f_a(a) = 2\pi a f_z(a) = 2\pi a f_z(a \cos \theta, a \sin \theta)$$

The inverse Hankel transform is defined

$$f_z(a) = \frac{1}{2\pi} \int_0^\infty \Phi_z(b) b J_0(ab) db \quad (\text{F.11})$$

Thus the amplitude has distribution

$$f_a(a) = a \int_0^\infty \Phi_z(b) b J_0(ab) db$$

Substituting the series expansion of $J_0(ab)$ given by (F.1) into (F.10) the intensity moments may be derived from the characteristic function of z using

$$\langle a^{2m} \rangle = \frac{(-4)^m (m!)^2}{(2m)!} \Phi_z^{(2m)}(0) \quad (\text{F.12})$$

where $\Phi_z^{(m)}(b)$ is the m th derivative of $\Phi_z(b)$ with respect to b

Independence and Correlation

Two random variables x and y are statistically independent if and only if their joint pdf $f_{x,y}(x,y)$ is equal to the product of their marginal pdf's $f_x(x)$ and $f_y(y)$

$$f_{x,y}(x,y) = f_x(x) f_y(y)$$

this is equivalent to the condition that for all moments

$$\langle x^m y^n \rangle = \langle x^m \rangle \langle y^n \rangle$$

When the above condition holds for the first moment (mean values)

$$\langle xy \rangle = \langle x \rangle \langle y \rangle$$

the random variables are said to be statistically uncorrelated. Thus independent random variables will also be uncorrelated; the reverse is only necessarily true if both random variables have **Gaussian distributions**.

Jacobian

Given two sets of n random variables $x = x_1, \dots, x_n$ and $y = y_1, \dots, y_n$ related by the transform $x_i = g_i(y)$, their joint pdf's $f_x(x)$ and $f_y(y)$ respectively are related by

$$f_y(y) = f_x(g_1(y), \dots, g_n(y)) |J|$$

where J , the Jacobian of the transform $g(y)$ is given by

$$J = \frac{\partial x}{\partial y} = \begin{vmatrix} \frac{\partial g_1}{\partial y_1} & \dots & \frac{\partial g_n}{\partial y_1} \\ \vdots & & \vdots \\ \frac{\partial g_1}{\partial y_n} & \dots & \frac{\partial g_n}{\partial y_n} \end{vmatrix}$$

Joint and marginal distributions

The joint probability that $x_1 < y_1, \dots, x_n < y_n$ is given by the joint cumulative distribution function $p_x(y)$, where $x = x_1, \dots, x_n$ and $y = y_1, \dots, y_n$. The joint probability density function of x is defined

$$f_x(y) = \frac{\partial^n p_x(y)}{\partial y_1 \dots \partial y_n}$$

Thus

$$p_x(y) = \int_{-\infty}^{y_1} \dots \int_{-\infty}^{y_n} f_x(x) dx_1 \dots dx_n$$

The marginal pdf of x_i is derived from the joint pdf of x by integrating it over $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n$

$$f_{x_i}(x_i) = \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} f_x(x) dx_1 \dots dx_{i-1} dx_{i+1} \dots dx_n$$

Kolmogorov-Smirnov test of fit

The Kolmogorov-Smirnov test of fit [33, 49, 56] is used to determine the degree of fit between the distribution of a set of data and a theoretic distribution. This is done by measuring the maximum absolute deviation (the L_∞ norm) between the cumulative frequency histogram of the data and the cumulative distribution function of the theoretic distribution. The distribution of the maximum absolute deviation d depends only on the number of independent data samples n . The degree of fit is given by the probability

$$p(n, d) = 2 \sum_{k=1}^{\infty} (-1)^{k-1} e^{-2k^2 nd^2}$$

As the deviation tends to zero the degree of fit tends to one which represents a good fit, and for large n as the deviation tends to one the degree of fit tends to zero which represents a poor fit. Figure F.1 shows a plot of the degree of fit versus \sqrt{nd} .

Matched filter

Given an input $r(x)$ consisting of a known signal $s(x)$ in additive noise $n(x)$

$$r(x) = s(x) + n(x)$$

the matched filter of the signal is defined as the filter $h(x)$ which maximises the signal to noise ratio of the output $r(x) * h(x)$ at $x = x_0$. When the noise has a constant power spectrum (i.e., white noise) the matched filter is given by $h(x) = s^*(x_0 - x)$.

Maximum likelihood estimate

Let x_1, \dots, x_n be a set of random variables whose joint pdf $F(x_1, \dots, x_n|y)$ depends on an unknown parameter y . Given a particular value of y the

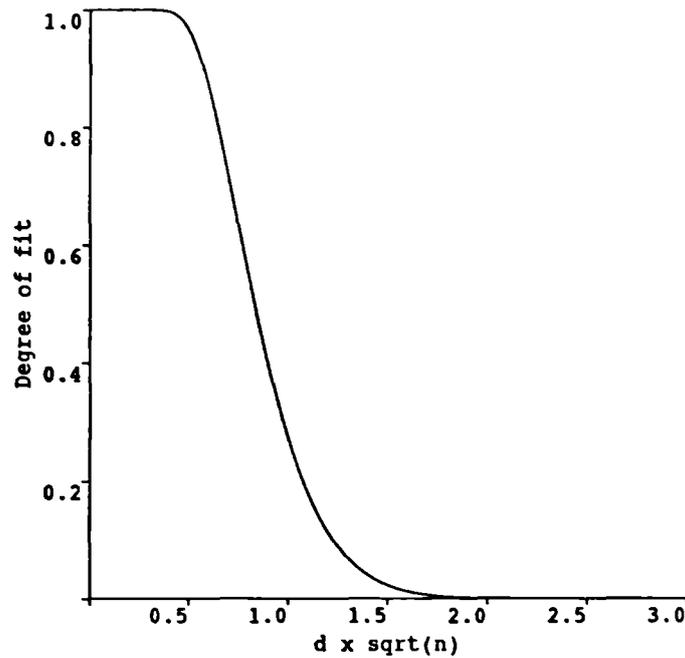


Figure F.1: Relation between maximum absolute deviation d , number of independent samples n and degree of fit $p(n, d)$.

joint pdf defines the likelihood of x_1, \dots, x_n then occurring and the maximum likelihood (ML) estimate of y is that value \hat{y} which maximises the joint pdf. Assuming the joint pdf is uni-modal the maximum likelihood estimate of y is found by determining where the first derivative of the likelihood with respect to y equals zero. As the maximum of the log of a positive function will occur at the same position as the maximum of the function ($\log(x)$ is a strictly monotonic increasing function for positive real x), the maximum likelihood estimate is often more easily found by differentiating the log-likelihood $L = \ln F$ with respect to y . When the random variables are independent this involves finding where the first derivative of the sum of the logs of the marginal pdf's of the x_i 's, with respect to y , equals zero.

Mean

The mean value μ_x (or first moment) of a random variable x gives its expected value and is defined by its ensemble average

$$\mu_x = \langle x \rangle$$

Moments

The m th moment of a random variable x is defined

$$\langle x^m \rangle = \int_{-\infty}^{\infty} x^m f_x(x) dx$$

where $f_x(x)$ is the probability density function of x . The moments of a random variable can also be generated from its characteristic function.

Normalised moments

The m th normalised moment of x is given by dividing its m th moment by its mean value raised to the power m

$$x^{(m)} = \frac{\langle x^m \rangle}{\langle x \rangle^m}$$

Power spectrum

The power spectrum, or power spectral density, of a stationary process $x(t)$ is defined as the Fourier transform of its autocorrelation function. As the autocorrelation function of a stationary process is defined by the convolution $x(t) * x^*(-t)$, from the convolution theorem the power spectrum of a stationary process is also given by the square of the magnitude of its Fourier transform.

Probability density function

The probability density function (pdf) of x is given by differentiating its cumulative density function

$$f_x(x) = \frac{dp_x(x)}{dx}$$

Stationarity

A random process $x(t)$ is considered to be strict-sense stationary if its statistical properties are invariant to a change of origin. The process is considered to be wide-sense stationary if the mean is constant and its autocorrelation function only depends on $t_1 - t_2$.

Variance and standard deviation

The variance σ_x^2 of a random variable x gives a measure of the width of its distribution and is defined

$$\sigma_x^2 = \langle (x - \mu_x)^2 \rangle = \langle x^2 \rangle - \langle x \rangle^2$$

where $\mu_x = \langle x \rangle$ is the mean value of x .

σ_x the square-root of the variance, gives the standard deviation of x .

Bibliography

- [1] M. Abramowitz and I. Stegun. *Handbook of Mathematical Functions*. Dover, New York, 1965.
- [2] H. H. Arsenault and G. April. Properties of speckle intergrated with a finite aperture and logrithmically transformed. *J. Opt. Soc. Am.*, 66:1160–1163, 1976.
- [3] M. S. Bartlett. *An introduction to stochastic processes*. Cambridge University Press, 2nd edition, 1966.
- [4] D. Blacknell, A. Freeman, S. Quegan, I. A. Ward, I. P. Finley, C. J. Oliver, R. G. White, and J. W. Wood. Geometric accuracy in airborne SAR images. *IEEE Trans. Aerosp. Electron. Syst.*, 25:241–258, 1989.
- [5] M. Born and E. Wolf. *Principles of Optics*. Pergamon, Oxford, 1980.
- [6] A. C. Bovik. On detecting edges in speckle imagery. *IEEE Trans. Acoust., Speech, Signal Proc.*, 36:1618–1627, 1988.
- [7] R. N. Bracewell. *The Fourier Transform and its Applications*. McGraw-Hill, New York, 1986.
- [8] R. Cook and J. Sandys-Renton. A parallel segmentation algorithm (merge using moments) for SAR images. *Applications of transputers*, 1:311–316, 1991.
- [9] J. C. Curlander, R. Kwok, and S. Pang. A post-processing system for automated rectification and registration of SAR images. *Int. J. Remote Sensing*, 8:621–638, 1987.
- [10] J. C. Curlander and R. N. McDonough. *Synthetic Aperture Radar: Systems and Signal Processing*. Wiley-Interscience, New York, 1991.
- [11] W. B. Davenport and W. L. Root. *An Introduction to the Theory of Random Signals and Noise*. IEEE Press, New York, 1987.
- [12] A. Della Ventura, A. Rampini, and R. Schettini. Image registration by recognition of corresponding structures. *IEEE Trans. Geosci. Remote Sens.*, 28:305–314, 1990.

- [13] L. M. Delves, R. Wilkinson, C. J. Oliver, and R. G. White. Comparing the performance of SAR image segmentation algorithms. *Int. J. Remote Sensing*, 13:2121–2149, 1992.
- [14] R. O. Duda and P. E. Hart. *Pattern classification and scene analysis*. Wiley-Interscience, New York, 1973.
- [15] C. Elachi. *Spaceborne Radar Remote Sensing: Applications and Techniques*. IEEE Press, New York, 1988.
- [16] European Space Agency. ERS-1 user handbook, 1992. ESA SP-1148.
- [17] R. L. Fante. Detection of multiscatter targets in K-distributed clutter. *IEEE Trans. Antennas and Propogat.*, 32:1358–1363, 1984.
- [18] V. S. Frost, K. S. Shaumugan, and J. C. Holtzman. Edge detection for synthetic aperture radar and other noisy images. In *Proc. IGARSS'82*, pages FA-2, 4.1–4.9, (Munich, W.Germany), 1982.
- [19] V. S. Frost, J. A. Stiles, K. S. Shanmugan, and J. C. Holtzmann. A model for radar images and its application to adaptive filtering of multiplicative noise. *IEEE Trans. Pattern Anal. Machine Intell.*, 4:157–166, 1982.
- [20] S. C. Geiss. Edge detection in SAR imagery using gradient operators. Technical Report 3743, Roy. Signals and Radar Estab., Malvern U.K., 1984.
- [21] S. C. Geiss. Aspects of SAR image segmentation. Technical Report 3789, Roy. Signals and Radar Estab., Malvern U.K., 1985.
- [22] J. W. Goodman. Statistical properties of laser speckle patterns. In J. C. Dainty, editor, *Laser speckle and related phenomena*, pages 9–75. Springer, New York, 1975.
- [23] A. Goshtasby. Registration of images with geometric distortions. *IEEE Trans. Geosci. Remote Sens.*, 26:60–64, 1988.
- [24] A. L. Gray, R. K. Hawkins, C. E. Livingstone, J. Campbell, E. Attema, and J. C. Morin. Towards a calibration of the CCRS airborne SAR's. In *Proc. IGARSS'88*, pages 707–709, (Edinburgh, Scotland), 1988.
- [25] A. Hendry, S. Quegan, and J. Wood. The visibility of linear features in SAR images. In *Proc. IGARSS'88*, pages 1517–1520, (Edinburgh, Scotland), 1988.
- [26] A. Hendry, J. Skingley, and A. J. Rye. Automated linear feature detection and its application to curve location in synthetic aperture radar. In *Proc. IGARSS'88*, pages 1521–1524, (Edinburgh, Scotland), 1988.
- [27] Hunting Technical Services and Marconi Research Centre. Study of land feature extraction from SAR images, 1984. Final report on ESA contract 5855/84/GP1(SC).

- [28] A. K. Jain. *Fundamentals of Digital Image Processing*. Prentice-Hall, Englewood Cliffs, NJ, 1989.
- [29] E. Jakeman. On the statistics of K-distributed noise. *J. Phys. A: Math. Gen.*, 13:31–48, 1980.
- [30] E. Jakeman. Statistics of integrated gamma-lorentzian intensity fluctuations. *Optica Acta*, 27:735–741, 1980.
- [31] E. Jakeman and P. N. Pusey. A model for non-Rayleigh sea echo. *IEEE Trans. Antennas and Propogat.*, 24:806–814, 1976.
- [32] B. K. Jang and R. T. Chin. Analysis of thinning algorithms using mathematical morphology. *IEEE Trans. Pattern Anal. Machine Intell.*, 12:541–551, 1990.
- [33] M. G. Kendall and A. Stuart. *The advanced theory of statistics*. Griffen, London, 1946.
- [34] S. Kingsley and S. Quegan. *Understanding Radar*. McGraw-Hill, New York, 1992.
- [35] H. Laur and G. M. Doherty. ERS-1 SAR calibration: history and results. In *CEOS workshop on SAR calibration*, (Ottawa, Canada), 1992.
- [36] T. Le Toan et al. Observation and modelling of pine forest responses and mapping of forest and land use types. In *Proc. MAESTRO Data Quality Workshop*, pages 183–188, (Noordwijk, The Netherlands), 1992.
- [37] F. W. Leberl. *Radargrammetric Image Processing*. Artech House, Norwood, MA, 1990.
- [38] S. N. Madsen. *Speckle theory modelling, analysis and applications related to synthetic aperture radar*. PhD thesis, Electromagnetics Institute, Technical University of Denmark, 1986.
- [39] P. Maragos and R. W. Schafer. Morphological systems for multidimensional signal processing. *Proc. IEEE*, 78:690–709, 1990.
- [40] D. Marr and E. Hildreth. Theory of edge detection. *Proc. Roy. Soc. London B*, 207:187–217, 1980.
- [41] G. Medioni and R. Nevatia. Matching images using linear features. *IEEE Trans. Pattern Anal. Machine Intell.*, 6:675–685, 1984.
- [42] C. J. Oliver. An analysis of template matching in image registration. *Optica Acta*, 31:233–248, 1984.
- [43] C. J. Oliver. A model for non-Rayleigh scattering statistics. *Optica Acta*, 31:701–722, 1984.

- [44] C. J. Oliver. The interpretation and simulation of clutter textures in coherent images. *Inverse Problems*, 2:481–518, 1986.
- [45] C. J. Oliver. Synthetic aperture radar imaging. *J. Phys. D: Appl. Phys.*, 22:871–890, 1989.
- [46] C. J. Oliver. Information from SAR images. *J. Phys. D: Appl. Phys.*, 24:1493–1514, 1991.
- [47] C. J. Oliver. Optimum texture estimators for SAR clutter. Personal communication (to be published in *J. Phys. D.*), 1993.
- [48] A. Papoulis. *Probability, Random Variables and Stochastic Process*. McGraw-Hill, New York, 1991.
- [49] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical recipes in C*. Cambridge University Press, 1988.
- [50] S. Quegan. Interpolation and sampling in SAR images. *IEEE Trans. Geosci. Remote Sens.*, 28:641–646, 1990.
- [51] S. Quegan, R. G. Caves, and P. Harley. Automated knowledge-based segmentation of SAR images. In P. Mather, editor, *TERRA-1: Understanding the terrestrial environment*, pages 221–232. Taylor and Francis, London, 1992.
- [52] S. Quegan, L. V. Dutra, and R. G. Caves. Does phase carry information in quadpolarised SAR data. In *Proc. IGARSS'92*, pages 83–85, (Houston, USA), 1992.
- [53] S. Quegan, A. J. Rye, A. Hendry, J. Skingley, and C. J. Oddy. Automatic interpretation strategies for synthetic aperture radar images. *Phil. Trans. R. Soc. Lond. series A*, 324:409–421, 1988.
- [54] R. K. Raney. Considerations for SAR image quantification unique to orbital systems. *IEEE Trans. Geosci. Remote Sens.*, 29:754–760, 1991.
- [55] R. Samadani and J. F. Vesecky. Finding curvilinear features in speckled images. *IEEE Trans. Geosci. Remote Sens.*, 28:669–673, 1990.
- [56] S. Siegel and N. J. Castellan. *Nonparametric Statistics*. McGraw-Hill, New York, 1988.
- [57] R. Touzi, A. Lopes, and P. Bousquet. A statistical and geometrical edge detector for SAR images. *IEEE Trans. Geosci. Remote Sens.*, 26:764–773, 1988.
- [58] F. T. Ulaby, R. K. Moore, and A. K. Fung. *Microwave Remote Sensing*, volume 1. Addison Wesley, Reading, MA, 1981.
- [59] F. T. Ulaby, R. K. Moore, and A. K. Fung. *Microwave Remote Sensing*, volume 2. Addison Wesley, Reading, MA, 1982.

- [60] F. T. Ulaby, R. K. Moore, and A. K. Fung. *Microwave Remote Sensing*, volume 3. Artech House, Dedham, MA, 1986.
- [61] R. G. White. Low level segmentation of noisy imagery. Technical Report 3900, Roy. Signals and Radar Estab., Malvern U.K., 1986.
- [62] R. G. White. Change detection in SAR imagery. *Int. J. Remote Sensing*, 12:339–360, 1991.
- [63] A. D. Woode, Y. L. Desnos, and H. Jackson. The development and first results from the ESTEC ERS-1 active radar calibration unit. *IEEE Trans. Geosci. Remote Sens.*, 30:1122–1130, 1992.
- [64] S. H. Yueh, J. A. Kong, J. K. Jao, R. T. Shin, H. A. Zebker, T. L. Toan, and H. Ottl. K-distribution and polarimetric terrain radar clutter. In J. A. Kong, editor, *Progress in electromagnetic research: polarimetric remote sensing*, pages 237–275. Elsevier, New York, 1990.
- [65] H. A. Zebker and R. Goldstein. Topographic mapping from interferometric SAR observations. *J. Geophys. Res.*, 91(B5):4993–4999, 1986.