

COHERENT CHANGE DETECTION UNDER A FOREST CANOPY

A thesis submitted for the degree of Doctor of Philosophy

Paul Pincus

School of Electrical and Electronic Engineering
University of Adelaide



THE UNIVERSITY
of ADELAIDE

January 21, 2020

Contents

Abstract	iii
Declaration	iv
Acknowledgements	iv
List of publications	v
List of symbols	vi
1 Introduction	1
1.1 Background and motivation	1
1.2 Structure and novelty	4
1.2.1 Image formation for maximum coherence	4
1.2.2 3D SAR coherent change detection for monitoring the ground under a forest canopy	5
1.2.3 Polarimetric calibration of circularly polarised synthetic aperture radar data	7
1.3 Airborne radar data	7
2 SAR image formation for maximum coherence	11
2.1 Introduction	11
2.2 Collection geometry	16
2.3 Plane-wave approximation	19
2.4 Spatial-frequency support	20
2.5 Transmitted chirp waveform	23
2.6 Scattering Model	24
2.7 Received signal	24
2.8 Baseband digital signal	25
2.9 Matched filtering	25
2.10 Focusing via beamforming	29
2.11 Pulse support in the 3D spatial-frequency domain	32
2.12 The slope-dependent projection yielding the 2D aperture	37
2.13 Aperture synthesis and resolution	40
2.14 Image formation on undulating terrain	40
2.15 Image spectrum	42
2.16 Spatially variant aperture trim	46
2.17 Application to real data	52
2.18 Conclusion	53
2.A Appendix	59
2.A.1 Statistical performance of the coherence estimator	59
2.A.2 Coherence in the spatial-frequency domain	61
2.A.3 Point-spread function for a uniform frequency support	62
2.A.4 Windows	63
2.A.5 Affine transformations and the Fourier transform	64
2.A.6 Projection-slice theorem	65
2.A.7 Hessian normal form	66
2.A.8 Analogue demodulation	66
2.A.8.1 Mix-down to IF	67
2.A.8.2 Dechirp	68
2.A.9 Sampling	70
2.A.10 Digital demodulation	71
2.A.10.1 Mix-down to baseband	72
2.A.10.2 Rechirp	73
2.A.10.3 Mitigation of RF spectral degradation	73

3	3D SAR coherent change detection for monitoring the ground under a forest canopy	75
3.1	Introduction	75
3.1.1	Foliage penetration	78
3.1.2	Forest scattering models	81
3.1.3	Hidden target detection	82
3.1.4	3D SAR	84
3.1.5	Change detection	87
3.1.6	Problem Statement	88
3.2	Radar design considerations	88
3.2.1	Wavelength selection	89
3.2.2	Effective dynamic range of coherence	90
3.3	Layover and interferometry	93
3.4	3D SAR beamforming	98
3.4.1	The steering vector	99
3.4.2	Types of weight vectors	100
3.4.3	Vertical beampattern	103
3.4.4	Height-range point-spread function	106
3.4.5	Imaging of a random volume over ground	108
3.4.6	Coherent change detection of the ground under a random volume	113
3.5	Multichannel coherence	116
3.6	Dual-layer coherence	117
3.7	Multichannel dual-layer coherence	121
3.8	3D SAR beamforming for volume suppression	126
3.8.1	Optimal volume suppression	126
3.9	Random-volume-over-ground model of a forest	128
3.9.1	Random volume coherence	128
3.9.2	Random-volume-over-ground coherence	132
3.10	RVOG beamformer performance	134
3.10.1	Optimal RVOG volume suppression	135
3.10.2	Coherence estimation error	137
3.10.3	Sensitivity to assumed model parameters	138
3.10.4	Sensitivity to perturbations using the matrix condition number	140
3.10.5	Sensitivity to ground height	143
3.11	3D SAR CCD simulation	145
3.11.1	Receiver operating characteristics	153
3.12	3D SAR CCD simulation using PolSARproSim	156
3.12.1	Polarimetric filtering of ground-trunk double bounce	164
3.12.2	Receiver operating characteristics for PolSARproSim	171
3.13	Conclusion	173
4	Polarimetric calibration of circularly polarised SAR data	175
4.1	Preface	175
4.2	Introduction	175
4.3	Polarimetric distortion models	178
4.4	Polarimetric basis change	182
4.5	Calibration via corner reflectors	184
4.6	Calibration via clutter	186
4.7	Data analysis	192
4.8	Conclusion	195
4.A	Appendix	200
4.A.1	Eigendecomposition of polarimetric SAR imagery	200
5	Conclusion	203
5.1	Summary	203
5.2	Future Work	204
	Bibliography	205

Abstract

Coherent change detection (CCD) is an established technique for remotely monitoring landscapes with minimal vegetation or buildings. By evaluating the local complex correlation between a pair of synthetic aperture radar (SAR) images acquired on repeat passes of an airborne or spaceborne imaging radar system, a map of the scene coherence is obtained. Subtle disturbances of the ground are detected as areas of low coherence in the surface clutter.

This thesis investigates extending CCD to monitor the ground in a forest. It is formulated as a multichannel dual-layer coherence estimation problem, where the coherence of scattering from the ground is estimated after suppressing interference from the canopy by vertically beamforming multiple image channels acquired at slightly different grazing angles on each pass. This 3D SAR beamforming must preserve the phase of the ground response. The choice of operating wavelength is considered in terms of the trade-off between foliage penetration and change sensitivity. A framework for comparing the performance of different radar designs and beamforming algorithms, as well as assessing the sensitivity to error, is built around the random-volume-over-ground (RVOG) model of forest scattering. If the ground and volume scattering contributions in the received echo are of similar strength, it is shown that an L-band array of just three channels can provide enough volume attenuation to permit reasonable estimation of the ground coherence. The proposed method is demonstrated using an RVOG clutter simulation and a modified version of the physics-based SAR image simulator PolSARproSim. Receiver operating characteristics show that whilst ordinary single-channel CCD is unusable when a canopy is present, 3D SAR CCD permits reasonable detection performance. A novel polarimetric filtering algorithm is also proposed to remove contributions from the ground-trunk double-bounce scattering mechanism, which may mask changes on the ground near trees.

To enable this kind of polarimetric processing, fully polarimetric data must be acquired and calibrated. Motivated by an interim version of the *Ingara* airborne imaging radar, which used a pair of helical antennas to acquire circularly polarised data, techniques for the estimation of polarimetric distortion in the circular basis are investigated. It is shown that the standard approach to estimating cross-talk in the linear basis, whereby expressions for the distortion of reflection-symmetric clutter are linearised and solved, cannot be adapted to the circular basis, because the first-order effects of individual cross-talk parameters cannot be distinguished. An alternative approach is proposed that uses ordinary and gridded trihedral corner reflectors, and optionally dihedrals, to iteratively estimate the channel imbalance and cross-talk parameters. Monte Carlo simulations show that the method reliably converges to the true parameter values. *Ingara* data is calibrated using the method, with broadly consistent parameter estimates obtained across flights.

Genuine scene changes may be masked by coherence loss that arises when the bands of spatial frequencies supported by the two passes do not match. Trimming the spatial-frequency bands to their common area of support would remove these uncorrelated contributions, but the bands, and therefore the required trim, depend on the effective collection geometry at each pixel position. The precise dependence on local slope and collection geometry is derived in this thesis. Standard methods of SAR image formation use a flat focal plane and allow only a single global trim, which leads to spatially varying coherence loss when the terrain is undulating. An image-formation algorithm is detailed that exploits the flexibility offered by back-projection not only to focus the image onto a surface matched to the scene topography but also to allow spatially adaptive trimming. Improved coherence is demonstrated in simulation and using data from two airborne radar systems.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint award of this degree.

I give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

14/8/19

Acknowledgements

The guidance of my supervisors, Professor Doug Gray and Dr Mark Preiss, was invaluable. My sincere thanks to them.

I would also like to thank the team behind the *Ingara* radar for their dedication in the laboratory, in the aircraft, and in the field.

This research was supported by the Defence Science and Technology Group.

To Kelly Hall, for your support and patience, I offer my love and several years of hot dinners.

List of publications

Pincus, P. & Preiss, M., “3D SAR coherent change detection for monitoring the ground under a forest canopy”, *IET Radar, Sonar & Navigation*, in press, accepted Apr. 2019.

Pincus, P. B. & Preiss, M., “Coherent change detection under a forest canopy”, in *International Conference on Radar (RADAR)*, Aug. 2018.

Pincus, P., Preiss, M., Goh, A. S. & Gray, D., “Polarimetric calibration of circularly polarized synthetic aperture radar data”, *IEEE Transactions on Geoscience & Remote Sensing*, 55(12):6824-6839, Dec. 2017.

Pincus, P. B., Preiss, M., Stacy, N. J. S. & Gray, D. A., “Image formation on undulating terrain using the upgraded *Ingara* L-band radar system”, in *IEEE Asia-Pacific Conference on Synthetic Aperture Radar (APSAR)*, Sep. 2015.

Pincus, P. B., Preiss, M. & Gray, D. A., “3D SAR beamforming under a foliage canopy from a single pass”, in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2015.

Pincus, P. B., Preiss, M., Goh, A., Stacy, N. J. S. & Gray, D. A., “Low frequency high resolution SAR imaging and polarimetric analysis of a Queensland tropical forest”, in *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, Jul. 2013.

List of symbols

Chapter 2:

SYMBOL	PAGE	DESCRIPTION
$A; A_c$	16	pulse acquisition position along flight-track; aperture centre
a	13	label for first (reference) pass
B	23	bandwidth
b	13	label for second (repeat) pass
c	21	propagation speed
D	31	image size
d	20; 31	distance; pixel spacing
f	21	frequency
g	24	ground complex reflectivity
h	18	height
I	13, 62	spatial-frequency support
J	39	set of spatial-frequency supports
k	21	spatial frequency
L	25	number of samples
l	25	sample index (fast-time)
l_{az}	18	azimuth axis for local tangent plane at O
l_{gr}	18	ground-range axis for local tangent plane at O
l_n	18	vertical axis for local tangent plane at O
N	31	number of pixels
O	16	origin of local coordinate system at image pixel
$r; \tilde{r}$	16; 18	range from O ; range from A
S	13	spectrum
s	13	image; pixel; signal; scattering response (context dependent)
T	23	duration
t	23	time
w	22, 63	resolution scaling factor due to spectral window
x	16	azimuth axis at O (nominal direction of flight-track is $-x$)
x	18	azimuth axis for global coordinate system at image centre
y	16	ground-range axis at O
y	18	ground-range axis for global coordinate system at image centre
z	16	vertical axis at O
z	18	vertical axis for global coordinate system at image centre
$\alpha_{az}; \alpha_{gr}$	18	azimuth slope angle; ground-range slope angle
β	23	chirp rate
γ	13	complex coherence
$\tilde{\delta}_{ant}$	19	antenna squint angle
$\tilde{\delta}_{ftk}$	19	flight-track squint angle
$\tilde{\delta}_{off}$	19	aperture centre offset squint angle
θ	16	azimuth angle at O
λ	21	wavelength
ρ	22, 63	spatial resolution
τ	24	delay
$\phi; \Phi$	21; 23	phase
ψ	16	grazing angle at O

Chapter 3:

SYMBOL	PAGE	DESCRIPTION
b	104	beampattern

d	118	distortion
f_v	129	two-way attenuation for volume propagation
i	104	image channel index
N_{ch}	99	number of image channels
n	118	noise
p	118	volume propagation factor
R	99	spatial covariance matrix
v	118	volume complex reflectivity
\mathbf{v}	99	steering vector
\mathbf{w}	76, 98	vector of weights
\mathbf{x}	98	vector of pixels
y	98	weighted sum of pixels
α_v	124	multichannel volume attenuation factor
β_v	124	ratio of cross-pass and single-pass multichannel volume coherences
Γ	116	matrix of coherences
ϵ	125	estimation error
κ	141	matrix condition number
λ	127	eigenvalue
μ	118	ratio of effective ground and volume scattering powers
$\sigma; \sigma^0$	91	radar cross-section; average radar cross-section per unit area
σ^2	99	signal power
$\sigma_e; \sigma_e^{\text{dB}}$	128	extinction coefficient; attenuation rate (decibels per unit length)

Chapter 4:

SYMBOL	PAGE	DESCRIPTION
A	181	transmit/receive channel imbalance matrix
a	178	antenna
C	180	polarimetric covariance matrix
D	180	distortion matrix
H	200	polarimetric entropy
h	177	horizontal polarisation (linear basis)
K	181	receive channel imbalance matrix
k	181	receive channel imbalance
l	177	left polarisation (circular basis)
M	181	cross-talk matrix
m	178	amplifier
N	180	noise matrix
n	179	noise
P	178	p -polarised antenna
P_s	200	total scattering power
p	177	a general polarisation (general basis)
Q	179	q -polarised antenna
q	177	a polarisation orthogonal to p (general basis)
R	179	receive distortion matrix
r	177; 178	right polarisation (circular basis); receive device distortion
$S; \tilde{S}$	177	true scattering matrix; observed scattering matrix
s	177; 178	scattering response; switch
T	179	transmit distortion matrix
t	178	transmit device distortion
U	182	unitary basis change matrix
u	181	receive cross-talk
v	177; 181	vertical polarisation (linear basis); transmit cross-talk
w	181	receive cross-talk
y	181	absolute gain and phase distortion
z	181	transmit cross-talk
α	181	ratio of transmit and receive channel imbalances
$\bar{\alpha}$	193, 200	covariance matrix eigendecomposition parameter
σ	188	covariance matrix element
τ	178	pulse index

Chapter 1

Introduction

1.1 Background and motivation

During the 1960s and early 1970s, the herbicide and defoliant Agent Orange was used to destroy vegetation cover in large areas of Southeast Asia [Buckingham 1982]. The effects on human health and the environment were devastating and long-lasting [Institute of Medicine 1994].

Radar offers a non-destructive way to observe the ground under a forest canopy from an airborne or spaceborne platform. Indeed, its use for ‘foliage penetration’ was first investigated during the the Vietnam War [Davis 2011 p. 4]. The key is to transmit a microwave signal at a wavelength that is longer than the leaves and twigs in the canopy, so that a non-negligible proportion of the incident energy propagates through the canopy and provides an echo from any objects at ground level, or from the ground itself [Bessette & Ayasli 2001].

The potential of foliage-penetrating radar for the detection of man-made targets, both stationary and moving, has been studied extensively [Winter *et al.* 1994, Hellsten *et al.* 1996, Nguyen *et al.* 1997, Soumekh 1997]. Alternatively, the ‘target’ may be the ground. The scattering response of the ground can be used to infer its environmental state, such as the level of moisture in the soil [Hajnsek *et al.* 2009a, McNairn *et al.* 2015]. Neither object detection nor geophysical estimation are the focus of this work.

The end goal pursued in this thesis is the reliable detection of disturbances of the ground in a forest. An example output might be a change map, consisting of an image of the landscape and some indication of where the ground has changed between two observations. To achieve this using radar, the echoes received from an illuminated scene must be resolved into an image, and images acquired at different times must be compared to locate any differences in the scene.

Synthetic aperture radar (SAR) is a standard technique for acquiring and processing radar echoes in order to generate a two-dimensional image of the illuminated scene [Cutrona *et al.* 1961]. Typically, a side-looking radar on a moving platform transmits and receives a sequence of pulses. Across-track resolution (in the direction of propagation) is achieved in inverse proportion to the bandwidth of the transmitted waveform [Klauder *et al.* 1960]. Along-track resolution (in the direction of travel) is achieved in inverse proportion to the length of track whose pulses are coherently combined; this coherent processing amounts to recovering the received echo that would have been acquired by a hypothetical antenna (the synthetic aperture) as long as the length of track [Munson *et al.* 1983].

Coherent change detection (CCD) is a standard technique for comparing SAR images by measuring the local complex correlation coefficient [Jakowatz *et al.* 1996 ch. 5.5]. In the SAR context, this coefficient is often referred to as the magnitude or degree of the interferometric coherence [Touzi & Lopes 1996]. Other methods of SAR change detection exist [Preiss *et al.* 2006, Oriot 2014], but CCD is particularly well-suited to detecting subtle changes caused by a rearrangement of scattering elements in a scene, such as pebbles on a gravel road. This is because a SAR image of natural clutter, such as rough ground, exhibits speckle: pixel-to-pixel fluctuations in the net observed scattering response (both magnitude and phase) due to coherent superposition of the individual unresolved responses from a particular arrangement of scattering elements viewed from a particular direction [Goodman 1975]. Although such fluctuations are

commonly treated as noise to be suppressed when analysing a single image [Lee 1981], the pattern formed by the speckle actually serves as a signature of the landscape for that acquisition, and the coherence magnitude is a measure of the local similarity of the complex speckle patterns in two SAR images [Rignot & van Zyl 1993]. If a part of the scene changed between the two acquisitions, then this area may be detected as an area of low coherence (in SAR parlance, decorrelation) between the two speckle patterns. CCD is usually applied to high-frequency (e.g. Ku-band, with wavelength $\sim 1.8\text{cm}$ [Tsunoda *et al.* 2000]) images of open ground with minimal vegetation.

The end goal of detecting ground disturbances in a forest could be satisfied by applying CCD to low-frequency (e.g. L-band, with wavelength $\sim 23\text{cm}$) SAR images of tree-covered ground. Numerous obstacles immediately present themselves, each of which, if not overcome, would have the effect of distorting the estimated coherence so that it is not an accurate measure of scene change on the ground. Four challenges in particular serve as initial motivation for this thesis. Firstly, the operating wavelength must be selected carefully: too short, and the level of foliage penetration may be insufficient; too long, and the coherence may be insensitive to scene changes. Given this trade-off, it is not possible to select a wavelength that avoids any scattering from the canopy. Secondly, therefore, scattering from the canopy will contribute to the observed coherence, with both volume decorrelation (due to multiple scattering elements at different heights) and temporal change (or lack of change) in the canopy obscuring the coherence of scattering from the ground. Thirdly, the scattering at ground level may arise from two main mechanisms: direct surface backscattering, which gives rise to the speckle pattern of interest, and ground-trunk double bounce, which is a sequence of two forward specular reflections that are unlikely to be affected by subtle ground disturbances and therefore will bias the coherence upwards. Lastly, for two conventional SAR images, the combination of non-identical flight-tracks and undulating terrain will lead to spatially varying surface or ‘baseline’ decorrelation that cannot be removed by simple filtering of the image spectra [Li & Goldstein 1990, Gatelli *et al.* 1994, Marechal 1995]. Note that trees have been widely cleared from flat ground to make way for agriculture and urban development, leaving forests more often in hilly areas.

3D SAR beamforming, also known as SAR tomography, is an experimental technique that could be used to suppress the scattering response from the canopy [Homer *et al.* 1996, Reigber & Moreira 2000]. By analogy with the conventional synthetic aperture obtained by combining pulses along the flight-track, multiple images acquired at different grazing angles can be coherently combined to construct a synthetic aperture in elevation, thereby supporting vertical resolution [Jakowatz & Wahl 2001, Pincus *et al.* 2009]. The output is a new image that is three-dimensional in the sense that the scattering elements whose responses contribute to each pixel are now constrained to those at a particular height (or rather, a narrow range of heights, given the vertical resolution). However, acquisition of a sufficient number of well-spaced images to support a useful level of resolution and an appropriate unambiguous extent for the scene typically requires several repeat passes of the platform, which is a significant burden, and the scene may change during data collection. To lessen this burden, it is common to recast SAR multi-image processing as a beamforming problem, and apply one of the many high-resolution reconstruction algorithms available in the array-processing literature that adaptively suppress sidelobe interference [Van Trees 2002, Lombardini *et al.* 2003]. Example algorithms applied to SAR imagery include Minimum-Variance Distortionless Response (MVDR) beamforming (also known as Capon beamforming, after the original author [Capon 1969]) and Multiple Signal Classification (MUSIC) [Reigber *et al.* 2005, Frey & Meier 2011a]. A common application of 3D SAR is the analysis of the vertical structure of a forest or an urban area by repeatedly beamforming the input 2D images to form a set of 3D images steered to a contiguous range of heights and plotting vertical profiles of the scattering intensity [Fornaro *et al.* 2005, Frey & Meier 2011b, Tebaldini & Rocca 2012]. 3D SAR has also been used for the detection of targets under canopy [Nannini *et al.* 2012, Huang *et al.* 2012]. Of particular interest here are beamforming algorithms that preserve phase, in order to obtain the complex speckle pattern due to scattering at ground level, with the scattering contributions from vegetation above the ground suppressed.

Performance assessment of different array designs and adaptive beamforming algorithms necessitates the use of models of the scene. Analytical models for the interferometric coherence of forest scattering have been developed in the SAR literature

[Cloude 2010 ch. 7]. Of particular note is the dual-layer random-volume-over-ground (RVOG) model [Treuhaft *et al.* 1996, Cloude 2010 ch. 5.2.4, 7.2-7.4], which treats the forest canopy as a homogeneous volume, consisting of many randomly positioned and randomly oriented scattering elements. The volume has only two parameters: height and attenuation rate (i.e. a fixed propagation loss in decibels per metre). The overall RVOG structure is further parameterised by a ground-volume scattering power ratio and the height of the ground surface, which is usually known only roughly. Key to the model's utility is that the scene coherence is given by a straightforward mathematical function. Even though the RVOG model is an extreme simplification of a forest, the fact that reasonably accurate parameter estimates have been obtained via model inversion suggests that it does capture the gross features of forest scattering [Praks *et al.* 2007, Hajnsek *et al.* 2009b].

As a complement to the spatial diversity that enables vertical resolution, polarisation diversity can be exploited to distinguish different scattering mechanisms, even when they manifest at the same location. Numerous polarimetric decompositions have been proposed that aim to determine the relative balance of some set of canonical mechanisms by which the scene may respond to the incident microwaves, such as rough surface scattering, specular double bounce, volumetric scattering from oriented particles, etc [van Zyl 1989, Cloude & Pottier 1996, Freeman & Durden 1998]. These mechanisms are usually derived and interpreted according to theoretical models of wave scattering, thus permitting the physical structure of the scene to be inferred to a limited extent. [Durden *et al.* 1989, Treuhaft & Siqueira 2000, Cloude 2010 ch. 3, 9.5]. For the CCD application considered here, polarimetric decompositions offer a way to separate the surface and double-bounce scattering mechanisms that may occur at ground level in a forest.

Beginning with the seminal papers of Shane Cloude and Konstantinos Papathanassiou in the late 1990s, there have been a large number of studies investigating how polarimetry and interferometry can be combined to characterise complex scenes, particularly the vertical structure of vegetation [Cloude & Papathanassiou 1997, 1998, Papathanassiou & Cloude 2001, Cloude & Papathanassiou 2003, Cloude 2010 (see references therein)]. This field is commonly called PolInSAR. The formulation and analysis of the forest change-detection problem in this thesis draws on the approach and models used in PolInSAR.

The polarisation diversity required for polarimetric decompositions is typically obtained by an imaging radar system with a dual-polarised antenna, some kind of switching control on transmit and a dual-channel receiver, together operating in an alternating-transmit, simultaneous-receive mode to measure the full scattering matrix of the scene [Zebker *et al.* 1987]. Importantly, for physical interpretation of the decomposed data to be valid, the system must be calibrated to correct for any polarimetric distortion i.e. imbalance or impurity (cross-talk) in the different polarimetric channels on transmit and receive [Freeman 1992]. The problem of reliably estimating the polarimetric distortion has received significant attention in the SAR literature [Freeman 1991, Quegan 1994, López-Martínez *et al.* 2007, Goh *et al.* 2007]. Initial approaches relied on specially constructed corner reflectors deployed in the scene to act as point targets with known scattering matrices, against which the collected data could be compared [Barnes 1986, Yueh *et al.* 1990]. Alternatively, clutter-based approaches have been developed that compare the data to the theoretical covariance matrix structure of distributed clutter that is assumed to exhibit reflection symmetry, whereby the average scattering response is mirrored on either side of the vertical plane containing the line-of-sight axis [van Zyl 1990, Klein 1992, Ainsworth *et al.* 2006]. Scattering theory and measurements indicate that the correlation between the co-polarised and cross-polarised responses in the linear basis (using horizontal and vertical polarisations) from reflection-symmetric clutter should be zero [Antar & Hendry 1985, Nghiem *et al.* 1992, Cloude & Pottier 1996]. Cross-talk causes the observed correlation to be greater than zero.

However, almost no polarimetric calibration techniques suitable for a circularly polarised imaging radar system on an airborne or spaceborne platform are described in the literature. (The author is aware of only one article that directly addresses this topic [Michelson *et al.* 1997].) In 2013, an interim version of DST Group's *Ingara* airborne imaging radar system was built using a pair of helical antennas that operated in the circular basis with left and right polarisations [Pincus *et al.* 2015b]. The need to calibrate this system was the impetus to develop polarimetric calibration techniques

for circularly polarised SAR data.

The last in the list of obstacles to obtaining an accurate estimate of the ground coherence is the problem of undulating terrain causing spatially varying surface decorrelation. Each pixel in a SAR image is supported by a band of spatial-frequencies in the focal plane, with the particular band determined by the waveform and the angular collection geometry at the pixel position. For corresponding pixels in two SAR images, if the spatial-frequency supports are not identical (due, say, to different grazing angles), then the non-overlapping portions of the bands will contribute uncorrelated energy that causes a loss of coherence at the pixel position. This decorrelation can be avoided by filtering the spatial-frequencies, or more precisely, trimming the spatial-frequency supports to their common band in the focal plane [Gatelli *et al.* 1994]. For flat ground, which is usually assumed to be the case, all pixels in an image will have approximately the same support, so a single global trim can be applied to each image spectrum directly [Jakowatz *et al.* 1996 ch. 5.2]. However, such a trim is accurate only for those areas where the scattering elements lie in the focal plane. Furthermore, the spatial-frequency supports depend on the local slope, because slope changes the effective angular collection geometry at the pixel position [Rodriguez & Martin 1992, Marechal 1995]. Thus, a global trim is not appropriate for undulating terrain, and decorrelation will occur in areas where the spatial-frequency supports are not matched.

An image-formation algorithm is needed that permits both focusing onto an undulating surface and controlling the spatial-frequency supports in a spatially varying manner. Given an interferometric pair of SAR acquisitions and knowledge of the scene topography, the focal surface could then be matched to the terrain, and the pair of spatial-frequency supports for each pixel position could be accurately trimmed to their common band. Back-projection in the ‘time’ (spatial) domain, unlike traditional frequency-domain algorithms, fits these requirements, although it should be noted that most proponents of this algorithm do not discuss the benefits it offers for interferometry [Desai & Jenkins 1992, Ulander *et al.* 2003, Jakowatz *et al.* 2008]. Spatially adaptive trimming of the spatial-frequency supports has been demonstrated by Blacknell and Andre (and coauthors) [Blacknell *et al.* 2010, 2011, André *et al.* 2013]. However, a precise formulation of how the spatial-frequency supports depend on slope, and how they should be trimmed, is not given in the literature.

The techniques, and associated gaps in knowledge, described in this section serve as the ingredients and the motivation for the thesis. The next section spells out the novel contributions made in each chapter.

1.2 Structure and novelty

The overarching goal of this thesis is to develop SAR processing techniques that enable the detection of disturbances of the ground in a forest. This application has not been proposed before. Techniques from the literature, outlined in Section 1.1, are extended and repurposed in pursuit of this goal.

The thesis consists of three main chapters covering quite different aspects of SAR processing. Given this, the relevant literature is reviewed in each chapter separately. Indeed, for the most part, the chapters stand alone. Nonetheless, solving the problems addressed in each chapter supports the overall goal.

1.2.1 Image formation for maximum coherence

The first of the three main chapters (i.e. Chapter 2) is concerned with SAR image formation, starting from the collection of raw radar pulses. The aim is to develop an image-formation algorithm that achieves maximum coherence between a pair of images covering undulating terrain by removing all surface decorrelation. The proposed algorithm is an extended version of back-projection (also referred to as beamforming in Chapter 2) that includes trimming of the spatial-frequency supports on a pixel-by-pixel basis. The concept of spatially adaptive trimming has been demonstrated in simulation and using ground-based laboratory data by Blacknell and Andre (and coauthors), however they did not formulate precisely how the support depends on slope or how it should be trimmed [Blacknell *et al.* 2010, 2011, André *et al.* 2013].

The following novel contributions are made in this chapter:

1. A precise formulation of the spatial-frequency support provided by each collected radar pulse for each pixel position, accounting for the local slope and collection geometry. See Section 2.12.
2. An analysis of the properties of SAR imagery formed by the back-projection algorithm onto an undulating focal surface. The spectrum of a SAR image formed by back-projection is shown to contain the superposition of the aliased spatial-frequency supports for all pixels. The shape of each support in the image spectrum, and the properties of the resulting point-spread function (e.g. mainlobe width, sidelobe orientation) are shown to depend strongly on the local slope. Comparison is made to conventional spotlight-mode imagery focused onto a flat plane using the polar-format algorithm. See Sections 2.14 and 2.15.
3. An image-formation algorithm detailing how the spatial-frequency supports should be trimmed given a pair of raw radar datasets and knowledge of the scene topography in order to remove the source of surface decorrelation and therefore maximise the interferometric coherence between the resulting SAR images. See Section 2.16.
4. Demonstration of coherence improvement when applying the image-formation algorithm to data acquired by two airborne imaging radar systems. See Section 2.17.

These points were summarised in the following paper:

Pincus, P. B., Preiss, M., Stacy, N. J. S. & Gray, D. A., “Image formation on undulating terrain using the upgraded *Ingara* L-band radar system”, in *IEEE Asia-Pacific Conference on Synthetic Aperture Radar (APSAR)*, Sep. 2015.

1.2.2 3D SAR coherent change detection for monitoring the ground under a forest canopy

In the second of the three main chapters (i.e. Chapter 3), the new idea of 3D SAR CCD for monitoring the ground under a forest canopy is proposed and analysed.

Relevant aspects of past foliage penetration experiments, forest scattering models, 3D SAR beamforming and change detection are first reviewed. Given the trade-off between foliage penetration and change sensitivity when selecting the operating wavelength, the relationship between wavelength, subtle scene change (randomly rearranged clutter elements) and observed coherence is investigated. Given the weak ground surface backscattering response expected at long wavelengths (i.e. for L-band and below), an upper limit on the system noise level is proposed.

Simulations are used to show that coherent interferometric processing of two complex 3D SAR images, each generated by MVDR (Capon) beamforming of multiple conventional 2D SAR images, does indeed work, with the output coherence magnitude image serving as a map of the changes in the speckle pattern associated with scattering at one height in the scene.

The following novel contributions are made in this chapter:

1. A formulation of the new idea of 3D SAR CCD as a multichannel dual-layer (ground surface and canopy volume) coherence estimation problem. Some of the resulting expressions are similar to those used in PolInSAR [Cloude 2010 ch. 6.2.2, 7.4]. However, in PolInSAR it is typically assumed that the ground coherence is unity (i.e. the ground is unchanged, so it exhibits no temporal decorrelation), whereas for the application here the ground coherence is the unknown to be estimated as the measure of change. Furthermore, where PolInSAR exploits polarimetric diversity to maximise and minimise the total (ground plus volume) coherence, here spatial diversity is exploited to minimise the above-ground interference from the volume. See Section 3.7.
2. A framework for analysing the performance of different radar designs, using the RVOG model of forest scattering. It is shown that an L-band array of just three channels, spaced by 0.05° in grazing angle, can attenuate the interference

due to a 20 m tall volume by 12 dB, and if the ground scattering intensity is no more than 2.5 dB below that of the volume, then the ground coherence can be estimated with an error of no more than 0.1 i.e. 10%. Such an array can be formed by a dual-antenna radar system operating in an alternating-transmit, simultaneous-receive mode to synthesise three effective phase centres at different across-track positions. The condition number of the matrix of volume coherences for all pairs of channels is proposed as a general measure of a radar design's sensitivity to error. See Section 3.10.

3. Demonstration of 3D SAR CCD for detecting disturbances of the ground beneath a canopy using two simulations. One simulation generates raw radar clutter for a large number of point scatterers which together form a scene that matches the RVOG model. The other simulation is a modified version of the open-source program PolSARproSim [Williams 2006], originally developed to directly generate polarimetric SAR images for testing PolInSAR algorithms, but extended and repurposed here to provide independent verification of the proposed processing techniques. Receiver operating characteristics are evaluated to show the improvement in detection performance; for the RVOG clutter simulation, when the ground and volume scattering intensities are equal, ordinary (single-channel) CCD achieves a detection probability of only 26% at a false-alarm rate of 5%, whereas 3D SAR CCD using a three-channel system achieves a detection probability of 76%. See Sections 3.11 and 3.12.
4. A polarimetric filtering algorithm that removes scattering contributions arising from the double-bounce mechanism. Ground-trunk double bounce is a source of high coherence that may obscure changes in the direct surface backscattering response from the ground near trees due to the spatial averaging involved in computing the coherence. The double-bounce signature is not assumed to be simply the theoretical dihedral response of two perpendicular metal plates, but is instead estimated from the data after applying the alpha-entropy-span polarimetric decomposition [Cloude & Pottier 1997, Pottier & Cloude 1997]. The novelty lies not in the decomposition itself but in cancelling one polarimetric component of the scattering response and then doing coherent interferometric processing (3D SAR beamforming and CCD) of the filter output. When applied to the fully polarimetric SAR imagery generated by PolSARproSim, the detection probability is improved by 7% compared to using vertical polarisation. See Section 3.12.1.

The main ideas in this chapter were published in the following papers:

Pincus, P. & Preiss, M., “3D SAR coherent change detection for monitoring the ground under a forest canopy”, *IET Radar, Sonar & Navigation*, in press, accepted Apr. 2019.

Pincus, P. B. & Preiss, M., “Coherent change detection under a forest canopy”, in *International Conference on Radar (RADAR)*, Aug. 2018.

Pincus, P. B., Preiss, M. & Gray, D. A., “3D SAR beamforming under a foliage canopy from a single pass”, in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2015.

Related ideas concerning 3D SAR processing and polarimetric decomposition of forest scattering at L-band were discussed in the following papers:

Pincus, P. B., Preiss, M. & Gray, D. A., “Aspects of 3D tomography for multiple-pass spotlight-mode airborne SAR”, in *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, Jul. 2009.

Pincus, P. B., Preiss, M., Goh, A., Stacy, N. J. S. & Gray, D. A., “Low frequency high resolution SAR imaging and polarimetric analysis of a Queensland tropical forest”, in *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, Jul. 2013.

1.2.3 Polarimetric calibration of circularly polarised synthetic aperture radar data

The last of the three main chapters (i.e. Chapter 4) addresses the problem of estimating and correcting for polarimetric distortion in an airborne imaging radar system.

Little attention has been given to this topic in the case when the radar operates in the circular basis, alternating between left and right polarisations on transmit and simultaneously receiving both polarisations. In the one relevant article known to the author, an estimation technique is proposed that uses corner reflectors, including two dihedrals oriented at different angles around the radar line-of-sight direction [Michelson *et al.* 1997]. However, the response of a dihedral is particularly sensitive to angular misalignment: for the strong specular double bounce from a dihedral to return in the backscatter direction, and therefore be received by a monostatic radar, the dihedral seam must be perpendicular to the radar line-of-sight. Achieving good alignment between a corner reflector deployed in the field and an airborne radar subject to wind is difficult.

Polarimetric calibration algorithms for airborne and spaceborne imaging radar systems that operate in the linear basis typically estimate the cross-talk by its first-order distortion of the covariance matrix of reflection-symmetric clutter [van Zyl 1990, Ainsworth *et al.* 2006]. To the author’s knowledge, no article considers how such clutter could be used to estimate distortion in the circular basis.

The following novel contributions are made in this chapter:

1. A unified framework for various radar system models. Past models have been the basis for calibrating particular types of systems, which permit different assumptions about the distortion [Sheen *et al.* 1989, Sarabandi & Ulaby 1990, Quegan 1994, Gau & Burnside 1995]. The general formulation proposed here allows for two generic transmitter designs: either a single amplifier followed by a high-power switch or a low-power switch followed by two amplifiers. In the latter case, it is shown that a particular factorisation of the polarimetric distortion matrix leads to a significant simplification of the cross-talk representation, from the standard four parameters to two reciprocal parameters, one for each of the antennas. See Section 4.3.
2. A distortion estimation technique using corner reflectors for circularly polarised SAR data. Dihedrals are not required, although they can be used. See Section 4.5.
3. An investigation of how clutter could be used to estimate the distortion in circularly polarised SAR data. It is shown that the distortion caused, to first-order, by cross-talk in the circular basis does not permit the individual cross-talk parameters to be estimated. See Section 4.6.
4. Demonstration of the proposed estimation techniques for the polarimetric calibration of data acquired by an airborne imaging radar system operating in the circular basis using a pair of helical antennas. See Section 4.7.

These points were published in the following article:

Pincus, P., Preiss, M., Goh, A. S. & Gray, D., “Polarimetric calibration of circularly polarized synthetic aperture radar data”, *IEEE Transactions on Geoscience & Remote Sensing*, 55(12):6824-6839, Dec. 2017.

1.3 Airborne radar data

Chapters 2 and 4 use L-band data acquired by the *Ingara* imaging radar system, built by the Radar Sensing and Exploitation Group within the Defence Science & Technology Group (DSTG, formerly DSTO, with “O” for Organisation), part of the Australian Department of Defence. Chapter 2 also uses P-band data acquired by the *Bright Sapphire* system, built by Astrium (now part of Airbus) and Systems Engineering and Assessment for the Defence Science and Technology Laboratory (Dstl), part of the United Kingdom Ministry of Defence. Both radars were flown in Australia on the Defence Experimentation Airborne Platform (DEAP), which is a modified Beechcraft 1900C aircraft owned by DSTG.

The *Ingara* system operated at X-band for many years [Stacy & Burgess 1994, Stacy *et al.* 2003]. In 2013, it was modified to operate at L-band using a pair of helical antennas; this is the version of *Ingara* that acquired the data used in this thesis [Pincus *et al.* 2015b]. In 2017, the helices were replaced by a four-element array of microstrip patches that enable transmit phase steering in azimuth [Preiss *et al.* 2018].

The *Bright Sapphire* system uses an ultra-wideband sinuous antenna and a repeating sequence of up to four chirp pulses at different centre frequencies to operate coherently from P-band to L-band [Pincus *et al.* 2013, Hayward *et al.* 2015]. During the trial from which the data used in thesis was acquired, only two pulses were used, one at P-band and one at L-band, due to limitations on the approved spectrum. Unfortunately, the L-band power amplifier failed, so only P-band data are available.


System parameters are listed in Table 1.1 and the antennas are pictured in Figure 1.1. The *Ingara* X-band and *Bright Sapphire* L-band channels are not used in this thesis, but their parameters are listed for interest.

Table 1.1: Typical operating parameters for the *Ingara* and *Bright Sapphire* radar systems. For the *Bright Sapphire* P-band channel, the bracketed values are the effective values after accounting for constraints imposed by (i) the spectrum approval authority on transmit and (ii) mitigation of RF interference on receive.

Parameter	Symbol	<i>Ingara</i>		<i>Bright Sapphire</i>	
		X-band	L-band	L-band	P-band
Dechirp on receive?		Y	N	Y	Y
Complex (I/Q) sampling?		N	N	Y	Y
Wavelength	λ	3 cm	23 cm	24 cm	88 cm (82 cm)
Transmit centre frequency	f_c (f_b)	10.1 GHz	1.32 GHz	1.2575 GHz	340 MHz (367.6 MHz)
Dechirp centre frequency	f_d	8.85 GHz	N/A	1.2591 GHz	344.1 MHz
Intermediate frequency	f_{if}	300 MHz	−100 MHz	N/A	N/A
Sampling frequency	f_s	1 GHz	500 MHz	53.85 MHz	107.7 MHz
Transmit bandwidth	B_c (B_b)	600 MHz	140 MHz	85 MHz	220 MHz (63.5 MHz)
Dechirp bandwidth	B_d	600 MHz	N/A	126.3 MHz	326.9 MHz
Transmit duration	T_c (T_b)	8–30 μ s	12 μ s	9.935 μ s	9.935 μ s (2.868 μ s)
Dechirp duration	T_d	8–30 μ s	N/A	14.763 μ s	14.763 μ s
Sampling duration	T_s	16.384 μ s	20.48 μ s	12.906 μ s	12.915 μ s
Effective PRF per channel	f_{prf}	0.45 kHz	0.45 kHz	4.247 kHz*	
Antenna type		patch array	helical pair	sinuous [Q-PAR]	
Peak input power		2 kW	0.5 kW	0.1 kW	
Polarisation		dual linear	dual circular	dual linear	
Polarimetric mode		full [†]	full [†]	full [†]	
Azimuth 3 dB beamwidth	θ_{bw}	1.5°	30°	72°	85°
Operating range	r	3–6 km	3 km	2 km	


*The transmitter cycles through the two frequency bands and two polarisations at 16.986 kHz.

[†]Alternating transmit and simultaneous receive. This halves the effective PRF per channel.



(a) *Ingara* X-band patch array antenna

(b) *Ingara* L-band helical antenna pair



(c) *Bright Sapphire* sinuous antenna (0.2–2 GHz)

Figure 1.1: Antennas for the *Ingara* and *Bright Sapphire* radar systems.

Chapter 2

SAR image formation for maximum coherence

2.1 Introduction

Synthetic aperture radar is a standard technique that generates two-dimensional images of an illuminated scene [Cutrona *et al.* 1961]. In typical operation, a side-looking radar on an airborne or spaceborne platform illuminates a patch of ground by a sequence of microwave pulses, with the echo from each pulse received and processed. After pulse compression, range resolution is achieved in inverse proportion to the bandwidth of the transmitted waveform [Klauder *et al.* 1960]. By coherently combining a block of pulses acquired at different positions along the platform track, along-track (azimuth) resolution is achieved in inverse proportion to the angle subtended on the ground by the processed extent of the track [Munson *et al.* 1983]. The block of pulses is termed the synthetic aperture because the azimuthal resolution is equivalent to that which would be obtained by a single antenna of the same length as the processed track extent.

Numerous algorithms have been developed to form synthetic aperture radar (SAR) images [Carrara *et al.* 1995, Cumming & Wong 2005]. Almost all of them start by modelling the received signal as the echo from a hypothetical point target fixed at a position in the scene. (Exceptions exist, whereby electromagnetic scattering models for extended or moving targets are derived and inverted to form the image; see references within Borden & Cheney [2013].) The task is then to coherently sum the contributions from this target across pulses, compensating for the change in range between platform and target so that the contributions add in-phase. The algorithms differ in how they achieve this motion compensation and the approximations they make, but they can be categorised according to whether they operate in the time (spatial) or frequency (spatial-frequency) domain, and if the latter case, can be further categorised by whether they do motion compensation to a point (spotlight mode) or a line (stripmap mode).

The received signal of course contains not one echo but a superposition of echoes from the continuum of scattering elements on the ground within the range gate of the system. Frequency-domain algorithms are designed to operate on all these responses in bulk, first compensating for the range delay to the reference point or line, and then making corrections to accommodate the responses from other locations, assuming a nominal collection geometry. For example, the polar-format algorithm (PFA) does motion compensation to a point, and then resamples the nonuniformly spaced spatial-frequency data to a regular grid, ready for Fourier inversion [Walker 1980, Jakowatz *et al.* 1996 ch. 3], while the range-migration algorithm (RMA) does motion compensation to a line, and then scales the data by a 2D phase function before doing a Stolt interpolation, again followed by Fourier inversion [Cafforio *et al.* 1991, Carrara *et al.* 1995 ch. 10]. Critically, the processing steps in all such algorithms are predicated on a straight and level flight-track and a planar focal surface, so that the collection geometry varies in a uniform and predictable way along the aperture and across the swath. When the actual geometry approximates this acquisition model moderately closely, the image formation works well, but the algorithms are inherently limited in their ability to handle highly nonlinear flight-tracks or a non-planar focal surface [Frey *et al.* 2009, Schwäbisch *et al.* 2010]. For example, before applying PFA,

a non-linear flight-track can be partially accommodated by projecting the data onto a selected processing plane, in a direction orthogonal to the desired focal plane. Now scatterers in the focal plane will be well-focused by PFA; however, other objects will defocus in proportion to both their offset from the focal plane and the magnitude of the deviations from the nominal flight-track [Jakowatz *et al.* 1996 ch. 3.8].

By contrast, the time domain offers the flexibility to handle arbitrary collection geometries which may vary substantially along the aperture or across the imaged scene [Jakowatz *et al.* 2008]. Pulses are no longer transformed in bulk, but are processed (and reprocessed) to suit different pixel positions, according to the actual platform-pixel range (time-delay). This suits scenarios where the range for different platform-pixel pairs varies significantly or somewhat unpredictably, such as for a light-weight or manoeuvring airborne platform operating at close-range, or a low-frequency radar system which requires a long synthetic aperture to achieve high azimuth resolution [Hellsten *et al.* 1996]. In addition, a digital elevation model can be used to define an undulating focal surface that matches the scene topography better than a horizontal plane [Jakowatz *et al.* 2008, Harcke *et al.* 2010]. Time-domain algorithms inherently allow for wavefront curvature, although this feature is not unique to them, since RMA does too [Milman 1993], and the plane-wave approximation made by PFA can be corrected by post-filtering [Doren *et al.* 1997, Preiss *et al.* 2002].

In essence, most time-domain algorithms are equivalent in attacking the image formation problem directly: they hypothesise a point target at each pixel position and coherently weight and sum the expected contributions across pulses. The weights are the phase delays needed to compensate for the differential range from platform to target. Each output pixel value therefore indicates the level of constructive interference achieved over the synthetic aperture by the actual scattering at the pixel position. Instead of motion-compensating to just a reference point or line, every pixel position serves as a point of phase stabilisation. Indeed, in this chapter it is argued that each pixel of a SAR image formed in the time domain is equivalent to the centre reference point of a spotlight-mode image formed in the frequency domain via PFA.

Time-domain algorithms are typically cast in terms of either beamforming output pixels [Jakowatz *et al.* 2008], where the synthetic aperture is treated as a sensor array, or back-projecting input pulses [Desai & Jenkins 1992], analogous to the method used in x-ray tomography [Kak & Slaney 1988]. The former terminology is favoured in this chapter, but the latter is used in this review section for consistency with the literature.

The most obvious obstacle to image formation in the time-domain is the heavy computational burden. A direct, geometrically exact implementation of the coherent addition process will be $O(N^3)$ [Wahl *et al.* 2008], where N is the computational size of the task (say, N input pulses and $N \times N$ output pixels) and $O(f(N))$ is the computational complexity to within an order-of-magnitude. Fast back-projection algorithms have been developed which trade-off precision for speed by recursively factorising standard back-projection to use subapertures and either image patches [McCorkle & Rofheart 1996, Xiao *et al.* 2000] or coarse pixel grids [Ulander *et al.* 2003], achieving $O(N^2 \log N)$, the same as PFA [Wahl *et al.* 2008]. However implemented, it is worth noting that time-domain processing fits a parallel structure quite naturally [Desai & Jenkins 1992] and is thus ideally suited to modern computer architectures [Rogan & Carande 2010].

A second downside to time-domain image formation is the difficulty in applying an autofocus routine to correct for defocusing caused by inaccurate motion compensation. Consider the phase-gradient algorithm, an established autofocus technique for PFA imagery which estimates and removes the residual phase error across pulses [Jakowatz & Wahl 1993]. The method requires that the synthetic aperture contributions to a selected point scattering response be directly accessible from the image via a Fourier transform of the point-spread function along azimuth. Images formed in the time domain do not satisfy this requirement [Jakowatz & Wahl 2009]; each pixel is formed by a curved contour of integration across range-compressed pulses such that the point-spread function is coupled in range and azimuth. Indeed, in this chapter it is shown that each pixel in a beamformed SAR image has a slightly different spatial-frequency support, by virtue of its unique position, and the 2D Fourier transform of such an image consists of the superposition of the supports for all pixels, which cannot be separated and directly related to the original pulse data. Autofocus routines for SAR images formed in the time-domain have been proposed [Zhang *et al.* 2013, Torgrimsson

et al. 2014], but this is an area of continuing research.

Given the airborne radar data which this thesis seeks to exploit—low-frequency, wide-beam, long-aperture, close-range [Pincus *et al.* 2013]—the time-domain beamforming approach is clearly an attractive image formation algorithm. However, the ultimate task here is not image formation but change detection. To that end, neither processing speed nor autofocus are the subject of this chapter. Instead, the flexibility afforded by beamforming is examined with the goal of optimising the subsequent performance of radar interferometry, which is the basis for coherent change detection.

In this chapter, a process of image formation is detailed which is designed to not only form a focused image, but also to maximise the potential repeat-pass coherence (similarity) between a pair of images by removing the decorrelation (differences) induced by differing collection geometries and undulating topography. This decorrelation could mask actual changes in the scene.

To elaborate on the concepts of coherence and decorrelation, observe that, for most natural and urban landscapes, microwave scattering can be treated as a stochastic process (in space) since each measurement consists of the coherent superposition of the responses from many scattering elements. Each pixel in a SAR image can be modelled as a complex Gaussian random variable; across pixels, the phase will be uniformly random and the intensity will exhibit local fluctuations akin to multiplicative noise [Lee 1981, Oliver & Quegan 1998 ch. 4.3]. Notwithstanding this statistical treatment, the measured scattering realisation is not arbitrary but is a product of the illuminated scene. The coherent superposition generates a complex speckle pattern which is characteristic of the scene when observed from one direction at one time using a particular frequency band [Goodman 1975].

Now consider two complex radar images, s_a and s_b , of the same scene viewed from similar directions, with band-limited 2D spectra S_a and S_b on spatial-frequency supports I_a and I_b , respectively. How can the images be compared? Given we have two complex, zero-mean, stochastic (in space), scattering processes, one suitable statistical metric is the complex correlation coefficient [Papoulis 1991 pp. 152, 188, 293–295]. In the radar (and laser) context, this is often called the complex coherence γ [Touzi & Lopes 1996, Goodman 1975 p. 38], defined as

$$\gamma_{ab} = \frac{E\{s_a s_b^*\}}{\sqrt{E\{s_a s_a^*\}E\{s_b s_b^*\}}} \approx \frac{\langle s_a s_b^* \rangle}{\sqrt{\langle s_a s_a^* \rangle \langle s_b s_b^* \rangle}} = \frac{\langle S_a S_b^* \rangle}{\sqrt{\langle S_a S_a^* \rangle \langle S_b S_b^* \rangle}} \quad (0 \leq |\gamma_{ab}| \leq 1). \quad (2.1)$$

The ensemble averages $E\{\dots\}$ in (2.1) are typically approximated, as shown, by spatial averages $\langle \dots \rangle$ over neighbouring or ‘similar’ resolution cells that are assumed to be independent realisations of the same wide-sense stationary scattering process [Touzi *et al.* 1999, Deledalle *et al.* 2015]. The statistical performance of the approximation for varying sample size is detailed in Appendix 2.A.1. Importantly, the spatial averages are equivalent to spectral averages—this last equality in (2.1) is derived in Appendix 2.A.2, where the complex coherence γ_{ab} is also linked to the more general concept of a coherence, or cross-correlation, function, allowing for a lag between the inputs.

The coherence magnitude $|\gamma_{ab}|$ is a measure of the similarity of the speckle patterns in the two images: when $|\gamma_{ab}| = 1$, the patterns are perfectly correlated, or equivalently, there is zero decorrelation, and when $|\gamma_{ab}| = 0$, the patterns are completely uncorrelated, or equivalently, there is complete decorrelation [Rignot & van Zyl 1993]. Because a speckle pattern is determined by the particular arrangement of scattering elements inside each resolution cell, it is highly sensitive to subtle changes in the scene [Preiss *et al.* 2006]. Therefore, it is used as the change detection metric in this thesis.

Given the Gaussian scattering model, the coherence magnitude $|\gamma_{ab}|$ is a scaled version of the maximum-likelihood estimator for change in the scattering response, disregarding changes in the overall brightness [Seymour & Cumming 1994, Jakowatz *et al.* 1996 ch. 5.5]. To see this latter point, observe that if $s_b = y s_a$ in (2.1) for any complex scalar y , then $|\gamma_{a(ya)}| = 1$ and $\arg \gamma_{a(ya)} = -\arg y$. To reiterate, the coherence magnitude is a correlation-based measure of change in the speckle patterns of two radar images.

For completeness, note that the interferometric phase $\arg \gamma_{ab}$ has long been used as a fine measure of topography, albeit wrapped at 2π intervals [Shapiro *et al.* 1972, Graham 1974, Bamler & Hartl 1998, Rosen *et al.* 2000].

There are several sources of decorrelation which could degrade the coherence magnitude [Zebker & Villasenor 1992, Cloude 2010 ch. 5.2]:

- temporal decorrelation due to changes in the scene between acquisitions,
- surface (a.k.a. baseline) decorrelation due to different angular views of the scene [Li & Goldstein 1990],
- volume decorrelation due to scattering elements at different elevations in the scene [Weber-Hoen & Zebker 2000],
- noise-induced decorrelation due to weak scattering or shadowing in the scene [Zebker & Villasenor 1992],
- processing-induced decorrelation due to poor image focusing or, most often, misregistration of the pair [Prats-Iraola *et al.* 2012].

The aim here is to maximise sensitivity to temporal decorrelation by removing surface decorrelation. Volume decorrelation will be addressed in a later chapter.

Surface decorrelation occurs due to the two images having mismatched spatial-frequency supports $I_a \neq I_b$ [Jakowatz *et al.* 1996 ch. 5.2]. Where the supports overlap ($I_a \cap I_b$), the spectra contain frequency components common to both images, but elsewhere ($I_a \oplus I_b$), the components are exclusive to one image. The latter subset do not contribute to the cross-spectral product in the numerator of (2.1), but do contribute to one of the individual image powers in the denominator, causing the coherence magnitude to be less than unity. Intuitively,

$$\gamma_{ab} \propto \frac{I_a \cap I_b}{I_a \cup I_b} \quad (2.2)$$

with $\gamma_{ab} = 1$ only if $I_a = I_b$ i.e. having equal supports is a necessary, but not sufficient, condition for maximum coherence. The amount of decorrelation is directly proportional to the relative size of the non-overlapping parts of the supports.

Crucially, the spatial-frequency support for a radar image depends on the angular collection geometry, principally the incidence and squint angles [Rodriguez & Martin 1992, Jakowatz *et al.* 1996 ch. 2.4]. If the angular collection geometries for a pair of images are not identical, then their supports will differ.

In fact, it will be seen in this chapter that the precise size and location of the spatial-frequency support for each point in the scene depends on the effective collection geometry at that point. In particular, the effective geometry depends on the local slope [Rodriguez & Martin 1992], so for undulating terrain, the support will vary spatially. For interferometry using a standard planar focal surface, the amount of non-overlapping support in some parts of the swath, and the subsequent surface decorrelation, can be significant [Marechal 1995, Lee & Liu 2001].

By trimming the supports (that is, filtering the images) to their common spectral region, this geometric source of decorrelation can be removed, at the expense of resolution [Gatelli *et al.* 1994]. However, this holds exactly only for scattering elements on the focal surface (the surface for which the trim is computed). If the scattering elements form a volume, then those above or below the focal surface will be trimmed inappropriately and will still cause decorrelation—such volume decorrelation cannot be avoided. If the scattering arises simply from a ground surface which is not matched to the focal surface, then limited slope adaptivity can still be achieved by iteratively estimating the interferometric phase and applying a spatially variant filter based on this phase which removes the uncorrelated parts of the spatial-frequency supports, thus giving estimates of the topography which are successively refined [Davidson & Bamler 1999].

Even when the terrain is a flat plane and the focal surface is matched to it, a global trim will not be precisely correct, because the incidence angle difference at the near (far) edge will be slightly bigger (smaller) than the nominal incidence angle difference at the scene centre. This effect becomes significant when the range swath is of a similar size to the nominal collection range, which is not uncommon for low-altitude airborne systems. A range-dependent filter is required to properly remove the surface decorrelation across the whole swath [Reigber 1999].

Hence, in order to remove all surface decorrelation and thereby maximise the probability of detecting genuine changes in a scene, (i) the focal surface should match the ground topography, and (ii) the spatial-frequency supports require a spatially variant trim which is tailored to suit the collection geometry local to each pixel. The

goal of this chapter is to develop a SAR image formation algorithm which forms well-focused images that satisfy both (i) and (ii), without any constraints on the flight-track or the terrain, except that both are known. Only image formation in the time domain is sufficiently flexible to facilitate this.

The need to control the spatial-frequency supports in this way was mentioned briefly by Ghiglia [1998], but the first (and only) dedicated discussion of this topic was by Blacknell and Andre (and coauthors) from the Defence Science and Technology Laboratory in the UK [Blacknell *et al.* 2010, 2011, André *et al.* 2013].

More clearly than previous authors, Blacknell and Andre distinguish between the spectrum of the image (obtained by a 2D FFT of the complex image) and the spatial-frequency support for each pixel (acquired by the radar and determined by the waveform and the effective collection geometry at that pixel’s position). The former shows the superposition or “amalgamation” of all of the latter [Blacknell *et al.* 2010]. They state that the pixel support is a “planar object” in the 3D spatial-frequency domain which projects down onto a 2D imaging plane in a direction which “will depend on the local tangent to the focal surface” [Blacknell *et al.* 2010]. (Actually, the support is not necessarily “planar”; it is better described as a “ribbon” [Jakowatz *et al.* 1996 p. 69].) For an undulating focal surface, the projection directions will be different for different pixels, leading to the supports being dispersed in the image spectrum. The authors describe a method of “topographic basebanding” whereby the constituent supports are shifted so that they are compactly centred (on top of each other) in the image spectrum, which is important, they argue, because it permits valid resampling of one image in order to register it onto another image [Blacknell *et al.* 2010].

In order to prevent surface decorrelation between a pair of radar images focused to an undulating surface, Blacknell and Andre outline a time-domain image formation process whereby the spatial-frequency supports are trimmed to their region of overlap on a pixel-by-pixel basis [André *et al.* 2013]. They note that the final “coherent” resolution will be spatially varying and may not be well-described by separate azimuth and range quantities, since the shape of a pixel’s support may no longer be approximately rectangular and aligned with the azimuth and range directions [André *et al.* 2013]. Compared to standard processing, they demonstrate an impressive improvement in coherent change detection over undulating terrain using both simulations [Blacknell *et al.* 2010] and real data acquired using a ground-based radar system in a laboratory-controlled experiment [André *et al.* 2013].

The published works of Blacknell and Andre are primarily conceptual and descriptive, rather than mathematical. In particular, they do not derive an expression for the projected location of a pixel’s spatial-frequency support in the image spectrum (although this would have been known to them to achieve their results). They also do not explain in detail how they trim a pair of supports, which might overlap in a complicated way.

This chapter elaborates on the published work of Blacknell and Andre by formulating the projection of a pixel’s spatial-frequency support mathematically in terms of the effective collection geometry and the azimuth and range slopes at the pixel position. The final location of the support in the image spectrum is then derived in terms of the azimuth and range pixel spacings. Their proposed method of topographic basebanding is not successfully replicated for a general flight-track, but it is found here that for valid resampling it is not necessary to achieve a fully compact spectrum; rather, it is sufficient to avoid wrapping of the supports at the image edges by selecting suitably small pixel spacings and applying a global shift. The trimming procedure implemented here is then described in detail. Overall, this chapter is intended to formulate and examine the concepts identified by Blacknell and Andre.

In fact, the chapter encompasses a full description of the signal processing chain for synthetic aperture radar, from transmitter to interferometric image pair. First the coordinate space is established and typical collection geometries are discussed in Section 2.2. Slope is formulated in terms of a local tangent plane at each pixel. After deriving the standard plane-wave approximation (not relied upon for image formation) in Section 2.3, the concept of spatial-frequency support is described in Section 2.4. This is intended to give an intuitive understanding of frequency support before specific signals are introduced. In Sections 2.5–2.8, the transmitted chirp, the received signal and the demodulated digital signal are formulated assuming a simplistic geometric model of scattering as a surface reflectivity function, which is nonetheless sufficient to

motivate the signal processing.

The key step of match-filtering the chirp echoes is formulated in Section 2.9, leading to the critical observation that the range-compressed output is a filtered, narrowband version of the projected scene reflectivity function. The beamforming algorithm for image focusing is set down and demonstrated in Section 2.10. Sections 2.11–2.13 then develop the concept of a 3D spatial-frequency domain within which the acquired pulse data undergo a slope-dependent projection onto the image frequency plane. The resulting 2D spatial-frequency support is the aperture which gives rise to resolution in the output image. In Section 2.14, the geometric interpretation of spatial-frequency is related to the standard formulation of Jakowatz *et al.*, and a simulated example compares the performance of PFA and beamforming over undulating terrain. The image spectrum is analysed in Section 2.15. The ingredients are then all put together in Section 2.16, where a modified algorithm for focusing an interferometric image pair is detailed, incorporating a spatially variant aperture trim which controls the apertures at a pixel-by-pixel level in order to minimise surface decorrelation induced between the pair by the combination of different collection geometries for the two passes and undulating terrain. The whole process is validated in Section 2.17 using data acquired by two airborne imaging radars: the *Ingara* L-band system and the *Bright Sapphire* P-band system.

Throughout this chapter, SAR image properties and processing techniques are demonstrated using simulated radar echo data generated for synthetic scenes composed of simple point scatterers. The coherent response of each point scatterer is an idealised geometric reflection at some fixed amplitude and phase delay; the response does not vary with aspect angle, frequency or polarisation. Clutter is realised using many point scatterers spread randomly over a surface or volume at a density of several scatterers per resolution cell. The raw pulse samples are demodulated, match-filtered and focused just as measured data would be. Pairs of simulated clutter images are used to test interferometric processing. This chapter studies the spatial properties of scattering, not its electromagnetic characteristics, so this simulation is sufficient.

2.2 Collection geometry

In this section, a coordinate system is established and the collection geometry for an airborne SAR acquisition over undulating terrain is formulated. This will serve as the basis for image formation and spatial-frequency trimming. Various collection geometry parameters are defined in Table 2.1 and depicted in Figure 2.1.

Begin with the nominal case of a planar focal surface and a broadside synthetic aperture, illustrated in Figure 2.1(a). Establish a Cartesian reference frame with origin O at a reference position in the illuminated ground scene, azimuth dimension x pointing antiparallel to the nominal flight-track, ground-range dimension y pointing (along the nominal ground) toward the aperture centre A_c , and vertical dimension z pointing up i.e. normal to the horizontal x - y plane. This gives a right-handed Cartesian coordinate system for a left-looking radar, which suits the radar systems used in this thesis.

Spherical coordinates will also be useful, with range r from the origin O toward the flight-track, azimuthal angle θ from the y axis toward the x axis and grazing angle ψ above the x - y plane (the grazing angle may also be called the elevation angle and is the complement of the incidence angle). As shown in Figure 2.1(a), a general aperture position A —that is, the position of the platform along the flight-track at the time of a pulse—can be specified using Cartesian coordinates as (x_A, y_A, z_A) or spherical coordinates as (r_A, θ_A, ψ_A) . Cartesian and spherical coordinates are related according to

$$x = r \cos \psi \sin \theta, \quad y = r \cos \psi \cos \theta, \quad z = r \sin \psi, \quad (2.3)$$

$$r = \sqrt{x^2 + y^2 + z^2}, \quad \theta = \text{atan2}(x, y), \quad \psi = \arctan z / \sqrt{x^2 + y^2}. \quad (2.4)$$

Cartesian coordinates can be represented as vectors $\mathbf{x} = [x, y, z]^T$ with length r pointing in direction (θ, ψ) . Of particular interest is the direction $\hat{\mathbf{r}}$ of the range axis,

$$\hat{\mathbf{r}} = \begin{bmatrix} \cos \psi \sin \theta \\ \cos \psi \cos \theta \\ \sin \psi \end{bmatrix}. \quad (2.5)$$

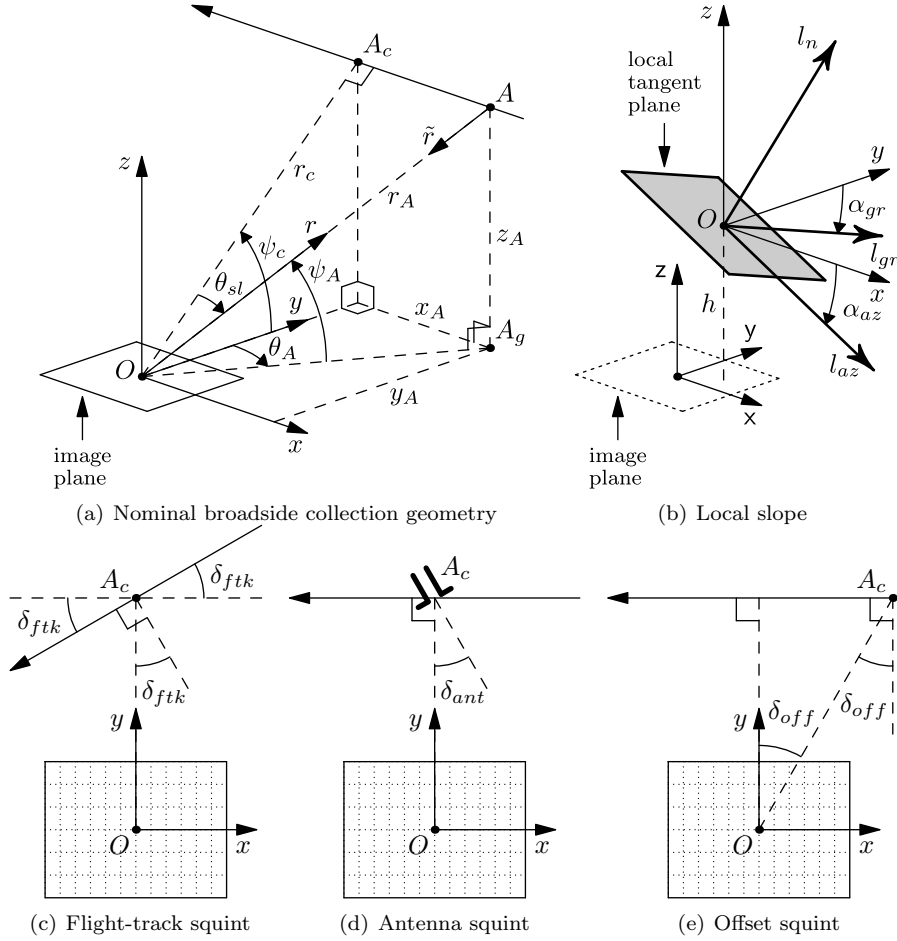


Figure 2.1: Diagrams showing typical SAR collection geometries. Refer to Table 2.1 for parameter definitions. The nominal case of a planar focal surface (coincident with the output image plane) and a broadside synthetic aperture ($\delta_{ftk} = \delta_{off} = 0$) is shown in (a). A local tangent plane to account for local slope is shown in (b). Top-down views in (c)-(e) show different types of squint.

Table 2.1: Geometry parameters in Figure 2.1.

SYMBOL DESCRIPTION

x	azimuth (<i>az</i>) axis of image plane (nominal along-track direction is $-x$)
y	ground-range (<i>gr</i>) axis of image plane (nominal across-track direction is y)
z	vertical axis (i.e. normal) of image plane
O	origin of coordinate system (reference position in scene)
A	general pulse acquisition point on the platform flight-track
A_g	projection of A along $-z$ onto the image plane
\square_c	indicates that \square is applicable to the aperture centre
r	range axis from O toward A
\tilde{r}	reverse range axis from A toward O
ψ	grazing (or elevation) angle above the image plane (N.B. the complement of grazing angle is incidence angle)
θ	azimuth angle off y in the image plane
θ_{sl}	azimuth angle off r_c in the slant plane
α_{az}	surface slope angle along azimuth (tilt below x)
α_{gr}	surface slope angle along ground-range (tilt below y)
l_{az}	azimuth axis (in x - z plane) of local tangent plane
l_{gr}	ground-range axis (in y - z plane) of local tangent plane
l_n	normal axis of local tangent plane
x, y, z	fixed coordinate frame for image plane (upright, sans serif font)
h	height offset of pixel from image plane
δ_{ftk}	squint angle due to flight-track heading (positive: forward off-broadside)
δ_{ant}	squint angle due to antenna steer or yaw (positive: forward off-boresight)
δ_{off}	squint angle due to aperture centre azimuthal offset (pos.: fwd off-broadside)




Figure 2.2: An undulating surface with local (x, y, z) coordinate frames and local tangent planes (grey rectangles) at different heights and slopes on the surface. For simplicity, each tangent plane axis l_{az} , l_{gr} and l_n is labelled only when it differs from the corresponding coordinate axis.

A more intuitive range direction from the perspective of the radar is the reverse range axis $\tilde{r} = r_A - r$, which is zero at the platform (at A) and increases toward the origin O .

The output SAR image will be a 2D grid of complex reflectivity measurements (i.e. pixels) at fixed azimuth and ground-range pixel spacings in the horizontal plane—call this the image plane.

It will be useful to distinguish two coordinate reference frames: a global image frame (x, y, z) with origin fixed to the centre of the image, denoted using upright, sans serif font, and a local pixel frame (x, y, z) with origin at a pixel. Their corresponding axes point in the same directions, as shown in Figure 2.1(b).

To accommodate varying terrain elevation, the planar focal surface can be replaced by an undulating focal surface $(x, y, z = h(x, y))$, where the height $h(x, y)$ of each pixel above the image plane is typically obtained from an external digital elevation model (DEM) [Gallant *et al.* 2011], which itself is usually derived from interferometric radar measurements [Farr *et al.* 2007]. In this chapter, the signal processing will proceed as if the height information is perfect and the focal surface is matched to the terrain. Realistically, the heights will have some error and the scattering objects will project onto the focal surface according to the principles of layover [Jakowatz *et al.* 1996 Appendix C].

Note that even though the constituent pixels may be focused to different heights, the output image will usually still be displayed in the horizontal image plane. Moreover, subsequent image-processing steps, including interferometric processing and Fourier transformation, will operate as if the input was a conventional planar image.

Image focusing via beamforming proceeds by computing and compensating for the propagation range of every pulse-pixel pair (see Section 2.10). This is facilitated here by successively shifting the local (x, y, z) frame such that O is positioned at each pixel in turn and the range r_A from O to each platform pulse position A is then recomputed for each pixel. Thus, every pixel serves as a reference origin for the pulses.

Spatially variant aperture trimming requires knowledge of both the local slope and the angular orientation (θ_A, ψ_A) of each collected pulse (see Section 2.16). As for range, these quantities are recomputed for every pixel.

To formalise the concept of slope, a tilted plane is defined such that it is tangent to the local surface at the reference point O . As shown in Figure 2.1(b), this local tangent plane has azimuthal axis l_{az} (in the x - z plane) tilted below the x axis by the azimuthal slope α_{az} , ground-range axis l_{gr} (in the y - z plane) tilted below the y axis by the ground-range slope α_{gr} , and normal axis l_n . These tilted axes can be expressed in vector form, as shown below.

$$\mathbf{l}_{az} = \begin{bmatrix} 1 \\ 0 \\ -\tan \alpha_{az} \end{bmatrix}, \quad \mathbf{l}_{gr} = \begin{bmatrix} 0 \\ 1 \\ -\tan \alpha_{gr} \end{bmatrix}, \quad \mathbf{l}_n = \mathbf{l}_{az} \times \mathbf{l}_{gr} = \begin{bmatrix} \tan \alpha_{az} \\ \tan \alpha_{gr} \\ 1 \end{bmatrix} \quad (2.6)$$

Note that \mathbf{l}_{az} and \mathbf{l}_{gr} are not necessarily orthogonal (consider the case when both slopes are non-zero), so the axes (l_{az}, l_{gr}, l_n) do not form a Cartesian coordinate system.

Observe too that the vectors in (2.6) do not have unit length. The normal vector \mathbf{l}_n leads directly to the following equation for the local tangent plane:

$$x \tan \alpha_{az} + y \tan \alpha_{gr} + z = 0. \quad (2.7)$$

This is analogous to the equation for the collection slant-plane in conventional spotlight-mode SAR [Jakowatz *et al.* 1996 p. 90 (compare Figure 2.1(b) to Figure 2.28 in this book)]. However, the slant plane specifies the acquisition geometry for a straight and level flight-track, with range tilt due to the depression angle ψ and azimuth tilt related to the squint angle δ_{ftk} , whereas the tangent plane here specifies the orientation of the ground surface with tilts due to the slope angles α_{az} and α_{gr} .

Figure 2.2 shows local (x, y, z) coordinate frames along with local tangent planes at different points on an undulating surface. Observe that the (x, y, z) axes are oriented in the same way but are positioned at different heights to suit the local surface height, whereas the (l_{az}, l_{gr}, l_n) axes are oriented differently to suit the local surface slope.

Instead of defining the local tangent plane via (2.6)–(2.7), it could be obtained using consecutive 3D matrix rotations of the original Cartesian axes by the two slope angles to give axes \tilde{x} and \tilde{y} in the tangent plane and normal \tilde{z} ; the effective platform orientation $(\tilde{\theta}, \tilde{\psi})$ could then be found in this rotated coordinate frame [Pincus *et al.* 2015b]. However, this approach has the disadvantages that the \tilde{x} and \tilde{y} axes would not necessarily be oriented in the intuitive azimuth and ground-range directions i.e. if both slopes were non-zero, then \tilde{x} would not lie in the x - z plane and \tilde{y} would not lie in the y - z plane. Moreover, there is an inherent ambiguity in choosing the order of the rotations, and because matrices do not commute, the two possible orderings give two different solutions. The ambiguity is avoided by specifying the tilted axes directly, as in (2.6).

The actual flight-track for a particular acquisition may deviate significantly from the nominal broadside case used to establish the global coordinate reference frame. The time-domain image formation allows for arbitrary non-linear motion, including variable along-track speed and across-track and above-track curvature, as long as it is accurately recorded by the motion sensors. Nonetheless, it will be useful to formulate the effect of squint for the case when the flight-track is still straight and level but has some azimuthal offset δ from the broadside direction. Squint affects the location and shape of the spatial-frequency support and affects how scattering objects above or below the focal surface lay over onto this plane [Jakowatz *et al.* 1996 ch. 2.4.5]. Three types of squint are distinguished here; in all cases, the angle δ is positive when the squint is forward of broadside (i.e. the radar is looking ahead of the track). Firstly, flight-track squint, depicted in Figure 2.1(c), occurs when the heading of the flight-track is offset by an angle δ_{ftk} from the $-x$ direction—this may be planned, or may be unintended. Observe that the aperture centre is unchanged so the scene is illuminated from the same nominal point of view, which is typically required for repeat-pass interferometry. Secondly, antenna squint, depicted in Figure 2.1(d), typically occurs when the platform has some yaw, so the track heading is broadside but the body alignment is not, and the radar is not mechanically or electrically compensated, resulting in the scene being illuminated at an angle δ_{ant} offset from the antenna boresight. Although this effect is usually ignored from a signal processing perspective, the quality of the output may be affected. Thirdly, offset squint, depicted in Figure 2.1(e), is equal to the azimuth angle of the aperture centre i.e. $\delta_{off} = \theta_c$. Offset squint can be useful when a platform yaw would cause an antenna squint at broadside, since an offset squint can be deliberately selected such that O is imaged from a different point-of-view which is on-boresight; this may be advantageous when, say, a calibration solution to be applied after image formation is valid for near-boresight angles only.

When describing an image in terms of resolution, grazing angle and other attributes which are truly pixel-specific, the parameters for the central pixel will be given as indicative values.

2.3 Plane-wave approximation

The beamformer will focus each pixel by compensating for the exact range r_A from O to A . However, to formulate the concept of spatial-frequency support, a linearised expression for the propagation distance to points near O is required. It will be seen that this approximation amounts to assuming plane waves.

Consider a point $P = (x_P, y_P, z_P)$ a small distance r_P from a pixel origin O , illuminated by a radar at $A = (x_A, y_A, z_A) = (r_A, \theta_A, \psi_A)$. A linearised expression for the distance d from A to P , in terms of the geometry of A relative to O , can be found using the binomial approximation $(1+x)^n \approx 1+nx$ for $x \ll 1$, assuming that $r_P \ll r_A$.

$$\begin{aligned} d &= \sqrt{(x_P - x_A)^2 + (y_P - y_A)^2 + (z_P - z_A)^2} \\ &= \sqrt{r_A^2 + r_P^2 - 2(x_P x_A + y_P y_A + z_P z_A)} \\ &= r_A \sqrt{1 + r_P^2/r_A^2 - 2(x_P x_A + y_P y_A + z_P z_A)/r_A^2} \\ &\approx r_A [1 + r_P^2/2r_A^2 - (x_P x_A + y_P y_A + z_P z_A)/r_A^2] \\ &\approx r_A - (x_P x_A + y_P y_A + z_P z_A)/r_A \\ &= r_A - (x_P \sin \theta_A \cos \psi_A + y_P \cos \theta_A \cos \psi_A + z_P \sin \psi_A) \end{aligned}$$


Rearranging and using (2.5), the distance offset $r_A - d$ can be approximated by the range offset r where

$$r = \mathbf{x}_P \cdot \hat{\mathbf{r}}_A = x_P \sin \theta_A \cos \psi_A + y_P \cos \theta_A \cos \psi_A + z_P \sin \psi_A \approx r_A - d. \quad (2.8)$$

The small distance $r = \mathbf{x}_P \cdot \hat{\mathbf{r}}_A$ is the projection of P onto the range axis from O toward A , and therefore r is the offset of P from O in the range direction. The distance approximation in (2.8) is illustrated in Figure 2.3.

The expression $r = \mathbf{x} \cdot \hat{\mathbf{r}}_A$ matches the Hessian normal form (2.103) of a plane oriented with normal $\hat{\mathbf{r}}_A$ and positioned a distance r from the origin (see Appendix 2.A.7); $\mathbf{x} = \mathbf{x}_P$ is a point on that plane. (This plane is distinct from the local tangent plane.) Approximating distances in this way therefore amounts to assuming plane waves for a small region around O .

Figure 2.3: Plane-wave approximation whereby the distance offset $r_A - d = \overline{OP_s}$ is approximated by the range $r = \overline{OP_p}$ obtained from the projection of P onto the range axis. (N.B. $d = \overline{AP} = \overline{AP_s}$)



For different acquisition positions A , the error in the plane-wave approximation for the distance offset to P will be different. Therefore, if the plane-wave approximation were relied on for image focusing, then the differential range error would cause a phase error in the coherent sum over the synthetic aperture, leading to defocusing. Intuitively, this defocusing would be worse for longer apertures (finer azimuth resolutions) and nearer ranges. For the polar-format algorithm (PFA), the phase error can be quantified analytically in terms of basic image parameters, and limiting the quadratic component of this error to avoid excessive defocus leads to a limit on the image size [Jakowatz *et al.* 1996 p. 96]. For example, given L-band data acquired at 3 km, a PFA-formed image with azimuth resolution of 1 m should be limited to a diameter of 325 m.

To reiterate, time-domain focusing, including via the beamforming method employed here, does not rely on the assumption of plane waves, because every pixel serves as a reference origin O ; nearby points P are of no concern. However, the plane-wave approximation is used here to formulate the concept of a spatial-frequency support.

2.4 Spatial-frequency support

Before developing a full signal model for the raw radar echo, the concept of a spatial-frequency support will be formulated simply by decomposing a pixel's range-dependent phase in terms of the Fourier transform kernel. This leads to the twin motivations for a large waveform bandwidth and a long synthetic aperture to achieve fine resolution.

Reconsider the case of a radar at $A = (x_A, y_A, z_A) = (r_A, \theta_A, \psi_A)$, operating with wavelength λ , illuminating a scattering object at $P = (x, y, z)$ (note dropped subscript). The echo will arrive with a phase offset $\phi = 2\pi \cdot 2d/\lambda$ relative to the transmitted signal, due to the two-way propagation distance $2d$. After image formation, including compensation for the range r_A to the pixel at O , the phase will contain a residual propagation component $\phi_r \approx -2\pi \cdot 2r/\lambda$, where $r \approx r_A - d$ is given by (2.8).

Let $k_r = 4\pi/\lambda = 4\pi f/c$ (for radar frequency f and propagation speed c), giving

$$\phi_r = -k_r r = -(k_x x + k_y y + k_z z) = -\mathbf{k} \cdot \mathbf{x} \quad (2.9)$$

where $\mathbf{k} = [k_x, k_y, k_z]^T$ and, analogous with (2.3),

$$k_x = k_r \cos \psi_A \sin \theta_A, \quad k_y = k_r \cos \psi_A \cos \theta_A, \quad k_z = k_r \sin \psi_A. \quad (2.10)$$

The k variables, with units of radians per metre, form the support for a scattering object in the sense that they scale its relative (x, y, z) position into a measured phase. They do not depend on this position themselves, due to the approximations leading to (2.8); they depend only on the radar wavelength λ and the angular geometry (θ_A, ψ_A) of the platform relative to the pixel, which are properties of the acquisition, not the scene.

Geometrically, (k_x, k_y, k_z) and (k_r, θ_A, ψ_A) from (2.10) can be interpreted as the Cartesian and spherical coordinates of a point in 3D k -space. The point is at radial distance $k_r = \sqrt{k_x^2 + k_y^2 + k_z^2}$ from the k -space origin; this origin corresponds to a radar operating at DC i.e. zero frequency.

When written out in full, the phase factor $e^{j\phi_r} = e^{-jk_r r}$ takes the form of the Fourier kernel. The components k_r and r therefore form a Fourier transform pair where k_r can be interpreted as a range spatial-frequency variable. Similarly, the pairs of corresponding Cartesian components $\{x, k_x\}$, $\{y, k_y\}$ and $\{z, k_z\}$ form Fourier transform pairs. Thus, k -space can be understood as the spatial-frequency domain.

Putting these interpretations together, a point (k_x, k_y, k_z) in 3D k -space is the spatial-frequency support for plane-wave scattering from any position (x, y, z) .

To account for the fact that the propagating waves are not actually planar, the reference origin can be shifted to different pixel positions, around which the plane-wave approximation can be considered locally valid. Given a particular acquisition, its orientation (θ_A, ψ_A) (computed using the local (x, y, z) coordinate frame from Section 2.2) will be different for different reference positions, so the precise support it provides in k -space for the associated pixels will be different. The support is interpreted as residing in a single, global spatial-frequency domain, since the Fourier spaces obtained by the transforms of all local (x, y, z) and global (x, y, z) frames are essentially equivalent—they have the same origin and their corresponding axes point in the same directions.

A realistic scattering surface will have a continuum of scattering elements causing a superposition of echoes at different delays. Their individual propagation phases, which share approximately the same spatial-frequency support, capture this identifying delay information, but a radar receiver necessarily samples the coherent sum of the scattering responses. An approach which enables these echoes to be separated is now sought.

A simple method would be to use a very short pulse so that the set of scattering objects contributing simultaneously is limited to a narrow range interval. However, from a signal-to-noise perspective, this is the exact opposite of what is needed: achieving a high average signal energy per pulse interval per scatterer would then require a very high transmit power which is difficult to achieve in practice, especially on an airborne or spaceborne platform [Jakowatz *et al.* 1996 ch. 1.2].

An alternative approach is to extend the spatial-frequency support from a point to have some finite extent. The shape of the support determines the shape of the point-spread function in the spatial domain via the Fourier transform. This point-spread function should have some nominal width such that scattering contributions from outside this width around a pixel will be excluded, or at least strongly attenuated, after image formation. A standard metric for this width is the extent over which the point-spread function is no less than half (or $10 \log_{10}(1/2) = -3$ dB) its peak power; this width is the effective resolution provided by the spatial-frequency support [Sullivan 2008 ch. 17.6]. For example, using the standard definitions and results in Appendices 2.A.3 and 2.A.5, if the support $I(k_x, k_y, k_z)$ is a rectangular prism in k -space centred at $k_0 = (k_{x0}, k_{y0}, k_{z0})$ with dimensions Δk_x , Δk_y and Δk_z , that is, from (2.82),

$$I(k_x, k_y, k_z) = \text{rect}\left(\frac{k_x - k_{x0}}{\Delta k_x}\right) \text{rect}\left(\frac{k_y - k_{y0}}{\Delta k_y}\right) \text{rect}\left(\frac{k_z - k_{z0}}{\Delta k_z}\right), \quad (2.11)$$

then the corresponding point-spread function $s(x, y, z)$ will be, from (2.90),

$$s(x, y, z) = \frac{\Delta k_x \Delta k_y \Delta k_z}{(2\pi)^{3/2}} e^{-j(k_{x0}x + k_{y0}y + k_{z0}z)} \text{sinc}\left(\frac{\Delta k_x}{2\pi}x\right) \text{sinc}\left(\frac{\Delta k_y}{2\pi}y\right) \text{sinc}\left(\frac{\Delta k_z}{2\pi}z\right) \quad (2.12)$$

with effective 3D resolution specified by ρ_x , ρ_y and ρ_z where, from (2.92),

$$\rho_x = w_{\text{rect}} \frac{2\pi}{\Delta k_x}, \quad \rho_y = w_{\text{rect}} \frac{2\pi}{\Delta k_y}, \quad \rho_z = w_{\text{rect}} \frac{2\pi}{\Delta k_z} \quad (2.13)$$

for spatial bandwidths Δk_x , Δk_y and Δk_z . Of course, the support may take other shapes, and can be deliberately weighted using a window, as described in Appendix 2.A.4, but it is a common feature of all spectral supports that the resolution is inversely proportional to the width of the support, so to achieve fine resolution, a wide support is required. This is a direct consequence of how the Fourier transform responds to affine coordinate transformations, as detailed in Appendix 2.A.5.

In conventional radar, range resolution is achieved by frequency diversity, whereby a wideband waveform, such as a chirp, is transmitted [Klauder *et al.* 1960]. This extends the support along a radial line in k -space (not simply along one of the Cartesian dimensions), since the wavelength λ is no longer a single value but takes on a range of values according to the waveform bandwidth. Note that this method assumes that the scattering response is constant over the frequency band.

In synthetic aperture radar, cross-range resolution, in both azimuth θ and elevation ψ , is achieved by spatial diversity, whereby responses from different angular views of the scene are coherently combined to form a synthetic aperture [Munson *et al.* 1983, Knaell & Cardillo 1995]. As noted previously, the angular location of the spatial-frequency support depends on the same (spatial) angles (θ_A, ψ_A) which characterise the acquisition, so by acquiring and combining data from many pulses along a flight-track, the total support will mark out an arc of constant radius in azimuth, leading to along-track resolution [Munson & Sanz 1984], and by acquiring and combining data from many flight-tracks (i.e. many passes) which illuminate the ground at different grazing angles, the total support will mark out an additional arc of constant radius in elevation, leading to vertical resolution [Reigber & Moreira 2000]. Resolution via spatial diversity can be recast as beamforming using a sensor array, except that here the ‘array’ is synthesised by one sensor (or occasionally, a few sensors) on a moving platform, and the ‘beam’ is a focused pixel at a particular position, as opposed to a wave arriving from a particular angle [Benitz 1997, DeGraaf 1998]. The foundational concept of aperture synthesis for both coherent and incoherent imaging is well-described by Hoor & Kassam [1990]. Note that this method assumes that the scattering response is constant over the subtended solid angle.

Combining these methods leads to a wedge of spatial-frequency support which enables 3D resolution [Chan & Farhat 1981, Jakowatz & Wahl 2001]. Figure 2.4 illustrates such a wedge, with radial width $\Delta k_r = 4\pi B_c/c$ provided by the waveform bandwidth B_c and angular width ($\Delta\theta, \Delta\psi$) provided by the (2D) synthetic aperture. The wedge is centred at k_0 , with radial offset $k_{r0} = 4\pi f_c/c$ because the waveform is transmitted at an RF carrier frequency f_c , and elevation offset ψ_c , as expected for side-looking airborne SAR systems which illuminate the ground. If the offset squint is non-zero, then an azimuthal offset $\theta_c = \delta_{off}$ would also be present.




Figure 2.4: Wedge of spatial-frequency support in 3D k -space, centred at k_0 , with radial width $\Delta k_r = 4\pi B_c/c$, azimuthal width $\Delta\theta$ (in the ground plane), elevation width $\Delta\psi$, radial offset $k_{r0} = 4\pi f_c/c$ and elevation offset ψ_c . An azimuthal offset $\theta_c = \delta_{off}$ may also be present (here, $\theta_c = 0$).

Note that the acquired data marking out the wedge will be discrete in every dimension: each received pulse is sampled, the azimuthal aperture is formed by

multiple pulses in one pass, and the elevation aperture is (usually) formed by multiple passes. Moreover, the pulse samples will be uniformly spaced in the radial direction (assuming a waveform with linear frequency modulation), but, without making any assumptions about the flight-track, there is no guarantee of uniformity across pulses or passes; the angular orientation of each radial line of pulse samples depends on the collection geometry for that pulse, as specified by (2.10).

In practice, the aperture in elevation is rarely formed because a multipass collection is too burdensome (except for a ground-based system [Yitayew *et al.* 2017]); for a single pass, the support reduces to a segment of an annulus with no extent in ψ .

The implicit assumptions that the scattering response is constant over both the frequency band and the subtended solid angle are equivalent to assuming that the scattering objects are isotropic point targets; for alternative approaches, see references within Borden & Cheney [2013].

The spatial-frequency support in 3D k -space will project into the 2D image spectrum during image formation. The manner of this projection, and the task of trimming a pair of projected supports to their common region, will be addressed in Sections 2.12 and 2.16.

Having established the principles by which resolution will be achieved and the kind of data required to provide the necessary spatial-frequency support, the details of data acquisition and image formation will now be formulated.

2.5 Transmitted chirp waveform

To generate the wide spatial-frequency support needed for range resolution, a chirp waveform is commonly transmitted [Klauder *et al.* 1960]. The radars considered in this thesis follow this tradition.

Specifically, let the transmitted signal $s_{tx}(t)$ be the real part of a chirp pulse $s_c(t)$ with centre frequency f_c , bandwidth B_c , duration T_c , start time $t=0$ and chirp rate $\beta = B_c/T_c$ giving a linear frequency modulation. The chirp can be formulated as

$$s_c(t) = e^{j[\Phi_c(t) + \Phi_{qc}(t)]} \text{rect}\left(\frac{t - T_c/2}{T_c}\right) \quad (2.14)$$

where

$$\Phi_c(t) = 2\pi f_c t \quad (2.15)$$

is the carrier phase and

$$\Phi_{qc}(t) = \pi\beta(t - T_c/2)^2 \quad (2.16)$$

is the quadratic chirp phase. The real transmitted signal is

$$s_{tx}(t) = \Re\{s_c(t)\} = \cos[2\pi f_c t + \pi\beta(t - T_c/2)^2] \text{rect}\left(\frac{t - T_c/2}{T_c}\right). \quad (2.17)$$

The instantaneous frequency $f(t)$ is the time-derivative of the phase.

$$f(t) = \frac{1}{2\pi} \cdot \frac{d}{dt} [\Phi_c(t) + \Phi_{qc}(t)] = f_c + \beta(t - T_c/2) \quad (2.18)$$

The chirp's frequency response is therefore a linear ramp in time, as illustrated in Figure 2.5. It is centred at $f(T_c/2) = f_c$ and has total frequency extent $f(T_c) - f(0) = \beta T_c = B_c$ i.e. the chirp bandwidth.





Figure 2.5: Chirp frequency ramp in time. The slope of the ramp is the chirp rate $\beta = B_c/T_c$.

Table 1.1 lists typical waveform parameters for the radar systems considered in this thesis. Figure 2.6 shows the spectrum of the chirp signal $s_c(t)$ in (2.14) using the

parameters of the *Ingara* L-band system. The extended bandwidth provides spatial-frequency support in the direction of propagation i.e. the range direction.

Figure 2.6: Chirp spectrum for *Ingara* L-band parameters $f_c = 1.32\text{GHz}$, $B_c = 140\text{MHz}$ and $T_c = 12\mu\text{s}$. The ringing at the ends of the band is an example of the Gibbs phenomenon, arising due the discontinuities at the ends of the time-domain signal enforced by the $\text{rect}()$ in (2.14) [Carlson 1998 p. 140].



2.6 Scattering Model

In this chapter, microwave scattering from the ground is modelled simply by a complex surface reflectivity function $g(x, y, z = h(x, y))$, where the magnitude corresponds to attenuation and the phase corresponds to sub-wavelength delay of the incident wave; the reflectivity depends on the electrical properties of the material and its shape [Jakowatz *et al.* 1996 ch. 1.2]. Realistically, the scattering response would vary with frequency, polarisation and aspect angle, but those dependencies are not accounted for at this stage—later chapters will consider polarisation. Volumetric multiple-scattering events and wave effects such as diffraction are ignored.

This amounts to modelling the scene as a continuum of isolated isotropic point targets, and modelling the incident electromagnetic wave as a simple ray, which is appropriate only in the geometric optics limit when the wavelength is small relative to the scattering elements [Tsang *et al.* 2000 p. 5]. At L-band ($\lambda \approx 23\text{cm}$) and below, it is certainly reasonable to question these models. However, in order to facilitate the signal processing required for image formation, they are sufficient. Indeed, the output image serves as a measure of how well the data conforms to the scattering model, and artefacts would indicate the presence of non-point-like structures and/or more complex wave interactions.

Amplitude factors due to propagation or radar hardware are neglected too, as they do not affect the signal processing but do clutter the formulation.

2.7 Received signal

The signal received by the radar antenna will be a superposition of echoes from the scattering elements forming the illuminated surface. Recall the collection geometry shown in Figure 2.1, with the radar platform at $A = (x_A, y_A, z_A) = (r_A, \theta_A, \psi_A)$ relative to a pixel origin O in the scene and range axis r pointing from O toward A . From the perspective of the radar, the delay of each echo depends directly on the propagation distance along the reverse range axis $\tilde{r} = r_A - r$ pointing from A toward O . The resultant signal will now be quantified in terms of the propagation distance $r_A - r$.

The echo $s_{echo}(t)$ from each element in isolation will be a version of the transmitted signal $s_{tx}(t)$ in (2.17) attenuated by the reflectivity magnitude $|g(x, y, z)|$, offset in phase by the reflectivity phase $\angle g(x, y, z)$ and delayed by the round-trip propagation time $\tau_r = 2(r_A - r)/c$ [Jakowatz *et al.* 1996 ch. 1.2].

$$\begin{aligned}
 s_{echo}(t) &= |g(x, y, z)| \cos \left[2\pi f_c(t - \tau_r) + \pi\beta(t - \tau_r - T_c/2)^2 + \angle g(x, y, z) \right] \\
 &\quad \times \text{rect} \left(\frac{t - \tau_r - T_c/2}{T_c} \right) \\
 &= |g(x, y, z)| \Re \left\{ e^{j[\Phi_c(t - \tau_r) + \Phi_{gc}(t - \tau_r) + \angle g(x, y, z)]} \right\} \text{rect} \left(\frac{t - \tau_r - T_c/2}{T_c} \right) \\
 &= \Re \{ g(x, y, z) s_c(t - \tau_r) \}
 \end{aligned} \tag{2.19}$$

For a continuum of scattering elements, the total response at range r will depend on the projection $g_p(r)$ of the reflectivity over the spherical wavefront at r [Redding &

Payne 2003].

$$g_p(r) = \int_z \int_y \int_x \delta \left(r_A - r - \sqrt{(x-x_A)^2 + (y-y_A)^2 + (z-z_A)^2} \right) g(x, y, z) dx dy dz \quad (2.20)$$

Putting (2.19) and (2.20) together, the received signal $s_{rx}(t)$ available at the antenna at time instant t will be the sum of the projected responses contributed from ranges whose propagation delay satisfies $t - T_c \leq \tau_r \leq t$, comprising the start of the echo from delay t , the end of the echo from delay $t - T_c$, and other parts of the echoes from delays in between.

$$\begin{aligned} s_{rx}(t) &= \int_r \Re \{ g_p(r) s_c(t - \tau_r) \} dr \\ &= \Re \left\{ \int_r g_p(r) s_c(t - \tau_r) dr \right\} \\ &= \Re \left\{ \int_r g_p(r) e^{j\Phi_{rx}(t,r)} \text{rect} \left(\frac{\tau_r - t + T_c/2}{T_c} \right) dr \right\} \end{aligned} \quad (2.21)$$

where, from (2.15) and (2.16),

$$\Phi_{rx}(t, r) = \Phi_c(t - \tau_r) + \Phi_{qc}(t - \tau_r) = \Phi_c(t) - \Phi_c(\tau_r) + \Phi_{qc}(t - \tau_r) \quad (2.22)$$

and, reexpressing the time support in (2.14) for varying τ_r , noting that $\text{rect}(t)$ in (2.83) is even,

$$\text{rect} \left(\frac{\tau_r - t + T_c/2}{T_c} \right) = \text{rect} \left(\frac{t - \tau_r - T_c/2}{T_c} \right). \quad (2.23)$$

2.8 Baseband digital signal

The received signal needs to be demodulated and sampled in preparation for digital image formation. These steps are detailed in Appendices 2.A.8–2.A.10.

The resulting baseband digital signal $s_b[l]$ is a superposition of chirp echoes centred on DC and minimally sampled at B_c :

$$s_b[l] = \int_r g_p(r) e^{j\Phi_b(t,r)} \text{rect} \left(\frac{\tau_r - t + T_c/2}{T_c} \right) \text{rect} \left(\frac{t - \tau_r - T_s/2}{T_s} \right) dr \quad (2.24)$$

where the phase

$$\Phi_b(t, r) = -\Phi_b(\tau_r) + \Phi_{qb}(t - \tau_r) = -2\pi f_c \tau_r + \pi\beta(t - \tau_r - T_c/2)^2 \quad (2.25)$$

is equal to the received phase in (2.22) except that the carrier component $\Phi_c(t)$ has been removed. Sample index $l = 0, 1, 2, \dots, L_b - 1$ corresponds to time

$$t = lt_b + \tau_s \quad (2.26)$$

where $t_b = 1/B_c$ is the time spacing between the L_b complex samples in $s_b[l]$ and τ_s is the time delay to the start of the sampling window, which has duration $T_s = L_b t_b$; this window acts as a range gate, limiting the ranges from which scattering responses contribute to the signal. In (2.24) and what follows, the variable t is retained for ease of interpretation, but it is implicitly discretised according to (2.26).

During demodulation, the received signal may be filtered such that the available portion of RF spectrum is some subset of the transmitted band. For example, the *Bright Sapphire* P-band data was rechirped and filtered to avoid RF interference and notches on transmit, as shown in Figure 2.32. In what follows, it is assumed that no filtering has been applied, so the transmit parameters f_c , B_c and T_c remain applicable.

2.9 Matched filtering

As is commonplace in radar signal processing, the demodulated signal is match-filtered in order to compress the delayed chirp echoes in time and thus estimate the projected reflectivity $g_p(r)$ as a function of range, while maximising the signal strength relative to

additive clutter and noise [North 1963, Curlander & McDonough 1991 ch. 3, 4.2.1, 5.1, Cumming & Wong 2005 ch. 3]. The matched filter effectively correlates the received data with a template echo successively shifted to different delays; the output indicates the similarity of the data to the template as a function of delay, or equivalently, it measures the magnitude and phase of the reflectivity relative to a unit-magnitude, zero-phase reference scatterer located at different ranges. The shape of the compressed chirp is the range point-spread function; it will be seen that its width—the resolution—is inversely proportional to the effective chirp bandwidth.

The matched filter $s_{mf}[l]$ is a conjugated and time-reversed digital version of the transmit waveform in (2.14), modified to account for the effects of demodulation i.e. the shift to baseband (necessitating complex samples) and the potential spectral crop:

$$s_{mf}[l] = e^{j\Phi_{mf}(t)} \text{rect}\left(\frac{t - \tau_s - T_c/2}{T_c}\right) \quad (2.27)$$

where the phase

$$\Phi_{mf}(t) = -\Phi_{qb}(\tau_s + T_c - t) = -\pi\beta(\tau_s + T_c - t - T_c/2)^2 = -\pi\beta(t - \tau_s - T_c/2)^2 \quad (2.28)$$

is matched to the quadratic phase variation $\Phi_{qb}(t)$ in (2.25). Note that discrete time $t = lt_b + \tau_s$ from (2.26) is referenced to the start of transmit, so the offset by the range-gate delay τ_s in (2.27) indicates that the filter is referenced to the start of sampling i.e. it is non-zero for time $\tau_s \leq t < \tau_s + T_c$, or equivalently, samples $0 \leq l < \lfloor T_c/t_b \rfloor$.

The filtered output $s_m[l]$ is given by the convolution of the demodulated data $s_b[l]$ in (2.24) with the matched filter in (2.27).

$$\begin{aligned} s_m[l] = s_b * s_{mf}[l] &= \sum_{\lambda} s_b[\lambda] s_{mf}[l - \lambda] = \sum_{\lambda} \int_r g_p(r) e^{j\Phi_m(t, r, r_l)} \text{rect}\left(\frac{\tau_r - t + T_c/2}{T_c}\right) dr \\ &\quad \times \text{rect}\left(\frac{t - \tau_s - T_s/2}{T_s}\right) \text{rect}\left(\frac{t - \tau_l - T_c/2}{T_c}\right) \end{aligned} \quad (2.29)$$

Now time $t = \lambda t_b + \tau_s$ is discretised by dummy sample index λ . Variables

$$\tau_l = lt_b + \tau_s - T_c \quad \text{and} \quad r_l = r_A - c\tau_l/2 \quad (2.30)$$

denote the time delay and associated range (from O) corresponding to output lag l . From (2.28), the shifted filter $s_{mf}[l - \lambda]$ has phase

$$\begin{aligned} \Phi_{mf}((l - \lambda)t_b + \tau_s) &= -\pi\beta[(l - \lambda)t_b + \tau_s - \tau_s - T_c/2]^2 \\ &= -\pi\beta[lt_b + \tau_s - T_c + T_c - (\lambda t_b + \tau_s) - T_c/2]^2 \\ &= -\pi\beta(\tau_l - t + T_c/2)^2 \\ &= -\pi\beta(t - \tau_l - T_c/2)^2 \\ &= \Phi_{mf}(t - \tau_l + \tau_s), \end{aligned} \quad (2.31)$$

and similarly considering the (even) $\text{rect}()$ in (2.27), $s_{mf}[l - \lambda]$ has time support $\tau_l \leq t < \tau_l + T_c$. The output phase $\Phi_m(t, r, r_l)$ is the sum of the demodulated echo phase $\Phi_b(t, r)$ in (2.25) and the matched filter phase $\Phi_{mf}(t - \tau_l + \tau_s)$ in (2.31).

$$\begin{aligned} \Phi_m(t, r, r_l) &= \Phi_b(t, r) + \Phi_{mf}(t - \tau_l + \tau_s) \\ &= \pi\beta[(t - \tau_r - T_c/2)^2 - (t - \tau_l - T_c/2)^2] - 2\pi f_c \tau_r \\ &= \pi\beta[t^2 + \tau_r^2 + T_c^2/4 - 2\tau_r t - T_c t + T_c \tau_r \\ &\quad - t^2 - \tau_l^2 - T_c^2/4 + 2\tau_l t + T_c t - T_c \tau_l] - 2\pi f_c \tau_r \\ &= \pi\beta[-(\tau_l + \tau_r)(\tau_l - \tau_r) + 2t(\tau_l - \tau_r) - T_c(\tau_l - \tau_r)] + 2\pi f_c(\tau_l - \tau_r) - 2\pi f_c \tau_l \\ &= [2\pi f_c + 2\pi\beta(t - (\tau_l + \tau_r)/2 - T_c/2)](\tau_l - \tau_r) - 2\pi f_c \tau_l \\ &= \left[\frac{4\pi f_c}{c} + \frac{4\pi B_c}{c} \left(\frac{t - (\tau_l + \tau_r)/2 - T_c/2}{T_c} \right) \right] [r_A - r_l - (r_A - r)] - \frac{4\pi f_c}{c} (r_A - r_l) \\ &= k_r(r - r_l) - k_{r0}(r_A - r_l) \end{aligned} \quad (2.32)$$

where $k_r = k_{r0} + k_t$ with

$$k_{r0} = \frac{4\pi f_c}{c} = \frac{4\pi}{\lambda_c} \quad (\lambda_c \text{ is wavelength}) \quad (2.33)$$

and

$$k_t = \frac{4\pi B_c}{c} \left(\frac{t - (\tau_l + \tau_r)/2 - T_c/2}{T_c} \right) \quad (2.34)$$

Hence, the convolution output is

$$s_m[l] = e^{-jk_{r0}(r_A - r_l)} \sum_{\lambda} \int_r g_p(r) e^{jk_r r} \text{rect}\left(\frac{\tau_r - t + T_c/2}{T_c}\right) dr e^{-jk_r r_l} \text{rect}\left(\frac{t - \tau_l - T_c/2}{T_c}\right) \times \text{rect}\left(\frac{t - \tau_s - T_s/2}{T_s}\right). \quad (2.35)$$

This is the key result. At each time instant t (each sample λ in the summation), the integral along range r —inherent to the data collection—is now an inverse Fourier transform of the actual scene reflectivity $g_p(r)$ over the range delay interval $t - T_c \leq \tau_r \leq t$ whose echoes contribute at this time instant. (The fact that the transform is inverse is only due to the definition of the range axis r as pointing towards the radar; defining range in the opposite direction would give the opposite transform.) The transform is evaluated at a single value of k_r . Therefore, k_r can be interpreted as a range spatial-frequency variable with units of radians per metre; r and k_r are a Fourier transform pair. From (2.33) and (2.34), k_r is simply a scaled and offset version of time which depends on the waveform parameters f_c and B_c but is not range-dependent, so reflectivities from different ranges transform to the same range spatial-frequency. (Contrast this with the range-dependent signal frequency after dechirp, discussed in Appendix 2.A.8.2.)

Furthermore, the summation across discrete time λ —due to the convolution—is now a discrete Fourier transform of the transformed reflectivity over a fixed, range-independent spatial-frequency interval

$$k_{r0} - \Delta k_r/2 \leq k_r < k_{r0} + \Delta k_r/2 \quad (2.36)$$

where

$$\Delta k_r = \frac{4\pi B_c}{c} \quad (2.37)$$

is the range spatial bandwidth. To see this, first observe that the maximum length of time that the filter and an echo can overlap is T_c (made explicit by the $\text{rect}()$ windows for τ_r and τ_l in (2.35)), so from (2.34), the spatial-frequency extent is limited to $k_{t+T_c} - k_t = 4\pi B_c/c = \Delta k_r$. Also observe that the spatial-frequency support for an echo at delay $\tau_r = \tau_l$ (i.e. the echo associated with a particular output lag, having time support $\tau_l \leq t < \tau_l + T_c$), will, from (2.34), start at $k_{r0} + k_{\tau_l} = k_{r0} - 4\pi B_c/2c = k_{r0} - \Delta k_r/2$ and end at $k_{r0} + k_{\tau_l + T_c} = k_{r0} + 4\pi B_c/2c = k_{r0} + \Delta k_r/2$. Therefore, by setting $\tau_r = \tau_l$, so that implicitly only the contributions from ranges $r = r_l$ are considered, (2.35) can be reexpressed using the deduced spatial-frequency support in (2.36), giving

$$s_m[l] = e^{-jk_{r0}(r_A - r_l)} \sum_{\lambda} \int_r g_p(r) e^{jk_r r} \text{rect}\left(\frac{\tau_r - t + T_c/2}{T_c}\right) dr e^{-jk_r r_l} \text{rect}\left(\frac{k_r - k_{r0}}{\Delta k_r}\right) \times \text{rect}\left(\frac{t - \tau_s - T_s/2}{T_s}\right). \quad (2.38)$$

Thus, the output of the matched filter, given by (2.38), is a filtered, narrowband version of the projected scene reflectivity function $g_p(r)$, evaluated at range r_l given by (2.30) for output lag l . Each sample estimate is scaled by a complex exponential whose phase encodes the corresponding propagation distance $2(r_A - r_l)$ counted in units of wavelength ($k_{r0} = 4\pi/\lambda_c$) modulo 2π —accounting for this phase when combining pulses is the essence of beamforming. From (2.30), sample $l = (2r_A/c - \tau_s + T_c)/t_b$ corresponds to $r_l = 0$ and propagation distance $2r_A$ between the pixel at O and the platform at A . The potential spatial-frequency support for the reflectivity estimates at all output ranges r_l is the same finite interval specified in (2.36), with offset k_{r0} from (2.33) determined by the effective radar centre frequency f_c , and length Δk_r in (2.37) limited by the effective radar bandwidth B_c .

Since the output consists of complex samples in the spatial domain with spatial bandwidth Δk_r , the required sampling rate is Δk_r radians per metre, or equivalently

in the time domain, B_c samples per second. The minimum spacing between samples at ranges r_l is therefore $2\pi/\Delta k_r$ metres.

Observe that (2.38) is linear and, for times within all of the $\text{rect}()$ windows, is time invariant. Therefore, ideal radar data collection together with matched filtering is a linear time-invariant system which is characterised completely by its impulse response [Carlson 1998 p. 98]. In this case, an impulse is a single point target $g_p(r) = \delta(r - r_0)$ at range r_0 inside the range gate with unit magnitude and zero phase, and the impulse response is the range point-spread function $s_{rpsf}[l]$, given by


$$\begin{aligned} s_{rpsf}[l] &= e^{-jk_{r0}(r_A - r_l)} \sum_{\lambda} \int_r \delta(r - r_0) e^{jk_r r} \text{rect}\left(\frac{\tau_r - t + T_c/2}{T_c}\right) dr e^{-jk_r r_l} \text{rect}\left(\frac{k_r - k_{r0}}{\Delta k_r}\right) \\ &= e^{-jk_{r0}(r_A - r_l)} \sum_{\lambda} e^{-jk_r(r_l - r_0)} \text{rect}\left(\frac{k_r - k_{r0}}{\Delta k_r}\right) \quad (\text{analogous to (2.82)}) \\ &= e^{-jk_{r0}(r_A - r_l)} \Delta k_r e^{-jk_{r0}(r_l - r_0)} \text{sinc}\left(\frac{\Delta k_r}{2\pi}(r_l - r_0)\right) \quad (\text{analogous to (2.90)}) \\ &= e^{-jk_{r0}(r_A - r_0)} \Delta k_r \text{sinc}\left(\frac{\Delta k_r}{2\pi}(r_l - r_0)\right). \end{aligned} \quad (2.39)$$

The expressions here for the range dimension match the form of the expressions presented in Section 2.4 for a generic rectangular spectral support and subsequent $\text{sinc}()$ point-spread function. Combining the expression for the range bandwidth in (2.37) with the expression for the 3 dB resolution of a $\text{sinc}()$ in (2.92), and allowing for a window with scaling factor w (see Appendix 2.A.4), the range resolution ρ_{rg} is

$$\rho_{rg} = w \frac{2\pi}{\Delta k_r} = w \frac{2\pi}{(4\pi B_c/c)} = w \frac{c}{2B_c}. \quad (2.40)$$

The range point-spread functions (without windowing) for the *Ingara* L-band and *Bright Sapphire* P-band radars are shown in Figure 2.7. Both curves follow the standard $\text{sinc}()$ sidelobe pattern, with first sidelobe 13.3 dB below the peak. Evaluating (2.40) using the parameters in Table 1.1, the *Ingara* L-band bandwidth of 140 MHz gives a range resolution of 0.95 m, whereas the *Bright Sapphire* P-band effective bandwidth of 63.5 MHz supports a range resolution of only 2.09 m.

Figure 2.7: Range point-spread functions for the *Ingara* L-band (blue) and *Bright Sapphire* P-band (green) radars, with resolutions (mainlobe 3 dB widths) 0.95 m and 2.09 m, respectively. Both curves follow the standard $\text{sinc}()$ sidelobe pattern, with first sidelobe 13.3 dB below the peak.



The range gate $\tau_s \leq t < \tau_s + T_s$, represented here by the $\text{rect}()$ window for τ_s in (2.38), will cut short the received echoes at near and far ranges. This reduction of the effective duration will cause a proportional reduction in (a) the length of time that the matched filter and such echoes overlap, (b) the consequent range bandwidth Δk_r (see the discussion after (2.37)) and therefore (c) the magnitude of the point-spread function in (2.39). To appreciate this last point, consider the effect of the range gate on the normalised peak power $|s_{rpsf}|^2$ of the point-spread function in (2.39): echoes whose propagation delay $\tau_s \leq \tau_r < \tau_s + T_s - T_c$ places them fully inside the range gate will have full range bandwidth, as given by (2.37), and unity normalised power; echoes from delays $\tau_r = \tau_s - T_c(1 - 1/\sqrt{2})$ (near-range) and $\tau_r = \tau_s + T_s - T_c/\sqrt{2}$ (far-range) will be cut-short by $1 - 1/\sqrt{2} \approx 29\%$, leading to a normalised power of $1/2$; echoes from delays $\tau_r = \tau_s - T_c/2$ (near-range) and $\tau_r = \tau_s + T_s - T_c/2$ (far-range) will be cut-short by 50%, leading to a normalised power of $1/4$; echoes from delays $\tau_r < \tau_s - T_c$ (near-range) and $\tau_r > \tau_s + T_s$ (far-range) will not be received in the range gate at all. This is illustrated in Figure 2.8.

To process the datasets for this project, matched filtering was implemented in the frequency domain by multiplying the spectra of the demodulated data $s_b[l]$ in




Figure 2.8: Normalised peak powers of the range point-spread function in (2.39) for point targets at different propagation delays. The delay to the scene centre is τ_0 . At near and far ranges, the range gate $\tau_s \leq t < \tau_s + T_s$ cuts off the echo, leading to reduced range bandwidth and peak power.

(2.24) and the matched filter in (2.27), and then inverse transforming. This effectively implements a circular convolution instead of the desired linear convolution, because a spectrum generated by the discrete Fourier transform is implicitly periodic [Carlson 1998 ch. 16.5–16.6]. In order to ensure that the consequent circular wrap (which is effectively an alias in time) does not corrupt the desired output extent, the time-domain input signals, of duration T_s and T_c , are both zero-padded to duration $T_s + T_c$ before transformation to the frequency domain, and the time-domain output is cropped to T_s , thus removing the aliased extent T_c . Following the proposal in Appendix 2.A.9, the cropped portion is selected such that only the output lags for which at least half the echo was inside the range gate are retained; this corresponds to delays $\tau_s - T_c/2 \leq \tau_l < \tau_s + T_s - T_c/2$ (length T_s), as expressed in (2.125). From the prior discussion of the range gate, the near-range and far-range portions of this range swath will have reduced spatial bandwidth Δk_r , leading to both degraded resolution by (2.40) and reduced power as indicated in Figure 2.8.

To summarise, matched filtering compresses the chirp echoes to give a range profile of the scattering responses from the scene. Point targets are represented according to the point-spread function given by (2.39); the mainlobe width depends inversely on the effective chirp bandwidth B_c , and is quantified by the range resolution ρ_{rg} in (2.40). Moreover, from (2.38), the radar and matched filter together produce a filtered, narrowband version of the projected scene reflectivity as a function of range.

A comparable derivation from the received signal in (2.21) to the match-filtered output in (2.38), with an equivalent interpretation in terms of spatial frequencies, could not be found in the standard SAR reference texts. Similar interpretations are arrived at in spotlight-mode SAR references, but they treat the case where the received signal is dechirped as part of the analogue demodulation before sampling (see Appendix 2.A.8.2), so the mathematical logic is different [Munson *et al.* 1983, Jakowatz *et al.* 1996]. Those texts which treat the chirp echoes directly, as is done here, do not explicitly formulate range compression in such spatial-frequency terms [Curlander & McDonough 1991, Cumming & Wong 2005].

The range compressed pulses can now be coherently combined to resolve the scene in azimuth i.e. along the flight-track.

2.10 Focusing via beamforming

To form the 2D radar image, resolved in azimuth (x) as well as ground-range (y), the range-compressed pulses are now coherently combined using a straightforward beamforming approach. An explanation for exactly how combining pulses achieves azimuth resolution is delayed until later sections, because it requires spatial-frequency concepts that are not explicitly used during beamforming. It is sufficient for now to note from (2.61) that the azimuth resolution ρ_{az} is inversely proportional to the angular extent $\Delta\theta_{sl}$ subtended by the synthetic aperture i.e. the length of flight-track from which the combined pulses were acquired. In this section, the beamforming algorithm is detailed.

Recall from (2.38), each match-filtered pulse sample $s_m[l]$ is a filtered estimate of the reflectivity at range r_l , scaled by $e^{-jk_{r0}(r_A - r_l)}$ whose phase encodes the corresponding propagation distance $2(r_A - r_l)$ counted in units of wavelength ($k_{r0} = 4\pi/\lambda$) modulo

2π . Focusing each pixel simply involves computing the weighted sum of the reflectivity estimates across all pulses in the selected aperture, where each estimate is just the pulse sample at the corresponding platform-pixel range r_A (i.e. $r_l = 0$), and each weight is just the conjugate of the expected phase for that range. Cancelling the range phase in this way for all pulses is equivalent to spotlight-mode motion compensation to a point [Carrara *et al.* 1995 ch. 2.5.3, 2.6.6]. Crucially, the beamforming process advocated here repeats the motion compensation for every pixel; recomputing the platform-pixel range like this is necessary to properly focus pulse data from a general, non-linear, flight-track onto a general, non-planar, focal surface. The quality of the focusing is limited by the accuracy of the range information; the platform position and pulse timing must be carefully measured and recorded during collection, in order that the pulse contributions from each pixel position are correctly extracted and compensated.

In beamforming terms, pulses in the synthetic aperture are equivalent to the data received by elements of a linear sensor array, and motion compensation is analogous to compensating for the differing path delays across the array given an incident wave [Van Trees 2002 ch. 2]. In standard beamforming, the phase delays are computed for a particular angle of arrival which is fixed for all array elements, while for (spotlight-mode) SAR processing, the phase delays are computed for ranges which vary pulse-to-pulse, and what is fixed is the focal position. Standard beamforming searches the angular space by repeatedly steering to different angles, whereas SAR beamforming focuses the image pixels by repeatedly steering to different pixel positions. The essential task of compensating for the differential propagation delays between pulses/elements, thereby enabling coherent addition and azimuthal/angular resolution, is the same.

Unlike the frequency-domain methods summarised in Section 2.1, this time-domain (or better, spatial-domain) approach to image focusing places no limits on the resolution, swath size, operating range or flight-track. As long as the pulse positions and timings are known, the beamforming algorithm will properly focus point targets over the whole scene, because the pulses are reprocessed (re-steered) for each pixel. In addition, because the algorithm compensates for the absolute range between the platform and the scene, the complex output pixels will have no phase contribution due to the propagation delay to the focal surface; by contrast, all of the frequency-domain methods effectively compensate for only the differential range across pulses, and therefore generate images with a residual ‘flat-earth’ phase component due to the propagation delay to the flat focal plane [Bamler & Hartl 1998, Jakowatz *et al.* 1996 p. 322].

Compare the beamforming approach of repeated motion compensation and summation to the polar format algorithm in particular, which does motion compensation to one reference point, the scene centre, (by applying a range-dependent phase ramp to each pulse in the frequency domain such that the echo from the reference point is centred in the receive window) before resampling and inverting the pulse data. Only the reference point is compensated for the (differential) range phase across pulses, so only this point is focused exactly; the focusing of all other pixels relies on the plane-wave approximation [Jakowatz *et al.* 1996 ch. 2.6], as discussed in Section 2.14. In addition, the Fourier inversion is commonly accompanied by half-length circular wraps (a.k.a. FFT-shifts) such that the scene centre is the first output sample (in a sense, it is the DC position), which means that the corresponding output value is simply the sum of the motion-compensated input data. Therefore, beamforming each pixel is equivalent to applying the polar-format algorithm using that pixel position as the centre reference point and then retaining only the centre value [Pincus *et al.* 2015b].

The flexibility offered by focusing the pixels independently permits their synthetic apertures to be different. From (2.61), the azimuth resolution depends on the subtended angular width $\Delta\theta_{sl}$, which will vary across the swath if the aperture is fixed. In the beamforming algorithm proposed and implemented here, a synthetic aperture is designed for each pixel which achieves the specified azimuth resolution by accommodating the collection geometry at that pixel position. This typically means that the aperture is longer for pixels at farther range, in order to keep the subtended angle constant. In a similar vein, the aperture design for each pixel can take into account a specified image bearing (relative to North) or one of the types of squint depicted in Figures 2.1(c)-2.1(e). In the common case of broadside image formation, the aperture centre is shifted along the flight-track for successive pixels along azimuth, thereby maintaining zero squint for all pixels in the swath—this feature is akin to traditional stripmap imaging, where the imaged swath runs beside the platform. Hence, for each pixel, not only

are the pulses processed differently, but the set of contributing pulses will be different.

Let the output image consist of $N_{az} \times N_{gr}$ pixels, with pixel spacings d_{az} and d_{gr} (in metres) along azimuth and ground-range, so the image dimensions are $D_{az} = N_{az}d_{az}$ and $D_{gr} = N_{gr}d_{gr}$. Assign each pixel a height h , either derived from external elevation data or set to a constant for the image; the set of heights defines the focal surface.

The beamforming algorithm for SAR image focusing is detailed in Algorithm 2.1 below. Note that each pulse is significantly upsampled to give an interpolated range profile. Windows can be applied to control range and azimuth sidelobes, as described in Appendix 2.A.4.

Algorithm 2.1 SAR image formation via beamforming

output data: A SAR image consisting of a set of focused complex pixels $\{s_f\}$, each at their specified position on the focal surface.

output specification: Image bearing, squint, azimuthal resolution, window type, dimensions and pixel spacings along azimuth and ground-range and the pixel heights.

input data: A set of demodulated and match-filtered pulses $\{s_m[l]\}$, each consisting of a set of complex samples of the range-compressed radar echo, as given by (2.38).

input metadata: The 3D position and the transmit and receive timing of every supplied pulse. The pulse positions should be specified in a coordinate frame whose origin is at the scene centre, as depicted in Figure 2.1(a).

- 1: **for all** pixels **do** #Aperture
- 2: Initialise $s_f \leftarrow 0$.
- 3: Design a synthetic aperture which achieves the specified bearing, squint and azimuthal resolution (by inverting (2.61)), allowing for the broadening effect of the window and noting the constraint of the beampattern.
- 4: **end for**
- 5: **for all** pulses **do** #Range interpolation
- 6: Obtain the pulse spectrum via an inverse FFT.
- 7: Apply a window w_{gr} to control range sidelobes.
- 8: Zero-pad the spectrum such that the total extent $\Delta k_{r,zp}$ supports a very fine pixel spacing $2\pi/\Delta k_{r,zp}$, say five times smaller than the smallest output pixel spacing.
- 9: Obtain the upsampled pulse $s_{m,zp}[l]$ via an FFT. The range profile is now interpolated onto a finely sampled range line.
- 10: **for all** pixels **do** #Beamforming
- 11: **if** this pulse contributes to the aperture for this pixel **then**
- 12: Compute the range r_A from pixel to platform.
- 13: Compute the corresponding index l_r into the upsampled range profile by inverting (2.30) using the new sample spacing.
- 14: Compute the range phase $\phi_r = k_{r0}r_A$. Other phase offsets such as the start-sampling delay can be included here if not accounted for during digital demodulation (see Appendix 2.A.10.1).
- 15: Compute the window value w_{az} to control azimuth sidelobes. This depends on the relative position of this pulse in this pixel's aperture.
- 16: Accumulate the weighted pulse sample into the complex pixel:
 $s_f \leftarrow s_f + w_{az}e^{j\phi_r} s_{m,zp}[l_r]$.
- 17: **end if**
- 18: **end for**
- 19: **end for**
- 20: Baseband the image spectrum by applying a 2D phase ramp to the focused image (see Section 2.15).

Observe that the pulses are processed independently, so steps 6-18 can be executed in parallel for different pulses, although each thread will need memory for its own version of the output image (to avoid the need for mutex locks), and these partial images must then be summed to obtain the output image. Alternatively, the upsampled pulses could be precomputed, and then the pixels processed in parallel, thus avoiding the need for partial images, at the expense of storing all the upsampled pulses; for high-resolution images, the memory trade-off usually favours parallel processing of the pulses.

The beamforming method of image focusing described here is exact in the sense that it compensates exactly for the platform-pixel range for each pulse-pixel pair. Indeed, only the range information is used, not the angular orientation, and the approximate (plane-wave dependent) concepts of 3D k -space and aperture synthesis, introduced in Sections 2.3–2.4 and further developed in Sections 2.11–2.13, are not relied upon, except at step 3 as a means to determine the aperture length. Nonetheless, the angular orientation of the locally valid spatial-frequency support is critical to understanding the image spectrum and the point-spread function, as will be seen in Section 2.15.

Figure 2.10 shows a beamformed SAR image of the radar calibration site in Adelaide, South Australia, using *Ingara* L-band data. Figure 2.9 shows corresponding optical views of the scene. Several corner reflectors are present: ordinary trihedrals, gridded trihedrals [Sheen *et al.* 1992] and dihedrals, although the latter are not visible because the displayed polarimetric channel is LR (left-circular transmit, right-circular receive), and double-bounce of circularly polarised waves generates no cross-polarisation [Cloude 2010 ch. 1.4.2]. No windowing has been applied, so the sidelobe leakage from the bright corner reflectors clearly marks out the point-spread function. The sidelobes are rotated due to the offset squint $\delta_{off} = -5.5^\circ$, which was selected so as to cancel out the flight-track squint $\delta_{ftk} = 1.5^\circ$ and the yaw of 4° , thereby forming the image on the antenna boresight. The azimuth sidelobes are split into an apparent bow tie because the beamformer integrates along a curved contour in the spatial-frequency domain, rather than a straight line (see Figure 2.19 for a simulated example).

Figure 2.12 shows a beamformed SAR image of the Tully Training Area, consisting of a grass-covered clearing surrounded by tropical forest in Northern Queensland (latitude 17.795 S), using *Bright Sapphire* P-band data. Figure 2.11 shows corresponding optical views of the scene. Three trihedral corner reflectors (all with edge length 2.45 m) are present: one in the open clearing (TH_c) and two partially hidden along a narrow jungle track (TH₁ and TH₂). No windowing has been applied, so the sidelobe leakage from the bright, unobscured corner reflector marks out the point-spread function. The scattering response from the clutter at P-band has low dynamic range: the cleared ground gives a very low response (on par with the noise even after the coherent gain from range compression and azimuth integration), and the forest appears as a homogeneous random medium whose individual tree components cannot be distinguished. Exceptions are some spikes due to ground-trunk double bounce from unobscured trees at the bottom-left edge of the clearing. Note that the displayed dynamic range for the *Bright Sapphire* image is only 40 dB, whereas for the *Ingara* image it is 60 dB. The hidden trihedrals are visible above the clutter, indicating a level of foliage penetration, although note that they are visible in images generated using the L-band channel as well [Pincus *et al.* 2013]. Unfortunately for this experiment, the forest canopy was not as tall and dense as usual due to damage from Cyclone Yasi fifteen months previous [Bureau of Meteorology 2011].

The eventual goal is to apply interferometric processing to a pair of SAR images. Of central concern, therefore, is the form of the image spectrum, and how this relates to the acquired spatial-frequency support. The analysis begins by extending the one-dimensional range formulation from Section 2.9 to three dimensions.

2.11 Pulse support in the 3D spatial-frequency domain

The model from Section 2.9 for the received, demodulated and match-filtered radar data is now converted from a one-dimensional formulation along range and range spatial-frequency to a three-dimensional Cartesian reference frame in space and k -space (the 3D spatial-frequency domain). This formulation is called the tomographic paradigm in the spotlight-mode SAR literature [Munson *et al.* 1983, Jakowatz *et al.* 1996 ch. 2.4]. Multiple pulses can then be interpreted as together providing extended spatial-frequency support in 3D k -space, in accordance with the generic model outlined in Section 2.4. Importantly, this also enables the underlying reflectivity $g(x, y, z)$ to be directly represented, instead of the projected version $g_p(r)$ from (2.20).

Begin by approximating the spherical wavefront, which was used to model the received signal in Section 2.7, as a plane. Using the Hessian normal form described in Appendix 2.A.7, the points $\mathbf{x} = (x, y, z)^T$ on the plane must satisfy $\mathbf{x} \cdot \hat{\mathbf{r}}_A = r$ where the unit normal $\hat{\mathbf{r}}_A$, given by (2.5), points in the range direction, from the pixel origin O toward the acquisition point A at (r_A, θ_A, ψ_A) , and the plane is offset from O by r .



(a) Satellite image (Google Earth, 20 June 2013). The offset squint to aperture centre is $\delta_{off} = -5.5^\circ$.



(b) Trihedral (TH) (edge length: 2.45 m) (c) Right-gridded trihedral (GT_r) (d) Left-gridded trihedral (GT_l) (zoom) (e) Dihedral (DH) (0.54 m) (f) Rotated dihedral (RD) (22.5°)

Figure 2.9: Optical views of the ground scene and corner reflectors in Figure 2.10.




Figure 2.10: *Ingara* L-band SAR image of the corner reflector calibration site (8 July 2013, 500×870 m, polarimetric channel LR, 60 dB dynamic range, no window). (Inset: 7×7 m, 30 dB dynamic range.) The nominal depression angle $\psi_c = 34.2^\circ$ gives a nominal ground-range resolution $\rho_{gr} = 1.15$ m. (The depression angle varies by 10° over the swath due to the close operating range $r_c = 2.8$ km.) The synthetic aperture for each pixel was designed to subtend an azimuthal angle $\Delta\theta_{sl} = 7.2^\circ$, giving a constant azimuthal resolution $\rho_{az} = 0.8$ m. Note the well-focused point responses from the buildings at far-range (bottom of image) and the four goal posts on the Australian-style football oval at near-range (top-right).



(a) Satellite image (Google Earth, 27 May 2011). Tropical forest surrounds the grassed clearing.



(b) Tower (TWR) (c) Clear trihedral (TH_c) (d) Hidden trihedral (TH₁) and the (e) Hidden trihedral jungle track (not visible in sat. img) (TH₂)

Figure 2.11: Optical views of the ground scene and corner reflectors in Figure 2.12.





Figure 2.12: *Bright Sapphire* P-band SAR image of the Tully Training Area (7 May 2012, 440×320 m, polarimetric channel HH, 40 dB dynamic range, no window). (Inset: 6×11 m, 30 dB dynamic range.) The nominal depression angle $\psi_c = 55.8^\circ$ gives a nominal ground-range resolution $\rho_{gr} = 3.72$ m. (The depression angle varies by 7.8° over the swath due to the close operating range $r_c = 1.95$ km.) The synthetic aperture for each pixel was designed to subtend an azimuthal angle $\Delta\theta_{sl} = 12.5^\circ$, giving a constant azimuthal resolution $\rho_{az} = 1.8$ m.

Such a plane was introduced in Section 2.3 and is shown again in Figure 2.13. The plane is tangent to the spherical wavefront at the point $r\hat{\mathbf{r}}_A$ (the point on the plane closest to O), so at this point, there is no approximation error.

Figure 2.13: Plane-wave approximation whereby the spherical wavefront of radius $r_A - r$, centred on the radar platform at A , is approximated by a planar wavefront tangential to the spherical wavefront at r on the range axis.



The projected reflectivity $g_p(r)$, expressed in (2.20) for a projection over the spherical wavefront, can be approximated using a projection over the tangent plane.

$$g_p(r) \approx \int_z \int_y \int_x \delta(r - \mathbf{x} \cdot \hat{\mathbf{r}}_A) g(x, y, z) dx dy dz \quad (2.41)$$

Note that this approximation enables the formulation of the concept of an extended spatial-frequency support in 3D k -space, but it is not necessary for time-domain image formation; beamforming will use the true range for each pulse-pixel pair.

Substituting (2.41) into the range inverse Fourier transform in (2.38) and dropping the $\text{rect}()$ window for brevity gives

$$\int_r g_p(r) e^{jk_r r} dr \approx \int_z \int_y \int_x \int_x \delta(r - \mathbf{x} \cdot \hat{\mathbf{r}}_A) g(x, y, z) e^{jk_r r} dr = G_{(\theta_A, \psi_A)}(k_x, k_y, k_z) \quad (2.42)$$

where the last equality follows by invocation of the projection-slice theorem, derived in Appendix 2.A.6. Thus, the one-dimensional Fourier transform along range of the projected reflectivity, where the projection is onto a line oriented at (θ_A, ψ_A) (indicated by $\hat{\mathbf{r}}_A$), is equal to the three-dimensional Fourier transform of the original reflectivity function evaluated along a line oriented at (θ_A, ψ_A) in k -space. Application of the projection-slice theorem requires that the summation in the projection be along straight lines or across planes, not over curved contours, hence the need for the plane-wave approximation.

From (2.38), the available support for the line $G_{(\theta_A, \psi_A)}(k_x, k_y, k_z)$ is limited to samples (indexed by λ) of the radial line segment $k_{r0} - \Delta k_r/2 \leq k_r < k_{r0} + \Delta k_r/2$.

Putting all this together, each collected, demodulated and match-filtered pulse, of the form in (2.38), has a spatial-frequency support which is a radial line segment at a particular orientation in 3D k -space. The radial extent Δk_r and offset k_{r0} are fixed by the radar waveform according to (2.33) and (2.37) respectively, but the orientation of the line segment depends on the instantaneous platform orientation (θ_A, ψ_A) with respect to each pixel origin, as indicated in (2.42); this orientation defines the orthogonal projection which generated $g_p(r)$.

Crucially, it is now apparent that collecting multiple pulses over the same scene along a flight-track would extend the spatial-frequency support by providing a variety of line segments at different orientations (θ_A, ψ_A) . For a general flight-track, the set of pulses will mark out a ribbon surface, as illustrated in Figure 2.14. If the flight-track were perfectly straight and level, then this surface would be planar [Ausherman *et al.* 1984, Jakowatz *et al.* 1996 pp. 71,187-188]; this is the impetus for the polar-format algorithm, discussed further in Section 2.14.

Note again that the geometry of the ribbon of spatial-frequency support depends on which pixel serves as the reference origin. In order to determine this geometry precisely for a given a set of pulses, their orientations (θ_A, ψ_A) would ideally be recomputed for every pixel; for adjacent pixels, the reorientation of the pulse supports in k -space would be very small, but across an image, the variation could be significant, especially if the swath extent and the operating range were of a similar order of magnitude, or if the focal surface had substantial changes in height.




Figure 2.14: Ribbon of spatial-frequency support in 3D k -space, marked out by multiple sampled lines of pulse supports. All lines have radial offset k_{r0} given by (2.33) and radial length Δk_r given by (2.37). Each line has a different orientation (θ_A, ψ_A) depending on the platform position along the flight-track relative to a reference pixel. The ribbon undulations have been exaggerated to emphasise that the pulse supports are usually not co-planar.

It can now be seen that the radar data does actually exhibit the kind of spatial-frequency support described in a general way in Section 2.4. A question remains as to exactly how the pulses in 3D k -space provide spatial-frequency support for the pixels in the 2D output image.

2.12 The slope-dependent projection yielding the 2D aperture

A 2D radar image will be displayed and post-processed in the horizontal image plane, even if the constituent pixels are focused to different heights, as discussed in Section 2.2. The k_x - k_y plane is the frequency space corresponding to the output image plane. The key quantity to be ascertained is the effective aperture, that is, the 2D spatial-frequency support which gives rise to resolution in the image plane, and more generally, determines the output point-spread function. This aperture results from an implicit projection of the pulse supports in 3D k -space onto the k_x - k_y plane. Crucially, the projection is in the direction normal to the local tangent plane at each pixel on the focal surface.

Recall from Section 2.2 that, given a pixel position with local azimuth and ground-range slopes α_{az} and α_{gr} , a local tangent plane was defined in (2.7) with normal \mathbf{l}_n specified by (2.6). Say that a pulse sample contributing to the aperture for this pixel has Cartesian and spherical coordinates (k_{xA}, k_{yA}, k_{zA}) and (k_r, θ_A, ψ_A) , respectively, in k -space. The sample projects onto the k_x - k_y plane at the intersection of the line

$$\begin{bmatrix} k_x \\ k_y \\ k_z \end{bmatrix} = \begin{bmatrix} k_{xA} \\ k_{yA} \\ k_{zA} \end{bmatrix} + \nu \begin{bmatrix} \tan \alpha_{az} \\ \tan \alpha_{gr} \\ 1 \end{bmatrix}, \quad \nu \in \mathbb{R} \quad (2.43)$$

and the plane $k_z = 0$, which occurs at

$$\begin{bmatrix} k_{xprj} \\ k_{yprj} \\ k_{zprj} \end{bmatrix} = \begin{bmatrix} k_{xA} - k_{zA} \tan \alpha_{az} \\ k_{yA} - k_{zA} \tan \alpha_{gr} \\ 0 \end{bmatrix} = \begin{bmatrix} k_r \cos \psi_A (\sin \theta_A - \tan \psi_A \tan \alpha_{az}) \\ k_r \cos \psi_A (\cos \theta_A - \tan \psi_A \tan \alpha_{gr}) \\ 0 \end{bmatrix}, \quad (2.44)$$

where the last equality was obtained using the coordinate conversions in (2.10).

To provide an intuitive feel for this projection, define the equivalent acquisition orientation (θ_{eq}, ψ_{eq}) which, for the same range spatial-frequency $k_r = 4\pi/\lambda$, would project onto the same point in the k_x - k_y plane if the slopes were zero.

$$k_{xprj} = k_r \cos \psi_{eq} \sin \psi_{eq} = k_r \cos \psi_A (\sin \theta_A - \tan \psi_A \tan \alpha_{az}) \quad (2.45)$$

$$k_{yprj} = k_r \cos \psi_{eq} \cos \psi_{eq} = k_r \cos \psi_A (\cos \theta_A - \tan \psi_A \tan \alpha_{gr}) \quad (2.46)$$

Dividing (2.45) by (2.46) gives

$$\tan \theta_{eq} = \frac{\sin \theta_A - \tan \psi_A \tan \alpha_{az}}{\cos \theta_A - \tan \psi_A \tan \alpha_{gr}}. \quad (2.47)$$

Squaring and summing (2.45) and (2.46) gives

$$\begin{aligned} k_r^2 \cos^2 \psi_{eq} &= k_r^2 \cos^2 \psi_A [(\sin \theta_A - \tan \psi_A \tan \alpha_{az})^2 + (\cos \theta_A - \tan \psi_A \tan \alpha_{gr})^2] \\ \therefore \cos \psi_{eq} &= \sqrt{(\sin \theta_A - \tan \psi_A \tan \alpha_{az})^2 + (\cos \theta_A - \tan \psi_A \tan \alpha_{gr})^2} \cos \psi_A. \end{aligned} \quad (2.48)$$

Now consider the special case when $\theta_A = 0^\circ$, which occurs for the centre of a broadside synthetic aperture. If, in addition, only one of the slopes is non-zero, the expressions for the transformed angles take simple forms, as shown below.

$$\theta_A = 0^\circ, \alpha_{gr} = 0^\circ \Rightarrow \tan \theta_{eq} = -\tan \psi_A \tan \alpha_{az}, \quad \cos \psi_{eq} = \sqrt{1 + (\tan \psi_A \tan \alpha_{az})^2} \cos \psi_A \quad (2.49)$$

$$\theta_A = 0^\circ, \alpha_{az} = 0^\circ \Rightarrow \quad \theta_{eq} = 0^\circ, \quad \cos \psi_{eq} = (1 - \tan \psi_A \tan \alpha_{gr}) \cos \psi_A \quad (2.50)$$

Hence, azimuthal slope is equivalent to a change in the azimuth collection angle and a small change in grazing angle, whereas ground-range slope is equivalent to only a change in grazing angle. Note that $\psi_{eq} \neq \psi_A + \alpha_{gr}$. In addition, note that an equivalent orientation (θ_{eq}, ψ_{eq}) satisfying (2.45)–(2.48) does not always exist. In particular, ψ_{eq} must satisfy $|\cos \psi_{eq}| \leq 1$ —of course, any combination of slope and orientation is possible in reality, the test here determines only whether a hypothetical equivalent orientation exists. Considering the special case in (2.49), α_{az} must therefore satisfy

$$(\tan \psi_A \tan \alpha_{az})^2 \leq \frac{1}{\cos^2 \psi_A} - 1 = \tan^2 \psi_A \Rightarrow |\tan \alpha_{az}| \leq 1 \Rightarrow |\alpha_{az}| \leq 45^\circ, \quad (2.51)$$

and considering the special case in (2.50), α_{gr} must therefore satisfy

$$\begin{aligned} -1 - 1/\cos \psi_A &\leq -\tan \psi_A \tan \alpha_{gr} \leq -1 + 1/\cos \psi_A \\ \Rightarrow (-1 + \cos \psi_A)/\sin \psi_A &\leq \tan \alpha_{gr} \leq (1 + \cos \psi_A)/\sin \psi_A \\ \Rightarrow \tan(-\psi_A/2) &\leq \tan \alpha_{gr} \leq 1/\tan(\psi_A/2) \\ \Rightarrow -\frac{\psi_A}{2} &\leq \alpha_{gr} \leq 90^\circ - \frac{\psi_A}{2} \quad \left(\arctan \frac{1}{x} = \begin{cases} 90^\circ - \arctan x & \text{if } x > 0 \\ -90^\circ - \arctan x & \text{if } x < 0 \end{cases} \right). \end{aligned} \quad (2.52)$$

Figure (2.15) illustrates how the equivalent grazing angle ψ_{eq} varies with the ground-range slope α_{gr} for the special case in (2.50). As expected, when $\alpha_{gr} = 0^\circ$, $\psi_{eq} = \psi_A$. The linear approximation $\psi_{eq} \approx \psi_A + \alpha_{gr}$ is valid only for small slopes i.e. $|\alpha_{gr}| \lesssim 5^\circ$, although the error decreases as ψ_A increases.

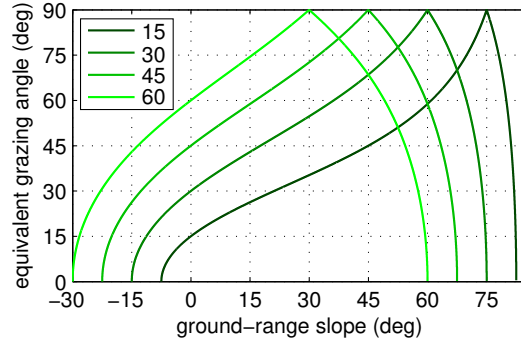


Figure 2.15: Equivalent grazing angle ψ_{eq} as a function of ground-range slope α_{gr} for different actual grazing angles ψ_A . This shows the special case in (2.50), when the acquisition is broadside and the azimuthal slope is zero. The curve for each value of ψ_A is restricted to the interval specified in (2.52).

From Section 2.11, a pulse provides a radial line of spatial-frequency support which starts and ends at radial distances k_{rs} and k_{re} ,

$$k_{rs} = k_{r0} - \Delta k_r / 2 \quad (2.53)$$

$$k_{re} = k_{r0} + \Delta k_r / 2, \quad (2.54)$$

where k_{r0} and Δk_r were defined in (2.33) and (2.37), respectively. Using (2.44), the 2D (k_x, k_y) coordinates I_{ps} and I_{pe} of the start and end of this line when projected onto the k_x - k_y plane are

$$I_{ps} = (k_{rs} \cos \psi_A (\sin \theta_A - \tan \psi_A \tan \alpha_{az}), k_{rs} \cos \psi_A (\cos \theta_A - \tan \psi_A \tan \alpha_{gr})) \quad (2.55)$$

$$I_{pe} = (k_{re} \cos \psi_A (\sin \theta_A - \tan \psi_A \tan \alpha_{az}), k_{re} \cos \psi_A (\cos \theta_A - \tan \psi_A \tan \alpha_{gr})) \quad (2.56)$$

Denote the projected line of pulse support by the ordered pair

$$I_p = (I_{ps}, I_{pe}) \quad (2.57)$$

and denote the effective aperture I as the set of all contributing pulse supports i.e.

$$I = \{I_p\}. \quad (2.58)$$

Formally, $I(k_x, k_y)$ is an indicator function over the 2D spatial-frequency domain; it is one where there is projected pulse support and zero everywhere else.

The orientation (θ_A, ψ_A) of the acquisition and the slopes $(\alpha_{az}, \alpha_{gr})$ are dependent on which pixel is serving as the reference point, so the effective aperture and the associated resolution will vary pixel-to-pixel. This spatial variation has long been recognised [Gatelli *et al.* 1994], although it is often ignored. When the image is focused onto an undulating surface (permitted by time-domain methods and modern computing power), the variation is amplified. Before forming the image, the position metadata provides sufficient information to compute the set J of effective apertures for all pixels in the image; denote

$$J = \{I\}. \quad (2.59)$$

Computing these apertures is the starting point for the proposed method of matching the spatial-frequency supports across two images, to be presented in Section 2.16.

The relevant published works of Blacknell and Andre discuss, at a conceptual level, the projection of the spatial-frequency supports given an undulating focal surface; see in particular Figure 4 in [Blacknell *et al.* 2010], Figure 6 in [Blacknell *et al.* 2011] or the brief description on page 5 in [André *et al.* 2013]. However, they do not formulate the projected support mathematically and therefore do not provide a method to quantify the support for a SAR dataset.

The initial published formulation of this projection used a local tangent plane obtained by twice-rotating the (x, y, z) coordinate frame, which not only led to ambiguity over the rotation order (as discussed in Section 2.2) but also provided little justification for the seemingly ad hoc projection [Pincus *et al.* 2015b]. The approach presented here specifies the tilted plane directly, and uses the associated normal \mathbf{l}_n to define the projection in a straightforward and intuitive way.

Figure 2.16 depicts the slope-dependent projection, given by (2.44), of one pulse support onto the k_x - k_y plane. Of particular interest is the effect of ground-range slope, shown in (b): compared to the projected support I_p in the nominal case (horizontal; black), if the local surface were tilted down toward the platform (positive slope; red), then I_p would be shifted down k_y and, more importantly, would be smaller, whereas if the local surface were tilted up toward the platform (negative slope; blue), then I_p would be shifted up k_y and would be larger.




Figure 2.16: Projection of the discretised, radial line of spatial-frequency support for one pulse onto the k_x - k_y plane, giving effective pulse support I_p . The projection is parallel to the local surface normal axis \mathbf{l}_n , first depicted in Figure 2.1(b). For the simple case of a broadside pulse and zero azimuth slope, (b) shows how different ground-range slopes give rise to different projections.

The projected 2D aperture I in (2.58) is the precise basis for resolution in the 2D output image.

2.13 Aperture synthesis and resolution

The essence of synthetic aperture radar is the coherent processing of a set of pulses, each with one-dimensional spatial-frequency support, to synthesise a two-dimensional strip of spatial-frequency support—an aperture—and thereby resolve the scattering responses from the illuminated scene in two dimensions i.e. along both azimuth and range. Here, the type of coherent processing is beamforming, as detailed in Section 2.10, and the slope-dependent form of the aperture was specified in Section 2.12. In general, aperture synthesis amounts to nothing more than combining many samples of the spectrum of an unknown function, in order to reconstruct that function [Hoctor & Kassam 1990]. Physically, the function defines a spatial distribution of (reflecting) sources, and the spectral samples are obtained from spatially distributed measurements. It is implicitly assumed that the different measurements sample the same underlying spectrum i.e. that the scattering response is uniform over the measurement space. In synthetic aperture radar, the spectral samples from successive pulses are measured at different times (along the flight-track), but once acquired, their time separation is irrelevant because the scene is assumed constant; it is the angular diversity along the synthetic aperture that gives rise to azimuthal resolution.

The spatial resolution ρ of the reconstructed reflectivity is inversely proportional to the spatial bandwidth Δk , as discussed in Section 2.4 and Appendix 2.A.3. The extent of the spatial-frequency support in the range direction (the width of the ribbon in Figure (2.14)) is limited by the chirp bandwidth, according to (2.37). The extent in the azimuth direction (the length of the ribbon in Figure (2.14)) is limited by the angular spread $\Delta\theta$ of the contributing pulses.

The resolution supported by an effective aperture I can be quantified approximately by inscribing a rectangle inside the area covered by I ; using Appendices 2.A.3–2.A.4, a $\Delta k_x \times \Delta k_y$ rectangle weighted in k_x and k_y by a window with scaling factor w would support azimuthal and ground-range resolutions $\rho_{az} = w \cdot 2\pi / \Delta k_x$ and $\rho_{gr} = w \cdot 2\pi / \Delta k_y$.

For a nominal broadside flight-track and planar focal surface, simple expressions for the resolutions can be obtained. The extent of the support along k_y is just the range bandwidth Δk_r scaled by $y_c/r_c = \cos\psi_c$ to project the centre pulse (acquired at A_c , where $\theta = 0^\circ$) onto the k_y axis. The range resolution ρ_{rg} was given in (2.40); the ground-range resolution ρ_{gr} is

$$\rho_{gr} = w \frac{2\pi}{\Delta k_y} = w \frac{2\pi}{\Delta k_r \cos\psi_c} = w \frac{c}{2B_c \cos\psi_c}. \quad (2.60)$$

The extent of the support along k_x depends on the length of the synthetic aperture, which is parameterised here by the angular width $\Delta\theta$ subtended in the k_x - k_y plane. Using the offset $k_{r0} \cos\psi_c$ to the centre of the projected region for simplicity, and noting from (2.33) that $k_{r0} = 4\pi/\lambda$, the azimuthal resolution is

$$\rho_{az} = w \frac{2\pi}{\Delta k_x} \approx w \frac{2\pi}{k_{r0} \cos\psi_c \Delta\theta} = w \frac{\lambda}{2\Delta\theta \cos\psi_c} = w \frac{\lambda}{2\Delta\theta_{sl}} \quad (2.61)$$

where $\Delta\theta_{sl} = \Delta\theta \cos\psi_c$ is the angular width subtended in the slant plane [Jakowatz *et al.* 1996 ch. 2.4.4]. If the radar operates in stripmap mode, with no steering of the antenna beam, then $\Delta\theta_{sl}$ is limited by the antenna beamwidth $\Delta\theta_{ant} = w_{rect} \lambda / D_{ant}$, where D_{ant} is the horizontal length of the (uniformly weighted) antenna, since a pulse contributes to the synthetic aperture for a pixel only if it illuminates the pixel position properly [Carrara *et al.* 1995 ch. 2.4.2]. The maximum azimuthal resolution that can be achieved in stripmap collection mode is therefore

$$\rho_{az, stripmap} = w \frac{\lambda}{2\Delta\theta_{ant}} = \frac{w}{w_{rect}} \cdot \frac{D_{ant}}{2}. \quad (2.62)$$

Observe that these nominal resolutions are independent of range, but they do vary with pixel position, principally due to variation in the effective grazing angle ψ .

An example of image formation on undulating terrain now follows.

2.14 Image formation on undulating terrain

The geometric interpretation of spatial-frequency developed in Sections 2.11 and 2.12 can be related to the formulation developed by Ausherman *et al.* [1984] and made

widely-known by Jakowatz *et al.* [1996]. Although their discussion assumes a planar focal surface and is tailored to the well-known polar format algorithm (PFA), which achieves image formation in the frequency domain, the principles of Fourier-domain projections they present are equally applicable to time-domain beamforming for a non-planar focal surface. (Regardless of how an image is focused, limitations imposed by the form of the data in the frequency domain cannot be avoided.) This analysis is illustrated by a simulated example of image formation on undulating terrain.

The ribbon formed in 3D k -space by the lines of support for the acquired pulses, shown in Figure 2.14, is the *collection surface*. If the flight-track were straight and level, then this surface would be planar, otherwise it will be non-planar [Ausherman *et al.* 1984, Jakowatz *et al.* 1996 pp. 71,187-188]. The k_x - k_y plane is the *processing plane* which the collected data are projected onto (and, for PFA, where the rectangular grid would be defined for resampling prior to Fourier inversion). Fourier transforming a SAR image, whether formed by beamforming or PFA, generates the image spectrum in the processing plane. The local tangent plane is a pixel-specific version of the *focus plane*, whose normal defines the direction of the projection onto the processing plane. The collection surface, the processing plane, the focus plane and the projection which links them are illustrated in Figure 23 in Ausherman *et al.* [1984] and Figures 3.48 and 3.49 in Jakowatz *et al.* [1996]; the extension to account for slope was shown here in Figure 2.16.

For PFA, the projection is an explicit processing step undertaken for “out-of-plane correction” [Jakowatz *et al.* 1996 ch. 3.8.1] (actually, this just amounts to accounting for the true platform positions instead of assuming a nominal straight and level track). For beamforming, however, the projection occurs implicitly during pixel focusing, and is used here to explain the spatial-frequency support observed in the image spectrum.

An important result of this spatial-frequency projection is that scattering objects which lie in the focus plane (or rather, its spatial equivalent) will be well-focused regardless of the flight-track, whereas objects which do not lie in the focus plane will be well-focused only if the flight-track is straight and level [Ausherman 1985, Jakowatz *et al.* 1996 ch. 3.8.1]. The projection of a non-planar collection surface preserves the relative phase of the scattering response across pulses only for those objects which lie in a plane whose normal matches the direction of the projection; this is the focus plane.

Operating in the time-domain offers the flexibility to change the focus height of each pixel to match the assumed terrain. (In this work, the set of focus heights across all pixels is called the focal surface.) During beamforming, the focus plane is implicitly shifted and tilted to match the assumed local tangent plane at each pixel position, thereby bringing undulating terrain into focus, regardless of the flight-track, if the assumed focal surface indeed matches the height of the true scattering surface. Where the assumption is wrong, and the flight-track is not straight and level, the scattering objects will be defocused in proportion to their height offset, regardless of the image formation technique.

It is worth reiterating that the geometric interpretation presented here assumes that the plane-wave approximation is (locally) valid and therefore the data can be analysed in the “tomographic paradigm” [Jakowatz *et al.* 1996 ch. 2.6].

Figure 2.17 compares PFA and beamformed SAR images of a synthetic scene consisting of nine point scatterers spread over a ridgeline which is viewed obliquely by two different flight-tracks, one straight and level and the other suffering a moderate sinusoidal altitude variation. Four main results are apparent. Firstly, PFA and beamforming generate almost exactly the same output when the beamformer employs the same fixed planar focal surface used for PFA: when the flight-track is straight and level (as in (e) and (f)), then all targets are well-focused, and when the flight-track is not level (as in (h) and (i)), then the targets not in the plane are azimuthally defocused in proportion to their height offset. Secondly, when the beamformer is supplied with the true terrain elevations (as in (g) and (j)), then all targets are well-focused, regardless of the flight-track. Thirdly, targets not at the focus height (in (e), (f), (h) and (i)) lay over onto the focus plane according to the same projection principles [Jakowatz *et al.* 1996 Appendix C], regardless of the image formation method and regardless of whether they are properly focused or not. Finally, the point-spread functions of the well-focused targets in the beamformed images (in (g) and (j)) nonetheless depend on the local slope of the focal surface, due to the slope-dependent (i.e. focus-plane-dependent) projection of the pulse supports onto the k_x - k_y plane, quantified by (2.55)–(2.56). For example, the terrain under the top-left target (at zero height) has very

steep positive ground-range slope (38°), leading to an effective aperture with small k_y extent (akin to the red case in Figure 2.16(b)) and therefore very coarse ground-range resolution in the image plane, whereas the terrain under the bottom-right target (also at zero height) has very steep negative ground-range slope (-38°), leading to an effective aperture with large k_y extent (akin to the blue case in Figure 2.16(b)) and therefore very fine ground-range resolution in the image plane.

The way in which the acquired and projected spatial-frequency support manifests in the observed image spectrum is studied next.

2.15 Image spectrum

The image spectrum, obtained by a 2D Fourier transform of the complex SAR image, represents the spatial frequencies in the image; this is obvious. The question is, how do these observed frequencies relate to the band of spatial frequencies measured by the radar? In particular, how does the line of spatial-frequency support provided by each pulse in 3D k -space map into the 2D spectrum? The answer, in a nutshell, is that the support projects onto the k_x - k_y plane and then aliases into the observed spectrum. To see this, the frequency space represented by the image spectrum is now characterised.

The image spectrum serves as a window onto the baseband region of the k_x - k_y plane. This is the plane onto which the pulse supports in 3D k -space are projected, as specified by (2.57), together providing an aperture whose size determines the output resolution, as presented in Section 2.12. The Fourier transforms along azimuth and ground-range that generate the image spectrum are accompanied by pre- and post-FFT-shifts (half-length circular wraps), with the latter ensuring that DC is in the middle of the image spectrum. The observed spectrum will cover the unambiguous spatial-frequency area bounded by $\pm\Delta k_{az}/2$ and $\pm\Delta k_{gr}/2$ where

$$\Delta k_{az} = 2\pi/d_{az} \quad \text{and} \quad \Delta k_{gr} = 2\pi/d_{gr} \quad (2.63)$$

and will have sample spacings $\delta k_{az} = 2\pi/D_{az}$ and $\delta k_{gr} = 2\pi/D_{gr}$. (k_{az} refers to k_x and k_{gr} refers to k_y ; these subscripts are used to emphasise that the quantities apply to the image spectrum.)

A point (k_{x_0}, k_{y_0}) will alias into the image spectrum at (k_{az_0}, k_{gr_0}) where

$$\begin{aligned} k_{az_0} &= k_{x_0} - a_{az}\Delta k_{az} \text{ for some integer } a_{az} \text{ such that } -\Delta k_{az}/2 \leq k_{az_0} \leq \Delta k_{az}/2 \\ k_{gr_0} &= k_{y_0} - a_{gr}\Delta k_{gr} \text{ for some integer } a_{gr} \text{ such that } -\Delta k_{gr}/2 \leq k_{gr_0} \leq \Delta k_{gr}/2. \end{aligned} \quad (2.64)$$


This corresponds to image indices in the DC-centred spectrum

$$n_{az} = \lfloor k_{az_0}/\delta k_{az} \rfloor + N_{az}/2 \quad \text{and} \quad n_{gr} = \lfloor k_{gr_0}/\delta k_{gr} \rfloor + N_{gr}/2. \quad (2.65)$$

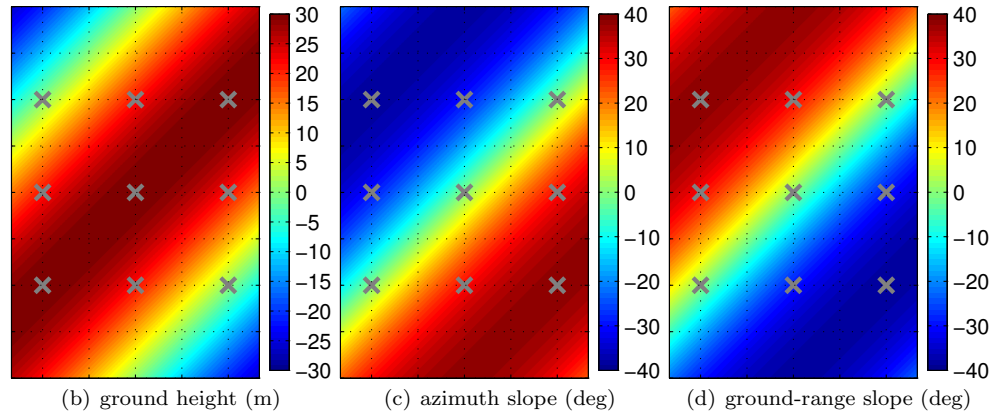
The image spectrum will show the superposition of the spatial frequencies for all pixels, and the observed spatial-frequency support will occupy the region covered by the set J of all pixel supports. Given the spatial dependence of the supports, discussed in Section 2.11, it should be expected that the content of the image spectrum will be dispersed i.e. the supports for different pixels will only partially overlap, giving total spectral coverage larger than the aperture for any single pixel. This spreading effect is amplified at closer ranges and for undulating focal surfaces.

Unfortunately, after aliasing into the image spectrum, it is quite possible for an aperture to wrap at the boundaries of the image spectrum. This means that the spatial variation in reflectivity is being sampled at the maximum observable spatial frequency $\pm\Delta k_{az}/2$ (if wrapped along azimuth) or $\pm\Delta k_{gr}/2$ (if wrapped along ground-range) i.e. the Nyquist limit. This is just due to different parts of the aperture aliasing down into different parts of the spectral window available for the image, and does not imply anything special about the underlying support. However, this will be a problem if the image is later resampled (in order, for example, to register it onto another image), because resampling stretches the coordinate space such that some areas are sampled more finely and other areas are sampled more coarsely, and the latter change will result in a sample rate (reciprocal pixel spacing) which is insufficient to properly support the aliased spatial frequencies.

The wrap can be avoided by basebanding the spectral content in the image spectrum—this is step 20 of the image formation process detailed in Algorithm 2.1.




(a) Altitude variation along the two 1.4 km flight-tracks. The blue flight-track is straight and level, whereas the green flight-track has a moderate sinusoidal variation (amplitude 10 m, period 700 m). There is no across-track variation (i.e. in y). The dots show the pulse positions.



(b) ground height (m)

(c) azimuth slope (deg)


(d) ground-range slope (deg)



(e) straight & level flight-track; PFA

(f) straight & level flight-track; BF (fixed horizontal focal plane)

(g) straight & level flight-track; BF (matched to terrain)



(h) sinusoidal altitude variation; PFA

(i) sinusoidal altitude variation; BF (fixed horizontal focal plane)

(j) sinusoidal altitude variation; BF (matched to terrain)

Figure 2.17: SAR images of a simulated scene consisting of nine point scatterers on terrain which models a ridgeline viewed obliquely, as depicted in (b)–(d) with grey crosses indicating the target positions. The images in (e)–(g) used the straight and level flight-track shown in (a), for which the k -space support is planar, whereas the images in (h)–(j) used the flight-track in (a) with sinusoidal altitude variation, for which the k -space support is non-planar. The images in (e) and (h) were focused via the polar format algorithm (PFA) using a fixed focal plane at $z=0$. The images in (f) and (i) were focused via beamforming (BF; see Section 2.10) using the same planar focal surface as for PFA. The images in (g) and (j) were also focused via beamforming, but using a focal surface matched to the terrain. For all images, $r_c=10$ km, $\psi_c=35^\circ$, $\delta_{ftk}=\delta_{off}=0^\circ$, $f_c=1.32$ GHz, $B_c=140$ MHz (same as the *Ingara* L-band waveform), no window, $\rho_{az}=75$ cm, $D_{az}\times D_{gr}=80\times 120$ m and displayed dynamic range 30 dB.

Denote the centre and the widths of the total occupied region in the k_x - k_y plane (i.e. from (2.59), the region covered by the set J of all apertures) as (k_{xc}, k_{yc}) and $\{\Delta k_{xi}, \Delta k_{yi}\}$, respectively. Applying a simple 2D phase ramp $e^{-j\Phi_{bb}}$ to the image, where

$$\Phi_{bb} = k_{xc}x + k_{yc}y, \quad (2.66)$$

will shift the central point (k_{xc}, k_{yc}) onto DC and thus baseband the spectrum. All the constituent apertures are shifted by the same amount. To avoid wrapping, the pixel spacings must also be sufficiently small such that the spectral content fits within the unambiguous extents Δk_{az} and Δk_{gr} of the image spectrum. From (2.63), this places upper limits on the pixel spacings:

$$d_{az} < 2\pi/\Delta k_{xi} \quad \text{and} \quad d_{gr} < 2\pi/\Delta k_{yi}. \quad (2.67)$$

Importantly, for interferometry, both images in the pair must be basebanded by the same 2D phase ramp, in order to ensure that no spatially varying differential phase is introduced. Since both images will presumably be focused onto the same pixel grid (and potentially trimmed to the same supports), their baseband shifts will naturally be the same.

Blacknell and Andre proposed a more sophisticated approach they called “topographic basebanding”, whereby the aperture I for each pixel is shifted so as to be centred on DC [Blacknell *et al.* 2010, 2011]. This eliminates the spread of different supports for different pixels, so all of the spectral content is compactly represented in the image spectrum. However, their technique was not able to be replicated here; it does not shift the supports correctly for all pixels on an undulating focal surface in the most general case of a non-linear flight-track. No matter, achieving such a compact spectrum is unnecessary; all that is needed is that the spectral content does not wrap, in order to permit image resampling.

The aperture observed in the image spectrum will be coupled in k_{az} and k_{gr} , since the beamformer coherently integrates along the natural curved contour in the spatial-frequency domain (i.e. along the ribbon in Figure 2.14), unlike the polar-format algorithm, which first resamples the spatial frequencies to a rectangular grid. Therefore, it is not possible to isolate the azimuthal support for a pixel to a single transformed azimuth row. This prevents application of many standard autofocus methods, including the phase gradient algorithm, which aim to estimate and correct for remnant phase errors across pulses after imperfect motion compensation [Jakowatz & Wahl 2009].

Figure 2.18 shows beamformed SAR images and their spectra for different collection geometries and focal surfaces, given a simulated scene consisting of a single point scatterer. Comparing (a) and (c) ($\psi_c = 35^\circ$) against (b) and (d) ($\psi_c = 50^\circ$), the increased grazing angle in the latter case causes the aperture to shift and shrink along k_y . Considering (j) and (l) too ($\alpha_{gr} = 15^\circ$), it can be seen that the ground-range slope has a very similar effect to the increased grazing angle (by (2.50), the equivalent grazing angle for this slope is 48.3°). Flight-track squint in (f) and (h) causes a shear of the aperture along k_y (and a consequent shear of the range sidelobes in the point-spread function), whereas offset squint in (e) and (g) causes a rotation and shift of the aperture (and a consequent rotation of the point-spread function). Azimuthal slope causes a shear of the aperture along k_x (analogous to the k_y shear induced by flight-track squint) and a shift along k_x (analogous to the shift induced by offset squint). Thus, simple differences in geometry can give rise to significant differences in the position and size of the projected aperture. All of these effects are precisely predicted by the formulation in Sections 2.2 and 2.12, and the boundary of the aperture in the image spectrum can be computed via (2.58) and (2.65).

The last row of Figure 2.18 demonstrates how the aperture may alias into the image spectrum awkwardly (the support in (o) is wrapped at the boundaries of the visible spectrum), and how this can be rectified by applying a 2D phase ramp to baseband the spectrum, without affecting the point-spread function.

Figure 2.19 redisplay the reference SAR image and spectrum from Figure 2.18 with a larger dynamic range (60 dB instead of 30 dB) such that the bow tie characteristic of the azimuth sidelobes is visible—this agrees with the point-spread functions observed for the corner reflectors in the *Ingara* L-band SAR image in Figure 2.10. This characteristic arises because the beamformer integrates over a curved contour in the




Figure 2.18: Beamformed SAR images and associated spectra of a simulated scene consisting of a single point scatterer. The images in columns one and two show the point-spread function, while the image spectra in columns three and four show the corresponding aperture projected onto the k_x - k_y plane ((c) is the spectrum of (a), (d) is the spectrum of (b), etc). In each case, the caption specifies how the collection geometry and/or focal surface differ from the reference case, which was a simple broadside synthetic aperture ($\delta_{ftk} = \delta_{off} = 0^\circ$) centred at grazing angle $\psi_c = 35^\circ$ and range $r_c = 1$ km, together with a flat focal plane ($\alpha_{az} = \alpha_{gr} = 0^\circ$). In all cases: $f_c = 1.32$ GHz, $B_c = 140$ MHz (same as the *Ingara* L-band waveform), no window, $\rho_{az} = 75$ cm, $d_{az} = d_{gr} = 15$ cm, $D_{az} = D_{gr} = 30$ m and 30 dB dynamic range for display. The basebanding step (step 20 of Algorithm 2.1) was not applied (except in (p)), in order to illustrate how the aperture projects onto different locations in the frequency ground plane.




Figure 2.19: The reference case from Figure 2.18 redisplayed with a larger dynamic range (60 dB instead of 30 dB). The bow tie characteristic of the azimuth sidelobes is now visible.

spatial-frequency domain; the divergence of the split sidelobes increases as the aperture angle is increased.

Figure 2.20 shows the case when the ground is undulating and the focal surface is matched to this (known) topography so that the image is draped over the ground. Observe that although the height variation in (a) is small, the slope variation in (b) and (c), and the consequent aperture variation in (d), is large. The pulse supports project onto the k_x - k_y plane in accordance with the local slope, so for undulating topography, the location, shape and size of the effective apertures are different for different pixel positions, leading to a spread spectrum where the constituent apertures are dispersed in the k_x - k_y plane, as seen in (g) and (h), and a spatially varying point-spread function, as seen in (e). In particular, the aperture size for each pixel, computed via (2.58) and shown in (d), determines the resolution apparent in (e): scatterer G has the worst ground-range resolution because the ground-range slope is at its maximum positive value so the projected aperture is at its smallest, whereas scatterer F has the best ground-range resolution because the ground-range slope is at its maximum negative value so the projected aperture is at its largest, and scatterer H has medium resolution because the slope is zero. The colour-coded boxes in (g) delineate the computed apertures I_F (biggest), I_G (smallest) and I_H (middle); it can be seen that they bound the main part of the actual spectral content for each scatterer—in this case, the apertures are somewhat distorted when observed via the image spectrum due to the fast variation of slope, particularly at H, causing the neighbouring pixels, onto which the sidelobes lay over, to have changing surface orientations. Importantly, the clutter in (f) also exhibits this aperture variation in terms of the resolution of the speckle pattern: the blob size varies inversely with the computed aperture size in (d), giving rise to alternating diagonal bands of fine and coarse resolution.

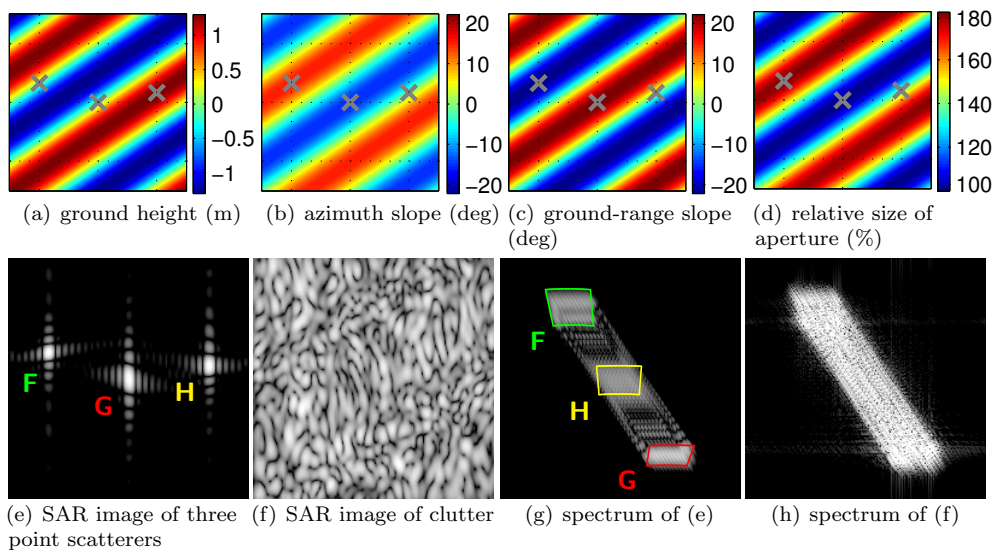


Figure 2.20: Beamformed SAR images and associated spectra of simulated scenes with undulating topography. The collection parameters are the same as for the reference case in Figure 2.18. The focal surface is matched to the ground. The upper row shows (a) the sinusoidal topography of the scene, (b) and (c) the sinusoidal slopes, and (d) the size of the effective aperture for every pixel (computed via (2.58)) relative to the size of the effective aperture for the centre pixel. Two scenes are simulated: (e) shows three point scatterers F , G and H , on the ground (no layover; their positions in (a)–(d) are marked by grey crosses), and (f) shows ground clutter, consisting of many point scatterers spread randomly on the ground. Their basebanded spectra (after applying step 20 of Algorithm 2.1) are shown in (g) and (h), respectively.

2.16 Spatially variant aperture trim

In Section 2.10 it was shown that a SAR image can be focused via beamforming. In Sections 2.12 and 2.15 it was shown how the effective spatial-frequency support for each pixel depends on the waveform and the relative acquisition geometry of the constituent

pulses in the synthetic aperture. Recall from (2.59) that J denotes the set of spatial-frequency supports i.e. the effective apertures for all pixels in an image. In Section 2.1 the problem of surface decorrelation was presented, whereby the coherence between a pair of images is degraded due to differing spectral supports. The task now is to control the supports such that, for two images s_a and s_b , all corresponding apertures are equal i.e. $J_a = J_b$. This is achieved not simply by post-processing two focused images, but by developing a modified image formation algorithm which operates on a pair of raw datasets and incorporates a spatially variant aperture trim.

The scenario of interest is when a radar platform makes two passes, denoted a and b , over the scene. The subsequent images, s_a and s_b , are focused to the same set of pixel positions. The two flight-tracks, however, will be different, either deliberately or simply due to realistic variations in air vehicle performance. Therefore, the projections onto the common image plane will be different, such that, for each pixel, the effective apertures will probably overlap but will not fully coincide, leading to uncorrelated spatial-frequency components which reduce the coherence (2.1) between the two observations of the pixel position. Conceptually, this is a standard problem in SAR interferometry, and it is fixed by trimming the apertures to their common area [Gatelli *et al.* 1994]. The complicating factor is that the required trim varies spatially, since the support varies spatially, as shown in Sections 2.12 and 2.15. When the focal surface is undulating, this spatial variation is amplified. A global trim would be inappropriate, and may actually degrade the coherence in areas which differ from the nominal geometry. A time-domain method of image formation, such as beamforming, is sufficiently flexible to permit the trim to be recomputed and reapplied on a pixel-by-pixel basis.

In this work, the aperture trim is undertaken in three steps. Firstly, the set $J_o = \{I_o\}$ of overlapping regions $I_o = I_a \cap I_b$ in the k_x - k_y plane is computed for all pairs of corresponding pixel apertures I_a and I_b given by (2.58) for the two passes. There may be no overlap for some pixels, which are left unformed with zero output value. Secondly, for each pulse, a set of start and end sample indices are computed where each (start, end) pair specifies the portion of the pulse which, when projected onto the k_x - k_y plane, fits in the overlap region I_o for a pixel. These are the trim indices. There may be no samples in the overlap region for some pulses, which are then dropped from the aperture. Many trimmed variants of each pulse may be specified to suit different pixels. Thirdly, the trimmed variants of each pulse are generated as needed during image formation. Thus, the first two steps use the metadata to precompute how the pulses are to be trimmed during the third step. These steps can be merged into the image formation process of Algorithm 2.1, giving a modified process which forms an interferometric pair of spatial-frequency-matched images by incorporating a spatially variant aperture trim, as detailed in Algorithm 2.2 below.

Algorithm 2.2 has been implemented with three features which trade-off accuracy for simplicity and memory efficiency. Firstly, the aperture and overlap region are represented as convex quadrilaterals or triangles inscribed within the true curved, irregular areas, thus sacrificing a small amount of potential overlap near the edges but greatly simplifying the overlap calculations at steps 7 and 12, which could be implemented using general-purpose 2D vector geometry functions. Secondly, only the most popular trim indices are stored for each pulse because, for a large image consisting of tens of thousands of input pulses and a similar number of output pixels, it is not practical to store all trim indices and then generate a unique trimmed pulse variant for every pulse-pixel pairing. Thirdly, the apertures and the trim indices can be computed for a subsampled pixel grid, since nearby pixels tend to have similar apertures. The coarseness of this pixel grid is user-selected: at one extreme, the grid could be selected to match the set of output pixels (no subsampling), and at the other extreme, the grid could be set to just a single pixel at the scene centre—this corresponds to the traditional global trim for a nominal collection geometry. Thus, for a given dataset, rerunning the image formation with different levels of spatial granularity enables the value of the spatially variant aperture trim to be assessed. Note that this third efficiency feature is not included in Algorithm 2.2, but involves simply replacing ‘for all pixels’ at steps 1, 11 and 16 with ‘for a subsampled set of pixels’.

Algorithm 2.2 extends naturally to three or more passes, giving a set of SAR images trimmed to their common spatial-frequency support. This will be used later for multichannel repeat-pass interferometry.

Example outputs of Algorithm 2.2 are shown in Figure 2.21 for a synthetic scene

Algorithm 2.2 Image-pair formation with spatially variant aperture trim

Input & output: As for Algorithm 2.1, except now given two passes a and b over the scene, two sets of pulses should be supplied on input, and two SAR images focused to the same set of pixel positions will be generated as output.

- 1: **for all pixels do** #Aperture and aperture overlap
- 2: **for all passes do**
- 3: Do steps 2–3 of Algorithm 2.1.
- 4: Compute the effective aperture I from (2.58) in the k_x - k_y plane.
- 5: Determine the four vertices of a trapezium (in American English, a trapezoid) inscribed within I . The trapezium need not be aligned with the k_x - k_y axes. A trapezium is general enough to capture the bulk of the energy in most apertures, allowing for a squinted collection and/or sloped terrain, but simple enough to make subsequent overlap calculations easy, since a trapezium is just one rectangle and two triangles.
- 6: **end for**
- 7: Determine the vertices of a convex quadrilateral (usual case) or triangle (rare) inscribed in the region where the two trapeziums overlap, if such an area exists. These vertices bound the overlap region I_o .
- 8: **end for**
- 9: **for all passes do**
- 10: **for all pulses do**
- 11: **for all pixels do** #Trim indices
- 12: Determine the portion of the line of pulse support I_p which lies inside the (convex) overlap region I_o .
 - If I_p lies completely inside I_o , then the whole pulse is usable.
 - If I_p lies completely outside I_o , then the whole pulse is unusable.
 - If I_p intersects the boundary of the overlap region once, then only part of the pulse, either from the start I_{ps} up to the intersection point, or from the intersection up to the end I_{pe} , is usable.
 - If I_p intersects the boundary of the overlap region twice, then only the part of the pulse between the two intersection points is usable.
- 13: Compute the start and end sample indices of the usable portion of the (discrete) pulse. Use the sample spacing after projection of the pulse onto the k_x - k_y plane.
- 14: **end for**
- 15: Store the most popular trim index pairs (say, the top 100).
- 16: **for all pixels do**
- 17: Record which pair of stored trim indices is closest to the ideal trim.
- 18: **end for**
- 19: Do step 6 of Algorithm 2.1.
- 20: **for all trim variants do** #Pulse trim
- 21: Make a copy of the pulse spectrum.
- 22: Trim the spectrum by setting to zero the samples before the trim start index and after the trim end index.
- 23: Do steps 7–9 of Algorithm 2.1. The range window in step 7 should be tailored for this variant to cover only the non-zero portion of the spectrum.
- 24: **end for**
- 25: **for all pixels do**
- 26: **if** this pulse contributes to the aperture for this pixel **then**
- 27: Select the trimmed pulse variant.
- 28: Do steps 12–16 of Algorithm 2.1.
- 29: **end if**
- 30: **end for**
- 31: **end for**
- 32: Do step 20 of Algorithm 2.1.
- 33: **end for**

consisting of a single point scatterer on sloped ground. Given data from four passes at different collection geometries, two are selected and processed to give an interferometric pair of SAR images which are matched in terms of their spatial-frequency support. By comparing the untrimmed spectra in (e)–(h) with the trimmed spectra in (i)–(l), it can be seen that the latter have been properly restricted to the region of spatial-frequency support which is common to the two selected passes—the goal of this algorithm is thus achieved. Of course, the act of trimming reduces the spatial bandwidth and the consequent spatial resolution, so the trimmed images in (m)–(p) have a much larger mainlobe than the untrimmed images in (a)–(d). Moreover, the common region may be irregularly shaped, leading to a distorted point-spread function which differs from the expected 2D sinc function; pair (a, c) , in particular, has distorted range sidelobes in (n) due to the skew top edge of the overlap quadrilateral in (j). Pair (c, d) in (l) and (p) illustrates how the trim output breaks down when the overlap area is very small—the point-spread function is now so large along azimuth that it is not fully contained in the image, and the Fourier transform of this cropped point response exhibits significant leakage of energy across the image spectrum. Note that the peak amplitudes of the responses in (m)–(p) are not the same (the 30 dB amplitude range was shifted for each image to show the sidelobes clearly), with reduced bandwidth causing reduced amplitude, in accordance with the Fourier transform properties shown in Appendix 2.A.5. Note too that the surprisingly large differences in grazing angle ($\Delta\psi = 5^\circ, 6^\circ$) between interferometric acquisitions are possible due to the large negative ground-range slope used for this example: from (2.48), the grazing angles 33° and 38° for passes a and b are equivalent to 17.2° and 22.5° over flat ground, and for flat ground, it can readily be shown via (2.57) that the overlap of the latter pair when projected onto the k_x - k_y plane is much greater than that of the former pair.

Further examples in Figure 2.22 show the real power of the spatially variant aperture trim for clutter on undulating terrain. The sinusoidal topography and the reference pass (a) are taken from Figure 2.20. Additional passes b and c at different grazing and squint angles are then separately matched to a . Given two passes, the key metric which distils the impact of their different collection geometries on the quality of the matched image pair is the relative size of the aperture overlap at each pixel i.e. the ratio of the areas of I_o and I . This metric is displayed as colour-mapped images in (a) and (e) for pairs (a, b) and (a, c) , respectively. It can be seen that the aperture overlap (just like the aperture itself) varies in proportion to the sinusoidal variation of the slopes of the focal surface in Figures 2.20(b)–(c), and that the overlap varies more for pair (a, c) than (a, b) due to the greater delta in grazing angle and squint.

First considering a scene consisting of isolated point scatterers, (b)–(d) show how the point-spread functions for pass b are changed when trimmed so as to match a , with the resolution degraded by different amounts at different pixel positions, depending on the local projection and consequent aperture overlap for that position (observe that in (d), scatterer F has much more overlap than G). Similarly, (f)–(h) show the same for pass c , with the trimmed response from scatterer G in (g) heavily distorted due to the small (22%), triangular overlap in (h).

Now considering clutter, (i)–(l) show how the resolution of the speckle patterns follows the local slope, and then is degraded (in proportion to the local slope) by the aperture trim. The coherence γ in (2.1) was computed between pairs of SAR images formed in four different ways (this calculation used a fixed rectangular averaging window, but ideally the window would vary with the resolution). When the apertures are not trimmed, as in (m) and (q), the coherence magnitude varies in proportion to the aperture overlap in (a) and (e), since where the overlap is low, there is a large proportion of uncorrelated spatial frequencies which lower the coherence—this demonstrates the essential problem of spatially varying surface decorrelation (loss of correlation) induced by differing collection geometries and undulating topography, introduced in Section 2.1. When a global trim is computed based on the apertures for the centre pixel position only, as in (n) and (r), the coherence magnitude still varies with the slope, but now regions with slopes similar to that at the centre have high coherence, and regions with different slopes have low coherence. When coarse spatial variation is permitted, as in the quadrant trim of (o) and (s), the coherence is improved in areas similar to the centre of each quadrant, but remains patchy overall. Essentially, the coherence level at each pixel depends on the similarity of the local terrain to the terrain of the reference pixel position for which the aperture trim was designed.

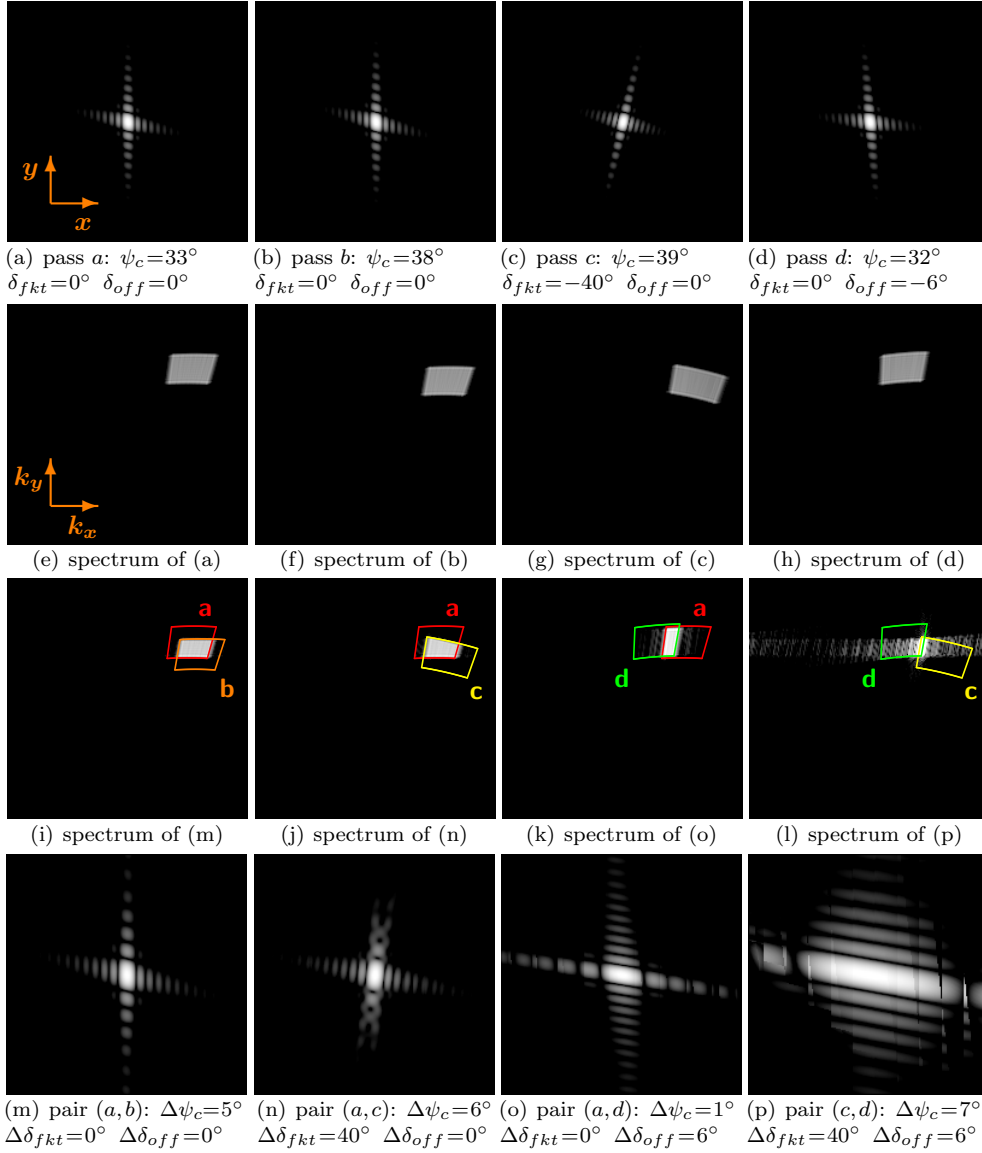


Figure 2.21: Beamformed SAR images and associated spectra of a simulated single point scatterer, showing the effect of the spatially variant aperture trim in Algorithm 2.2. The ground scene (and matching focal surface) has slopes $\alpha_{az} = -20^\circ$ and $\alpha_{gr} = -10^\circ$. Echo data was synthesised for four passes with grazing and squint angles specified in (a)–(d). The other data collection and image display parameters are the same as in Figure 2.18 (the 30 dB dynamic range is referenced to the peak value in each image). Aperture-trimmed images are displayed in (m)–(p) for different pairs of passes (only the image from the first pass in each pair is shown since the two output images have the same support). Their (non-basebanded) spectra in (i)–(l) are annotated with colour-coded boxes delineating the original (untrimmed) apertures from the constituent passes.

Finally, when pixel-to-pixel variation of the aperture trim is implemented, as proposed in Algorithm 2.2 and shown in (p) and (t), the resulting coherence is high over the whole image ($|\bar{\gamma}_{ab}| = 0.91$, $|\bar{\gamma}_{ac}| = 0.85$), indicating that the spatially varying surface decorrelation has been mostly removed—the goal of this chapter is thus achieved.

The remnant surface decorrelation in these examples is due to the fast variation in slope relative to the point-spread function: the total scatterer response (mainlobe and sidelobes) is spread over an area with changing surface orientation, so no single trim, designed assuming a point scatterer on a tilted tangent plane, can suffice to remove all differences in the spatial-frequency supports. Indeed, the whole formulation of spatial-frequency support in Sections 2.4 and 2.12 presumed a point-scatterer model which requires the total response to be on a plane at fixed azimuth and ground-range tilts. Faster-varying slope and lower resolution both contribute to a point-spread function being distributed over a greater range of terrain, breaking the spatial-frequency model





Figure 2.22: Beamformed SAR images and coherence maps of simulated scenes with undulating topography, showing the effect of the spatially variant aperture trim. Both the sinusoidal ground heights and the scattering elements in the scene (three isolated point scatterers for rows 1–2, continuous clutter for rows 3–5) are the same as those in Figure 2.20. Echo data was synthesised for three passes: pass *a* is the same as that in Figures 2.20(e)–(h) ($\psi_c = 35^\circ$, $\delta_{f_{tk}} = 0^\circ$), pass *b* is shown in (b) and (i) and pass *c* is shown in (f) and (k). The percentage aperture overlap at every pixel for the pair (a,b) is shown in (a) and for the pair (a,c) in (e). The SAR image in (c) shows pass *b* formed via Algorithm 2.2 so as to be spatial-frequency-matched to pass *a*. Similarly, (g) shows pass *c* matched to pass *a*. The colour-coded boxes in the (non-basebanded) spectra in (d) and (h) delineate the original (untrimmed) apertures for the point scatterers *F*, *G* and *H* from passes *a* (red), *b* (yellow) and *c* (green). In the clutter case, (j) and (l) show passes *b* and *c*, respectively, matched to pass *a* (using a full pixel trim). The maps of the coherence magnitude ($|\gamma| = 0 \rightarrow$ black, $|\gamma| = 1 \rightarrow$ white) between clutter images from passes *a* and *b* (row 4) and passes *a* and *c* (row 5) show the effect of different aperture-control choices during image formation: no trim (i.e. Algorithm 2.1), a global trim (Algorithm 2.2 using only the centre pixel to design a constant aperture trim), a quadrant trim (coarse spatial variation over just four pixels across the four image quadrants) and a true pixel-by-pixel trim (the finest possible spatial variation).

employed here. Following the same argument, there is more remnant decorrelation for pair (a, c) in (t) because the overlap in some areas is much lower than for pair (a, b) (compare (e) to (a)), so after aperture trimming, the mainlobe of the point responses in those areas is much wider, covering more terrain variations. Repeated simulations indicate that, as the amplitude of the sinusoidal topography is increased from zero (a flat plane), the mean coherence decreases, as shown in Figure 2.23, due to the slope changing more rapidly and the aperture trim becoming less effective. (For a flat plane, simulation imperfections limited the coherence to 0.98, instead of the expected unity.) Note that for an amplitude of 1.3 m, as was used in Figures 2.20 and 2.22, the variation in slope is $4.4^\circ/\text{m}$ along ground-range, which is much higher than would typically be exhibited by natural landscapes, so the remnant decorrelation observed here is greater than what would be expected in more realistic imaging scenarios.

Figure 2.23: Variation of the mean coherence for pass pairs (a, b) and (a, c) with the amplitude of the sinusoidal topography. The topography follows the same pattern as shown in Figure 2.20(a), but the amplitude of the sinusoidal undulations was increased from zero (a flat plane), causing the slopes to not only become steeper but also change faster. The scene for Figures 2.20, 2.22 used 1.3 m.



2.17 Application to real data

Figures 2.24–2.27 show the results of applying the spatially variant aperture trim described in the previous section to repeat-pass data acquired by both the *Ingara* L-band radar (over rolling hills near Cape Jervis, South Australia) and the *Bright Sapphire* P-band radar (over the Tully Training Area in northern Queensland). In both cases, elevation data were extracted from a smoothed 1 arc second digital elevation model Gallant *et al.* [2011] and interpolated onto the pixel grid via a 2D spline fit. The elevation model itself was derived from height measurements obtained from interferometric radar data collected during the shuttle radar topographic mission [Farr *et al.* 2007]; these point measurements are sparse and somewhat approximate, with errors of the order of several metres, especially in the vicinity of trees and buildings (but sometimes even when there are no such complications) Gallant *et al.* [2011]. For the *Ingara* acquisitions, the terrain slopes are large ($\alpha > 20^\circ$), and the two passes have an extremely large separation in grazing angle ($\Delta\psi_c = 5.1^\circ$), so the size of the aperture overlap varies enormously over the swath (between zero and 77%), as shown in Figure 2.24(f), with some patches in near-range having no overlap at all. For the *Bright Sapphire* acquisitions, the terrain is relatively flat, but the passes are made at close-range ($r_c = 2\text{ km}$), such that their grazing angles vary considerably over the swath (by 10° from near to far range), and their nominal difference is large ($\Delta\psi_c = 1.8^\circ$), leading to a significant, though smooth, variation in the size of the aperture overlap (from 54% at near-range to 81% at far-range), as shown in Figure 2.26(f).

The key result is shown by the maps of the coherence magnitude ($|\gamma| = 0 \rightarrow$ black, $|\gamma| = 1 \rightarrow$ white) in Figures 2.25(c)–(f) for the *Ingara* pair and Figures 2.27(c)–(f) for the *Bright Sapphire* pair: as the aperture trim becomes more localised (from a global trim to a pixel trim), the coherence improves (the mean coherence $|\bar{\gamma}|$ increases and the map becomes whiter), because more of the uncorrelated spatial frequencies which cause surface decorrelation are removed. The *Ingara* results were first presented in 2015 [Pincus *et al.* 2015b].

The unrealistic step changes in coherence along straight lines in Figure 2.25(e) show the kind of processing artefacts that can arise when the spatially variant aperture trim is implemented using a subsampled pixel grid, discussed in Section 2.16. If the spatial sampling of the topographic data and the subsequent aperture calculations is too coarse to capture the true topographic variation encoded in the spatial frequency supports of the focused pixels, then the apertures for some pixels will be trimmed inaccurately, and furthermore, nearby pixels that happen to lie in adjacent grid cells may have their apertures trimmed very differently, leading to quite different coherence

levels, depending on how the local terrain differs from that at the grid centre locations used to design the trims. For the quadrant trim, topographic data at just four locations is used to design four aperture trims, one per quadrant (see Figures 2.22(o) and (s) for simulated examples). Comparing Figures 2.25(e) and 2.27(e), both obtained using a quadrant trim, the latter does not exhibit processing artefacts, simply because this terrain is relatively flat—compare the slopes in Figures 2.24(d)–(e) for the site near Cape Jervis to the slopes in Figures 2.26(d)–(e) for the Tully Training Area—so the applied aperture trims do not vary dramatically across the quadrants. The greater the terrain undulations, the finer the grid used to calculate the aperture trims needs to be to avoid processing artefacts.

Admittedly, the final coherence maps in Figures 2.25(f) and 2.27(f) for the pixel-trimmed images have significant remnant decorrelation. This is probably partly due to inaccuracies in the height information: height errors lead to slope errors which lead to incorrect k -space projections and trims. For the *Ingara* pair, processing-induced decorrelation is also likely due to imperfect co-registration of the two images, especially because, wherever the focus height differs (mistakenly) from the true ground height, the undulating terrain will lay over onto the focal surface differently for the two passes, making for a spatially varying warp function (of one image relative to the other) which can be difficult to estimate. For the *Bright Sapphire* pair, volume decorrelation is certainly present in the forested areas (everywhere except the clearing) due to the above-ground branches (not the foliage, given the 82 cm wavelength at P-band). Indeed, using a random volume over ground model, typically employed for interferometric analysis of forest data [Treuhaft *et al.* 1996, Treuhaft & Siqueira 2000, Cloude *et al.* 2004], with a volume height of 10 m (expect a rather low average canopy height due to the damage from Cyclone Yasi the previous year), an exponential extinction coefficient of 0.05 dB/m [Cloude *et al.* 2000] and a ground-to-volume scattering power ratio of 1 dB (expect slightly higher ground response due to significant penetration and strong ground-trunk double bounce at P-band), the total predicted coherence is 0.6, which approximately agrees with the mean coherence observed in Figure 2.27(f). This suggests that, in this case, the surface decorrelation has been almost totally removed, leaving only the volume decorrelation.

To reiterate, the fact that the coherence improved in Figures 2.25 and 2.27 as the trim of the spatial-frequency supports was made more localised demonstrates that the spatially variant aperture trim described and implemented here is an effective means of reducing surface decorrelation induced by the differing collection geometries of repeat-pass airborne radar systems imaging an undulating scene.

2.18 Conclusion

A signal processing chain has been described which takes two raw radar datasets acquired over undulating terrain and forms an interferometric pair of SAR images focused to this terrain whose spatial-frequency supports are matched for all pixels, such that the surface decorrelation induced by their differing collection geometries is minimised.

The signal transformations involved in transmission, reception, analogue demodulation, sampling and digital demodulation to baseband were formulated, considering both simple mix-down and dechirp receivers. Matched filtering was shown to give a filtered, narrowband version of the projected scene reflectivity function. A beamforming algorithm which forms a SAR image focused to an undulating surface with constant azimuth resolution was detailed and demonstrated using both *Ingara* L-band and *Bright Sapphire* P-band data.

The concepts of a 3D spatial-frequency domain (k -space) and 2D spatial-frequency support (the effective aperture, which determines the point-spread function) were formulated and examined. The aperture for each pixel was quantified precisely for every constituent pulse using a projection which accounts for the local slope and local collection geometry. The idea of an equivalent collection geometry for a flat plane which gives the same projected support was developed as an intuitive way to understand the projection. The spectrum of a beamformed SAR image was shown to contain the superposition of the aliased apertures for all pixels.

Putting all this together, a modified beamforming algorithm was presented which

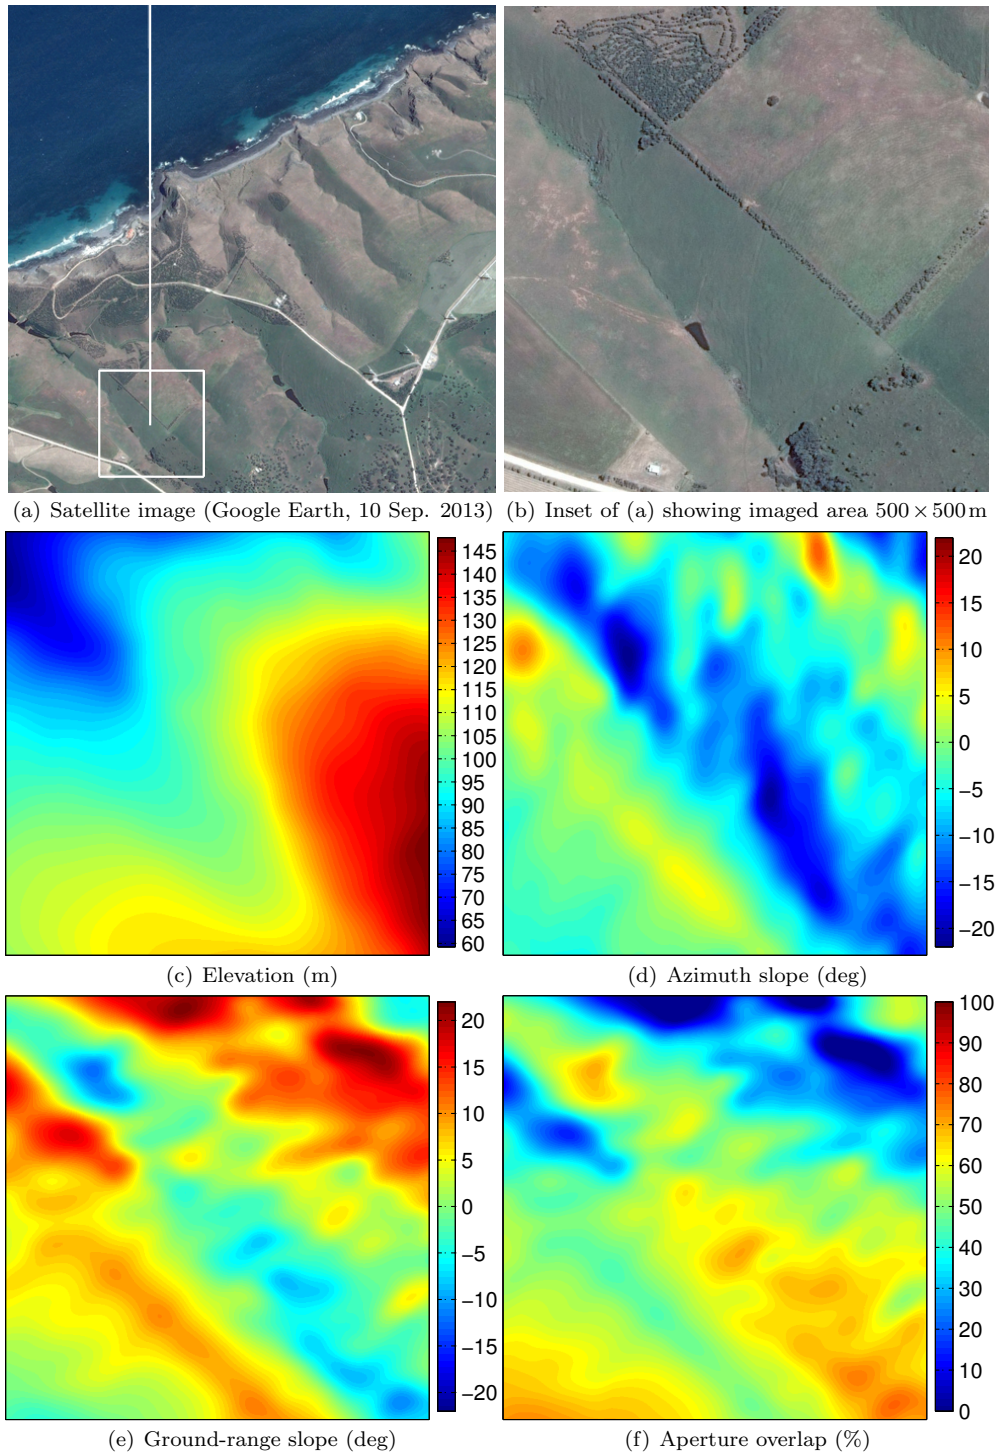


Figure 2.24: Optical views and metadata for the terrain near Cape Jervis, South Australia. The rolling hills along the coast are apparent in (a), while (b) shows the area imaged by the *Ingara* radar in Figure 2.25. The varying elevation over the imaged area, shown in (c), gives rise to steep positive and negative slopes along azimuth (left-to-right) and ground-range (bottom-to-top), shown in (d) and (e) respectively, which in turn cause a very wide variation in the aperture overlap, shown in (f), for the two passes presented in Figure 2.25. At the dark-blue patches in near-range, the size of the common spatial-frequency support falls to zero.




Figure 2.25: *Ingara* L-band SAR images and coherence maps focused onto undulating terrain near Cape Jervis, South Australia (10 July 2013, 500×500 m, polarimetric channel LR, 60 dB dynamic range, Hamming window, integration angle $\Delta\theta_{sl} = 11.3^\circ$, resolutions $\rho_{az} = 0.75$ m and $\rho_{gr} = 1.63$ m). Two passes (*a*, *b*) were acquired thirty minutes apart at $r_c = (3.3, 4.1)$ km, $\delta_{ftk} = (-6.6^\circ, -6.9^\circ)$ and $\psi_c = (26.6^\circ, 21.5^\circ)$ —note the large difference in grazing angle $\Delta\psi_c = 5.1^\circ$. Compared to the full-resolution image in (a), the pixel-trimmed image in (b) exhibits varying resolution in accordance with the computed aperture overlap in Figure 2.24(f); where there is no overlap, the pixels are left unformed and displayed as black voids. Four pairs of images were generated with different levels of aperture control: no trim, a global trim, a quadrant trim and a pixel trim. For each pair, the images were co-registered and their coherence computed using a fixed averaging window 3.25×4.75 m ($az \times gr$) to give ~ 12 equivalent looks [Touzi *et al.* 1999]; (c)–(f) show the coherence magnitude of each pair.

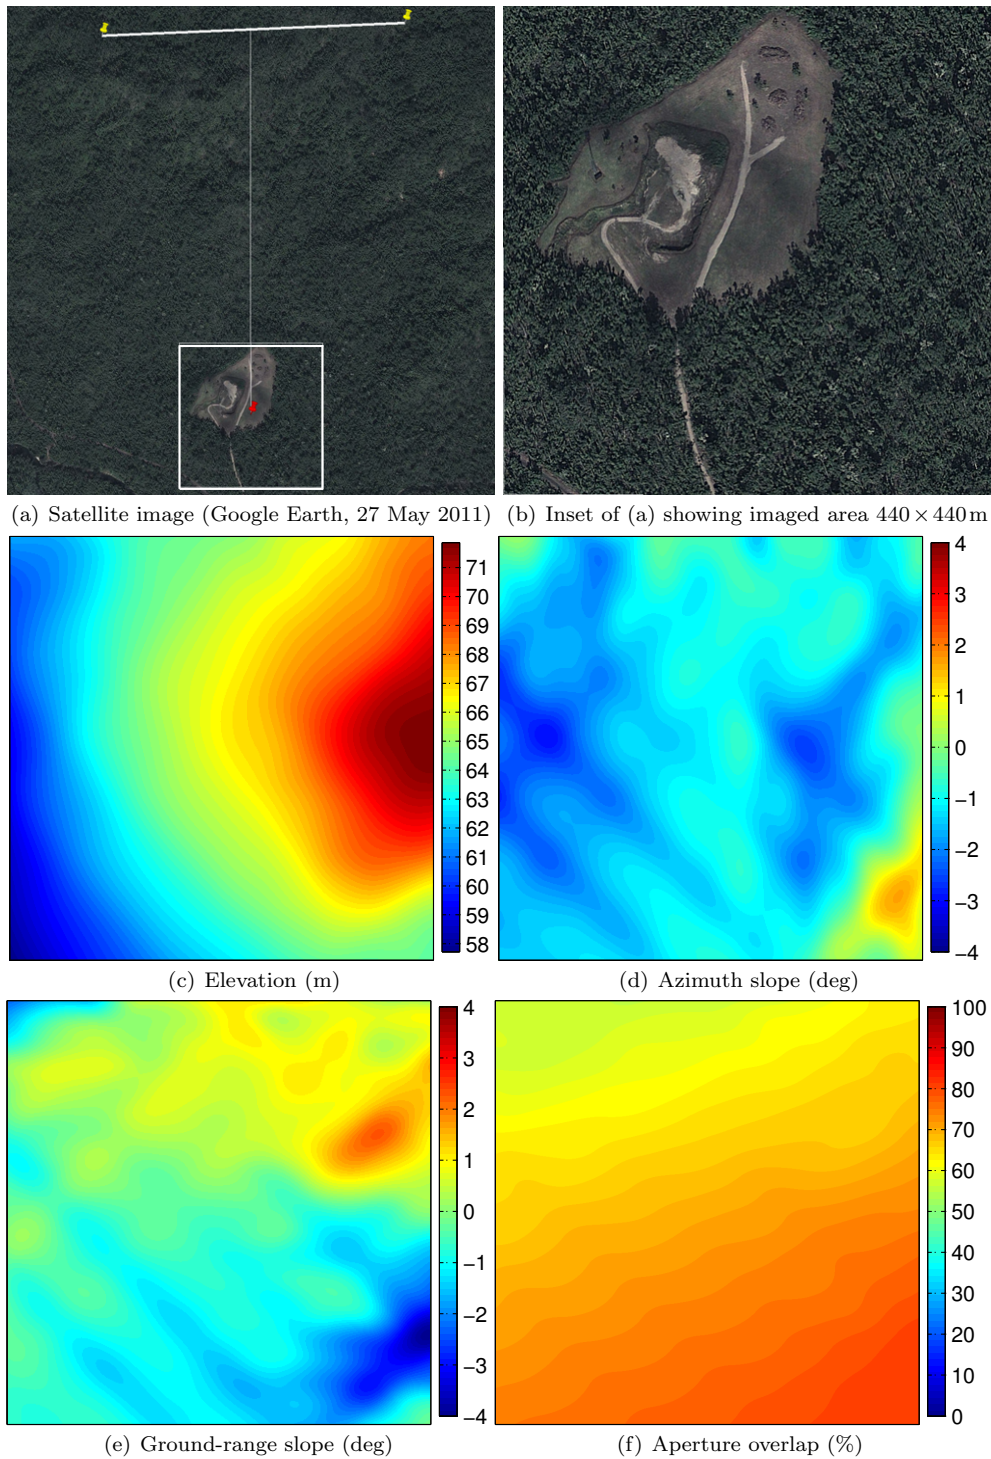


Figure 2.26: Optical views and metadata for the terrain of the Tully Training Area in northern Queensland (the same area shown in Figure 2.12). The dense and continuous vegetation cover surrounding the clearing is apparent in (a), while (b) shows the area imaged by the *Bright Sapphire* radar in Figure 2.27. The slowly varying elevation over the imaged area, shown in (c), gives rise to gentle slopes along azimuth (left-to-right) and ground-range (bottom-to-top), shown in (d) and (e) respectively. The variation in aperture overlap, shown in (f), is dominated by the varying grazing angle, which changes by 10° over the swath for the two passes presented in Figure 2.27.




Figure 2.27: *Bright Sapphire* P-band SAR images and coherence maps of the Tully Training area, northern Queensland (7 May 2012, 440×440 m, polarimetric channel H, 40 dB dynamic range, no window, integration angle $\Delta\theta_{st} = 12.5^\circ$, resolutions $\rho_{az} = 1.8$ m and $\rho_{gr} = 3.58$ m (before trimming)). Two passes (a, b) were acquired thirty minutes apart at $r_c = (2.00, 1.97)$ km, $\delta_{ftk} = (0.6^\circ, 3.3^\circ)$ and $\psi_c = (54.2^\circ, 56.0^\circ)$. Compared to the full-resolution image in (a), the pixel-trimmed image in (b) exhibits varying resolution in accordance with the computed aperture overlap in Figure 2.26(f)—observe the degraded ground-range resolution at near-range (top of image). Four pairs of images were generated with different levels of aperture control: no trim, a global trim, a quadrant trim and a pixel trim. For each pair, the images were co-registered and their coherence computed using a fixed averaging window 6.5×11.5 m ($az \times gr$) to give ~ 9 equivalent looks; (c)–(f) show the coherence magnitude of each pair.

forms an interferometric pair of images given two acquisitions over known undulating terrain by applying a spatially variant aperture trim to restrict the aperture at each pixel position to the area which is common to both passes. It was shown in simulation and using airborne radar data acquired by both the the *Ingara* L-band and *Bright Sapphire* P-band systems that the coherence between such a pair was significantly higher compared to a pair formed with no trim or a global trim, because uncorrelated spatial frequencies causing surface decorrelation were removed.

Remnant surface decorrelation may be observed if the slopes change quickly relative to the point-spread function. The aperture trim is designed presuming a point-scatterer model whereby the total scattering response lies on a plane at fixed azimuth and ground-range tilts; if the response lies on a surface with changing orientation, then the aperture trim will not remove all of the non-overlapping spatial-frequency supports. In addition, the trim is dependent on the accuracy of the height information.

2.A Appendix

2.A.1 Statistical performance of the coherence estimator

Reconsider the definition of coherence in (2.1). A more careful formulation would distinguish between the true coherence γ and an estimator $\hat{\gamma}$ for this coherence. The statistical performance of the estimator $|\hat{\gamma}|$ for the coherence magnitude is now presented in terms of its bias and standard deviation.

Let s_a and s_b be two complex, zero mean random variables. The complex (degree of) coherence

$$\gamma_{ab} = \frac{\mathbb{E}\{s_a s_b^*\}}{\sqrt{\mathbb{E}\{s_a s_a^*\} \mathbb{E}\{s_b s_b^*\}}} \quad (2.68)$$

is a parameter describing the joint distribution of s_a and s_b [Touzi & Lopes 1996, Goodman 1975 p. 38]. In more general statistical treatments, it is called the complex correlation coefficient [Papoulis 1991 pp. 152, 188, 293-295]. Its magnitude $0 \leq |\gamma_{ab}| \leq 1$ is a measure of the mutual linear dependence of s_a and s_b , where zero indicates that they are uncorrelated (i.e. there is no linear relationship between them), and unity indicates that they are perfectly correlated (i.e. they are linearly related, so there is some complex scalar y such that $s_b = y s_a$, resulting in $|\gamma_{a(ya)}| = |y^* \mathbb{E}\{s_a s_a^*\} / \sqrt{|y|^2 \mathbb{E}\{s_a s_a^*\} \mathbb{E}\{s_a s_a^*\}}| = 1$).

The sample complex coherence or sample complex correlation coefficient

$$\hat{\gamma}_{ab} = \frac{\langle s_a s_b^* \rangle}{\sqrt{\langle s_a s_a^* \rangle \langle s_b s_b^* \rangle}} \quad (2.69)$$

is an estimator for γ_{ab} , where the data averages $\langle \dots \rangle$ in (2.69) are intended to approximate the ensemble averages $\mathbb{E}\{\dots\}$ in (2.68). Because $\hat{\gamma}_{ab}$ is a function of random variables, it too is a random variable, and therefore may be characterised by a probability distribution [Wackerly *et al.* 1996 ch. 7.1]. Of particular interest are metrics of this distribution as they relate to the true parameter being estimated. Standard metrics are the bias, the standard deviation (std) and the root-mean-squared error (RMSE), defined and interpreted as follows for a general parameter θ and associated estimator $\hat{\theta}_N$ using N samples [Wackerly *et al.* 1996 ch. 8.2]:

$$\text{bias}(\hat{\theta}_N) = \mathbb{E}\{\hat{\theta}_N\} - \theta \quad (\text{estimator accuracy}) \quad (2.70)$$

$$\text{std}(\hat{\theta}_N) = \sqrt{\mathbb{E}\{(\hat{\theta}_N - \mathbb{E}\{\hat{\theta}_N\})^2\}} \quad (\text{estimator precision}) \quad (2.71)$$

$$\text{RMSE}(\hat{\theta}_N) = \sqrt{\mathbb{E}\{(\hat{\theta}_N - \theta)^2\}} \quad (\text{average estimation error}) \quad (2.72)$$

where

$$\text{RMSE}^2 = \text{std}^2 + \text{bias}^2. \quad (2.73)$$

These metrics characterise the statistical performance of the estimator. Crucially, the values of these metrics vary with the number of samples N used to compute the estimate.

To continue, assume that the random variables s_a and s_b follow normal distributions, which is commonplace (although not at all universal) for SAR imagery.

The estimator $\hat{\gamma}_{ab}$ in (2.69) is a scaled version of the maximum-likelihood estimator for γ_{ab} ; the scale is very often close to one, such that $\hat{\gamma}_{ab}$ retains maximum-likelihood properties. [Seymour & Cumming 1994, Jakowatz *et al.* 1996 ch. 5.5]. Two important properties of a general maximum-likelihood estimator follow:

- As N increases, its distribution tends towards a normal distribution (regardless of the distribution of the original random variables) with mean equal to the true parameter value and variance equal to the Cramér-Rao lower bound on the variance of all possible unbiased estimators [Kay 1993 ch. 7.5].
- As N increases, the estimator converges (in a probabilistic sense) to the true parameter value (i.e. it is consistent). Most often, this convergence manifests as $\lim_{n \rightarrow \infty} \{\text{bias}(\hat{\theta}_N)\} = 0$ and $\lim_{n \rightarrow \infty} \{\text{std}(\hat{\theta}_N)\} = 0$, but due to subtleties in the definition of convergence, these are technically only sufficient conditions for consistency, not necessary conditions [Wackerly *et al.* 1996 ch. 9.3, 9.8].

Despite these attractive large-sample properties, the statistical performance of a maximum-likelihood estimator is not necessarily good for small N [Kay 1993 ch. 7.5].

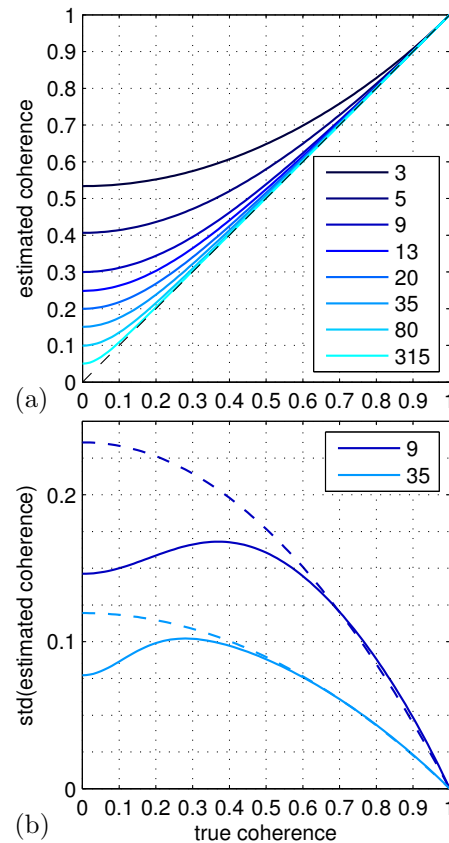
For a finite number of samples, the estimator $|\hat{\gamma}_{ab}|$ of the coherence magnitude $|\gamma_{ab}|$ is biased upwards [Carter *et al.* 1973, Joughin *et al.* 1994, Touzi & Lopes 1996, Touzi *et al.* 1999]. Figure 2.28(a) illustrates the bias by plotting the estimated coherence magnitude as a function of the true coherence for different sample sizes. The estimates were computed via (2.69) using a simple random number simulation; the curves produced by this method precisely match those in the literature (for the same sample size), where the latter were instead generated using the rather complicated analytical expressions for the moments of the PDF of the estimator [Carter *et al.* 1973, Touzi *et al.* 1999, Bamler & Hartl 1998]. It can be seen that the bias is greater at lower coherences, but decreases as the number of samples is increased, although the relationship between sample size and bias is highly nonlinear. The effect of the bias is to limit the dynamic range of the observed coherence such that low coherences are not visible.

Figure 2.28(b) plots the standard deviation of the estimator $|\hat{\gamma}_{ab}|$. It can be seen that the spread of estimates will be largest for low to moderate coherences, but decreases as the number of samples is increased. Observe that at high coherences, where the estimator is unbiased, the precision is very close to the minimum prescribed by the Cramér-Rao lower bound,

$$\text{std}_{\text{CRLB}}(|\hat{\gamma}_{ab}(N)|) = (1 - |\gamma_{ab}|^2) / \sqrt{2N} \quad [\text{Seymour \& Cumming 1994}], \quad (2.74)$$

whereas at low coherences, where the estimator is biased, the precision is actually much better than the bound, which is applicable to unbiased estimators only.

Figure 2.28: Statistical performance of the coherence magnitude estimator for different sample sizes. In (a), the bias of the estimator is shown by the deviation from the dashed line; the bias is larger at lower true coherences and for smaller sample sizes. In (b), the standard deviation of the estimator is shown as a solid line, and the square-root of the Cramér-Rao lower bound is shown as a dashed line. The curves were obtained by a simple random number simulation using many (one million) realisations of a pair of complex, zero-mean, normally distributed random variables generated with a particular correlation (the quantity on the horizontal axis). The magnitude of the complex coherence was repeatedly estimated via (2.69) using a particular number of samples (the quantity in the legend). Thus, in (a), each point on each curve is the mean of many magnitude estimates across different random realisations, all with the specified level of true coherence and all the result of averaging over the specified number of samples. In (b), each point on each solid curve is the standard deviation of those magnitude estimates.



Throughout this section, it is assumed that the samples used in the estimator are independent realisations of the random variables. Because SAR imagery is typically oversampled, neighbouring pixels are correlated, and it is actually resolution cells that correspond to the statistical notion of independent samples. The level of oversampling can be characterised by the equivalent number of looks (ENL) per pixel, which is the ratio of the pixel area to the effective resolution cell size after all processing. Assuming complex Gaussian statistics, so the average intensity can be modelled by a gamma distribution, the ENL per pixel is equal to the square of the ratio of mean to standard deviation in intensity i.e. the square of the reciprocal of the coefficient of variation [Oliver & Quegan 1998 ch. 4.4].

2.A.2 Coherence in the spatial-frequency domain

The equivalence of the spatial and spatial-frequency expressions for the complex coherence γ_{ab} in (2.1) is derived here.

Begin with the 2D Fourier transform and its inverse, defined here over the 2D spatial domain (x, y) and the associated spatial-frequency domain (k_x, k_y) .

$$S(k_x, k_y) = \frac{1}{2\pi} \int_y \int_x s(x, y) e^{-j(k_x x + k_y y)} dx dy \quad (2.75)$$

$$s(x, y) = \frac{1}{2\pi} \int_{k_y} \int_{k_x} S(k_x, k_y) e^{j(k_x x + k_y y)} dk_x dk_y \quad (2.76)$$

Denote the 2D sum by $\mathbb{S}\{\dots\}$.

$$\mathbb{S}\{s\} = \int_y \int_x s(x, y) dx dy \quad (2.77)$$

$$\mathbb{S}\{S\} = \int_{k_y} \int_{k_x} S(k_x, k_y) dk_x dk_y \quad (2.78)$$

Given two 2D spatial signals $s_a(x, y)$ and $s_b(x, y)$, with associated spectra $S_a(k_x, k_y)$ and $S_b(k_x, k_y)$, the sum of their conjugate product is

$$\begin{aligned} \mathbb{S}\{s_a s_b^*\} &= \int_y \int_x s_a(x, y) s_b^*(x, y) dx dy \\ &= \int_y \int_x \frac{1}{2\pi} \int_{k_y} \int_{k_x} S_a(k_x, k_y) e^{j(k_x x + k_y y)} dk_x dk_y s_b^*(x, y) dx dy \\ &= \int_{k_y} \int_{k_x} S_a(k_x, k_y) \cdot \frac{1}{2\pi} \int_y \int_x s_b^*(x, y) e^{j(k_x x + k_y y)} dx dy dk_x dk_y \\ &= \int_{k_y} \int_{k_x} S_a(k_x, k_y) \left[\frac{1}{2\pi} \int_y \int_x s_b(x, y) e^{-j(k_x x + k_y y)} dx dy \right]^* dk_x dk_y \\ &= \int_{k_y} \int_{k_x} S_a(k_x, k_y) S_b^*(k_x, k_y) dk_x dk_y \\ &= \mathbb{S}\{S_a S_b^*\} \quad [\text{Bracewell 2003 p. 162}]. \end{aligned} \quad (2.79)$$

Similarly, $\mathbb{S}\{|s_a|^2\} = \mathbb{S}\{s_a s_a^*\} = \mathbb{S}\{S_a S_a^*\} = \mathbb{S}\{|S_a|^2\}$ and the same for $\mathbb{S}\{|s_b|^2\}$. Thus, the sum (the total energy) can be computed in either domain. This is just Parseval's theorem in two dimensions [Carlson 1998 p. 159].

The average $\langle \dots \rangle$ is simply a sum $\mathbb{S}\{\dots\}$ normalised by a length (an area in 2D). Since γ_{ab} in (2.1) is a ratio of spatial averages, the normalisation can be neglected, leaving a ratio of sums, to which the result in (2.79) applies. Therefore, the complex coherence γ_{ab} can be expressed in either the spatial or spatial-frequency domain, justifying the last equality in (2.1).

Note that Cloude [2010 ch. A3, p. 428] actually defines $\langle \dots \rangle$ to be the sum $\mathbb{S}\{\dots\}$, but this seems to be non-standard.

Allowing for a spatial offset (lag) (x', y') between s_a and s_b gives the coherence (or cross-correlation) function $\Gamma_{ab}(x', y')$, where

$$\Gamma_{ab}(x', y') = \langle s_a(x + x', y + y') s_b^*(x, y) \rangle. \quad (2.80)$$

Its spectrum is the cross-spectral density, obtained by appropriately averaging the conjugate product $S_a(k_x, k_y) S_b^*(k_x, k_y)$ [Born & Wolf 1999 ch. 10.3]. Normalising the function by the constituent powers, that is, the autocorrelations at zero lag, gives the complex degree of coherence $\gamma'_{ab}(x', y')$, where

$$\gamma'_{ab}(x', y') = \frac{\Gamma_{ab}(x', y')}{\sqrt{\Gamma_{aa}(0, 0) \Gamma_{bb}(0, 0)}} \quad [\text{Born & Wolf 1999 ch. 10.3}]. \quad (2.81)$$

The complex coherence γ_{ab} in (2.1) is effectively $\gamma'_{ab}(0, 0)$, that is, the normalised cross-correlation at zero lag.

2.A.3 Point-spread function for a uniform frequency support

The point-spread function is derived for a discrete, uniform (unweighted), finite frequency support. As the number of samples is increased, the support approaches a continuous $\text{rect}()$ indicator function, and the corresponding point-spread function approaches a $\text{sinc}()$ function.

Begin by expressing an indicator function $I(k)$ centred on k_0 with width Δk as

$$I(k) = \text{rect}\left(\frac{k - k_0}{\Delta k}\right) \quad (2.82)$$

where

$$\text{rect}(x) = \begin{cases} 1 & \text{if } |x| \leq 1/2 \\ 0 & \text{otherwise.} \end{cases} \quad (2.83)$$

Say $I(k)$ is discretised using L samples $l = 0, 1, 2, \dots, L-1$ uniformly spaced such that

$$k = l\delta k + k_0 - \Delta k/2 \quad (2.84)$$

where

$$\delta k = \Delta k/L. \quad (2.85)$$

N.B. $\delta k \neq \Delta k/(L-1)$. Observe that

$$\begin{aligned} l = 0 & \Rightarrow k = k_0 - \Delta k/2 \\ l = L-1 & \Rightarrow k = k_0 + \Delta k/2 - \delta k. \end{aligned}$$

Implicitly, each sample represents the function over the small extent $\delta k = \Delta k/L$, and therefore the L samples together cover $\Delta k = L\delta k$.

The point-spread function $s(u)$ is the Fourier transform of the indicator function $I(k)$ i.e.

$$s(u) = \frac{1}{\sqrt{2\pi}} \int I(k) e^{-jku} dk. \quad (2.86)$$

Switching to the discrete version by substituting (2.84) and implementing the integral as a normalised sum gives

$$s(u) = \frac{\delta k}{\sqrt{2\pi}} e^{-j(k_0 - \Delta k/2)u} \sum_{l=0}^{L-1} e^{-jl\delta k u}. \quad (2.87)$$

The summation is a geometric series with a known general solution, as shown below.

$$z = e^{-j\delta k u} \rightarrow \sum_{l=0}^{L-1} z^l = \frac{1 - z^L}{1 - z} \quad (2.88)$$

Therefore, the point-spread function simplifies as follows.

$$\begin{aligned} s(u) &= \frac{\delta k}{\sqrt{2\pi}} e^{-j(k_0 - \Delta k/2)u} \frac{1 - e^{-jL\delta k u}}{1 - e^{-j\delta k u}} \\ &= \frac{\delta k}{\sqrt{2\pi}} \frac{e^{-jk_0 u} e^{j(\Delta k/2)u} - e^{-j(\Delta k/2)u}}{e^{-j(\delta k/2)u} e^{j(\delta k/2)u} - e^{-j(\delta k/2)u}} \\ &= \frac{\delta k}{\sqrt{2\pi}} \frac{e^{-jk_0 u} \sin\left(\frac{\Delta k}{2}u\right)}{e^{-j(\delta k/2)u} \sin\left(\frac{\delta k}{2}u\right)} \\ &= \frac{\Delta k}{\sqrt{2\pi}} e^{-j(k_0 - \Delta k/2L)u} \frac{\sin\left(\pi \frac{\Delta k}{2\pi} u\right)}{L \sin\left(\pi \frac{\Delta k}{2\pi L} u\right)} \quad (\text{using } \delta k = \Delta k/L \text{ in (2.85)}) \quad (2.89) \end{aligned}$$

$$\begin{aligned} &\approx \frac{\Delta k}{\sqrt{2\pi}} e^{-jk_0 u} \frac{\sin\left(\pi \frac{\Delta k}{2\pi} u\right)}{\pi \frac{\Delta k}{2\pi} u} \quad (\text{use } \sin \theta \approx \theta \text{ if } L \text{ is large}) \\ &= \frac{\Delta k}{\sqrt{2\pi}} e^{-jk_0 u} \text{sinc}\left(\frac{\Delta k}{2\pi} u\right) \quad (\text{where } \text{sinc}(x) = \frac{\sin(\pi x)}{\pi x}) \quad (2.90) \end{aligned}$$

This is the standard result expected for a continuous function [Carlson 1998 app. C] (see also Appendix 2.A.5). Intuitively, as the number of samples L in the discrete

indicator function is increased, its point-spread function more closely approximates the $\text{sinc}()$ point-spread function of the continuous $\text{rect}()$ indicator function. If the number of samples is small, then (2.89) must be used. For a given $\Delta k u$ product, the error in using (2.90) in place of (2.89) can be computed.

The point-spread function in (2.89) takes extreme values when

$$u = u_n = \frac{2\pi}{\Delta k} n, \quad \forall n \in \mathbb{Z}. \quad (2.91)$$

- At $u_0 = 0$, both the numerator and denominator are zero, leading to a peak response by L'Hôpital's rule. This is the centre of the main lobe.
- At all u_n where $n \neq 0$ is not a multiple of L , only the numerator is zero, leading to a null response. The peak-to-null width is $\Delta u_n = u_1 - u_0 = 2\pi/\Delta k$.
- At all u_n where $n \neq 0$ is a multiple of L , both the numerator and denominator are again zero, leading to more peak responses. These are undesired grating lobes i.e. aliases.

Figure 2.29 shows the reference $\text{sinc}()$ from (2.90) and the exact discrete sum from (2.89) for a small value of L . Both curves have a mainlobe centred on $u = 0$, nulls at multiples of $2\pi/\Delta k$ (as specified by (2.91), and sidelobes whose peaks occur near, but not quite at, odd multiples of $\pi/\Delta k$. Importantly, the $\text{sinc}()$ does not exhibit the grating lobes of the discrete sum, which manifest at $n = L$ (and multiples thereof).




Figure 2.29: Point-spread functions for $\Delta k = 2\pi$ and $L = 7$: reference $\text{sinc}()$ from (2.90) (blue) and exact discrete sum from (2.89) (green). A common magnitude offset $20\log_{10}|s(0)|$ has been removed from both curves so that their peak is at zero.

The power function $|\text{sinc}(x)|^2$ has a maximum of unity at $x = 0$ and falls to $1/2$ (i.e. -3 dB) at $x \pm 0.443$ (found numerically). Therefore, the effective half-power resolution ρ of the point-spread function in (2.90) obtained for a uniform bandwidth of Δk is

$$\rho = w_{\text{rect}} \frac{2\pi}{\Delta k} \quad (w_{\text{rect}} = 0.886). \quad (2.92)$$

2.A.4 Windows

The high sidelobes in the point-spread functions in Figure 2.29 are due to the abrupt discontinuities at the ends of the uniform support in (2.82). It is commonplace to apply a weighting function to an aperture so that the ends taper towards zero; the sidelobes of the point-spread function are thus reduced at the expense of a wider mainlobe, because the effective aperture is smaller. Numerous weighting functions, often called windows, have been derived and compared in the signal processing literature [Harris 1978].

Figure 2.30 depicts three windows: uniform (i.e. a $\text{rect}()$ indicator function, as used in the previous section), Hamming and Taylor. The Taylor window is parameterised by (i) the number of sidelobes \bar{n} adjacent to the mainlobe that have almost constant magnitude (the sidelobes farther out will fall-off in magnitude), and (ii) the peak sidelobe level SLL relative to the mainlobe peak [Carrara *et al.* 1995 ch. D.2]. Extending (2.86), the point-spread function is the Fourier transform of the weighting function. It can be seen that, as the weighting function applies more tapering, the sidelobe magnitude reduces, but the mainlobe widens.

The mainlobe performance of a window can be quantified in terms of the factor w by which the reciprocal of the bandwidth is scaled to give the half-power (i.e. -3 dB) mainlobe width. The values for the windows in Figure 2.30 were determined numerically and are listed in Table 2.2.

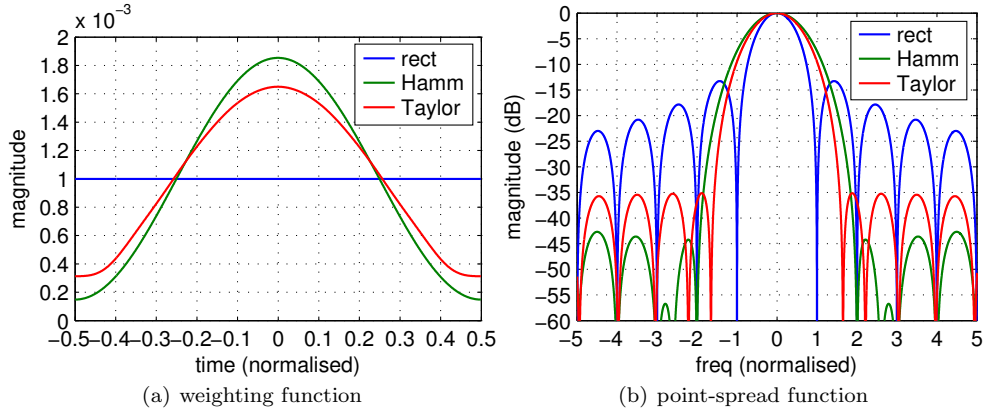


Figure 2.30: Uniform (blue), Hamming (green) and Taylor (red) windows and their point-spread functions. Each window has 1000 values and is normalised so that it sums to one. The parameters of the Taylor window are $\bar{n} = 9$ and $SLL = -35$ dB.

Table 2.2: Window scaling factor w for the half-power (-3 dB) mainlobe width.

Window	w	w/w_{rect}
Uniform (i.e. <code>rect()</code>)	0.886	1.00
Hamming	1.304	1.47
Taylor ($\bar{n} = 9$, $SLL = -35$ dB)	1.177	1.33

2.A.5 Affine transformations and the Fourier transform

If a function is defined over a coordinate space which undergoes some kind of affine transformation, how is the Fourier transform of that function affected? It undergoes a related affine transformation in frequency space which is derived below. The approach of Bracewell [2003 pp. 154-161] is followed, with an extension from two dimensions to a general N -dimensional space.

Let $f(\mathbf{x}) \in \mathbb{C}$ be a complex-valued function over an N -dimensional Cartesian coordinate space $\mathbf{x} = (x_1, x_2, \dots, x_N) \in \mathbb{R}^N$ where x_n is the n^{th} dimension. In two-dimensional space, $N = 2$ and $\mathbf{x} = (x, y)$, while in three-dimensional space, $N = 3$ and $\mathbf{x} = (x, y, z)$.

The N -dimensional Fourier transform of $f(\mathbf{x})$ is

$$F(\mathbf{k}) = \frac{1}{(2\pi)^{N/2}} \int_{\mathbf{x}} f(\mathbf{x}) e^{-j\mathbf{k}^T \mathbf{x}} d\mathbf{x} \quad (2.93)$$

where \mathbf{k} is the N -dimensional frequency space $\mathbf{k} = (k_1, k_2, \dots, k_N) \in \mathbb{R}^N$ arranged such that k_n corresponds to x_n i.e. the Fourier transform along the n^{th} dimension of \mathbf{x} on input is evaluated at frequencies along the n^{th} dimension of \mathbf{k} on output. The angular-frequency form is used here in order to match the conventions established in radar imaging [Jakowatz *et al.* 1996, Cloude 2010], although general imaging often uses ordinary frequency $\mathbf{u} = \mathbf{k}/2\pi$ [Bracewell 2003]. In two dimensions, the Fourier transform in (2.93) takes the form given by (2.75) in Appendix 2.A.2 when considering the spectrum of an image. Three-dimensional transforms are used when developing the concept of spatial-frequency support in Section 2.4.

An affine transformation

$$\mathbf{x}' = A\mathbf{x} + \mathbf{x}_0 \quad (2.94)$$

is a type of mapping between coordinate spaces \mathbf{x} and \mathbf{x}' which allows for a linear transformation A and a translation \mathbf{x}_0 . Element a_{nn} on the diagonal of A serves to expand/zoom-out (if $a_{nn} > 1$) or contract/zoom-in (if $a_{nn} < 1$) the n^{th} dimension. Element a_{mn} off the diagonal of A effects a shear whereby the m^{th} dimension is expanded (if $a_{mn} > 1$) or contracted (if $a_{mn} < 1$) in proportion to the n^{th} dimension. If $\det A = 1$ and $A^{-1} = A^T$, then A is a rotation [Murray *et al.* 1994 p. 24]; if $\det A = -1$ and $A^{-1} = A^T$, then A effects both a rotation and a reflection (together they may be called an improper rotation). Assuming the matrix inverse exists, the inverse transformation is $\mathbf{x} = A^{-1}(\mathbf{x}' - \mathbf{x}_0)$. The Jacobian of the transformation is the determinant, so $d\mathbf{x} = d\mathbf{x}' / \det A$.

Let $f_t(\mathbf{x}) = f(\mathbf{x}')$ be the original function f evaluated on a coordinate space $\mathbf{x}' = A\mathbf{x} + \mathbf{x}_0$ subject to an affine transformation. The Fourier transform of $f_t(\mathbf{x})$ is

$$\begin{aligned}
F_t(\mathbf{k}) &= \frac{1}{(2\pi)^{N/2}} \int_{\mathbf{x}} f_t(\mathbf{x}) e^{-j\mathbf{k}^T \mathbf{x}} d\mathbf{x} \\
&= \frac{1}{(2\pi)^{N/2}} \int_{\mathbf{x}'} f(\mathbf{x}') e^{-j\mathbf{k}^T A^{-1}(\mathbf{x}' - \mathbf{x}_0)} d\mathbf{x}' / \det A \\
&= \frac{e^{j\mathbf{k}^T A^{-1} \mathbf{x}_0}}{(2\pi)^{N/2} \det A} \int_{\mathbf{x}'} f(\mathbf{x}') e^{-j\mathbf{k}^T A^{-1} \mathbf{x}'} d\mathbf{x}' \\
&= \frac{e^{j\mathbf{k}'^T \mathbf{x}_0}}{(2\pi)^{N/2} \det A} \int_{\mathbf{x}'} f(\mathbf{x}') e^{-j\mathbf{k}'^T \mathbf{x}'} d\mathbf{x}' \quad (\mathbf{k}' = A^{-T} \mathbf{k} \Rightarrow \mathbf{k}'^T = \mathbf{k}^T A^{-1}) \\
&= \frac{e^{j\mathbf{k}'^T \mathbf{x}_0}}{\det A} F(\mathbf{k}'), \tag{2.95}
\end{aligned}$$

which is simply a scaled version of the original Fourier transform F evaluated on the transformed frequency space

$$\mathbf{k}' = A^{-T} \mathbf{k}. \tag{2.96}$$

If A is a rotation ($\det A = 1$ and $A^{-1} = A^T$), then $\mathbf{k}' = A\mathbf{k}$ i.e. the original coordinate space \mathbf{x} and the associated frequency space \mathbf{k} are both rotated by A . Hence, if $f(\mathbf{x})$ and $F(\mathbf{k})$ are a Fourier transform pair, then, for a rotation A_θ , $f(A_\theta \mathbf{x})$ and $F(A_\theta \mathbf{k})$ are a Fourier transform pair as well.

If A is diagonal, then expansion (contraction) of the n^{th} dimension in \mathbf{x} by a_{nn} causes contraction (expansion) of the n^{th} dimension in \mathbf{k} by $1/a_{nn}$. Moreover, if $f_n(x_n)$ and $F_n(k_n)$ are a Fourier transform pair along the n^{th} dimension, then, for a stretch a_{nn} and shift x_{0n} ,

$$f_n(a_{nn}x_n + x_{0n}) = f_n(a_{nn}(x_n + x_{0n}/a_{nn})) \quad \text{and} \quad \frac{1}{a_{nn}} e^{jk_n x_{0n}/a_{nn}} F_n\left(\frac{k_n}{a_{nn}}\right) \tag{2.97}$$

are a Fourier transform pair as well.

For the inverse Fourier transform, affine transformations have the same effect as derived above, except that the sign of the complex exponential induced by a shift is swapped. In this case, if $f_n(x_n)$ and $F_n(k_n)$ are a Fourier transform pair along the n^{th} dimension, then, for a stretch a_{nn} and shift k_{0n} of the spectrum,

$$\frac{1}{a_{nn}} e^{-jx_n k_{0n}/a_{nn}} f_n\left(\frac{x_n}{a_{nn}}\right) \quad \text{and} \quad F_n(a_{nn}k_n + k_{0n}) = F_n(a_{nn}(k_n + k_{0n}/a_{nn})) \tag{2.98}$$

are a Fourier transform pair as well.

2.A.6 Projection-slice theorem

The projection-slice theorem states that the Fourier transform of a projected function is equal to a particular slice of the Fourier transform of the original function; this slice has an orientation which matches the orientation of the projected function [Bracewell 1956, 2003 pp. 498-500]. In three dimensions, the theorem takes two forms: if the projected function is linear (after summation over two dimensions), then the slice will be a line, whereas if the projected function is planar (after summation over one dimension), then the slice will be a plane [Jakowatz *et al.* 1996 ch. 2.4.2]. The linear-slice version of the theorem is now derived for a general N -dimensional space.

Let $f(\mathbf{x}) \in \mathbb{C}$ be a complex-valued function over an N -dimensional Cartesian coordinate space $\mathbf{x} = (x_1, x_2, \dots, x_N) \in \mathbb{R}^N$ where x_n is the n^{th} dimension. In two-dimensional space, $N = 2$ and $\mathbf{x} = (x, y)$, while in three-dimensional space, $N = 3$ and $\mathbf{x} = (x, y, z)$.

The projection of $f(\mathbf{x})$ onto the x_1 axis is

$$f_1(x_1) = \int_{\mathbf{x}'} \delta(x_1 - x'_1) f(\mathbf{x}') d\mathbf{x}' \tag{2.99}$$

where \mathbf{x}' is a dummy variable for summation across \mathbf{x} and the delta function $\delta(x_1 - x'_1)$ ensures that only the $N - 1$ dimensions x_2, \dots, x_N contribute to the summation.

As in (2.93), the N -dimensional Fourier transform of $f(\mathbf{x})$ is

$$F(\mathbf{k}) = \frac{1}{(2\pi)^{N/2}} \int_{\mathbf{x}} f(\mathbf{x}) e^{-j\mathbf{k}^T \mathbf{x}} d\mathbf{x} \quad (2.100)$$

where \mathbf{k} is the N -dimensional frequency space $\mathbf{k} = (k_1, k_2, \dots, k_N) \in \mathbb{R}^N$ arranged such that k_n corresponds to x_n i.e. the Fourier transform along the n^{th} dimension of \mathbf{x} on input is evaluated at frequencies along the n^{th} dimension of \mathbf{k} on output.

Along the line k_1 , for which $k_2 = \dots = k_N = 0$, the Fourier transform becomes

$$F(k_1, 0, \dots, 0) = \frac{1}{(2\pi)^{N/2}} \int_{\mathbf{x}} f(\mathbf{x}) e^{-jk_1 x_1} d\mathbf{x} = \frac{1}{(2\pi)^{N/2}} \int_{x_1} \underbrace{\left(\int_{\mathbf{x}'} \delta(x_1 - x_1') f(\mathbf{x}') d\mathbf{x}' \right)}_{f_1(x_1)} e^{-jk_1 x_1} dx_1 \quad (2.101)$$

The right-hand side of (2.101) is the one-dimensional Fourier transform of the function after it has been projected onto x_1 , and this equals the N -dimensional Fourier transform of the original function evaluated along k_1 on the left-hand side. This establishes the projection-slice theorem for one particular line: the x_1 axis.

It remains to prove that the theorem holds for a projection onto any line. Let $\hat{\mathbf{n}}$ denote the direction of a general line. Instead of directly projecting onto this line, first rotate $f(\mathbf{x})$ to give a new function $f_\theta(\mathbf{x})$ such that the orientation of $f_\theta(\mathbf{x})$ with respect to $\hat{\mathbf{n}}$ is the same as the orientation of $f(\mathbf{x})$ with respect to the x_1 axis (i.e. this would rotate the x_1 axis onto $\hat{\mathbf{n}}$, but the axes are fixed, it is the function that rotates here). Projecting $f_\theta(\mathbf{x})$ onto x_1 is equivalent to projecting $f(\mathbf{x})$ onto $\hat{\mathbf{n}}$. Clearly, (2.101) is applicable to the rotated function $f_\theta(\mathbf{x})$ and its Fourier transform $F_\theta(\mathbf{k})$: $f_\theta(\mathbf{x})$ can be projected onto x_1 and Fourier transformed, and this will equal $F_\theta(k_1, 0, \dots, 0)$. It was shown in Appendix 2.A.5 that the Fourier transform of a rotated function is equal to the same rotation of the Fourier transform of the original function. Therefore, $F_\theta(k_1, 0, \dots, 0)$ is equivalent to the Fourier transform $F(\mathbf{k})$ of the original function evaluated along $\hat{\mathbf{n}}$. Thus, projecting onto any direction $\hat{\mathbf{n}}$ and applying the Fourier transform is equal to evaluating the Fourier transform of the original function along $\hat{\mathbf{n}}$.

Hence, the one-dimensional Fourier transform of a projected function, where the projection is onto a line with orientation $\hat{\mathbf{n}}$, is equal to the N -dimensional Fourier transform of the original function evaluated along a line oriented at $\hat{\mathbf{n}}$ in frequency space; this is the projection-slice theorem.

2.A.7 Hessian normal form

The Hessian normal form is a convenient way to define a plane in three dimensions (or more generally, a hyper-plane in N dimensions). This section establishes the form.

Note that $\mathbf{a} \cdot \mathbf{b} = \mathbf{a}^T \mathbf{b}$.

Consider a plane containing a point \mathbf{p} and having a normal $\hat{\mathbf{n}}$. A point \mathbf{x} lies in the plane if and only if

$$(\mathbf{x} - \mathbf{p}) \cdot \hat{\mathbf{n}} = 0, \quad (2.102)$$

that is, the vector from \mathbf{p} to \mathbf{x} is orthogonal to the normal $\hat{\mathbf{n}}$. Expanding gives

$$\mathbf{x} \cdot \hat{\mathbf{n}} = r \quad (2.103)$$

where $r = \mathbf{p} \cdot \hat{\mathbf{n}}$ is the distance from the origin to the closest point on the plane. To see this, expand the dot product in (2.103) as $r = \mathbf{x} \cdot \hat{\mathbf{n}} = |\mathbf{x}| \cos \theta$. Observe that the length $|\mathbf{x}| = r / \cos \theta$ achieves a minimum value of r when $\theta = 0$. Therefore, the point on the plane closest to the origin is $r\hat{\mathbf{n}}$.

2.A.8 Analogue demodulation

The received signal will usually have very low power due to the two-way propagation loss [Sullivan 2008 ch. 17.6]. Significant gain is required to bring the signal up to a level suitable for sampling, and this is usually achieved by a chain of amplifiers. However, this could lead to oscillation of RF leakage via unintended feedback loops. One way to avoid this is to block the feedback loops by shifting the signal frequency band at some point (or multiple points) in the chain [Couch 2005 ch. 12.1.7]. Furthermore, the

new centre frequency can be chosen to be a moderate ‘intermediate frequency’ (IF), well-below RF but above DC, which permits the use of fixed high-Q filters with good band-pass selectivity. The *Ingara* L-band radar system implements this frequency shift using a conventional superheterodyne receiver design whereby the received signal is mixed with a local oscillator tone; this method is described from a signal-processing perspective in Section 2.A.8.1 below.

It may also be desirable to reduce the signal bandwidth in order to lower the required sampling frequency. The *Ingara* X-band variant and the *Bright Sapphire* radar system achieve this by mixing the received signal with a version of the transmitted chirp; this method, called ‘dechirp’ or ‘deramp’, is described in Section 2.A.8.2.

Regardless of whether the demodulation process is a simple mix-down or dechirp, the receiver design can have the mixer output centred at either a convenient non-zero IF (as described above), so that the signal band does not overlap with DC, or at DC, so that the demodulated sidebands are effectively converted to positive and negative frequencies. In the latter case, the mixing operation must be implemented using a dual-channel in-phase and quadrature (I/Q) mixer, in order to properly represent the whole demodulated frequency span without aliasing, and subsequent stages, notably the sampling stage, must handle both channels. As indicated in Table 1.1, the *Ingara* system uses a non-zero IF, whereas the *Bright Sapphire* system mixes down (via dechirp) to DC using I/Q demodulation.

The following formulation presumes ideal device behaviour in order to convey the motivation behind the demodulation techniques. Practical effects such as amplifier noise and mixer intermodulation products are ignored, although of course they cannot be avoided in a real system.

2.A.8.1 Mix-down to IF

In this demodulation technique, the received RF signal $s_{rx}(t)$ centred at f_c is mixed with a local oscillator tone $s_{lo}(t)$ at f_{lo} ,

$$s_{lo}(t) = \cos(2\pi f_{lo}t + \phi_{lo}) = \frac{1}{2} \left(e^{j(2\pi f_{lo}t + \phi_{lo})} + e^{-j(2\pi f_{lo}t + \phi_{lo})} \right) \quad (2.104)$$

and filtered to give a replica $s_{mx}(t)$ of the input signal shifted down to a convenient intermediate frequency $f_{if} = f_c - f_{lo}$. Note that the local oscillator tone is ‘free-running’ in the sense that it is not stopped and restarted pulse-to-pulse, but continues across pulses, so the phase offset ϕ_{lo} may vary pulse-to-pulse. In particular, for pulse index $m = 0, 1, 2, \dots, M - 1$,

$$\phi_{lo} = \phi_{lo0} + 2\pi f_{lo}mT_{prf} \quad (2.105)$$

where $T_{prf} = 1/f_{prf}$ is the inter-pulse period for a pulse-repetition frequency (PRF) f_{prf} . In the usual case, f_{lo} is a multiple of f_{prf} , so ϕ_{lo} would be constant, but if f_{lo}/f_{prf} is not an integer, then ϕ_{lo} would vary according to (2.105)—this is inconvenient but can be remedied digitally (see Section 2.A.10).

Mixing the received signal $s_{rx}(t)$ in (2.21) with the local oscillator tone $s_{lo}(t)$ in (2.104) and suppressing amplitude factors gives

$$s_{rx}(t)s_{lo}(t) = \Re \left\{ \int_r g_p(r) \left(e^{j\Phi_{mx}(t,r)} + e^{j\Phi_{mxs}(t,r)} \right) \text{rect} \left(\frac{\tau_r - t + T_c/2}{T_c} \right) dr \right\} \quad (2.106)$$

where, from (2.22),

$$\Phi_{mx}(t,r) = 2\pi(f_c - f_{lo})t - \phi_{lo} - \Phi_c(\tau_r) + \Phi_{qc}(t - \tau_r) \quad (\text{difference component}) \quad (2.107)$$

$$\Phi_{mxs}(t,r) = 2\pi(f_c + f_{lo})t + \phi_{lo} - \Phi_c(\tau_r) + \Phi_{qc}(t - \tau_r) \quad (\text{sum component}). \quad (2.108)$$

Filtering out the sum component $\Phi_{mxs}(t,r)$ in (2.106) gives the demodulated signal

$$s_{mx}(t) = \Re \left\{ \int_r g_p(r) e^{j\Phi_{mx}(t,r)} \text{rect} \left(\frac{\tau_r - t + T_c/2}{T_c} \right) dr \right\} \quad (2.109)$$

centred at the intermediate frequency $f_{if} = f_c - f_{lo}$.

Observe that the bandwidths of the IF signal $s_{mx}(t)$ in (2.109) and the RF signal $s_{rx}(t)$ in (2.21) are the same, so the sampling requirements are unchanged by the demodulation to IF.

The *Ingara* L-band system uses $f_c = 1.32$ GHz and $f_{lo} = 1.42$ GHz, which gives a negative IF i.e. $f_{if} = f_c - f_{lo} = -100$ MHz. The positive spectral band of the real signal is shifted from $+1.32$ GHz to -100 MHz whereas the negative spectral band is shifted from -1.32 GHz to $+100$ MHz. As a consequence, the signal is swapped from an up-chirp to a down-chirp, so that before (after) demodulation the frequency increases (decreases) with time. On the other hand, any amplitude modulation retains the original variation with time.

2.A.8.2 Dechirp

Instead of mixing with a simple tone, an alternative demodulation technique is to mix the received signal with a chirp. This effectively dechirps the received signal because each chirp echo is converted to a tone whose frequency is proportional to the propagation delay of the echo [Jakowatz *et al.* 1996 ch. 1.3, app. D.2.2]. The resultant bandwidth depends on the processed swath rather than just the transmitted signal.

The dechirp chirp $s_d(t)$ must have the same chirp rate β as the transmitted signal, but all other parameters—centre frequency f_d , bandwidth B_d , duration T_d and start time $t = \tau_d$ —are flexible. In particular, an offset centre frequency ($f_d \neq f_c$) can be used to downshift the RF signal to a practical IF where $f_{if} = f_c - f_d$, and a longer signal ($T_d > T_c$) permits echoes from different ranges to be fully demodulated. Mirroring (2.14)–(2.17),

$$s_d(t) = e^{j[\Phi_d(t) + \Phi_{qd}(t)]} \text{rect}\left(\frac{t - T_d/2}{T_d}\right) \quad (2.110)$$

where

$$\Phi_d(t) = 2\pi f_d t \quad (2.111)$$

is the carrier phase and

$$\Phi_{qd}(t) = \pi\beta(t - T_d/2)^2 \quad (2.112)$$

is the quadratic dechirp phase with $\beta = B_d/T_d$. The real dechirp signal is

$$s_{dr}(t) = \Re\{s_d(t)\} = \cos[2\pi f_d t + \pi\beta(t - T_d/2)^2] \text{rect}\left(\frac{t - T_d/2}{T_d}\right) \quad (2.113)$$

$$= \frac{1}{2} \left(e^{j[\Phi_d(t) + \Phi_{qd}(t)]} + e^{-j[\Phi_d(t) + \Phi_{qd}(t)]} \right) \text{rect}\left(\frac{t - T_d/2}{T_d}\right). \quad (2.114)$$

The form of (2.114) is analogous to the local oscillator tone $s_{lo}(t)$ in (2.104), but unlike that demodulation signal, the applied dechirp chirp $s_{dr}(t - \tau_d)$ has a defined start time $t = \tau_d$ relative to the start of pulse transmission at $t = 0$.

Mixing the received signal $s_{rx}(t)$ in (2.21) with the delayed dechirp chirp $s_{dr}(t - \tau_d)$ from (2.114) and suppressing amplitude factors gives

$$\begin{aligned} & s_{rx}(t) s_{dr}(t - \tau_d) \\ &= \Re \left\{ \int_r g_p(r) s_c(t - \tau_r) dr \right\} \Re \{ s_d(t - \tau_d) \} \\ &= \Re \left\{ \int_r g_p(r) (e^{j\Phi_{dch}(t,r)} + e^{j\Phi_{dchs}(t,r)}) \text{rect}\left(\frac{\tau_r - t + T_c/2}{T_c}\right) \text{rect}\left(\frac{t - \tau_d - T_d/2}{T_d}\right) dr \right\} \end{aligned} \quad (2.115)$$

where, from (2.22), the dechirp difference component is

$$\begin{aligned} \Phi_{dch}(t,r) &= \Phi_c(t - \tau_r) + \Phi_{qc}(t - \tau_r) - \Phi_d(t - \tau_d) - \Phi_{qd}(t - \tau_d) \\ &= 2\pi f_c(t - \tau_r) + \pi\beta(\tau_r - t + T_c/2)^2 - 2\pi f_d(t - \tau_d) - \pi\beta(t - \tau_d - T_d/2)^2 \\ &= 2\pi(f_c - f_d)t - 2\pi f_c \tau_r + 2\pi f_d \tau_d + \pi\beta [t^2 + \tau_r^2 + T_c^2/4 - 2\tau_r t - T_c t + T_c \tau_r \\ &\quad - t^2 - \tau_d^2 - T_d^2/4 + 2\tau_d t + T_d t - T_d \tau_d] \\ &= 2\pi f_{if} t - 2\pi f_c(\tau_r - \tau_d) - 2\pi f_{if} \tau_d + \pi\beta [-2(\tau_r + T_c/2 - (\tau_d + T_d/2))t \\ &\quad + (\tau_r + T_c/2)^2 - (\tau_d + T_d/2)^2] \\ &= 2\pi f_{dch} t + \phi_{dch} \end{aligned} \quad (2.117)$$

with range-dependent dechirp output frequency

$$f_{dch} = f_{if} - \beta[\tau_r + T_c/2 - (\tau_d + T_d/2)] \quad (2.118)$$

and range-dependent dechirp output phase

$$\phi_{dch} = -2\pi f_c(\tau_r - \tau_d) - 2\pi f_{if}\tau_d + \pi\beta[(\tau_r + T_c/2)^2 - (\tau_d + T_d/2)^2], \quad (2.119)$$

and the dechirp sum component is

$$\begin{aligned} \Phi_{dchs}(t, r) &= \Phi_c(t - \tau_r) + \Phi_{qc}(t - \tau_r) + \Phi_d(t - \tau_d) + \Phi_{qd}(t - \tau_d) \\ &\propto 2\pi(f_c + f_d)t \quad (\text{analogous to (2.117)}). \end{aligned}$$

Filtering out the sum component $\Phi_{dchs}(t, r)$ in (2.115) gives the demodulated signal

$$s_{dch}(t) = \Re \left\{ \int_r g_p(r) e^{j\Phi_{dch}(t, r)} \text{rect}\left(\frac{\tau_r - t + T_c/2}{T_c}\right) \text{rect}\left(\frac{t - \tau_d - T_d/2}{T_d}\right) dr \right\} \quad (2.120)$$

centred at the intermediate frequency $f_{if} = f_c - f_d$. The bandwidth B_{dch} of the filter limits the available range swath, since from (2.118),

$$\begin{aligned} f_{dch, min} &= f_{if} - B_{dch}/2 = f_c - f_d - \beta[\tau_{r, max} + T_c/2 - (\tau_d + T_d/2)] \\ f_{dch, max} &= f_{if} + B_{dch}/2 = f_c - f_d - \beta[\tau_{r, min} + T_c/2 - (\tau_d + T_d/2)] \\ \Rightarrow B_{dch} &= f_{dch, max} - f_{dch, min} = \beta(\tau_{r, max} - \tau_{r, min}) = \frac{\tau_{r, max} - \tau_{r, min}}{T_c} B_c. \end{aligned} \quad (2.121)$$

From (2.121), dechirp achieves bandwidth reduction only if $\tau_{r, max} - \tau_{r, min} < T_c$ i.e. the swath is relatively small [Jakowatz *et al.* 1996 ch. 1.3]. The receiver should be designed to handle the signal bandwidth corresponding to the desired range swath.

It is worth considering how $\Phi_{dch}(t, r)$ is limited by the dechirp chirp, before the filter is applied. Using the time window of the dechirp signal in (2.115), and noting that $\beta = B_c/T_c = B_d/T_d$, the following limits on the dechirp output frequency in (2.118) can be obtained:

$$\begin{aligned} \text{Far-range echo starting at dechirp end: } \tau_{r, max} &= \tau_d + T_d \Rightarrow f_{dch, min} = f_{if} - (B_c + B_d)/2 \\ \text{Near-range echo ending at dechirp start: } \tau_{r, min} &= \tau_d - T_c \Rightarrow f_{dch, max} = f_{if} + (B_c + B_d)/2 \\ &\Rightarrow B_{dch} = B_c + B_d. \end{aligned}$$

In order to avoid generating DC output, the dechirp chirp should therefore be designed such that it does not overlap in frequency with the transmit chirp i.e.


$f_{dch, min} = f_{if} - (B_c + B_d)/2 > 0 \Rightarrow f_c - B_c/2 > f_d + B_d/2$ —the *Ingara* X-band system follows this approach. Alternatively, if the two chirps do overlap, then I/Q demodulation is required, since according to (2.118), near-range (far-range) echoes are effectively converted to positive (negative) frequencies—the *Bright Sapphire* radar follows this approach. Comparison of the waveform parameters in Table 1.1 bears out this fundamental difference in design. These practical constraints are not apparent from the standard textbook formulations of Jakowatz *et al.* [1996 ch. 1.3] and Carrara *et al.* [1995 app. C], which implicitly fix $f_d = f_c$ and assume quadrature demodulation.

Note that the residual quadratic phase components of ϕ_{dch} in (2.119) will distort SAR images formed using frequency-domain methods applied directly to dechirped data, particularly when fine-resolution, wide-area imagery is generated from low-frequency radar data [Jakowatz *et al.* 1996 ch. 2.6]. A range-varying ‘deskew’ filter can be applied digitally before image formation to remove the quadratic phase [Carrara *et al.* 1995 app. C].

To summarise, the dechirped output $s_{dch}(t)$ in (2.120) is a weighted sum of range-dependent tones where the weighting is the projected reflectivity $g_p(r)$, the frequency and phase of the constituent tones are given by (2.118) and (2.119), and the timing of the tones is limited by the intersection of the $\text{rect}()$ windows in (2.120).

The demodulation process is illustrated in Figure 2.31, where chirp echoes from different ranges are dechirped by a longer, wider-bandwidth, offset chirp, generating tones whose frequency depends on their respective propagation delay. Observe four features of this example: firstly, the signal bandwidth has been reduced i.e. $B_{dch} < B_c$.

Secondly, since $f_c - B_c/2 > f_d + B_d/2$, the dechirp output does not cross DC (i.e. B_{dch} covers positive frequencies only), so quadrature demodulation is not required. Thirdly, the tones are delayed in proportion to their range; they can be aligned after sampling by either applying the aforementioned deskew filter or, as discussed in Section 2.A.10.2, rechirping the data. Fourthly, the early (late) part of the echoes from near (far) ranges are cut-off because of the finite duration of the dechirp chirp; the available echo shrinks with increasing distance from the scene centre (beyond a central region for which the echoes and the dechirp chirp overlap in time completely), leading to both reduced power and degraded resolution.



2.A.9 Sampling

After analogue transformation to a suitable power level and frequency band, the received signal is sampled. The continuous-time function notation $f(t)$ should now be replaced by the discrete-time function notation $f[l]$ such that

$$t = lt_s + \tau_s \quad (2.122)$$

where $l = 0, 1, 2, \dots, L_s - 1$ is the sample index, L_s is the number of samples per raw pulse, t_s is the time-spacing between samples for a sampling frequency $f_s = 1/t_s$ and τ_s is the time delay to the start of sampling; the duration of the sampling window is $T_s = L_s t_s$. In what follows, the variable t is sometimes retained for ease of interpretation, but it is implicitly discretised according to (2.122).

If analogue I/Q demodulation was applied, as in the *Bright Sapphire* system, then dual-channel ‘complex’ sampling of the in-phase and quadrature components is required at a rate no less than the bandwidth. Otherwise, single-channel ‘real’ sampling of the IF signal is sufficient, as for the *Ingara* system, but the minimum sampling rate is twice the bandwidth (this is the standard Nyquist/Shannon sampling theorem) [Carlson 1998 p. 356]. Relevant system parameters are listed in Table (1.1).

For real sampling, a further restriction on the sampling rate must be enforced: aliases of the positive and negative spectral bands must not overlap. Sampling a continuous-time band-pass signal, centred on f_0 and band-limited to B , at f_s , will generate a set of aliases of the positive spectral band with envelopes $\text{rect}[(f - (f_0 + m_p f_s))/B]$ and a set of aliases of the negative spectral band with envelopes $\text{rect}[f - (-f_0 + m_n f_s)]/B$, where $m_p, m_n \in \mathbb{Z}$; a true representation of the continuous-time signal requires that none of these aliased bands overlap in the unambiguous extent $[-f_s/2, f_s/2]$. For example, the *Ingara* IF of -100 MHz prevents sampling at the minimum rate of $2B_c = 280$ MHz because the band at $[-170, -30]$ MHz would alias to $[110, 250]$ MHz for $m_n = 1$ which overlaps with the opposite band $[30, 170]$ MHz for $m_p = 0$. In fact, the sampling rate needs to be at least $f_{if} + B_c/2 - (-f_{if} - B_c/2) = 2f_{if} + B_c = 340$ MHz to ensure zero overlap. If the IF had been -210 MHz, then it would be possible to sample at the minimum rate because the band at $[-280, -140]$ MHz would alias to $[0, 140]$ MHz

for $m_n = 1$ and the opposite band at $[140, 280]$ MHz would alias to $[-140, 0]$ MHz for $m_p = -1$, just avoiding overlap.

If simple mix-down demodulation was used, then from (2.109), the sampled (complex) signal $s_s[l] = s_s(lt_s + \tau_s)$ would be

$$s_s[l] = \int_r g_p(r) e^{j\Phi_{mx}(t,r)} \text{rect}\left(\frac{\tau_r - t + T_c/2}{T_c}\right) \text{rect}\left(\frac{t - \tau_s - T_s/2}{T_s}\right) dr \quad (2.123)$$

centred at f_{if} with bandwidth B_c and $\Phi_{mx}(t,r)$ given by (2.107), whereas if dechirp demodulation was employed, then from (2.120), the sampled (complex) signal would be

$$s_s[l] = \int_r g_p(r) e^{j\Phi_{dch}(t,r)} \text{rect}\left(\frac{\tau_r - t + T_c/2}{T_c}\right) \text{rect}\left(\frac{t - \tau_d - T_d/2}{T_d}\right) \text{rect}\left(\frac{t - \tau_s - T_s/2}{T_s}\right) dr \quad (2.124)$$

centred at f_{if} with bandwidth B_{dch} and $\Phi_{dch}(t,r)$ given by (2.117). If real sampling is used, then this should be indicated by wrapping the above expressions with $\Re\{\}$.

The sampling window $\tau_s \leq t < \tau_s + T_s$, indicated by the $\text{rect}()$ windows for τ_s in (2.123) and (2.124), acts as a range gate, whereby echoes may be fully, partially or not at all sampled, depending on their propagation delay relative to the sampling window. Say that an echo counts as received if at least half of its total duration T_c is sampled; the range gate of supported delays is therefore

$$\tau_s - T_c/2 \leq \tau_r < \tau_s + T_s - T_c/2. \quad (2.125)$$

In the *Ingara* radar system, the range-gate delay τ_s to the start of the sampling window is dynamic: it is adjusted pulse-to-pulse according to $\tau_s = \tau_0 + T_c/2 - T_s/2$ so that the echo from a reference ground position is approximately centred in the receive window using the available platform position information, even as the instantaneous range $r_0 = c\tau_0/2$ to that position changes. (The dechirp delay τ_d is similarly dynamic, when applicable.) Furthermore, in standard *Ingara* stripmap mode, the reference position is continually updated in order to mark-out a reference line beside the flight-track. The radar thus images a planned strip of ground without the range gate being steered off-target by perturbations in platform motion.

For the *Ingara* L-band radar operating at typical range 3 km and using the timing parameters given in Table 1.1, the range-gate delay would be $\tau_s = \tau_0 + T_c/2 - T_s/2 \approx 16 \mu\text{s}$, so from (2.125) the minimum and maximum ranges for received echoes would be $r_{min} = 1.5 \text{ km}$ and $r_{max} = 4.57 \text{ km}$. However, this minimum range is actually less than the nominal altitude 1.72 km (5650 ft) at a typical depression angle around 35° ; the range gate will unfortunately include $[1720 \times 2/c - (\tau_s - T_c)]/T_c = 62\%$ of the strong nadir echo from under the aircraft. Increasing the operating range to, say, 5 km, leads to $\tau_s \approx 29 \mu\text{s}$, giving $r_{min} = 3.45 \text{ km}$ (and $r_{max} = 6.52 \text{ km}$), which is greater than the new altitude 2.87 km (9410 ft); this new range gate would pass only $[2870 \times 2/c - (\tau_s - T_c)]/T_c = 18\%$ of the nadir echo, at the expense of a significant power loss due to the increased propagation distance.

Analogue demodulation is designed to make the received signal convenient for sampling. After sampling, digital demodulation, described in the next section, is designed to extract the underlying baseband signal which facilitates image formation.

2.A.10 Digital demodulation

It is desirable to totally remove the carrier variation with time and represent only the baseband phase variation $-\Phi_c(\tau_r) + \Phi_{qc}(t - \tau_r)$ in the received signal in (2.21), as this phase contains the range information which is the basis for image focusing. If analogue I/Q demodulation was done, then the stored complex samples are already at baseband. Otherwise, the real, sampled IF signal can be further demodulated, this time in software, as described in Section 2.A.10.1, in order to obtain the underlying complex signal.

It is also convenient to obtain a single form for the baseband signal, regardless of the analogue demodulation. This is achieved here by rechirping the signal if dechirp had been applied, as described in Section 2.A.10.2, so that the output after digital demodulation is always a superposition of complex chirp echoes at baseband. Rechirping the signal is not necessary—standard SAR processing chains typically work with the dechirped signal directly [Jakowatz *et al.* 1996 ch. 1.3]—but it has the

important feature of recovering the original chirp spectrum common to all echoes regardless of their delay, which offers three advantages otherwise precluded by the undesirable time-frequency characteristic of the dechirped echoes observed in Figure 2.31, where the responses had both a range-dependent frequency offset and a range-dependent time delay. Firstly, RF spectral degradation, including discontinuities in the transmit waveform and interference contaminating the received data, can now be mitigated as part of the digital demodulation—the mitigation processing applied to the *Bright Sapphire* P-band channel is detailed in Section 2.A.10.3. Secondly, a window can later be applied to shape the point-spread function equally for all responses. Thirdly, spatially variant (range-dependent) spatial-frequency trimming can be implemented more easily.

To allow for the digital filtering involved in mitigating the RF spectral degradation, let the baseband chirp echoes have final useable bandwidth $B_b \leq B_c$ centred at f_b and corresponding duration $T_b \leq T_c$ such that the chirp rate is unchanged i.e. $\beta = B_c/T_c = B_b/T_b$. For the *Bright Sapphire* P-band channel, the values of B_b and T_b are given in brackets in Table 1.1. In all other cases, no filtering is applied, so $f_b = f_c$, $B_b = B_c$ and $T_b = T_c$.

After digital demodulation, the baseband signal $s_b[l]$ will be a superposition of chirp echoes centred on DC and minimally sampled at B_b using L_b complex samples per demodulated pulse such that $T_s = L_b t_b$ where $t_b = 1/B_b$ is the new time-spacing between samples. From (2.123),

$$s_b[l] = \int_r g_p(r) e^{j\Phi_b(t,r)} \text{rect}\left(\frac{\tau_r - t + T_b/2}{T_b}\right) \text{rect}\left(\frac{t - \tau_s - T_s/2}{T_s}\right) dr \quad (2.126)$$

with phase (from (2.15) and (2.16))

$$\Phi_b(t, r) = -\Phi_b(\tau_r) + \Phi_{qb}(t - \tau_r) = -2\pi f_b \tau_r + \pi\beta(t - \tau_r - T_b/2)^2. \quad (2.127)$$

If (analogue) dechirp is followed by (digital) rechirp, as proposed here, then the output signal can still be expressed using $s_b[l]$ in (2.126), assuming that the dechirp window $\text{rect}[(t - \tau_d - T_d/2)/T_d]$ in (2.124) does not limit the echoes any more than the sampling window and therefore can be neglected.

Section 2.A.10.1 describes how the baseband signal in (2.126) is obtained from the mixed-down, IF signal in (2.123). Section 2.A.10.2 describes how (2.126) is obtained from the dechirped signal in (2.124).

2.A.10.1 Mix-down to baseband

Digital demodulation of a real sampled signal to baseband is analogous to the mixing process described in Section 2.A.8.1, but now $s_{lo}(t)$ can be complex. In particular, complex samples of the second complex exponential term in (2.104) are computed for $f_{lo} = f_{if}$ and $\phi_{lo} = 2\pi f_{if} \tau_s$ over the interval $[0, T_s)$ for samples $l = 0, 1, 2, \dots, L_s - 1$, which is equivalent to the time interval $[\tau_s, \tau_s + T_s)$ with zero phase offset. Multiplying the sampled pulse by this tone shifts the spectral band at f_{if} down to DC, thus basebanding the signal. The phase offset accounts for possible pulse-to-pulse variations in the range-gate delay. In addition, if the received signal was mixed down using a local oscillator tone at f_{lo} such that f_{lo}/f_{prf} is not an integer (see Section 2.A.8.1), then the consequent pulse-to-pulse phase variation can also be removed at this point by suitable modification of ϕ_{lo} to match (2.105).

The digital mix-down must be followed by a filter to remove the sum component at $2f_{if}$. This could be easily implemented in the frequency domain as a brick-wall low-pass filter with cut-off at f_{if} i.e. by just zeroing the high-frequency part of the spectrum. However, since the baseband spectrum will have known bandwidth B_b , it is more convenient to just crop the spectrum to this bandwidth; the resulting complex baseband signal (with L_b samples) is thus down-sampled by the factor f_s/B_b to its minimal sampling rate.

It is tempting to implement the frequency shift in the frequency domain as a circular wrap, but such a shift would be limited to an integer number of samples only, which requires that f_{if} be a multiple of the frequency increment f_s/L_s . Also, this technique would not account for any pulse-to-pulse phase variations.

One complication remains. The basebanded spectrum should correspond to the positive band of the original RF spectrum [Curlander & McDonough 1991 p. 185], but

the raw sampled signal may have the positive and negative bands swapped, caused by either a negative IF, as for the *Ingara* L-band receiver discussed in Section 2.A.8.1, or (deliberate) aliasing due to the choice of sampling rate, as discussed in Section 2.A.9. Numerical software libraries such as FFTW often produce only half the spectrum of a real signal, so in order to recover the opposite half, the initial half-spectrum should be conjugated and reversed, since the magnitude and phase of the spectrum of a real signal are even and odd functions of frequency, respectively [Carlson 1998 p. 167].

2.A.10.2 Rechirp

Rechirping the sampled signal is essentially the opposite of the dechirp process described in Section 2.A.8.2. Multiplying the sampled signal by a digital replica of $s_d^*(t - \tau_d)$ i.e. the delayed conjugate of the dechirp chirp from (2.110), will cancel the dechirp components of the signal phase in (2.116).

Implementing this directly will generate a chirp phase back up at RF. To support this, the digital signal should first be upsampled in order to properly represent the upper chirp frequency $f_c + B_c/2$. Observing the RF spectrum after rechirp can be informative about the RF environment.

In fact, since the desired output is at baseband, the signal should then be mixed down and down-sampled according to the procedure in Section 2.A.10.1. Alternatively, the rechirp and mix-down steps could be combined by forming a single digital signal containing both the rechirp chirp and the mix-down tone such that both the dechirp and carrier phase components in (2.116) are cancelled, leaving only the range-dependent components shown in (2.127).

2.A.10.3 Mitigation of RF spectral degradation

The modern, congested microwave spectrum can cause RF interference [Griffiths *et al.* 2015]. In reaction to this, spectrum management authorities may impose severe restrictions on wideband transmission, potentially necessitating a discontinuous waveform. Both of these effects degrade radar performance, and both were suffered by the *Bright Sapphire* radar operating over the Tully Training Area in northern Queensland (far northern Australia).

Various techniques to suppress RF interference have been proposed for low-frequency radar [Lamont-Smith *et al.* 2006, Meyer *et al.* 2013]. In this work, two simple, intuitive methods were applied to the *Bright Sapphire* P-band data. Firstly, unusually high-power pulses in the synthetic aperture were set to zero. This involved computing and sorting the average sample power per pulse and then thresholding the difference in power between adjacent sorted pulses; a step-change of 0.5 dB was set as the threshold above which pulses were classified as outliers. It is not known what caused the occasional isolated high-power pulses observed in the data, but given the high PRF (see Table 1.1), it is very unlikely that they are valid echoes from glinting objects in the scene. Secondly, interferences manifesting at a constant delay across pulses were removed. This involved coherently summing all range-compressed pulses to give an ‘average’ range-compressed pulse which was then coherently subtracted from the pulses; signal components that constructively added in-phase across the aperture will thus be cancelled. Again, it is not known what caused the low-power, constant-delay signals observed in the data (perhaps a spurious source inside the radar), but given the long stripmap synthetic apertures processed here, it is very unlikely that they are valid echoes, since the phase response from fixed ground objects will change pulse-to-pulse, leading to destructive interference (without motion compensation).

Besides RF interference, two transmit notches mandated by the spectrum approval authority led to a discontinuous waveform. Despite various experiments, no robust method was found to coherently combine the total response over the whole, nominally large, band ($B_c = 220$ MHz). Bandwidth extrapolation across the notches is effectively required, and that implies some scattering model which may match particular point targets but is unlikely to suit natural clutter. Instead, the notches were simply avoided by selecting, cropping and basebanding a notch-free segment ($B_b = 63.5$ MHz) of the RF spectrum after rechirp.

Figure 2.32 illustrates the above mitigation techniques for the *Bright Sapphire* P-band channel in terms of the RF spectrum of the radar signal as it is successively

transformed. The theoretical notched transmit chirp (a)-(blue) and wider-band dechirp chirp (b)-(green) are plotted for reference; the remaining plots (c)–(f) show the average measured power spectral density. The RF environment given no transmission, shown in (c)-(red), indicates moderate in-band noise in the range 240–300MHz and strong out-of-band signals corresponding to television channels Ten, ABC, Seven and Nine. Furthermore, observe that whilst there are strong external signals in the larger 400.0–429.9MHz notch, there is no significant RF source in the smaller 328.7–335.3MHz notch, suggesting that this latter notch was not actually necessary. Comparing the data before (d)-(cyan) and after (e)-(magenta) interference suppression, it can be seen that the perturbations due to external noise in the range 240–300MHz have been suppressed, as desired. Finally, the selected notch-free segment of the RF spectrum (which also has comparatively low external noise) is depicted in (f)-(black).




Figure 2.32: RF spectra for the *Bright Sapphire* P-band channel at HH polarisation over the Tully Training Area in northern Queensland. (a) The transmit chirp ($f_c = 340$ MHz, $B_c = 220$ MHz) with notches 328.7–335.3 MHz and 400.0–429.9 MHz. (b) The dechirp chirp ($f_d = 344.1$ MHz, $B_d = 326.9$ MHz). (c) The received data given no transmission i.e. the external interference, including four strong signals corresponding to television channels. (d) The received signal before interference suppression. (e) The received signal after interference suppression. (f) The selected portion of the received signal used for image formation ($f_b = 367.6$ MHz, $B_b = 63.5$ MHz). The vertical power scale is not calibrated in an absolute sense, so only relative comparisons between measurements (c)–(f) are valid (theoretical plots (a) and (b) are positioned for visual clarity only).

Chapter 3

3D SAR coherent change detection for monitoring the ground under a forest canopy

3.1 Introduction

Coherent change detection (CCD) is a radar image processing technique able to detect subtle changes in natural landscapes [Rignot & van Zyl 1993, Jakowatz *et al.* 1996 ch. 5.5]. It has shown great utility for the persistent surveillance of remote areas [Sandia National Laboratories 2014]. Given two radar images of the same scene collected at different times, CCD involves nothing more than generating a map of the coherence magnitude $|\gamma|$, where the complex coherence γ was defined in (2.1). As discussed in Section 2.1, the coherence magnitude is a measure of the similarity of the complex speckle patterns in the two images, where each speckle pattern is the observed set of pixel-to-pixel fluctuations in the net scattering response determined by coherent superposition of the individual responses from a particular arrangement of scattering elements viewed from a particular direction [Goodman 1975]. Even a small rearrangement of the scattering elements (e.g. pebbles on a gravel road) may be enough to change the net response in a resolution cell, and over a wider area this would manifest as a change in the observed speckle pattern [Andre *et al.* 2015]. The coherence maps in Figures 2.22, 2.25 and 2.27 in Chapter 2 are all examples of CCD. The technique is typically applied to high-frequency (e.g. Ku-band [Tsunoda *et al.* 2000]), fine-resolution images of open ground with minimal vegetation.

In this chapter, the possibility is studied of extending the CCD technique to detect changes on the ground under a forest canopy [Pincus *et al.* 2015a]. Figure 3.1 illustrates the twin challenges imposed by the canopy: attenuation of the propagating signal and layover (projection) onto the ground. In addressing the first challenge, an operating wavelength must be selected that, on the one hand, is long enough to at least partially penetrate the canopy, but on the other hand, is short enough to be sensitive to any changes of interest on the ground. Given a suitable wavelength, the layover interference could be reduced by vertically beamforming multiple image channels acquired (ideally simultaneously) at different grazing angles; the output would be a ‘3D’ image in the sense that the contributing scattering is (coarsely) resolved in height. The magnitude of the coherence between two such 3D images, acquired from repeat passes, would serve as a measure of scene change on the ground under the canopy between the two acquisition times. To achieve coherent change detection, the beamformer must preserve phase, giving a 3D image which is complex; various beamforming methods will be analysed. This concept of 3D SAR CCD is the key proposal in this chapter—it is summarised in Figure 3.2.

To permit sufficient attenuation of the canopy, the radar system, collection geometries and forest scene must satisfy the restrictions of a rather small, but not unachievable or unrealistic, operating envelope. This envelope will be characterised through a novel analysis of a theoretical covariance matrix populated with channel-pair coherences predicted using the random-volume-over-ground (RVOG) model of a forest, which is well-established in the radar remote-sensing literature. A key result is that




Figure 3.1: The difficulties of imaging the ground under a forest canopy, and a proposed radar system to overcome them. The canopy is both a lossy propagation medium, attenuating the ground signal, and a source of scattering, giving rise to clutter that lays over onto the ground. A dual-antenna, across-track interferometer would synthesise three effective phase centres separated in grazing angle ψ by alternating the transmit antenna pulse-to-pulse and receiving every pulse on both antennas. The middle, bistatic phase centre [Willis 2008] is sampled at twice the PRF of the two outer, monostatic phase centres. By operating at a foliage-penetrating wavelength, and vertically beamforming the three SAR images corresponding to the three phase centres, the desired ground response can be separated from the canopy interference.

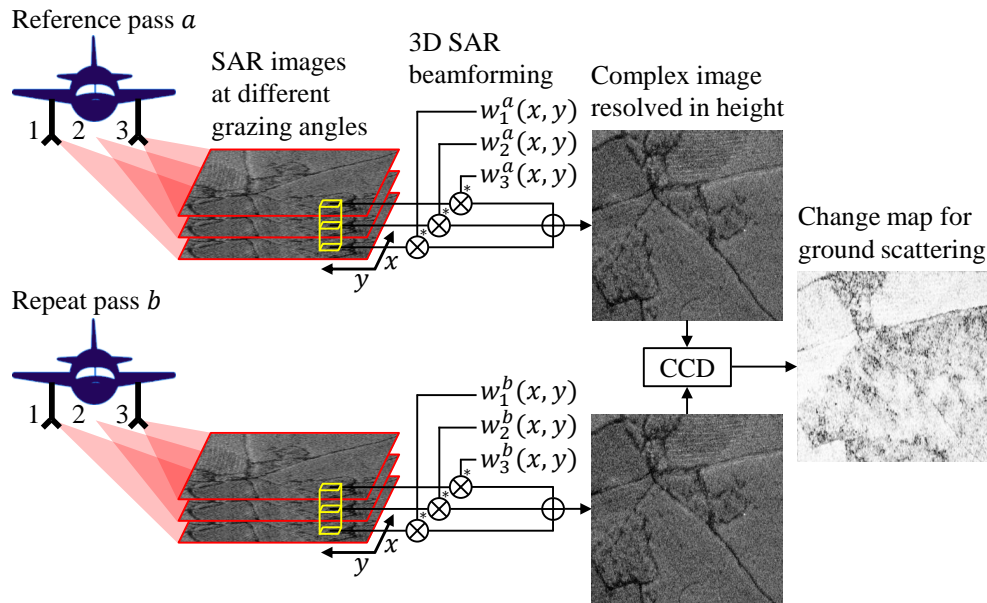


Figure 3.2: The 3D SAR CCD concept proposed in this chapter. The radar platform makes two passes over a forest scene, hours or days apart. On each pass, the multichannel radar illuminates the scene at different grazing angles using a foliage-penetrating wavelength, and a set of ordinary (i.e. 2D) SAR images (all focused to the same ground surface) are formed, one per phase centre; for example, given the across-track interferometer shown in Fig. 3.1, three images would be formed. The set of images from each pass are then vertically beamformed using complex weights $\mathbf{w} = [w_1, w_2, \dots]^T$ (which may vary spatially) to produce a complex ‘3D’ SAR image coarsely resolved in height—that is, laid-over scattering from the above-ground canopy is attenuated. Finally, the magnitude of the coherence γ_{ab} between the two 3D images gives a map of the changes on the ground below the canopy.

combining just three channels provides sufficient attenuation of the scattering response from a typical model volume to permit the true ground coherence to be estimated with reasonable accuracy.

Therefore, a suitable radar system would be a dual-antenna across-track interferometer that operates in an alternating-transmit, simultaneous-receive mode to synthesise three effective phase centres at slightly different grazing angles, as shown in Figure 3.1, thus collecting the required multichannel data in a single flight past the scene. Such a mode was first demonstrated by Intermap in Canada [Schwäbisch *et al.* 2008, Zhang *et al.* 2012]. Multichannel acquisition in a single-pass has three important advantages over multi-pass collection: firstly, single-pass acquisition is obviously much less of a collection burden; secondly, it avoids the possibility of scene changes between channels; thirdly, it gives rise to a much easier motion-compensation problem, particularly for airborne platforms.

This chapter begins with a wide discussion of the literature in order to appreciate the phenomenology of forest scattering and the established processing tools available for the proposed application. The goal of this chapter is then restated more precisely, and the novel aspects of this chapter are made clear.

The feasibility of CCD under a forest canopy relies on the radar being sensitive enough to measure the direct ground echo despite lossy propagation through the canopy, and also being sensitive to subtle changes in the ground clutter. In general, as the wavelength increases, penetration improves, but the surface backscatter response itself becomes weaker and is less sensitive to change. Therefore, the required radar noise performance is quantified, and the choice of radar operating wavelength is considered by analysing the variation of coherence with changes in clutter at different wavelengths. Operation at L-band (23 cm wavelength) would seem to offer the greatest sensitivity to change whilst still penetrating the canopy.

The main signal-processing task of 3D SAR beamforming is then described in detail. To fully understand beamforming in a SAR context, layover and interferometry are first discussed, followed by the concepts of a steering vector and a vertical beam pattern. Three beamforming techniques are presented: conventional, minimum-variance distortionless response (i.e. MVDR or Capon), and null-steer; all preserve phase. A SAR simulation is used to illustrate properties of the point-spread function in the height-range plane, along with example height-range images of a volumetric scene. Most importantly, the simulation is used to demonstrate that the complex coherence can be computed between two 3D images, each the output of an MVDR beamformer.

Measurement of ground change under a canopy is then formulated as a multichannel dual-layer coherence estimation problem, where the two layers are the ground surface and a generic volume permitting lossy propagation. The effect of the beamformer is captured by a multichannel volume attenuation factor, which raises the effective ground-volume scattering power ratio, so the output ground coherence estimate is less sensitive to bias from the volume. The beamformer that optimises the volume attenuation factor is derived and found to be equivalent to the MVDR beamformer.

The RVOG model is introduced and used to assess the performance of the beamformer in terms of volume attenuation and coherence estimation. The sensitivity to (i) deviations from the expected RVOG model and (ii) general phase-noise perturbations are investigated. Importantly, it is shown that the sensitivity to error is strongly determined by the radar design (in terms of the number of channels and their spacing), and that this sensitivity is captured by the condition number of the matrix of volume coherences between all pairs of channels.

The most significant source of error in practice is likely to be inaccurate knowledge of the height of the ground surface onto which the constituent SAR images should be focused. In standard repeat-pass change detection of bare ground, accurate knowledge of this height is not required; if the two images are focused to the wrong height (a common case), then the resulting differential layover (i.e. different horizontal shifts of the scene content, which may vary spatially due to differences in the local grazing angles and local terrain height) can be accommodated by registering one image onto the other according to some warp function derived from the images [Jakowatz *et al.* 1996 ch. 5.3.2]. By contrast, in the proposed processing scheme to achieve repeat-pass change detection of forest-covered ground, a 3D SAR beamformer is first used to isolate scattering at the focus height (limited by the vertical resolution and the array ambiguity), so subsequent change detection of the ground requires that the focus height matches the ground height,

otherwise changes will be detected for some other scattering layer. Thus, knowledge of the true ground height (in the frame of reference that the position of the radar platform is determined, and at least within the unambiguous extent of the beamformed array) is needed. This is a very demanding requirement, and indeed, determining the topography of forested terrain is a research problem in its own right [Wheeler & Hensley 2000]. For most of this chapter, it is assumed that this height is known exactly, but realistically, it would be known only approximately. It is shown that the multichannel coherence is extremely sensitive to even small offsets (of the order of a few metres) between the SAR focal surface and the ground scattering surface, but that this sensitivity can be turned into an advantage by using it to determine the ground height.

The proposed processing scheme is tested using two separate simulations (unfortunately, no suitable real data was available during this study). Firstly, RVOG clutter is synthesised at the level of raw pulse samples by accumulating the echoes from many point scatterers randomly distributed in 3D space. Various radar designs, collection geometries and beamformers are then compared. Even with only three channels per pass, it is found that as long as the sample covariance matrix supplied to the MVDR beamformer on each pass is an accurate estimate of the underlying covariance matrix, then the volume is strongly suppressed by this beamformer, and the resulting repeat-pass coherence accurately reveals subtle changes in the ground clutter.

Secondly, the open-source software program PolSARproSim is utilised. It was developed by Williams [2006] and has been used by several authors [Zhang *et al.* 2008, Lavalle *et al.* 2009, Ballester-Berman & Lopez-Sanchez 2010, Ainsworth & Williams 2011, Hensley *et al.* 2014]. It directly generates complex polarimetric SAR images using a “physics-based approach” [Williams 2006 ch. 3.1.3] derived from electromagnetic theory (but it is not a full-wave electromagnetic field solver). In order to test repeat-pass multichannel change detection, the source code was modified by adding support for multiple channels per pass and adding a mechanism to make controlled changes in the scattering elements that constitute the ground. Unlike the RVOG model, this simulation accounts for ground-trunk double bounce, a scattering mechanism that is highly correlated across passes, so it masks potential changes on the ground near trees, and whose effective phase centre is at ground height, so it cannot be suppressed by vertical beamforming.

To overcome this, a polarimetric filtering algorithm is proposed that first estimates the average scattering mechanism corresponding to ground-trunk double bounce by using an eigendecomposition of the polarimetric covariance matrix, and then projects the data onto an orthogonal scattering mechanism. The coherence contrast between changed and unchanged areas is thereby increased.

It is worth reiterating that the goal here is to detect and form images of subtle changes on the ground under a forest canopy. The goal is not to characterise or monitor the forest structure itself. Nor is the goal to directly detect man-made targets, whether static or moving, under the canopy. Nonetheless, the proposed approach draws on techniques originally intended for these other goals. Therefore, in the following subsections, relevant background is presented from the areas of foliage penetration, hidden target detection, 3D SAR and general change detection.

3.1.1 Foliage penetration

The practicality of the ideas put forth in this chapter relies on the physical phenomenon of foliage penetration by radio waves, which has long been exploited by radar systems [Davis 2011]. Development of foliage penetrating radar was motivated by the U.S. experience during the Vietnam War [Davis 2011 p. 4]. The first systems were mounted on temporary masts in order to look down through the nearby canopy and detect motion by the Doppler response across pulses [Davis 2011 ch. 1.1]. An airborne imaging system, IMFRAD, was built for the US Airforce in the mid 1970s, prompting one officer (rank: 1st Lt) to announce that “aerial detection of targets in dense foliage won’t be a problem anymore ... [long] wavelengths provide a natural filtering effect that rejects the echoes from small objects, while highlighting tactical sized targets under foliage.” [Seipel 1976]. However, the low resolution (e.g. 50×50 ft [Seipel 1976]), large size and non-real-time processing of foliage penetrating radar systems in general discouraged further investment for several years [Davis 2011 p. 21].

A renaissance in the late 1980s saw the development of a variety of airborne imaging

radar systems with a foliage penetration capability; several are listed in Table 3.1. Many experimental studies followed, seeking to evaluate this capability, and its impact on target detection, as a function of frequency, polarisation, incidence angle and forest type [Giglio 1994, Bessette & Ayasli 2001]. Key results are now given in terms of the observed foliage attenuation rate in decibels per metre, which is a useful metric for comparing different types of measurements by different sensors at different sites, although it is somewhat misleading because it implies that the canopy is homogeneous, which is generally not the case. The experimental methodologies are also briefly described in order to show the difficulty and variability when characterising foliage penetration.

Table 3.1: Early airborne SAR systems with a foliage penetration capability.

Name	Main organisation	Operating bands [†]	Features	References
P-3 SAR	Environmental Institute of Michigan (ERIM)	L, C, X, UHF	<ul style="list-style-type: none"> • Polarimetric • Installed in a U.S. Navy P-3 aircraft • Ultrawideband UHF (515 MHz bandwidth in selectable range 215–900 MHz) added later 	[Sullivan <i>et al.</i> 1988] [Sheen <i>et al.</i> 1996]
AIRSAR	Jet Propulsion Laboratory (JPL)	L, C, P	<ul style="list-style-type: none"> • Polarimetric • Simultaneous tri-band operation 	[Held <i>et al.</i> 1988]
FOLPEN II	SRI* International	P	<ul style="list-style-type: none"> • Ultrawideband (200–400 MHz) • Waveform: 5 ns impulse transmitted at 100 kHz 	[Vickers <i>et al.</i> 1988]
E-SAR	German Aerospace Center (DLR)	L, C, X, P	<ul style="list-style-type: none"> • X, P and polarimetric operation added later 	[Horn 1988] [Horn 1996]
CARABAS	Swedish Defence Research Establishment (FOA)	VHF	<ul style="list-style-type: none"> • Ultrawideband (20–90 MHz) • Antenna: two 5 m dipoles in inflatable canvas sleeves trailing behind aircraft 	[Hellsten 1992] [Hellsten <i>et al.</i> 1996]

*Formerly the Stanford Research Institute.

[†]Operating band: centre frequency f_c , wavelength λ (nominal designations) [IEEE Std 521-2002]

VHF:	100 MHz,	3 m
P:	350 MHz,	85 cm
UHF:	667 MHz,	45 cm
L:	1.3 GHz,	23 cm
C:	5.4 GHz,	5.6 cm
X:	10 GHz,	3 cm

Ulaby *et al.* [1990a] used a vector network analyser and an elevated (19 m high) parabolic reflector antenna operating at 1.6 GHz to illuminate a short, dense pine forest (13.7 m average height, 2300 trunks per hectare) at 40° incidence. By comparing the polarimetric response with and without a trihedral corner reflector (height 85 cm $\approx 4.5\lambda$) under the canopy, foliage attenuation rates of 0.52 dB/m for horizontal polarisation (HH) and 0.51 dB/m for vertical polarisation (VV) were inferred. These rates are very high for L-band—the authors describe the forest as “a somewhat extreme case in terms of the foliage penetration problem” due to its high density and thick foliage.

Durden *et al.* [1991] used AIRSAR data to characterise two conifer stands in northern California: one tall and sparse (30 m average height, 200 trunks per hectare), the other much denser (25 m average height, 700 trunks per hectare). Comparing the magnitude of the responses of several trihedrals (2.44 m) in the open and under the canopy, worst-case foliage attenuation rates of 0.15 dB/m at L-band and 0.27 dB/m at C-band were observed at 40° incidence (both polarisations) for the sparse stand, although the attenuations were highly variable across reflectors due to the particular arrangement of trees surrounding each one. For the dense stand, the reflectors were not even visible.

Sheen *et al.* [1994] used an ultra-wideband transmitter (300–1300 MHz) mounted on an elevated horizontal rail (15 m high, 10 m long) and a ground-based receiver to measure the foliage attenuation through four different forest canopies. At 800 MHz (37.5 cm wavelength) and 45° incidence, the (HH, VV) results were 0.41, 0.40 dB/m for red pine (9.8 m average height, 1900 trunks per hectare), 0.51, 0.33 dB/m for plantation scotch pine (6.8 m average height, 1800 trunks per hectare), 0.11, 0.40 dB/m for hardwood sugar maple (18.6 m average height, 2000 trunks per hectare) and 0.20, 0.10 dB/m for aspen (13.5 m average height, 1600 trunks per hectare). Observe that tree type is a much stronger determinant of the attenuation rate than tree height. Furthermore, different tree types can have very different polarimetric responses, suggesting a strong dependence on branch orientation at this wavelength.

Fleischman *et al.* [1996] studied AIRSAR data acquired in 1990 over a short, dense conifer forest in northern Maine (11.5 m average height, 1850 trunks per hectare [Durden *et al.* 1993]) consisting mainly of balsam fir with densely-packed foliage. Comparing the magnitude of the responses of trihedrals (2.44 m) in the open and under the canopy, median foliage attenuation rates at 45° incidence for polarisations (HH, VV) were 0.93, 0.85 dB/m at C-band, 0.39, 0.41 dB/m at L-band and 0.17, 0.29 dB/m at P-band. Observe that the rates at L-band and C-band are much higher than those from the earlier AIRSAR measurements over the much taller but sparser Californian forest, indicating that the key influencing factor here is not tree height but tree density. In addition, only at P-band was there a clear differential polarimetric response and a large HH-VV phase difference, which implies that the scattering mechanism of ground-trunk double bounce was dominant only at P-band [Freeman & Durden 1998]. Nonetheless, at L-band, the double-bounce mechanism may be stronger than the direct ground backscatter, as is assumed in PolSARproSim for HH [Williams 2006 ch. 8.14.25].

Binder *et al.* [1995] applied the same technique using data acquired in 1993 by the FOLPEN II (UHF) and CARABAS (VHF) sensors to obtain attenuation rates (HH only) of 0.06 dB/m (UHF) and 0.03 dB/m (VHF) for a 25 m tall tropical forest in Panama (using eleven 4.88 m trihedrals), and 0.10 dB/m (UHF) and 0.04 dB/m (VHF) for a ~12 m tall, mixed coniferous-deciduous forest in northern Maine (using thirty-four 2.44 m and 4.88 m trihedrals!). Comparing to the earlier AIRSAR measurements, the authors conclude that “foliage attenuation decreases with decreasing frequency, beginning at prohibitive levels at C-band and reaching negligible amounts at VHF”.

Bessette & Ayasli [2001] similarly analysed P-3 and CARABAS data acquired between 1995 and 1999 from five different forest sites around America. They observed that the variation in the median attenuation across sites (~5 dB) was usually significantly smaller than the variation across corner reflectors within a site (~15 dB), indicating the strong influence of local heterogeneity. Of particular significance was their empirical model for the total attenuation caused by a forest canopy, developed by pooling these and past datasets from AIRSAR, FOLPEN II, P-3 and CARABAS:

$$\text{two-way attenuation (dB)} = \alpha f_c^\beta / \cos \theta \quad (3.1)$$

where f_c is the centre frequency in MHz, θ is the incidence angle, and α and β are regression parameters. For a 50% fit to all of the available data (i.e. the middle trend line), $\alpha = 0.18$ and $\beta = 0.53$ for HH, and $\alpha = 0.30$ and $\beta = 0.47$ for VV. Note that this model does not account for tree height, which is not a strong determinant of attenuation. Nonetheless, for ease of comparison, the total attenuation can be converted to an attenuation rate for a given tree height; choosing a reference height of 20 m and an incidence angle of 45°, the predicted (HH, VV) attenuation rates are 0.20, 0.22 dB/m at L-band and 0.10, 0.12 dB/m at P-band. To reiterate, these are rough indicators of average canopy penetration across the sites considered by the authors, and the actual measurements varied significantly.

Alternative nominal attenuation rates are given by Cloude *et al.* [2000]: 0.5 dB/m at C-band, 0.1 dB/m at L-band and 0.05 dB/m at P-band. Considering these values together with those derived from Bessette & Ayasli [2001] suggests that, on average, the attenuation rate at P-band would be half that at L-band.

The variation of penetration with frequency leads to a corresponding variation in the height of the effective phase centre observed via radar interferometry [Dall 2007], where the primary observable is not the backscatter power (or the scattering matrix for polarimetric data) but the complex coherence between two observations at

slightly different incidence angles [Bamler & Hartl 1998]. Exploiting this variation, VandenBerg *et al.* [1996] generated two interferometric height maps of a forest scene, one using P-3 SAR operating at P-band (approximately) and the other using an X-band system. The heights were found to have a vertical offset approximately equal to the height of the trees; the P-band phase centre was on the forest floor, whereas the X-band phase centre was near the top of the canopy. GeoSAR, a dual P- and X-band single-pass interferometric radar system, was subsequently developed by JPL in order to map the topography of forested terrain [Wheeler & Hensley 2000]. An analogous relationship between penetration and frequency has been measured interferometrically for ice, with penetration depths of 12–35 m at C-band [Weber-Hoen & Zebker 2000] and approximately 30 cm at Ka-band ($\lambda \sim 8.5$ mm) [Hensley *et al.* 2016].

Briefly, high-resolution imaging under foliage has also been demonstrated using airborne lidar (laser radar) systems ($\lambda \sim 0.5 \mu\text{m}$) [Marino & Davis 2005], but the illuminated swath is an order of magnitude smaller than that for SAR systems, making it impractical for wide-area surveillance.

3.1.2 Forest scattering models

In parallel with experimental studies, models were developed that characterise forest scattering in terms of scattering mechanisms, notably the direct canopy (volume backscatter) response, the direct ground (surface backscatter) response and ground-trunk double bounce (a sequence of two forward reflections) [Durden *et al.* 1989, 1991, 1993, Moghaddam & Saatchi 1995, Freeman & Durden 1998]. The mechanisms are derived using approximate expressions for electromagnetic scattering from canonical shapes (e.g. cylinders and spheroids), surfaces and volumes [Cloude 2010 ch. 3 and references therein]. This approach is a compromise between full electromagnetic simulation of a forest, which is computationally impractical, and reliance solely on experiment, which would lack wide applicability. Moreover, the relative balance of the different mechanisms provides insight into the effective structure of a forest at different frequencies and polarisations.

Durden *et al.* [1991] inferred the level of foliage penetration according to the extent to which different scattering mechanisms manifested in polarimetric SAR data: a canopy consisting of randomly oriented elements should scatter with approximately equal and in-phase co-polarisations (i.e. $|\text{HH}/\text{VV}| \sim 0$ dB and $\arg\{\text{HH}/\text{VV}\} \sim 0^\circ$), whereas pure ground-trunk double bounce should exhibit a horizontally polarised response that is stronger and out-of-phase with the vertically polarised response (i.e. $|\text{HH}/\text{VV}| > 0$ dB and $\arg\{\text{HH}/\text{VV}\} \sim 180^\circ$) [Ulaby *et al.* 1987, van Zyl 1989]. Thus, the greater the differential in the observed polarimetric response, the greater the penetration through the canopy must have been. Note that if the canopy consists of elements with some dominant orientation, and/or the surfaces have dielectric properties such that the Brewster effect is significant, then this logic may not be wholly applicable [Guinvarc’h & Thirion-Lefevre 2015, Cloude 2010 ch. 3].

A rich area of study has been the extension of these forest scattering models for radar interferometry [Hagberg *et al.* 1995, Askne *et al.* 1997, 2003, Treuhaft *et al.* 1996, Treuhaft & Cloude 1999, Treuhaft & Siqueira 2000, Ballester-Berman & Lopez-Sanchez 2007, 2011]. The key observable in this case is the complex interferometric coherence between two radar observations of a scene made at slightly different incidence angles. This coherence depends on the product of the individual coherences associated with various physical processes and features of SAR image formation, notably system noise, scene change and volume decorrelation, where the latter is the coherence loss due to the coherent superposition of laid-over (i.e. uncompensated) scattering from objects at different heights [Zebker & Villasenor 1992, Cloude 2010 ch. 5.2.2, 5.2.5]. In this context, scene change is often called temporal decorrelation to indicate a loss of coherence over time.

Analytical models for the interferometric coherence of different polarimetric channels have been developed which rely on only a few physical parameters [Cloude 2010 ch. 7]. Given radar data that partially penetrates the canopy, inversion of these models permits approximate estimation of these parameters and thus the forest structure [Cloude 2006, Garestier *et al.* 2008, Neumann *et al.* 2010]. The same idea can be applied to crops [Ballester-Berman *et al.* 2005, Preiss & Stacy 2005, Erten *et al.* 2016]. Such inversion methods, collectively termed PolInSAR, have been a key focus of much

of the radar imaging literature for the past two decades. In several influential publications, Cloude and Papathanassiou detailed how the polarimetric variation of the interferometric coherence could be exploited in order to maximally separate the phase centres corresponding to different scattering mechanisms, thus enabling more accurate and more robust parameter estimation [Cloude & Papathanassiou 1997, 1998, Cloude *et al.* 2000, Papathanassiou & Cloude 2001, Cloude & Papathanassiou 2003] (see also Flynn *et al.* [2002] and Colin *et al.* [2005a] for key contributions).

Of particular note is the dual-layer random-volume-over-ground (RVOG) model [Treuhaft *et al.* 1996, Cloude 2010 ch. 5.2.4, 7.2-7.4], which treats the forest canopy as a homogeneous volume, consisting of many randomly positioned and randomly oriented scattering elements, through which the wave propagates before and after scattering off the ground surface at the bottom of the volume. The volume has only two parameters: its height and the (constant) attenuation rate for propagation. The overall RVOG structure is further parameterised by a ground-volume scattering power ratio and the height of the ground surface, which is usually known only roughly. Inversion using P-, L-, C- and X-band SAR data has repeatedly been shown to permit moderately accurate estimates of ground topography and canopy height over boreal (coniferous), temperate, tropical and mangrove forests, with RMS errors (compared to LIDAR measurements or in-situ ground truth) of 1.5–3.5 m [Praks *et al.* 2007, Hajnsek *et al.* 2009b, Mercer *et al.* 2009, Praks *et al.* 2012, Kugler *et al.* 2014, Lee & Fatoyinbo 2015, Kugler *et al.* 2015, Khati *et al.* 2018]. Hence, even though the RVOG model is of course an extreme simplification of a forest, it does broadly capture the gross features of forest scattering.

In the simplest form of the RVOG model, the only source of coherence loss is volume decorrelation. This is suitable for dual-antenna, single-pass interferometry, but for single-antenna, repeat-pass interferometry, the scene may change between the two observations. Generalisations and extensions of the RVOG model have been developed that include extra coherence factors and/or structural parameters to account for temporal decorrelation of the ground and volume layers, with an emphasis on modelling natural effects such as wind (which causes structural change in the volume) and moisture variation (which causes electromagnetic changes in both layers) [Askne *et al.* 2003, Lavalle *et al.* 2012, Lee *et al.* 2013, Lavalle & Hensley 2015].

In this chapter, the RVOG model is repurposed in order to predict the performance of the proposed change-detection algorithm efficiently over a wide range of scenes. For this application, the model is formulated to include separate coherences for the ground and volume layers, with the volume coherence assumed to be determined solely by volume decorrelation, and the ground coherence the key unknown to be estimated as a measure of change. Temporal decorrelation parameters that model particular physical processes are not included.

Trading-off simplicity and invertibility for greater realism, various methods of simulating the electromagnetic response of model forest scenes have also been developed [Ulaby *et al.* 1990b, Lin & Sarabandi 1999, Thirion *et al.* 2006, Williams 2006, Dhar *et al.* 2012]. Comparison of these simulations with real data reveals the dominant sources of scattering at different frequencies: small branches and large leaves in the upper canopy at C-band, branches in the middle and lower canopy, plus some ground-trunk double bounce, at L-band, and very large branches together with ground-trunk double bounce at P-band [Sarabandi & Lin 2000, Thirion-Lefevre & Colin-Koeniguer 2007, ESA Report 2012 p. 45]. In addition, ground-trunk double bounce (whose phase centre is at ground level) is stronger for HH compared to VV, so the effective phase centre for HH is lower.

3.1.3 Hidden target detection

Having considered the forest itself in terms of measurement, modelling and estimation, now consider the problem of detecting a target hidden under the forest canopy. This is an obvious application for a foliage-penetrating imaging radar. The target-to-clutter ratio in the final SAR image is determined not just by the level of foliage penetration, but also by the strength of foliage backscatter [Winter *et al.* 1994, Fleischman *et al.* 1996, Binder *et al.* 1995, Bessette & Ayasli 2001] and by the effect of foliage-induced fluctuations on the observed impulse response [Durden *et al.* 1993, Sheen *et al.* 1994, Touns *et al.* 1996]. The latter degradation arises when the modulation (in both magnitude and phase) of the target response by the canopy varies azimuthally, so the

target point-spread function is distorted and does not benefit from the full coherent gain expected over the synthetic aperture. Toups *et al.* [1996] used transmitters under the canopy together with AIRSAR (as receiver) to measure the magnitude and phase of the modulation pattern imposed by the canopy, and then quantified the effect of this pattern in terms of the integrated sidelobe ratio of the predicted point-spread function; a high ratio would cause low image contrast which then limits detectability [Carrara *et al.* 1995 ch. 8.3.3]. They found that this ratio increased (worsened) in proportion to the level of foliage attenuation; the effect was particularly strong at C-band. Hence, higher frequencies suffer not only from higher foliage attenuation, which reduces the average target power, but also greater fluctuations in magnitude and phase, which reduce the focus quality. The azimuthal variation in penetration can also significantly modulate the observed polarimetric response of the target [Pincus *et al.* 2013, Sletten & Brozena 2016]. These results further highlight the local heterogeneity of real forest scattering, which can be difficult to model.

Results from the study of forest phenomenology fed into the development of a variety of target detection algorithms Nguyen *et al.* [1997]. If the canopy response is strong, then polarimetric [Marino *et al.* 2010] or polarimetric-interferometric [Cloude *et al.* 2004] techniques may be able to distinguish man-made targets (particularly vehicles) hidden underneath. If the penetration is high (the more common case in the literature), then the main task is to distinguish between the point-like response of a target and the numerous bright points due to ground-trunk double bounce, which in this context are false alarms [Colin *et al.* 2005b, Lundberg *et al.* 2006]. In a single SAR image, the two types of responses can look very similar, but while the ground-trunk clutter is usually azimuthally and temporally stable, the target signal is likely to vary with aspect angle (although this variation decreases with frequency [Hellsten *et al.* 1996 Fig. 2 & 3]) and the target itself may be present in or absent from the illuminated scene. Hence, some kind of comparative processing applied to images acquired at different squint angles or times may highlight anomalies that are potential targets. Halversen *et al.* [1994] compared these two approaches using FOLPEN II data collected over a Maine forest containing concealed trucks. To exploit the azimuthal variation, the available pulses from a single pass were split into sections (ideally three, not two, to capture the potential broadside flash) and used to form lower-resolution SAR images at different squints. To capture the temporal variation, repeat-pass datasets—acquired before and after controlled vehicle arrivals—were used to form full-resolution SAR images which then had to be carefully co-registered. Applying the same adaptive difference processing to the split-aperture and repeat-pass image sets, the target detection performance was comparable, and much better than that achieved using a constant false-alarm rate (CFAR) filter applied to a single image.

The angular dependence of the target response has recently motivated the use of a circular flight-track and a full 360° synthetic aperture, which not only ensures that the brightest target glints are captured (for detection) [Fröling *et al.* 2012] but may permit better characterisation of the target (for recognition) when combined with other techniques DeGraaf [2005]. If the angular dependence is known and the resolution is fine enough, then the predicted target signature could potentially be discriminated from the forest clutter [Sletten & Brozena 2016]. Of course, circular flight-tracks entail a very large penalty in coverage efficiency (area per unit time), so this approach is not pursued here.

Repeat-pass comparative processing for target detection is commonly referred to as change detection [James & Hendrickson 1994]. In this context, change detection is a method of target detection whose key advantage is the suppression of false alarms generated by constant clutter. What is really being detected is a change in the status of a target: from absent to present (appearance), or present to absent (disappearance). Secondary ground disturbances associated with arrival or departure, such as vehicle tracks, are irrelevant in this context. Good examples of this approach are the target detection algorithms developed for the CARABAS sensor, which rely on repeat-pass change detection to suppress the bright, point-like false alarms due to ground-trunk double bounce by up to two orders of magnitude [Ulander *et al.* 1999, 2004, 2005, 2011a]. The trade-offs for this clutter removal are the added difficulties of image registration Silk & Hunter [1999] and differential RFI mitigation [Ulander *et al.* 1999], where the latter problem may arise due to different RF interferences at the different acquisition times. In the VHF/UHF case, the need for even greater clutter suppression

has motivated the use of a bistatic acquisition geometry, whereby the transmitter and receiver are at different incidence angles such that the ground-trunk double bounce is substantially avoided [Ulander & Martin 2005, Ulander *et al.* 2010, 2011b]. This bistatic approach has been combined with change detection [Ulander *et al.* 2012]. An isolated change of the polarimetric response within a forest may also be indicative of a hidden target [Marino *et al.* 2013, Rosa *et al.* 2016].

Briefly, if the target is moving under the canopy, then along-track interferometry, using two antenna elements at different along-track positions under the platform, permits detection of the target via its Doppler signature [Soumekh 1997, Davis 2011 ch. 8.3]. Commercial systems have been developed based on this [Robinson 2007].

3.1.4 3D SAR

When forming a conventional SAR image resolved in range and azimuth, scattering responses from objects above and below the focal surface will lay over onto the surface, thus obscuring the actual response at the surface height [Jakowatz *et al.* 1996 ch. 2.4.5, 3.7.5, Appendix C]. The projection occurs over the wavefront i.e. along a curve of constant range—see Figure 3.1. For penetration through a three-dimensional medium such as a forest, vertical resolution would clearly be of great benefit in separating scattering contributions from different heights. Together with conventional azimuth and range resolution, the processed SAR output would resolve the scene in all three dimensions, hence the term ‘3D SAR’ used in this chapter. An equivalent term in the literature is ‘SAR tomography’, so-called because the vertical resolution is achieved by coherently combining multiple SAR images acquired at slightly differently incidence angles; each image is a ‘slice’ or $\tau\acute{o}\mu\omicron\varsigma$ in Ancient Greek [Chan & Farhat 1981, Reigber & Moreira 2000]. The initial motivation for combining multiple acquisitions was simply to improve the interferometric estimation of ground height [Homer & Longstaff 1995, Lombardini & Lombardo 1996], but the potential of vertical resolution followed immediately [Knaell & Cardillo 1995, Homer *et al.* 1996].

In its basic concept, 3D SAR simply extends the principles of coherent aperture synthesis from the azimuth dimension (for ordinary SAR) to the elevation dimension [Hocor & Kassam 1990]. Given pulse measurements of the scattering response of the scene at different azimuth and elevation angles, the total spatial-frequency support forms a wedge in k -space, as depicted in Figure 2.4, with radial extent determined by frequency diversity and angular extents (in azimuth and elevation) provided by spatial diversity [Jakowatz & Wahl 2001, Pincus *et al.* 2009]. The resulting 3D point-spread function (including 3D resolution and grating lobes) depends on the nominal widths and sample spacings of the wedge along the cardinal directions, as discussed in Section 2.4 and Appendix 2.A.3.

In practice, for airborne SAR, the k -space wedge is rarely acquired or processed directly. Instead, as first demonstrated by Reigber & Moreira [2000] using E-SAR, a monostatic radar system is repeatedly flown past a scene at slightly different incidence angles. A conventional 2D SAR image is formed for each pass. The set of images must be mutually registered such that a 2D pixel index (m, n) corresponds to the same real world position in all images—this may not be the case initially due to uncompensated platform position offsets, which may be relatively constant pulse-to-pulse (so do not cause defocus) but vary pass-to-pass. Critically, an additional phase calibration (an autofocus across passes) may be required to effectively register the images to a sub-pixel level and correct for any phase distortions caused by system variations e.g. different parts of the antenna beam pattern being used on different passes [Kragh 2009, Boss *et al.* 2010, Tebaldini & Guarnieri 2010, Tebaldini *et al.* 2016a,b, Ponce *et al.* 2016, Cantalloube *et al.* 2017]. As a further complication, scattering objects above or below the focal surface will layover onto this surface differently in different images, depending on the incidence angle, so any scattering sources used for calibration must be stable across passes. The ground itself may layover in this way, if the focal surface does not match the local topography, which is usually known only approximately. Once registered and calibrated, the 2D images can be coherently combined using a conventional delay-and-sum beamformer, essentially equivalent to that described in Section 2.10 for azimuth focusing, giving a new image focused to a single height. Off-height scattering will be attenuated according to the point-spread function. By repeatedly refocusing the 2D images to a contiguous range of heights and stacking the

outputs, a 3D representation of the scene is obtained.

However, having to make multiple passes over the scene is an immense burden, so the number of images providing elevation support is typically two orders of magnitude less than the number of pulses providing azimuth support for each image (~ 10 vs ~ 1000) [Nannini *et al.* 2009]. In addition, these passes are usually nonuniformly spaced—see the set of TerraSAR-X satellite images used by Reale *et al.* [2011]. The vertical point-spread function obtained with conventional synthetic aperture processing is therefore often too poor to properly distinguish scattering elements vertically in forest and urban scenes, with (i) low resolution (due to a small total elevation extent across all passes), (ii) grating lobes that severely limit the unambiguous height (due to a large elevation separation between adjacent passes), and (iii) high sidelobes (due to nonuniform spacing between passes)—see the point-spread functions studied by Lombardini & Pardini [2008]. To overcome these limitations, more exotic processing techniques have been adapted from the fields of beamforming and sparse signal estimation.

A common approach is to recast SAR multipass processing as a beamforming problem in its own right [Lombardini *et al.* 2003]. For each pixel position, the corresponding pixel values in the available images together serve as the vector input to an established array processing algorithm [Van Trees 2002], as if the 2D-focused scattering responses had been measured directly and simultaneously by a generic spatially distributed sensor array. The output ‘beam’ is a new pixel whose value is an estimate of the reflectivity at a particular height, as selected by a phase-only steering vector. Now instead of just doing conventional phase-delay-and-sum beamforming across images, a wide variety of higher-resolution, sidelobe-suppressing reconstruction algorithms can be applied.

One such method is Minimum-Variance Distortionless-Response (MVDR) beamforming, which weights and sums the inputs so as to minimise the output power subject to unity gain at the steered height, thereby nulling interferences that would otherwise leak through sidelobes or even the (outer) mainlobe of the vertical beampattern [Van Trees 2002 ch. 6.2,6.3,6.6]. This improved performance comes at the expense of greater sensitivity to array calibration errors that may cause the signal response from the desired height to be cancelled [Lombardini *et al.* 2009, Van Trees 2002 ch. 6.6]; variants are available which selectively trade-off resolution for robustness [Lorenz & Boyd 2005, Gershman *et al.* 2006]. The weights are adaptive, as they depend on the inverse of the local sample covariance matrix; given a finite number of independent looks contributing to the computed matrix, the output power estimate will be biased downwards, although the bias goes to zero as the number of looks tends to infinity [Capon & Goodman 1970, Pardini *et al.* 2014]. Note that in the SAR literature this approach is often called the Capon method, after the original author [Capon 1969, Lombardini *et al.* 2003].

Alternatively, the parametric Multiple Signal Classification (MUSIC) method partitions the eigendecomposition of the sample covariance matrix into (orthogonal) signal and noise subspaces according to a ‘known’ (i.e. previously estimated) number of scattering sources, and then projects all possible steering vectors onto the noise subspace to find the heights (but not the actual reflectivities) that give the smallest projection magnitude and therefore most likely correspond to the scattering sources [Van Trees 2002 ch. 9.3.2.1]. The resolution for MUSIC is even finer than that for MVDR, if the noise subspace is accurately estimated. However, MUSIC is arguably not well-suited to the case of SAR imaging of a forest volume, because there is no clear way to determine the number of notional scattering sources, and the heterogeneous canopy will layover differently pass-to-pass (due to the different incidence angles), so the clutter interference will manifest as non-stationary noise, causing subspace leakage i.e. an eigenvalue spectrum that is not clearly separated into signal and noise components [Guerci 2002]. Admittedly, this differential layover can lessen the ability of MVDR to suppress interference too.

Nonetheless, these and other beamforming methods have been widely used for 3D SAR imaging of forests. Lombardini & Reigber [2003] compared conventional beamforming and MVDR using L-band data from 14 passes, demonstrating a dramatic improvement in resolution and sidelobe performance in the latter case. Reigber *et al.* [2005] compared vertical structure estimation by (i) model-free 3D SAR (both MVDR and MUSIC) using 14 passes and (ii) model inversion via PolInSAR using two (fully polarimetric) passes, with the two sets of estimates of ground and canopy heights found to be in good agreement, although the PolInSAR results had a coarser

horizontal resolution due to greater averaging and they were very sensitive to the angular separation between the two passes. Frey & Meier [2011a,b] characterised forest scattering mechanisms as a function of height, frequency and polarisation by beamforming various polarimetric combinations of L-band (12 passes) and P-band (17 passes) data. Comparing conventional beamforming, a robust form of MVDR, and MUSIC, the latter two methods gave broadly similar results. Albinet *et al.* [2012] and Minh *et al.* [2013, 2014] used a P-band vertical antenna array fixed to a 55 m tower in French Guiana to repeatedly form range-height images using conventional beamforming (since fine resolution was not required) and assess how the coherence of the tropical forest varied over hours, days and months. Ferro-Famil *et al.* [2012] proposed a way to estimate 3D polarimetric covariance matrices using either conventional or MVDR beamforming, and demonstrated the application of traditional polarimetric decompositions, such as the alpha-entropy-span decomposition [Cloude & Pottier 1997], to show how the decomposition parameters varied with height in a forest. Nannini *et al.* [2012] and Huang *et al.* [2012] compared several high-resolution beamforming techniques applied to polarimetric L-band data (21 passes) in order to try to detect large vehicles concealed in a German forest. In a great feat of data processing, Ponce *et al.* [2016] coherently combined 19 fully polarimetric circular collections (i.e. $19 \times 360^\circ$ apertures) at L-band. A variety of 3D SAR processing techniques were illustrated and qualitatively compared by Reigber *et al.* [2015].

In order to quantify forest structure, Tebaldini [2009, 2010] formulated a general framework by which polarimetric multipass data can be decomposed into a set of independent scattering mechanisms (as opposed to scattering objects) characterised by their power and their polarimetric and interferometric covariance matrices. The decomposition relies on inverting a physical or mathematical model, such as the RVOG model, together with covariance matrix matching, and therefore can be understood as a multipass extension of PolInSAR [Cloude & Papathanassiou 2003], rather than an instance of 3D SAR beamforming. However, the decomposition is typically analysed by using MVDR to generate vertical intensity profiles for the different scattering mechanisms. Analysis of six passes over a boreal forest in northern Sweden showed that a simple dual-layer model of a forest as a volume above a ground surface is reasonably accurate at both L-band and P-band, noting that at P-band a significant portion of the scattering response apparent at ground level arises due to ground-volume interactions (even for the cross-polarised response) [Tebaldini & Rocca 2012]. Using this decomposition and conventional (not MVDR) beamforming, d’Alessandro & Tebaldini [2012, 2013] studied the phenomenology of ground scattering in a tropical forest at P-band given six passes by forming single-look complex polarimetric images at the nominal ground level and using the relative phase of the co-polarised channels to distinguish between direct and double-bounce scattering mechanisms. Pardini & Papathanassiou [2017] refined the decomposition technique for the specific case of separating surface and volume layers, where the vertical structure of the surface is given by a single steering vector to an assumed or estimated height. This method, together with MVDR beamforming (and significant multi-looking), was used recently to conduct an extensive analysis of scattering from tropical forests in Gabon, with average ground-volume power ratios for HH found to be -15 dB at L-band and -5 dB at P-band; the ratio could be increased by up to 3 dB at both frequencies by using the optimisation techniques of PolInSAR to select a different polarisation [Pardini *et al.* 2018].

These examples indicate the different motivations which have shaped how 3D SAR has been used in the past, and therefore indicate usage modes which have not yet been explored. In many cases, the intent is to determine the forest vertical structure, notably the ground and canopy heights, by scanning beams in height and plotting the vertical variation in beamformer output power; this is analogous to traditional direction-of-arrival estimation, where resolution is of foremost importance, and the beamformer output need not represent a physical signal (as in MUSIC). In rare cases, the complex polarimetric responses from different layers are estimated in order to fully characterise the scattering mechanisms. For target detection, the canopy is interference to be nulled. Across the literature for 3D SAR, considering both urban and natural landscapes, no examples could be found where the speckle pattern obtained as the complex output from a high-resolution beamformer, such as MVDR, was used. Indeed, Reigber *et al.* [2015] mistakenly states that MVDR generates power estimates only, whereas MVDR can actually be used to estimate the complex signal. Finally,

there is little interest in change detection, except in terms of characterising temporal decorrelation of the forest itself. To the author’s knowledge, there are no published examples of coherent change detection being applied to 3D SAR images.

The implication thus far has been that multiple SAR images for 3D processing can be acquired only by making multiple passes over the scene, entailing a very large penalty in coverage efficiency. However, single-pass acquisition is possible. Held & O’Brien [1992] described a gimbaled Ku-band system consisting of a single transmitter and three receivers which, although intended for along-track interferometry, could be rotated 90° to give three across-track channels, thus supporting “limited adaptive beamforming” to separate targets at different heights. Lombardini *et al.* [2004] applied conventional beamforming, MVDR and MUSIC to three-channel data acquired by an airborne X-band system consisting of one active element and two receive-only elements in a nonuniformly spaced, across-track formation; they demonstrated that the scattering responses from a bridge and a valley underneath could be separated vertically, with their measured height difference (~100m) approximately matching ground truth. Most relevant to this study, Zhang *et al.* [2012] used just three across-track channels to produce extremely coarse vertical profiles of a forest via MVDR, and correctly estimated the ground and canopy heights via MUSIC. The data was collected in a single pass by a dual-antenna L-band system operating in a ping-pong mode to generate fully polarimetric data at three across-track phase centres [Schwäbisch *et al.* 2008]—this was the inspiration for the system proposed in Figure 3.1. [Duque *et al.* 2015] demonstrated vertical beamforming using the three channels generated by the TanDEM-X (X-band) pair of SAR satellites operating in ‘alternating bistatic’ mode [Krieger *et al.* 2010]. Most unusually, Schmitt & Zhu [2016] produced vertical profiles of a forest using a Ka-band system with four across-track receivers.

It can be seen that 3D SAR beamforming, applied to multichannel image data from an across-track array operating with a foliage-penetrating wavelength, offers a way to attenuate the canopy interference, permitting the complex ground reflectivity to be estimated in a single pass. This is what is needed for coherent change detection.

3.1.5 Change detection

Different types of change are revealed by different types of change metrics [Rignot & van Zyl 1993, Oriot 2014].

Incoherent ratio metrics indicate change in the backscatter power, and have found wide application to the target detection problem, particularly at low frequencies (UHF and below) for hidden targets like vehicles [Halversen *et al.* 1994, Ulander *et al.* 2005, Ulander & Lundberg 2008], and at high frequencies (X-band and above) for very small targets such as landmines [Ranney & Khatri 2008]. Note that simple subtraction of the image intensities is statistically inappropriate for SAR imagery [Rignot & van Zyl 1993].

Coherent metrics, notably $|\gamma|$ in (2.1), indicate change (i.e. decorrelation) in the complex speckle pattern, and therefore are sensitive to subtle ground disturbances where the scattering elements have been rearranged but the mean power is not necessarily affected; at high frequencies (X-band and above), vehicle tracks on dirt roads stand out clearly and foot tracks may be faintly visible, depending on the softness of the ground [Jakowatz *et al.* 1996 ch. 5.5, Tsunoda *et al.* 2000]. Note that change associated with target detection is typically blob-like and therefore amenable to automatic detection, whereas anomalous ground disturbances often (but not always) mark out extended curves which may be better interpreted via an image that provides context. Coherent metrics suffer from false alarms due to decorrelation in areas of low SNR (e.g. shadow regions and very smooth surfaces with low backscatter) since the noise is uncorrelated across images; these can be mitigated either by comparing against a reference change image acquired previously which contains the false alarms only [Barber & Kogon 2012], or by explicitly accounting for noise in the change metric [Wahl *et al.* 2016].

Preiss *et al.* [2006] proposed a single likelihood-ratio metric which captures changes in power or correlation or both; Newey *et al.* [2012] extended this to account for noise.

The above incoherent, coherent and joint metrics have been extended for polarimetric data [Novak 2005a,b, Preiss & Stacy 2008, Barber 2015].

As a related but distinct category, polarimetric change metrics detect and classify wide-area change in the polarimetric signature of a landscape, due to such causes as

crop growth, logging, flooding or oil slicks [Conradsen *et al.* 2003, Erten *et al.* 2012, Marino *et al.* 2013, Migliaccio *et al.* 2015].

Most of these change metrics follow from likelihood ratios or statistical distance measures derived in formal hypothesis-testing frameworks assuming underlying Gaussian statistics. By contrast, Sabry [2009] proposed a basis set of change metrics, analogous to the Stokes parameters for electromagnetic waves, which covers different types of change.

To the author’s knowledge, none of the change metrics has been applied to the problem of detecting changes in the speckle pattern of the ground under a forest canopy.

3.1.6 Problem Statement

The goal in this chapter is to detect and form images of subtle changes on the ground under a forest canopy. To reiterate, the scene of interest is the foliage-obscured ground usually associated with hidden-target detection, but the changes of interest are the subtle ground disturbances (such as vehicle tracks) usually associated with open, sparsely vegetated landscapes.

The emphasis on subtlety motivates the use of a coherent change metric which is sensitive to decorrelation (a loss of coherence) of the speckle pattern in SAR images. The standard coherence metric $|\gamma|$ in (2.1) is used here. In typical SAR terminology, this is repeat-pass coherent change detection (CCD), the output of which is a change map i.e. an image of the coherence resolved in range and azimuth.

In order to see under the canopy, a wavelength which at least partially penetrates the canopy is required. However, the wavelength must also be sensitive to changes on the ground via its backscatter response. Considering both the literature on foliage penetration and an analysis of coherent sensitivity to change in Section 3.2.1, this chapter will focus on radar operation at L-band (23 cm wavelength).

To attenuate interference caused by layover of the canopy response onto the ground, 3D SAR beamforming applied to a set of 2D SAR images (acquired at slightly different incidence angles) is investigated. Only phase-preserving methods are considered, since the coherence metric will be applied to the beamformer output. For practical airborne operation, the set of SAR images needs to be acquired in a single pass, which severely limits the number of images. Attention is therefore focused on the potential of applying MVDR (a.k.a. Capon) beamforming to just three images.

The signal of interest is the direct ground backscatter, but a strong component of the forest scattering response may be ground-trunk double bounce. This scattering mechanism is constant across passes, so it gives rise to high coherence in the area around each tree, masking potential changes on the ground. Its interferometric phase centre is located at ground height, so vertical beamforming cannot remove it. Given fully polarimetric SAR data, an adaptive polarimetric filter that rejects ground-trunk double bounce is constructed via an eigendecomposition of the polarimetric covariance matrix.

Although the individual signal-processing steps (the ‘ingredients’) are well-established in the literature, the proposed processing chain (the ‘recipe’) is entirely novel. Coherent change detection applied to the ground under a forest canopy has not been attempted before. Coherent change metrics have not previously been applied to the output from a general 3D SAR beamforming process. The complex speckle pattern generated as the output of an MVDR (i.e. Capon) beamformer in particular has not been utilised at all in the 3D SAR literature. Using a polarimetric decomposition to remove a scattering mechanism prior to coherent change detection has not previously been proposed. These novel applications will be studied in this chapter.

3.2 Radar design considerations

The feasibility of measuring change under a forest canopy using radar is not immediately apparent. Before establishing a theoretical framework for this problem, it is useful to bring to the forefront the main difficulties that must be overcome. These difficulties are now considered from the point of view of the radar designer.

3.2.1 Wavelength selection

The selection of the operating wavelength involves several important trade-offs, as listed in Table 3.2. For the most part, these are standard factors to be weighed-up by the radar system designer [Cutrona 1990]. However, sensitivity to change is uniquely important to the application of coherent change detection, and has seen less quantitative study, so it is worthy of further investigation.

Table 3.2: Factors determining wavelength selection.

Backscatter intensity	Shorter wavelengths provide a stronger backscatter response from natural surfaces (grass, gravel, etc), because the surface appears rougher [Ulaby & Long 2014 ch. 10].
Foliage penetration	Longer wavelengths better penetrate the canopy (see (3.1)).
Change sensitivity	Shorter wavelengths are more sensitive to scene changes.
Range resolution	Shorter wavelengths (i.e. higher centre frequencies) more easily support a wider bandwidth and therefore finer range resolution (see (2.60)). As the fractional bandwidth (ratio of bandwidth to centre frequency) increases, the antenna may be more difficult to construct and calibrate.
Azimuth resolution	Shorter wavelengths more easily support a finer azimuth resolution (see (2.61)). As the wavelength increases, the integration angle to achieve a particular resolution increases too, and this requires either a wider antenna beamwidth (i.e. less antenna gain) or a steering capability (for spotlight-mode collection). Furthermore, at a fixed range, a wider angle implies a longer aperture (in space and time) over which the airborne platform should remain stable and the motion measurements remain accurate, although this is partly alleviated by the fact that longer wavelengths are less sensitive to perturbations in platform motion.
Antenna size	A shorter wavelength permits the use of a smaller antenna, which would be easier to mount on the airborne platform.
Spectrum	Broadly speaking, higher-frequency portions of the RF spectrum are less congested, so it easier to obtain approval for wide-band transmission, and less RF interference will be received [Vickers 2002, Griffiths <i>et al.</i> 2015, Davis & Reis 2015].

Intuitively, it may be expected that shorter wavelengths would be more sensitive to scene change. To investigate this, a SAR clutter simulation was used to quantify the sensitivity of the coherence metric to scene change as a function of both wavelength and the level of change. Multiple sets of raw radar pulse data were synthesised for a flat homogeneous clutter scene consisting of many randomly distributed point scatterers ($\sim 10/\text{m}^2$). For each dataset, the scatterers were randomly (and independently) shifted relative to their original position, giving rise to a different speckle pattern. The shifts were uniformly distributed in angle and normally distributed in distance with zero mean and a particular standard deviation for each dataset. A SAR image was formed from each dataset and its coherence with the original (zero-shift) image computed via (2.1) i.e. the standard coherence estimator. The whole exercise was repeated at different wavelengths; to permit comparison, the bandwidth (140 MHz) was fixed to give a constant ground-range resolution (1.7 m at 35° graze with a Hamming window), and the aperture lengths were varied to give a constant azimuth resolution (1 m).

Fig. 3.3 shows CCD images at different wavelengths for the same random shifts of the scattering elements. Each pixel in each CCD is the local coherence magnitude between a pair of SAR images formed before and after the shifts, evaluated using a rectangular averaging window that gave 35 effective looks. Clearly, shorter wavelengths are more sensitive to change in the speckle pattern, as expected. As the wavelength becomes shorter, the change in propagation range of each shifted element of the clutter causes a larger change in the phase of its scattering response, leading to a greater change in the coherent superposition of the responses from all elements in each (fixed-size) resolution cell, thus generating a speckle pattern that differs from the original pattern more dramatically, which is measured as a greater decorrelation.





Figure 3.3: CCDs at different wavelengths after the elements of a clutter scene were randomly shifted with a standard deviation of 6 cm. Black corresponds to zero coherence (the SAR image pixels before and after the change are uncorrelated) and white corresponds to unit coherence (the pixels are perfectly correlated).

Note that in each simulated repeat-pass scenario, the change is statistically homogeneous across the scene, but in practice, the coherence estimation window may cover areas that have changed to differing extents, including areas that have not changed at all. Averaging over the window thus degrades the resolution of the CCD and makes discrimination between changed and unchanged areas more difficult.

Fig. 3.4 shows how the coherence at different wavelengths varies with the level of change, specifically, the standard deviation of the random shifts. Each point on each curve is the coherence computed over the whole scene using 1032 effective looks. Complete decorrelation (say, coherence 0.15 for this simulation) occurs when the

Figure 3.4: Coherence magnitude at different wavelengths as a function of the level of change in a clutter scene. Each curve was obtained by repeated execution of a SAR simulation where change was implemented as random shifts in the positions of point scatterers. The resolution was fixed.



standard deviation of the shift is roughly half of the wavelength. Note that the exact shape of these curves holds only for this simulation; in a real imaging scenario, a particular shift of a scattering object may cause a coherence loss different to that indicated by Figure 3.4. Nonetheless, the curves indicate the relative performance of the different wavelengths when scene changes can be modelled as a random rearrangement of point-like scattering elements. Their shape broadly agrees with the plot by Zebker & Villasenor [1992 Fig. 4] for a different simulation.

Sensitivity to change is what is sought in this CCD application, so a steeper fall-off of coherence with shift magnitude is better. However, there is clearly a trade-off between change sensitivity, which drives a shorter wavelength, and foliage penetration, which drives a longer wavelength. Weighing-up all the factors listed in Table 3.2, together with the literature review in Section 3.1.1, it is tentatively proposed that L-band (~ 23 cm wavelength) offers the best compromise, primarily because it is the shortest wavelength that still penetrates the canopy sufficiently to permit measurement of the ground echo. L-band is also a convenient window for spectrum approval; it can be difficult to obtain approval for wideband transmission at lower frequencies [Pincus *et al.* 2013].

3.2.2 Effective dynamic range of coherence

Coherence, as defined in (2.1), has a magnitude that ranges between zero (uncorrelated) and one (perfectly correlated). However, it is usually not possible to measure coherence over its full range due to practical limitations at the two extremes.

At the low extreme (coherence near zero), the sample coherence estimator in (2.1), which involves a spatial average, is biased such that low coherences are not observable, as discussed in Appendix 2.A.1. There is a trade-off between averaging more, to lessen the bias and extend the observable range of coherence, and averaging less, to

avoid degrading the original resolution of the radar image. For the purpose of change detection, the true coherence magnitude need not be known exactly, whereas resolution is critical, so it is reasonable to lean towards less averaging. A convenient compromise is to average over nine resolution cells, which, from Figure 2.28, sets the coherence floor at 0.3.

At the high extreme (coherence near unity), the coherence is limited by noise. A standard result in SAR interferometry is that coherence can be expressed as a product of decorrelation factors, each due to a different source [Zebker & Villasenor 1992]. For the moment, it is sufficient to decompose the observed coherence as

$$\gamma_{ab} = \gamma'_{ab} \gamma_{SNR} \quad (3.2)$$

where

$$\gamma_{SNR} = (1 + SNR^{-1})^{-1} \quad (0 \leq \gamma_{SNR} \leq 1) \quad [\text{Zebker \& Villasenor 1992}], \quad (3.3)$$

captures the decorrelation due to system noise and γ'_{ab} is the net coherence due to all other sources. In (3.3), noise is encapsulated in the signal-to-noise ratio (SNR), which depends on the radar system, the range to the scene, and the scene itself as observed at the angular geometry of the radar. It can be seen that the maximum observable coherence is $\gamma_{max} = \gamma_{SNR} < 1$, obtained when the actual coherence (neglecting noise) is unity i.e. $\gamma'_{ab} = 1$. An extended decomposition of coherence is derived in Section 3.6; see (3.76) and (3.77) in particular.

The scene modulates the received power, and therefore the SNR, via its apparent ‘brightness’ i.e. the extent to which it scatters the incident energy in the direction of the receiver, which obviously depends on the structural and material features of the scene. For discrete targets (e.g. a vehicle), this brightness is typically quantified by the radar cross-section σ (unit: m^2), which is the area of an equivalent isotropic scatterer i.e. a hypothetical object that would scatter the same amount of energy in every direction as the actual target did in the observed direction [Long 2001 ch. 2.1]. For spatially distributed clutter (e.g. the ground), the brightness is instead quantified by the normalised radar cross-section σ^0 , which is the average radar cross-section per unit area [Ulaby & Long 2014 ch. 5.5]; it is implicit that the scattering elements giving rise to the clutter be uniformly distributed over the area contributing to the average, so the underlying scattering process is spatially stationary, even if measurements fluctuate sample to sample [Long 2001 ch. 6.1.1]. Since σ^0 is dimensionless, it is often called the scattering coefficient, and in the case of a monostatic radar, the backscattering coefficient. Importantly, the received power, and therefore the SNR, are directly proportional to the (normalised) radar cross-section [Long 2001 ch. 2.1.1].

The sensitivity of an imaging radar system can be characterised by its noise-equivalent sigma-nought σ_{NE}^0 , which is the value of σ^0 that yields unity SNR [Ulaby & Long 2014 ch. 14.6.2]. Since SNR depends on range, σ_{NE}^0 does as well; when considering radar system performance in general, it is necessary to assess the radar hardware together with the planned operating geometry. Lower values of σ_{NE}^0 are better, in that less bright scenes can still be imaged with reasonable fidelity, or the same scenes can be imaged at longer ranges.

Typically, a maximum (worst-case) σ_{NE}^0 is determined for a particular application, and the radar designer attempts to achieve this. Say that the application requires that a scene with brightness σ_{OP}^0 be imaged with signal-to-noise ratio SNR_{OP} , where the subscript OP indicates that this is the chosen operating point. Because σ_{NE}^0 is just a value of σ^0 (for which $SNR = 1$), and SNR and σ^0 are directly proportional, one can immediately write-down:

$$\frac{SNR_{OP}}{1} = \frac{\sigma_{OP}^0}{\sigma_{NE}^0} \Rightarrow \sigma_{NE}^0 = \frac{1}{SNR_{OP}} \sigma_{OP}^0. \quad (3.4)$$

In this way, the radar performance can be specified without delving into the full radar range equation, which contains numerous hardware-specific parameters that together determine SNR [Cutrona 1990].

The aim here is to determine a value for σ_{NE}^0 in (3.4) that permits adequate estimation of the coherence of the ground under a forest canopy. Somewhat arbitrarily, let $\gamma_{max} = 0.8$ i.e. say that ‘no change’ is adequately represented by a coherence measurement of 0.8. From (3.3), this sets the SNR operating point at $SNR_{OP} = 6$ dB.

For the scattering operating point σ_{OP}^0 , it is necessary to turn to the literature. Table 3.3 lists backscattering coefficients at L-band for different types of land cover, averaged over many historical measurements made at different sites using different radar systems. The grazing angle is not the same for all scene types just due to limited data availability. Two important points stand out. Firstly, for a given type of land cover, the variability across measurements, as indicated by the standard deviation, is usually greater than the shift in the mean that results from changing polarisation or changing grazing angle by several degrees (the latter aspect is clear from plots of the full dataset [Ulaby & Dobson 1989]). Thus, the precise nature of the scene is (not surprisingly) the most important factor that determines the backscattering coefficient; even within each category in Table 3.3, there is great scope for heterogeneous scattering responses. Secondly, tree-covered land is approximately 10 dB brighter than all other types of terrain, and exhibits approximately half the variability. As a rough average of averages, excluding the tree category (since it is only the ground response that is sought), say that the scattering coefficient at L-band of the ground in a forest is nominally -22 dB.

Table 3.3: Statistics of measured backscattering coefficients σ^0 for different types of land cover at L-band, collated from numerous sources and summarised by Ulaby & Dobson [1989]. In the table, angle refers to grazing angle and N refers to the number of σ^0 data points contributing to the collated statistics (not the number of samples contributing to each reported data point). The mean and standard deviation (std) are each given for co-polarisations HH and VV as $\sigma_{HH}^0/\sigma_{VV}^0$.

type (ch. 5.2)	angle	N	mean (dB)	std (dB)	page
soil & rock surfaces	40°	207/ 65	-20.7/-19.5	4.6/4.9	120
grasses (rangeland, pasture & small crops)	30°	105/106	-24.0/-25.1	5.8/6.1	168
short vegetation (grass, shrubs & wetlands)	30°	347/347	-23.9/-21.6	4.9/5.5	232
shrubs (bushy plants & large crops)	30°	242/241	-23.8/-20.1	4.5/4.4	200
trees (forests & orchards)	35°	235/235	-10.7/-11.3	2.2/1.8	148

However, in the application considered here, the ground is not observed directly, but through the canopy. Propagation through the canopy will be lossy, causing a reduction in SNR for the ground echo. One way to account for this is to use an effective scattering coefficient that combines the effects of ground scattering and canopy propagation. The propagation loss can be predicted using the random-volume model of a forest canopy presented in Section 3.9.1: setting a volume height of 20 m, an attenuation rate of 0.1 dB/m, and a grazing angle of 30°, the two-way attenuation given by $f_v(0)$ in (3.112) is -8 dB. Thus, the effective backscattering coefficient at L-band for the ground under a forest canopy is nominally $\sigma_{OP}^0 = -22 - 8 = -30$ dB—this is the scattering operating point.

Hence, from (3.4), a suitable radar system must achieve

$$\sigma_{NE}^0 = \sigma_{OP}^0 / SNR_{OP} = -30 - 6 = -36 \text{ dB}. \quad (3.5)$$

This is a demanding requirement, but achievable by airborne systems: the interim version of the *Ingara* L-band radar (which used a single helical antenna for each polarisation) managed -31 dB [Pincus *et al.* 2015b]; DLR reports that the multi-frequency F-SAR system achieves -42 dB for L-band at 30° grazing [Reigber *et al.* 2012]; Intermap aims for -40 dB, although to achieve this they have to fly low (1 km altitude) and make do with a small swath (1 km) [Mercer *et al.* 2009]. Spaceborne systems, by contrast, typically operate with much worse noise performance: at L-band, Japan’s PALSAR-2 manages -24 dB [Kankaku *et al.* 2013], and at C-band, the European Space Agency’s Sentinel-1 [ESA Sentinel-1 2017 ch. 4.2.8,6.2.8] and Canada’s RADARSAT-2 [MDA 2016 Table 2-2] achieve roughly -25 dB in high-resolution modes.

Figure 3.5 illustrates how the noise decorrelation γ_{SNR} varies with SNR according to (3.3). In addition, from (3.4), the radar performance metric σ_{NE}^0 is related to SNR by $\sigma_{NE}^0 = \sigma_{OP}^0 - SNR$ (all quantities in decibels). At the operating point, $\gamma_{max} = 0.8$, $SNR_{OP} = 6$ dB, $\sigma_{OP}^0 = -30$ dB and $\sigma_{NE}^0 = -36$ dB.

In summary, the dynamic range of the observable coherence is limited at the low end by estimation bias and at the high end by noise. Figure 3.6 illustrates how the observed coherence varies with the true coherence in two different estimation




Figure 3.5: Noise decorrelation γ_{SNR} as a function of SNR. At the chosen operating points, $\gamma_{max} = 0.8$, so $SNR_{OP} = 6$ dB, and $\sigma_{OP}^0 = -30$ dB, so $\sigma_{NE}^0 = -36$ dB—this is the noise performance metric that the radar must satisfy in order to measure the expected ground echo with sufficient SNR such that the noise decorrelation is no more than 0.8.

scenarios. The blue curve corresponds to the operating points selected previously i.e. averaging over nine samples to set the coherence floor at 0.3 and achieving an SNR of 6 dB to set the coherence ceiling at 0.8; this provides a minimum level of estimation performance, where the dynamic range is reduced by half compared to the ideal. This plot was obtained using the same random number simulation that was used in Appendix 2.A.1 to plot the bias of the coherence estimator in Figure 2.28, except that here uncorrelated complex Gaussian noise was added to the correlated complex Gaussian random variables before their coherence was estimated via (2.1).




Figure 3.6: Observed coherence as a function of true coherence in different estimation scenarios. The blue curve shows a minimal-performance scenario, where the coherence is estimated using 9 independent samples and the SNR is 6 dB, whereas the cyan curve shows a high-performance scenario, where the coherence is estimated using 35 independent samples and the SNR is 9.5 dB. Assuming that the scattering coefficient σ^0 of the scene is -30 dB, an SNR of 6 dB (9.5 dB) corresponds to $\sigma_{NE}^0 = -36$ dB (-39.5 dB). The black dashed line shows the ideal where the estimated coherence always matches the truth. Note that the sample size N_{avg} and the SNR are not related; the chosen values just show typical cases.

The next section moves on from phenomenology and radar design to establish the signal-processing basis for interferometry and 3D SAR.

3.3 Layover and interferometry

This section addresses the basic but important question of how a pixel in a SAR image represents the scattering response of an object. An understanding of the pixel phase, in particular, will serve as the foundation for the development of 3D SAR processing techniques, which rely on coherent summation of corresponding pixels from multiple SAR images. A critical result derived in this section is an expression for the interferometric wavenumber k_{z0} , which relates the observed phase difference ϕ between corresponding pixels in two aperture-trimmed SAR images to the height z of a scattering object above the focal surface.

One of the key features of the time-domain beamforming method of SAR image formation, as described in Section 2.10, is that it does motion compensation to a point, at every pixel position. Conceptually, for each pixel position, the scene is assumed to be empty except for a single point target at the pixel position that reflects the incident pulses with some fixed magnitude fraction and phase delay. The received pulses are compensated for all known modulations and coherently summed in order to obtain estimates of the magnitude and phase, which are recorded as the complex pixel value.

The most obvious and most important effect to compensate for is propagation delay. In particular, for every pulse-pixel pair (i.e. a pairing of (i) a platform position at the time of a pulse and (ii) a pixel position), the delay due to the propagation range r is compensated separately and explicitly, in order to remove the propagation phase offset $2\pi(2r/\lambda)$ (i.e. the two-way propagation distance counted in units of the radar operating wavelength λ). This phase varies pulse-to-pulse and therefore would prevent coherent addition across pulses if not compensated.

Consider a point-like scattering object in the scene at a distance d from the platform. Recall that it will be represented in the image according to some point-spread function due to the finite spatial-frequency support provided by the radar waveform and the synthetic aperture, as discussed in Sections 2.4, 2.9 and 2.13. More specifically, after beamforming, the object will give rise to a scattering response at many pixels in the image, scaled in magnitude by the point-spread function and offset in phase by $2\pi \cdot 2(d-r)/\lambda$ due to the residual propagation distance $d-r$ after compensation for the range r to each pixel. (See Figure 2.3 for a diagram and Figures 2.1–2.2 for coordinate reference frames.) The peak response will appear at the pixel whose compensated range is closest to the actual propagation distance for the object.

The description above is applicable regardless of whether the scattering object lies on the focal surface, where the pixels are evaluated, or is vertically offset from the surface by some amount Δz . In the latter case, however, the scattering response is subject to a horizontal shift, called layover. The range component of layover arises due to the lack of vertical resolution along the (curved) wavefront incident on the object at the local grazing angle, leading to the projection of the scattering response along an isorange contour onto the focal surface, as depicted in Figure 3.1. An azimuthal component of layover arises if the collection geometry is squinted (see Figures 2.1(c) and 2.1(e)): the pulse-to-pulse propagation compensation for the range-only layover position will leave a residual phase on the scattering response that varies (approximately) linearly along the aperture in proportion to the size of the squint, thus acting as a phase ramp that causes an along-track shift when the pulses are summed. For the textbook case of plane waves and a flat focal plane, the layover can be quantified by simple expressions that capture the linear projections from the slanted and squinted acquisition geometry to the upright image geometry [Jakowatz *et al.* 1996 ch. 2.4.5, 3.7.5, Appendix C]. In particular, the shift Δy due to the range component of layover is

$$\Delta y = \Delta z \tan \psi \quad [\text{Jakowatz } et al. 1996 \text{ p. 91}]. \quad (3.6)$$

For the more general case of spherical wavefronts and an undulating focal surface, it is sufficient to say that an object will layover onto the focal surface at a position whose set of propagation distances for all pulses in the synthetic aperture most closely matches the set of propagation distances to the object's true position. This was verified by first simulating SAR images for various imaging geometries and focal surfaces where the scene consisted of a single elevated point target, and then, for every pixel, computing the sum of the squared residual propagation distances $d-r$ for all pulses in the aperture; in each SAR image, the peak response of the laid-over target occurred at the pixel whose associated sum was smallest.

Now consider two observations, a and b , of the scene, made from broadside positions A and B . The distances to the scattering objects and the pixels will be different across the observations, so the residual propagation phases will be different, and the layover of objects above or below the focal surface will be different as well. Figure 3.7 illustrates the differential layover of an elevated object, together with the differential range to a nearby pixel. For a scattering object at distances d_a and d_b , and a pixel position at ranges r_a and r_b , the net phase ϕ is just the difference between the residual propagation phases at the corresponding pixels in the two SAR images:

$$\phi = 2\pi \frac{2(d_b - r_b)}{\lambda} - 2\pi \frac{2(d_a - r_a)}{\lambda} = \frac{4\pi}{\lambda} (d_b - d_a) - \frac{4\pi}{\lambda} (r_b - r_a). \quad (3.7)$$

The phase in (3.7) is the basis for SAR interferometry, by which surface topography can be measured. Interferometric processing typically involves making two observations of a scene at slightly different grazing angles, forming two SAR images with mutually trimmed spatial-frequency supports, registering them spatially, computing their complex coherence via (2.1), unwrapping the phase of the coherence, and relating this




Figure 3.7: Differential layover for two broadside collection geometries shown in the range-height plane, with platform positions A and B , a scattering object at P , and a focal surface (here a horizontal plane) where pixels are evaluated at the crosses. The long-dashed lines indicate the propagation distances d_a and d_b from A to P (blue) and B to P (red), respectively; the solid coloured lines indicate the curved wavefronts incident on the object at P ; the short-dashed lines indicate the ranges r_a and r_b to the pixel at O . The elevated object will layover onto the focal plane, appearing at P'_A (P'_B) in the SAR image acquired at A (B), such that $d_a = \overline{AP} = \overline{AP'_A}$ ($d_b = \overline{BP} = \overline{BP'_B}$). Although the precise layover position may be between pixel positions, the radar point-spread function, which is sketched above for the case of no aperture window, ensures that nearby pixels will contain the scattering response. During image formation, the pixel at O is compensated for the range $r_a = \overline{AO}$ ($r_b = \overline{BO}$), so the uncompensated propagation distance $d_a - r_a = \overline{AP} - \overline{AO}$ ($d_b - r_b = \overline{BP} - \overline{BO}$) gives rise to a phase offset. Thus, the interferometric phase difference between the pixel at O in image A and the corresponding pixel in image B is proportional to the differential propagation distance $d_b - r_b - (d_a - r_a) = d_b - d_a - (r_b - r_a) = \overline{BP} - \overline{AP} - (\overline{BO} - \overline{AO})$. In this diagram, the layover is drawn to scale, but the platform separation is unrealistically large in order to clearly depict the differential layover $P'_B - P'_A$.

phase—termed the interferometric phase—to the vertical offset of the scattering surface above or below the focal surface [Shapiro *et al.* 1972, Graham 1974, Bamler & Hartl 1998, Rosen *et al.* 2000]. An additional step of ‘flat-earth’ phase removal is required for SAR images formed using frequency domain methods, but in the case of time-domain beamforming, the act of motion compensation to every pixel position conveniently removes the range variation over the focal surface, as shown by the $r_b - r_a$ negative phase term in (3.7). In general, the interferometric phase consists of two components: a geometric part determined by the difference in the propagation distances, as given by (3.7), and a noise part arising from any differences in the scattering process itself [Just & Bamler 1994]. For simplicity, in what follows the noise part is ignored and the interferometric phase is attributed solely to propagation delay. Typically, only the spatial variation of the terrain height can be estimated reliably, because absolute height measurement requires both precise calibration for any cross-channel distortion and determination of the absolute phase, which are difficult [Madsen & Zebker 1992, Rosen *et al.* 2000].

To see that the phase in (3.7) can be exploited in this way, first note that the residual propagation distance $d - r$ can be separated into Cartesian components and linearised via the plane-wave approximation, as shown in Section 2.3, leading to the expression for the range offset given in (2.8) and repeated here:

$$d - r \approx x \sin \theta \cos \psi + y \cos \theta \cos \psi + z \sin \psi \quad (3.8)$$

The propagation difference now depends on the 3D position (x, y, z) of the object relative to the pixel position, and the angular geometry (θ, ψ) of the acquisition relative to the pixel position. Note that if the plane-wave approximation is reapplied at every pixel position by recomputing the angular geometry (effectively setting each pixel position in turn as the reference origin), then the error caused by the approximation will be very small. As before, scaling the difference in (3.8) by $4\pi/\lambda$ will generate the (approximate) residual propagation phase. Given two observations, acquired at (θ_a, ψ_a) and (θ_b, ψ_b) using potentially different wavelengths λ_a and λ_b , the propagation

phase difference ϕ in (3.7) becomes

$$\phi = \frac{4\pi}{\lambda_b}(x \sin \theta_b \cos \psi_b + y \cos \theta_b \cos \psi_b + z \sin \psi_b) - \frac{4\pi}{\lambda_a}(x \sin \theta_a \cos \psi_a + y \cos \theta_a \cos \psi_a + z \sin \psi_a). \quad (3.9)$$

The immediate goal is to make ϕ insensitive to the horizontal components of the scattering object's position, which is unknown, so that ϕ serves as a measure of vertical offset z only. In particular, when the object is on the focal surface ($z=0$), the phase difference should be zero. This can be achieved by selecting the two wavelengths such that the effect of the angular difference at $z=0$ is cancelled. If the azimuthal angles are equal, then the constraint on the wavelengths is simple:

$$\left. \begin{array}{l} \phi = 0 \\ z = 0 \\ \theta_a = \theta_b \end{array} \right\} \Rightarrow \frac{\lambda_a}{\lambda_b} = \frac{\cos \psi_a}{\cos \psi_b}. \quad (3.10)$$

Applying this solution to the case where the object is not on the focal surface, the observed phase difference would be

$$\begin{aligned} \phi &= \frac{4\pi}{\lambda_a} \left(\begin{array}{l} x \sin \theta_a \cos \psi_a + y \cos \theta_a \cos \psi_a + z \cos \psi_a \tan \psi_b \\ -x \sin \theta_a \cos \psi_a - y \cos \theta_a \cos \psi_a - z \sin \psi_a \end{array} \right) \quad (\theta_a = \theta_b \text{ and } \lambda_b = \frac{\cos \psi_b}{\cos \psi_a} \lambda_a) \\ &= \frac{4\pi}{\lambda_a} \cos \psi_a (\tan \psi_b - \tan \psi_a) z. \end{aligned} \quad (3.11)$$

Thus, the observed phase would depend on the vertical position component only.

This idea can be extended using a more abstract formulation where the radar acquisition variables of angle and wavelength are combined into the concept of spatial-frequency support, as described in Section 2.4. This support is then expressed geometrically in wavenumber space or k -space using the spatial-frequency variables defined in (2.10), which are repeated here for a general wavelength λ and general acquisition angles (θ, ψ) :

$$k_r = \frac{4\pi}{\lambda}, \quad k_x = k_r \sin \theta \cos \psi, \quad k_y = k_r \cos \theta \cos \psi, \quad k_z = k_r \sin \psi = k_y \frac{\tan \psi}{\cos \theta}. \quad (3.12)$$

One point of spatial-frequency support (i.e. one wavelength at one angular geometry) determines the value of propagation phase generated by a scattering object, referenced to a particular pixel position. Given two observations with point supports $(k_{x_a}, k_{y_a}, k_{z_a})$ and $(k_{x_b}, k_{y_b}, k_{z_b})$, the phase difference ϕ in (3.9) can be expressed using (3.12) as

$$\phi = k_{x_b} x + k_{y_b} y + k_{z_b} z - (k_{x_a} x + k_{y_a} y + k_{z_a} z). \quad (3.13)$$

In order to ensure that the same value of propagation phase is generated for all possible object positions $(x, y, 0)$ on the focal surface, the supports must be at the same horizontal position in k -space, as shown below:

$$\left. \begin{array}{l} \phi = 0 \\ z = 0 \end{array} \right\} \Rightarrow \begin{array}{l} k_{x_a} = k_{x_b} \\ k_{y_a} = k_{y_b} \end{array}. \quad (3.14)$$

In synthetic aperture radar, each 'observation' of the scene (in the sense of an image pixel) is, of course, the coherent summation of many wideband pulses acquired at different azimuthal angles. Therefore, the spatial-frequency support is extended from a point to a ribbon surface, as depicted in Figure 2.14. For the current purpose, this means that one frequency component acquired at one angular view in a pulse contributing to observation a , can be matched against a different frequency component acquired at a different angular view in a pulse contributing to observation b , such that the generated phase contributions are the same. Extending the logic from the point case in (3.14), the constraint on the two synthetic apertures to ensure that this phase matching is always satisfied is that their spatial-frequency supports must cover the same region in the horizontal k_x - k_y plane. Recall the definition in (2.57) and (2.58) of the indicator function $I = f(k_x, k_y)$ that specifies, for one pixel position, the coverage provided by the synthetic aperture when projected onto the spatial-frequency ground

plane $k_z = 0$, and recall too the definition in (2.59) of the set $J = \{I\}$ of indicator functions for all pixel positions in an image. The requirement that must be met in order to ensure equal propagation phases for all possible object positions on the focal surface can be stated succinctly as

$$J_a = J_b. \quad (3.15)$$

That is, the spatial-frequency supports for corresponding pixels must be the same. This is precisely the requirement identified in Chapter 2 to achieve maximum coherence magnitude; in that case, the non-overlapping portions of the spatial-frequency supports were undesirable because they degraded the magnitude of the coherence, but here, the problem is that the non-overlapping supports distort the phase of the coherence by generating different propagation phases for scattering from the focal surface. The requirement in (3.15) is met by filtering the supports such that they match [Gatelli *et al.* 1994, Cloude 2010 ch. 5.1.1.1]; in the general case of an undulating focal surface, the spatially variant aperture trim presented in Section 2.16 can be applied.

For each pixel position, the centre (k_{x0}, k_{y0}) of the trimmed region of spatial-frequency support can be associated with equivalent radar acquisition parameters as follows:

$$\theta_0 = \arctan \frac{k_{x0}}{k_{y0}}, \quad \psi_0 = \frac{\psi_a + \psi_b}{2}, \quad k_{r0} = \frac{k_{y0}}{\cos \theta_0 \cos \psi_0}, \quad \lambda_0 = \frac{4\pi}{k_{r0}}. \quad (3.16)$$

Note that the grazing angles are not actually constrained by the trim, which is applied in the spatial-frequency ground plane— ψ_0 above is just a nominal value—so the vertical wavenumbers remains unequal for the two observations. Using (3.12) and (3.16),

$$\begin{aligned} k_{z_a} &= (k_{y0} / \cos \theta_0) \tan \psi_a = k_{r0} \cos \psi_0 \tan \psi_a \\ k_{z_b} &= (k_{y0} / \cos \theta_0) \tan \psi_b = k_{r0} \cos \psi_0 \tan \psi_b. \end{aligned} \quad (3.17)$$

Note too that λ_0 is not quite equal to the actual radar operating wavelength; given the definitions of ψ_0 and k_{r0} in (3.16), the trimmed support effectively corresponds to a frequency band that is both cropped and slightly down-shifted relative to the transmitted band.

Putting all this together, if the spatial-frequency supports for corresponding pixels in two SAR images are trimmed to their common region, as specified in (3.15), then the propagation phase difference ϕ would depend on the vertical position component only, as can be seen by combining (3.13), (3.14) and (3.17):

$$\begin{aligned} \phi &= k_{x_a} x + k_{y_a} y + k_{z_b} z \\ &\quad - k_{x_a} x - k_{y_a} y - k_{z_a} z \quad (\text{aperture trim} \Rightarrow k_{x_a} = k_{x_b} \text{ and } k_{y_a} = k_{y_b}) \\ &= (k_{z_b} - k_{z_a}) z \\ &= k_{r0} \cos \psi_0 (\tan \psi_b - \tan \psi_a) z. \end{aligned} \quad (3.18)$$

Let $\Delta\psi = \psi_b - \psi_a$. If $|\Delta\psi| \ll 1$, then

$$\begin{aligned} \tan \psi_b - \tan \psi_a &= \tan(\psi_0 + \Delta\psi/2) - \tan(\psi_0 - \Delta\psi/2) \\ &= \frac{\tan \psi_0 + \tan \Delta\psi/2}{1 - \tan \psi_0 \tan \Delta\psi/2} - \frac{\tan \psi_0 - \tan \Delta\psi/2}{1 + \tan \psi_0 \tan \Delta\psi/2} \\ &= \frac{2(1 + \tan^2 \psi_0) \tan \Delta\psi/2}{1 - \tan^2 \psi_0 \tan^2 \Delta\psi/2} \\ &= \frac{2 \tan \Delta\psi/2}{(1 - \tan^2 \psi_0 \tan^2 \Delta\psi/2) \cos^2 \psi_0} \\ &\approx \frac{\Delta\psi}{\cos^2 \psi_0} \end{aligned}$$

where the error in the approximation is less than 1% for $\psi_0 \leq 65^\circ$ and $|\Delta\psi| \leq 5^\circ$. Hence,

$$\phi = k_{z0} z \quad (3.19)$$

where

$$k_{z0} = k_{r0} \cos \psi_0 (\tan \psi_b - \tan \psi_a) \approx k_{r0} \frac{\Delta\psi}{\cos \psi_0} \quad (3.20)$$

is the interferometric wavenumber, which linearly transforms height (above the focal surface) into interferometric phase. As a further approximation, k_{r0} is typically computed using the original centre wavelength, ignoring the small frequency shift induced by the trim [Jakowatz *et al.* 1996 p. 318].

Inverting (3.19), it can now be seen that the height z of the true ground surface above the focal surface (at $z = 0$) can be estimated by measuring, unwrapping and scaling the interferometric phase. Before unwrapping, the 2π phase wrap limits the unambiguous height extent to

$$h_{\text{amb}} = 2\pi/k_{z0}. \quad (3.21)$$

An implicit assumption of SAR interferometry is that the scene acts as a surface, with scattering elements at a single height in each 2D resolution cell. If instead the scattering elements were at different heights, collectively acting as a volume, then the height measurement would be ambiguous. The aperture trim removes the differential propagation phase at one height only; the elements at all other heights would lay over differently in the two observations, as depicted in Figure 3.7, and contribute different propagation phases, which cannot be separated. The volume could instead be characterised by an effective phase centre at height

$$h_{v,\text{eff}} = \phi_v/k_{z0} \quad (3.22)$$

where ϕ_v is the interferometric phase measured for the volume. Even better would be to image the volume by resolving its structure vertically, as well as in range and azimuth. This is the goal of 3D SAR.

Before moving on, it is important to note that practical implementation of interferometric techniques requires that the constituent SAR images be first spatially registered so that corresponding pixels represent the same location in the scene [Jakowatz *et al.* 1996 ch. 5.3]. The initial images may be misregistered for several reasons; three causes stand out for airborne repeat-pass acquisition. Firstly, the antenna phase centre may not have been tracked accurately (relative to the operating wavelength) over one or both of the apertures, leading to uncompensated propagation [Prats *et al.* 2004]. Secondly, the system reference oscillator may suffer from drift, causing unknown variations in the pulse timing across passes. These two sources of phase error may give rise to range and azimuth shifts during image formation. Thirdly, even if the system hardware performs perfectly, differential layover will occur in areas where there is both a vertical offset between the focal surface and the scattering object and a difference in the angular geometry of the acquisitions (see Figure 3.7) [Jakowatz *et al.* 1996 ch. 2.4.5, 3.7.5, Appendix C] [Prats *et al.* 2005]. Misregistration as small as one-eighth of a resolution cell will significantly raise the variance of the interferometric phase, thus degrading the coherent relationship between the SAR images [Just & Bamler 1994]. Hence, spatial registration is almost always required before coherent processing of fine-resolution airborne repeat-pass SAR images.

3.4 3D SAR beamforming

Consider a set of conventional SAR images of a scene resolved in azimuth and range. Scattering elements above or below the focal surface will layover into each 2D image and contribute residual propagation phase components, as discussed in the previous section; both the layover and the phase are functions of both the height of each scatterer and the grazing angle of each acquisition. If the images were acquired at slightly different grazing angles, then collectively they will be sensitive to the true heights of the scattering elements via the differential layover and differential propagation phase. It should be possible to combine the images in a way that preserves the responses from one height and suppresses the responses from other heights; such 3D processing is the subject of this section.

The images will be combined as a weighted sum by applying a complex weight vector \mathbf{w} to a set \mathbf{x} of complex pixels, one per image, all corresponding to the same pixel position. This process was depicted in Figure 3.2. The complex output is

$$y = \mathbf{w}^H \mathbf{x} \quad (3.23)$$

with expected power

$$\sigma_y^2 = E\{yy^*\} = E\{\mathbf{w}^H \mathbf{x}(\mathbf{w}^H \mathbf{x})^*\} = \mathbf{w}^H E\{\mathbf{x}\mathbf{x}^H\} \mathbf{w} = \mathbf{w}^H R \mathbf{w} \quad (3.24)$$

where R is the covariance matrix. The weight vector \mathbf{w} is treated as deterministic, although in practice it may be selected in a way that depends on the radar data. Different weight vectors may be used at different pixel positions. (Note that in this context, y is used to indicate the output signal, not a position.)

Crucially, the coherent gain of the weighted sum varies with height, by virtue of the ‘memory’ of the scattering height embedded in the pixel phase. After applying a particular weight vector to a set of pixels that contain the scattering responses and propagation phases from objects at different heights, the contributions from one height may add in-phase to give a relatively large value, and the contributions from a different height may add out-of-phase to give a relatively small value. This height selectivity can be exploited by appropriately chosen weights, such that the summed output is a 3D SAR image, coarsely resolved to a particular height.

3D SAR processing is akin to beamforming [Lombardini *et al.* 2003]. At each pixel position, the set of input pixels can be considered to be the vector output from a sensor array spaced in grazing angle, and the task of weighting and summing is effectively beam-steering along the vertical dimension. The field of beamforming provides a framework in which weight vectors can be selected and compared.

3.4.1 The steering vector

An important concept in beamforming is the steering vector \mathbf{v} . Given a hypothetical signal source at a particular angle or location relative to the sensor array, the steering vector contains the phasors required to compensate for the differing propagation delays to the sensors. When applied to the array data, the coherent gain of the sum is then maximised for signals from this source.

For the textbook case of angle-only beamforming, where a uniform linear array of narrowband sensors is used to sample plane waves arriving from a particular direction, the phases in the steering vector depend directly on the effective array spacing along the steered direction, counted in units of the operating wavelength [Van Trees 2002 ch. 2.3–2.5]. In the broadside direction, for which the (planar) wavefronts are parallel to the array, there are no propagation delays between sensors, so the steering vector would be just a vector of ones.

For 3D SAR beamforming in height, the ‘sensor outputs’, i.e. the input data to the beamformer, are the pixels in N_{ch} 2D SAR images at a particular range-azimuth position on a focal surface. The height of the focal surface itself at $z = 0$ is the SAR-equivalent of the ‘broadside direction’, and the steering vector for this height is a vector of ones i.e.

$$\mathbf{v}_0 = \mathbf{1}_{N_{ch} \times 1}. \quad (3.25)$$

This is a consequence of the fact that the image formation process described in Section 2.10 does motion compensation to every pixel position, whereby the range-dependent propagation phase for every pulse-pixel pair is explicitly cancelled. This point was explained in detail at the start of Section 3.3. A scattering object located at a pixel position does not contribute any residual propagation phase to that pixel value. Since this applies for all images, the spatial phase difference between the image channels is zero, and a steering vector of ones effectively steers the input pixels to the focal surface at $z = 0$.

In the SAR case, steering to different heights requires more than just changing the steering vector: the images must also be corrected for the differing layover at the different heights, because the horizontal layover shift depends on the vertical separation between the steered height and the focal surface. Moreover, the change in layover will be different for images acquired at different angular geometries, because layover also depends on the precise grazing and squint angles that the radar phase centre makes with the pixel position at the steered height. Not only should the change in layover be accounted for, but also the change in the residual propagation phase on the pixels. If equations can be derived that characterise all of the effects of height-dependent scattering, then corrections can be applied for each height of interest by first resampling the images to accommodate the predicted horizontal shift—this shift

is unlikely to be a whole number of pixels—and then applying an associated phase modulation [Jakowatz & Wahl 2001].

A different approach is taken in this work. In order to steer to different heights, the whole image formation process is repeated for all input datasets, with the height of the focal surface adjusted each time. This is more suited to the case where the focal surface is not flat, so the layover would be much more difficult to predict. In addition, this avoids any errors due to resampling or due to approximations made in deriving explicit expressions for the layover projection, such as assuming plane waves or neglecting azimuthal shifts induced by squint. Of course, this increases the computation time, but in the context of change detection over a typical interval of hours, a processing delay of minutes is unimportant. Hence, at the 3D processing stage, the steered height is always the height of the focal surface at $z = 0$, and the steering vector is always a vector of ones, as given in (3.25). The equivalent concept in the traditional angle-only beamforming case would be to redefine the input to the beamformer as the array data after compensation for the propagation differences, so the steering vector used by the beamformer itself would always be just a vector of ones, and ‘steering’ to different angles would involve supplying different processed versions of the sensor outputs. This might seem to be just a matter of semantics in the standard beamforming case, but in the 3D SAR case it makes sense because precisely this type of processing is undertaken for image formation.

3.4.2 Types of weight vectors

Beamforming applications may be roughly divided into three categories: direction-of-arrival estimation, signal detection and signal estimation. In the context of 3D SAR, the first category corresponds to estimating the heights of scattering objects that are known (or assumed) to be in the scene, with a prime example being the dual measurement of ground topography and canopy height for forest scenes [Reigber *et al.* 2005, Frey & Meier 2011b]. The second category corresponds to detecting the presence of anomalous targets, notably vehicles hidden under a forest canopy [Nannini *et al.* 2012, Huang *et al.* 2012]. The third category involves estimating either the intensity or the full complex reflectivity at multiple heights of interest, such as for the generation of vertical forest profiles that roughly characterise the structure and density of the vegetation [Cloude 2006, Tebaldini & Rocca 2012, Ferro-Famil *et al.* 2012]. In the first two categories, the need for fine resolution despite the small number of image channels has driven the investigation of a variety of exotic beamforming schemes, many of which maximise resolution at the expense of signal fidelity, as was discussed in Section 3.1.4.

The application here of coherent change detection on the forest floor demands that the complex reflectivity of the ground be estimated with reasonable accuracy. Thus, the beamforming task fits in the third category of signal estimation, and the processing must preserve phase. This estimation problem has been addressed in the literature only rarely: one study used conventional beamforming of six airborne passes to characterise ground scattering in a tropical forest at P-band [d’Alessandro & Tebaldini 2012, d’Alessandro *et al.* 2013]. In this work, the adaptive MVDR (a.k.a. Capon) beamformer is the main technique employed. The conventional and null-steer beamformers are also used to characterise the beamforming problem and provide a reference point for comparison. Later, a clairvoyant beamformer matched to the RVOG model will be derived; it is equivalent to the MVDR beamformer if the sample covariance matrix equals the matrix of coherences predicted by the RVOG model.

The conventional beamformer \mathbf{w}_{conv} does nothing more than apply the steering vector to compensate for the differential propagation phases (if any) between channels. This maximises the coherent gain at the steered height. Given the steering vector \mathbf{v}_0 in 3.25,

$$\mathbf{w}_{\text{conv}} = \frac{1}{N_{ch}} \mathbf{v}_0. \quad (3.26)$$

Coherent integration of SAR images using the conventional beamformer is analogous to coherent integration of radar pulses, discussed in Section 2.13; both are examples of aperture synthesis. Along-track resolution depends essentially on the azimuthal angle $\Delta\theta_{sl}$ subtended in the slant-plane; vertical resolution depends essentially on the grazing angle interval $\psi_{N_{ch}} - \psi_1$ covered by the set of N_{ch} acquisitions in the direction perpendicular to the nominal slant-plane (see (3.39) and (3.42)). Although 3D SAR

processing is discussed here in beamforming terms, the uncompressed multichannel data effectively forms a wedge in the spatial-frequency domain, as depicted in Figure 2.4, and this is what supports 3D resolution, as discussed in Section 2.4. The conventional beamformer will be used to provide a performance baseline for comparison with the other beamformers.

The adaptive MVDR beamformer \mathbf{w}_{mvdr} is the solution to the optimisation problem of minimising the output power from all heights, subject to the constraint that the complex signal from the steered height be passed undistorted [Van Trees 2002 ch. 6.2, 6.3, 6.6]. In the array-processing literature, it is often called the Minimum-Variance Distortionless-Response (MVDR) beamformer, and in the SAR literature, it is often called the Capon method, after the original author [Capon 1969, Lombardini *et al.* 2003]. For data covariance matrix \hat{R} , the data-dependent weight vector is

$$\mathbf{w}_{\text{mvdr}}(\mathbf{x}, \mathbf{y}) = \frac{\hat{R}^{-1}(\mathbf{x}, \mathbf{y})\mathbf{v}_0}{\mathbf{v}_0^H \hat{R}^{-1}(\mathbf{x}, \mathbf{y})\mathbf{v}_0}, \quad (3.27)$$

where the dependence on the pixel position (\mathbf{x}, \mathbf{y}) (counted in metres) has been made explicit. Recall from Section 2.2 that $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is the global image reference frame, shown in Figure 2.1(b), where \mathbf{x} corresponds to azimuth (subscript *az*) and \mathbf{y} corresponds to ground-range (subscript *gr*).

The sample covariance matrix \hat{R} serves as an estimate for the covariance matrix R of the underlying random process that gave rise to the observations. In the case of multichannel SAR imagery of natural landscapes, the random process is electromagnetic wave scattering from distributed clutter, and each observed realisation \mathbf{x}_m is a vector of N_{ch} complex reflectivities manifest as pixel values at a particular pixel position (here, pixel position is generically indexed by m). Due to the physical phenomenon of speckle, each vector is distributed as an N_{ch} -variate complex zero-mean Gaussian random variable, and each resolution cell is an independent realisation of the scattering process [Lee *et al.* 1994, Oliver & Quegan 1998 ch. 4]. The maximum-likelihood estimator in this case is implemented as a spatial average of N_{avg} unit-rank outer products [Goodman 1963]:

$$\hat{R}(\mathbf{x}, \mathbf{y}) = \frac{1}{N_{\text{avg}}} \sum_{m=1}^{N_{\text{avg}}} \mathbf{x}_m \mathbf{x}_m^H. \quad (3.28)$$

As with all maximum-likelihood estimators, it is asymptotically unbiased and asymptotically attains the Cramér-Rao lower bound on the variance. However, for small N_{avg} , its statistical performance is not necessarily good [Kay 1993 ch. 7.5]. It has also been characterised as “excessively sensitive” to outliers [Huber & Ronchetti 2009 p. 200].

Combining (3.24) and (3.27), the expected output power using the adaptive MVDR beamformer is

$$\sigma_{y, \text{mvdr}}^2 = \mathbf{w}_{\text{mvdr}}^H \hat{R} \mathbf{w}_{\text{mvdr}} = \frac{\mathbf{v}_0^H \hat{R}^{-H}}{\mathbf{v}_0^H \hat{R}^{-H} \mathbf{v}_0} \hat{R} \frac{\hat{R}^{-1} \mathbf{v}_0}{\mathbf{v}_0^H \hat{R}^{-1} \mathbf{v}_0} = \frac{\mathbf{v}_0^H \hat{R}^{-H} \mathbf{v}_0}{\mathbf{v}_0^H \hat{R}^{-H} \mathbf{v}_0 \cdot \mathbf{v}_0^H \hat{R}^{-1} \mathbf{v}_0} = \frac{1}{\mathbf{v}_0^H \hat{R}^{-1} \mathbf{v}_0}. \quad (3.29)$$

For a finite number of independent looks contributing to the sample covariance matrix \hat{R} , $\sigma_{y, \text{mvdr}}^2$ will be biased downwards compared to the power that would have been obtained using the underlying covariance matrix R ; the bias goes to zero as the number of looks tends to infinity [Capon & Goodman 1970, Pardini *et al.* 2014]. Moreover, for a fixed number of looks, $\sigma_{y, \text{mvdr}}^2$ will suffer greater variance than the power $\sigma_{y, \text{conv}}^2 = \mathbf{w}_{\text{conv}}^H \hat{R} \mathbf{w}_{\text{conv}}$ obtained using the conventional weights [Capon & Goodman 1970, Van Veen 1991]. Qualitatively, there is a trade-off between adaptively suppressing interference, which is good for resolution, and minimising statistical variability, which is good for detection performance [d’Assumpcao & Gray 2007].

Selecting the pixel positions to average over is not trivial, because the observations \mathbf{x}_m should be realisations of the same random scattering process. Neighbouring pixels may be assumed to be realisations of the same process, but of course, this does not hold at the boundaries between different types of landscapes (urban, grassland, forest, roads, etc), or if there are isolated specular targets (e.g. a vehicle hidden in a forest). Elaborate averaging schemes have been devised recently that emphasise statistical similarity over locality [Deledalle *et al.* 2015]. Nonetheless, in this work, (3.28) is computed simply

over a rectangular window of $N_{\text{avg}} = N_{\text{avg},az} \times N_{\text{avg},gr}$ pixels centred at (x, y) . Given pixel spacings d_{az} and d_{gr} and resolutions ρ_{az} and ρ_{gr} , the window should cover enough resolution cells to ensure that the covariance matrix \hat{R} has full rank, as indicated below:

$$\frac{N_{\text{avg},az} d_{az} \times N_{\text{avg},gr} d_{gr}}{\rho_{az} \times \rho_{gr}} \geq N_{ch}. \quad (3.30)$$

Note that in this implementation, the estimate $\hat{R}(x, y)$ in (3.28) is recomputed at every pixel position, with each window centred on each position, so neighbouring pixels have slightly different weight vectors. Also note that the spatial averaging needed to obtain $\hat{R}(x, y)$ and therefore $\mathbf{w}_{\text{mvdr}}(x, y)$ does not cause any loss of horizontal resolution if the underlying scattering process is stationary, because \mathbf{w}_{mvdr} is still separately applied at every pixel position as $y(x, y) = \mathbf{w}_{\text{mvdr}}^H(x, y) \mathbf{x}(x, y)$; the output value y is not an average of the inputs pixels in the local neighbourhood around (x, y) .

The effect of the adaptive beamformer is to steer up to $N_{ch} - 1$ nulls onto interferences that would otherwise leak through sidelobes or even the (outer) mainlobe of the beampattern of the conventional beamformer. The nulls are automatically steered to the heights of the strongest interferences via the inverse covariance matrix. The fine resolution consequently achieved by the adaptive beamformer makes it the main focus of this chapter—the operating envelope for the proposed 3D SAR CCD application depends on the predicted performance of this beamformer.

This improved performance compared to the conventional beamformer comes at the expense of greater sensitivity to errors. Most seriously, the signal response from the desired height may be mistaken for interference and cancelled [Lombardini *et al.* 2009]. These errors can arise from many different sources, both physical and statistical, that either degrade the accuracy of the covariance matrix estimate \hat{R} or invalidate the assumption that the array geometry is properly captured by the steering vector \mathbf{v}_0 [Gershman *et al.* 2006]:

- The clutter scene, as observed in the averaging window, may be non-stationary or may contain outlier scattering responses that degrade \hat{R} .
- The number of looks in the averaging window may be inadequate (a finite sampling limitation), leading to an imprecise and biased estimate \hat{R} .
- Knowledge of the terrain height may be poor, causing the SAR images to be focused onto a surface that does not match the true ground surface; \mathbf{v}_0 in (3.25) steers to the height of the focal surface, so the ground scattering response may be treated as interference and nulled.
- Knowledge of the relative geometry of the radar collections may be poor, leading to uncompensated phase differences between the SAR image channels that are not captured by using \mathbf{v}_0 in (3.25) as the steering vector.
- The radar system may be miscalibrated such that the RF signal paths for the different channels have different amplitude gains or phase delays that again are not captured by using \mathbf{v}_0 in (3.25) as the steering vector.

Variants of the MVDR beamformer are available which selectively trade-off resolution for robustness [Lorenz & Boyd 2005, Gershman *et al.* 2006]. These algorithms are not considered in this chapter, as the focus is on presenting the principles and operating envelope of 3D SAR CCD, rather than preparing a particular processing scheme for real-world implementation.

A useful property of the MVDR beamformer in (3.27) is that if it is steered onto a unit rank signal component of the covariance matrix, then it is invariant to the power of this component, and depends only on the other components (the interference). To see this, decompose the matrix as $R = R_1 + R_2$ where $R_1 = \sigma_1^2 \mathbf{v}_1 \mathbf{v}_1^H$ (for any complex vector \mathbf{v}_1). The inverse of R is given by the Sherman-Morrison-Woodbury formula [Horn & Johnson 2013 ch. 0.7.4] as follows:

$$\begin{aligned} R^{-1} &= (R_2 + \sigma_1^2 \mathbf{v}_1 \mathbf{v}_1^H)^{-1} \\ &= R_2^{-1} - \frac{\sigma_1^2 R_2^{-1} \mathbf{v}_1 \mathbf{v}_1^H R_2^{-1}}{1 + \sigma_1^2 \mathbf{v}_1^H R_2^{-1} \mathbf{v}_1} \\ &= \frac{1}{1 + \sigma_1^2 \mathbf{v}_1^H R_2^{-1} \mathbf{v}_1} (R_2^{-1} + \sigma_1^2 \mathbf{v}_1^H R_2^{-1} \mathbf{v}_1 R_2^{-1} - \sigma_1^2 R_2^{-1} \mathbf{v}_1 \mathbf{v}_1^H R_2^{-1}) \end{aligned} \quad (3.31)$$

Note that $\mathbf{v}_1^H R_2^{-1} \mathbf{v}_1$ is a scalar. Given (3.31), the MVDR weight vector when steered in the direction corresponding to \mathbf{v}_1 is

$$\frac{R^{-1} \mathbf{v}_1}{\mathbf{v}_1^H R^{-1} \mathbf{v}_1} = \frac{R_2^{-1} \mathbf{v}_1 + \sigma_1^2 \mathbf{v}_1^H R_2^{-1} \mathbf{v}_1 R_2^{-1} \mathbf{v}_1 - \sigma_1^2 R_2^{-1} \mathbf{v}_1 \mathbf{v}_1^H R_2^{-1} \mathbf{v}_1}{\mathbf{v}_1^H R_2^{-1} \mathbf{v}_1 + \sigma_1^2 \mathbf{v}_1^H R_2^{-1} \mathbf{v}_1 \mathbf{v}_1^H R_2^{-1} \mathbf{v}_1 - \sigma_1^2 \mathbf{v}_1^H R_2^{-1} \mathbf{v}_1 \mathbf{v}_1^H R_2^{-1} \mathbf{v}_1} = \frac{R_2^{-1} \mathbf{v}_1}{\mathbf{v}_1^H R_2^{-1} \mathbf{v}_1}. \quad (3.32)$$

This result is applicable to the multichannel dual-layer model developed in Section 3.7, where the beamformer is steered vertically towards the unit-rank ground component, with the aim of suppressing the above-ground volume interference.

To reiterate, the MVDR beamformer is optimal in the sense that it maximises the signal-to-interference ratio (SIR), where the signal is distinguished from the interference by the steering vector. If the signal and interference conform to Gaussian probability distributions, then under certain conditions, the same beamformer may also be optimal in a maximum-likelihood (ML) sense; the fact that different optimisation criteria lead to the same result is well-known [Van Trees 2002 ch. 6.2, Capon *et al.* 1967]. Recent papers [Joerg *et al.* 2017, Pardini & Papathanassiou 2017] concerned with separating surface and volume layers have highlighted two observations by Marzetta [1983]: firstly, ML estimation of the signal power is ill-posed, because there is no unique combination of signal power and interference covariance matrix that maximises the likelihood function, and secondly, ML estimation of the complex signal amplitude is not ill-posed, and the MVDR estimate converges asymptotically to the unique ML estimate as the number of looks increases i.e. $\hat{R} \rightarrow R$. For the application here, it is the complex amplitude of the ground response, not just its power, which we seek to estimate for each pass, because the complex signal (i.e. the complex speckle pattern) determines the coherence between passes. At any rate, in this chapter, no probability distributions are assumed; the MVDR beamformer is proposed because it maximises the SIR.

The null-steer beamformer \mathbf{w}_{null} is a solution to the design problem of finding a weight vector that most closely exhibits a desired beampattern [d'Assumpcao & Mountford 1984]. The beampattern is specified as a set of control points i.e. a set of desired responses \mathbf{a}_c at particular heights \mathbf{z}_c . Define the matrix V_c whose columns are the steering vectors to the controlled heights. The design problem can be formulated as finding a solution to the system of linear equations

$$V_c^H \mathbf{w} = \mathbf{a}_c. \quad (3.33)$$

The solution that minimises the Euclidean norm $\mathbf{w}^H \mathbf{w}$ is

$$\mathbf{w}_{\text{null}} = \text{pinv}(V_c^H) \mathbf{a}_c \quad (3.34)$$

where

$$\text{pinv}(V) = V^H (V V^H)^{-1} \quad (3.35)$$

is the Moore-Penrose pseudo-inverse [d'Assumpcao & Gray 2007 p. 57]. If the interference field is unknown, then the adaptive beamformer in (3.27) is ideal. However, if information about the interference field is available, then the null-steer beamformer in (3.34) offers a way to make use of this knowledge, without relying on covariance matrix estimation, which may be problematic if the data is highly variable or if the amount of data to average over is limited. The null-steer beamformer provides a way to predict the performance of the adaptive MVDR beamformer, without having to generate a covariance matrix.

3.4.3 Vertical beampattern

A weight vector is characterised by its vertical beampattern, which specifies how it attenuates scattering from different heights. The beampattern is conceptually equivalent to the point-spread function, which was derived in Section 2.A.3 for a uniform frequency support, but in the 3D SAR case, where there are usually only a few nonuniformly spaced samples (i.e. the channels at different grazing angles), the beampattern formulation arguably provides a more intuitive understanding that does not depend on the Fourier transform. The beampattern directly computes how the weight vector, which is steered to the focal surface at $z=0$ via the steering vector \mathbf{v}_0

in (3.25), interacts with the interferometric phases that would result from scattering from heights $z \neq 0$. Thus, the beampattern $b(z)$ for a weight vector \mathbf{w} is given by

$$b(z) = |\mathbf{w}^H \mathbf{v}(z)|^2 \quad (3.36)$$

where the i^{th} element of the propagation vector $\mathbf{v}(z) = [e^{j\phi_i(z)}]$ contains the interferometric phase $\phi_i(z) = k_{z_i} z$ (from (3.19)) that would be observed by the pairing of the first and i^{th} channels, acquired at grazing angles ψ_1 and ψ_i , with interferometric wavenumber (from (3.20))

$$k_{z_i} = k_{r0} \frac{\psi_i - \psi_1}{\cos(\psi_i + \psi_1)/2} \quad (i = 1, 2, \dots, N_{ch}). \quad (3.37)$$

Here, the first channel is arbitrarily chosen as the reference channel, with wavenumber $k_{z_1} = 0$. Setting a different channel as the reference channel would change the $\{k_{z_i}\}$ values, but not change the beampattern.

Observe from (3.37) that the channels will not be uniformly spaced in k_z , even if they are uniformly spaced in grazing angle, which is itself highly unlikely unless the channels are acquired from a single platform in a single pass, as depicted in Figure 3.1. In general, it is not possible to accurately summarise the beampattern in terms of only one or two metrics, such as resolution and unambiguous extent, which are well-defined only for uniformly sampled data; the beampattern should be considered in its entirety.

However, if the channels are approximately uniformly spaced in grazing angle, then an approximate characterisation of the beampattern of the conventional beamformer may be reasonable, thereby providing a baseline indication of the potential performance of the array. One approach is to use average acquisition parameters as follows:

$$\bar{\psi} = \frac{1}{N_{ch}} \sum_{i=1}^{N_{ch}} \psi_i, \quad (3.38)$$

$$\overline{\Delta\psi} = \frac{1}{N_{ch} - 1} \sum_{i=2}^{N_{ch}} \psi_i - \psi_{i-1} = \frac{\psi_{N_{ch}} - \psi_1}{N_{ch} - 1} \quad (\psi_1 < \psi_2 < \dots < \psi_{N_{ch}}) \quad (3.39)$$

$$k_{z_{\text{nom}}} = k_{r0} \frac{\overline{\Delta\psi}}{\cos \bar{\psi}} \quad (\text{from (3.20)}) \quad (3.40)$$

$$h_{\text{amb}} = \frac{2\pi}{k_{z_{\text{nom}}}} \quad (\text{from (3.21)}) \quad (3.41)$$

$$\rho_z = w_{\text{rect}} \frac{2\pi}{N_{ch} k_{z_{\text{nom}}}} = w_{\text{rect}} \frac{\lambda}{2N_{ch} \overline{\Delta\psi}} \cos \bar{\psi} \quad (\text{analogous to (2.61)}) \quad (3.42)$$

where computation of the average channel spacing in (3.39) requires the grazing angles to be sorted in increasing order. The beampattern

$$b_{\text{conv}}(z) = \frac{1}{N_{ch}^2} |\mathbf{v}_0^T \mathbf{v}(z)|^2 \quad (3.43)$$

of the conventional beamformer \mathbf{w}_{conv} in (3.26) will approximate the geometric series in (2.88), and therefore will take the sinc-like form expressed in (2.89) (and illustrated in Figure 2.29), with 3-dB resolution given by ρ_z in (3.42) and grating lobes (aliases) at multiples of the unambiguous height extent given by h_{amb} in (3.41). Note that the averaging approach in (3.38)–(3.40) works well if the (small) nonuniformity is well-modelled as random perturbations around a mean spacing, but if the nonuniformity is instead caused by one or two outliers, then the average quantities may be quite misleading; for example, h_{amb} (if meaningful) may be better computed using $\min(\Delta\psi)$.

Figure 3.8 shows example beampatterns of the conventional beamformer given nine channels that are either uniformly spaced in grazing angle (blue) or nonuniformly spaced (green) due to random perturbations such that the standard deviation of their separations is 13 per cent of the nominal spacing. It can be seen that the beampatterns follow the sinc-like shape of a standard point-spread function, but that the nonuniformity causes raised sidelobes. One way to mitigate the nonuniformity is to interpolate over the angular sector of interest [Lombardini & Pardini 2008], although the results are still not ideal. To reiterate, one of the advantages of the proposed

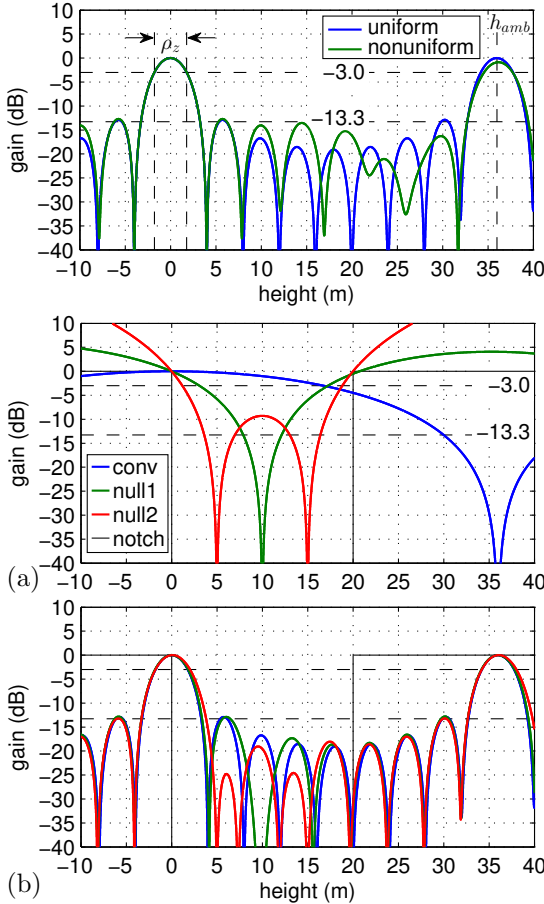


Figure 3.8: Vertical beam patterns for the conventional beamformer given uniformly spaced (blue) and nonuniformly spaced (green) channels, where the latter set of channels was obtained from the former by adding zero-mean random perturbations such that the standard deviation of the grazing angle separations $\{\Delta\psi_i\}$ was 0.02° . Parameters: $\lambda = 23$ cm, $N_{ch} = 9$, $\bar{\psi} = 35^\circ$, $\overline{\Delta\psi} = 0.15^\circ$, $h_{amb} = 36$ m, $\rho_z = 3.5$ m.

Figure 3.9: Vertical beam patterns for the conventional and two null-steer beamformers using uniformly spaced channels. For (a), $N_{ch} = 3$ and $\overline{\Delta\psi} = 0.05^\circ$, whereas for (b), $N_{ch} = 9$, and $\overline{\Delta\psi} = 0.15^\circ$ (in both cases, $\bar{\psi} = 35^\circ$ and $\lambda = 23$ cm). In (a), the main lobe of the conventional beamformer (blue) covers heights up to 36 m, with $\rho_z = 31.9$ m and $h_{amb} = 108$ m. The null-steer beamformers were constructed using control points $\{\mathbf{z}_c = [0, 10], \mathbf{a}_c = [1, 0]\}$ (green) and $\{\mathbf{z}_c = [0, 5, 15], \mathbf{a}_c = [1, 0, 0]\}$ (red).

single-platform single-pass collection mode is that it ensures that the grazing angle spacing will be approximately uniform.

However, having nine across-track channels on an aircraft is probably impractical; it would require five antennas (two active, three passive) spread out in a line from wing-tip to wing-tip, with the outer two alternating on transmit, and all antennas operating on receive, thereby generating two monostatic phase centres and seven distinct bistatic phase centres. A more realistic set-up would have just three, closely spaced, channels, as depicted in Figure 3.1. For example, the Intermap radar system operates in this way, using two antennas 3.5 m apart (attached to the ends of a rigid beam passing through the unpressurised part of the fuselage), so that when operating at their (deliberately low) altitude of ~ 1000 m (and a nominal grazing angle of $\sim 45^\circ$), the grazing angle interval subtended on the ground by the two antennas is 0.1° , and the grazing angle spacing between the three phase centres is 0.05° [Schwäbisch *et al.* 2008].

Using this more realistic three-channel set-up, Figure 3.9(a) compares the vertical beam patterns of the conventional beamformer (blue) and two null-steer beamformers, one with one null (green), the other with two nulls (red). Given such few, closely spaced, channels, the conventional resolution is terrible ($\rho_z = 31.9$ m)—the whole vertical extent of a typical forest would be inside the main lobe. However, the ability to steer nulls, notably inside the main lobe, suggests that it would be possible to attenuate the canopy response without disturbing the ground response. Since the intended CCD application requires imaging of the ground only, the ideal beam pattern would look like a notch filter (black), with the notch width matched to the height of the forest (say, 20 m tall), and the notch edge aligned to the vertical position of the ground; it can be seen that the two-null beam pattern (red) approaches this goal. The price to pay when steering nulls inside the main lobe is large gain outside the expected scene extent i.e. below 0 m and above 20 m, which means that the 3D processing will be highly sensitive to errors in knowledge of the ground position that cause incorrect placement of the notch. Nonetheless, the results obtained by Intermap using the adaptive MVDR beamformer, which is just an adaptive version of the null-steer beamformer, indicate that it is possible to scan

beams vertically using just three, closely-spaced, channels, and use this to coarsely distinguish between the ground and the canopy [Zhang *et al.* 2012, Huang *et al.* 2013].

Figure 3.9(b) shows the three beampatterns for the corresponding beamformers given nine channels. Although all three curves broadly follow the generic shape of a sinc, the sidelobes of the null-steer beamformers are moderately rearranged in order to place nulls at the specified locations, which are now outside the main lobe.

3.4.4 Height-range point-spread function

The discussion of the vertical beampattern thus far is a little misleading, because it ignores the coupling between height and range induced by the grazing angle $\bar{\psi}$, which is inherent to all side-looking imaging systems. A more complete understanding of the 3D performance requires consideration of the 2D point-spread function in the height-range plane. In what follows, azimuth processing is assumed but not explicitly discussed, because the along-track (i.e. azimuth) direction is mutually perpendicular to height and range, and is not affected by the grazing angle.

Resolution is obtained natively along the slant-range (subscript rg) and slant-height (subscript ns) directions, where slant-height is normal to the slant-plane (i.e. the nominal plane containing the propagation vectors between the scene centre and the pulse positions along the azimuthal aperture). For conventional processing that gives a sinc-like point-spread function, the 2D (y, z) coordinates of the first nulls along the slant-range and slant-height directions are

$$(\rho'_{rg} \cos \bar{\psi}, \rho'_{rg} \sin \bar{\psi}) \quad (\text{first null along slant-range}) \quad (3.44)$$

$$(\rho'_{ns} \sin \bar{\psi}, \rho'_{ns} \cos \bar{\psi}) \quad (\text{first null along slant-height}) \quad (3.45)$$

where the peak-to-null widths ρ' of the 2D sinc mainlobe are given by

$$\rho'_{rg} = c/2B_c \quad (\text{from (2.40)}) \quad (3.46)$$

$$\rho'_{ns} = \lambda/2N_{ch}\bar{\Delta\psi} \quad (\text{from (3.42)}). \quad (3.47)$$

Thus, the waveform bandwidth B_c and the total effective subtended angle $N_{ch}\bar{\Delta\psi} = \frac{N_{ch}}{N_{ch}-1}(\psi_{N_{ch}} - \psi_1)$ determine the size of the point-spread function, while the nominal grazing angle $\bar{\psi}$ determines its orientation. Note that the first range null intersects with the ground-range (y) axis at

$$(\rho'_{rg}/\cos \bar{\psi}, 0) \quad (3.48)$$

so the ground-range resolution ρ_{gr} is dilated by $1/\cos \bar{\psi}$ compared to ρ_{rg} , which accords with (2.60).

The vertical beampatterns plotted in Figures 3.8 and 3.9 in the previous section actually showed the point-spread functions in the slant-height direction, with the ‘height’ quantity on the horizontal axis corresponding to the z -coordinate here. Hence, the vertical resolution ρ_z in (3.42) corresponds to the z -coordinate of the first null along slant-height in (3.45) (leaving aside the window scale). Similarly, the unambiguous height extent h_{amb} in (3.41) corresponds to the z -coordinate of the grating lobe at

$$(h_{amb} \tan \bar{\psi}, h_{amb}). \quad (3.49)$$

To visualise the 2D point-spread function, height-range images of a single point target have been generated according to the method described in Algorithm 3.1. There is nothing special about this algorithm; the constituent steps are listed here just for clarity. The algorithm makes extensive use of the SAR simulation employed in Chapter 2, in order to synthesise the scattering responses of targets in a simulated scene and do all the processing to form SAR images.

Various radar configurations and platform geometries are considered here. In all cases, the aperture phase centres are positioned at the same height but uniformly spaced in the across-track dimension with some horizontal baseline, in order to achieve a particular grazing angle spacing ($\bar{\Delta\psi}$). Three platform (altitude, nominal grazing angle) geometries are considered: (500m, 30°), (1000m, 45°) and (1500m, 60°); three multichannel (horizontal, grazing angle) spacings are considered: (5.25m, 0.15°), (3.50m, 0.10°) and (1.75m, 0.05°); two array lengths are considered: three and nine

Algorithm 3.1 Height-range image formation

output data: A height-range image consisting of a stack of range-lines at one azimuth index, each steered to a different height.

input definitions:

- A scene, consisting of point targets in 3D space, each with some complex reflectivity.
- N_{ch} collection apertures, each at a particular geometry with respect to the scene centre: grazing angle, squint angle, altitude, aperture length, pulse-spacing, position and timing perturbations, etc.
- A set of steer heights. (Usually, each focal surface is just a horizontal plane at some height, which is the steer height.)

- 1: **for all** collection apertures **do** #Acquisition
- 2: Synthesise the pulse echoes by computing and summing the scattering responses of the elements of the scene.
- 3: **end for**
- 4: **for all** steer heights **do** #Processing
- 5: **for all** collection apertures **do**
- 6: Form a 2D SAR image.
- 7: **end for**
- 8: Form a 3D SAR image steered to this height by coherently combining the N_{ch} 2D SAR images using one of the beamforming methods described in Section 3.4.2.
- 9: Extract from the 3D SAR image a line along range at one azimuth index and insert it into the output height-range image.
- 10: **end for**

channels. Nine channels spaced at 5.25 m would require a platform with a 42 m wing span (a Boeing 747 has a 60 m wing span), while three channels spaced at 1.75 m would require a 3.5 m wing span. The combination of $(1000 \text{ m}, 45^\circ)$, $(1.75 \text{ m}, 0.05^\circ)$ and three channels matches the Intermap design [Mercer *et al.* 2009]. Note that a grazing angle spacing of 0.05° is very small. In all cases, the simulated waveform parameters are $f_c = 1.32 \text{ GHz}$ ($\lambda = 22.7 \text{ cm}$) and $B_c = 140 \text{ MHz}$, to match the *Ingara* L-band system.

The height-range images in Figures 3.10–3.12 show the point-spread function. It is tilted by the nominal grazing angle $\bar{\psi}$ because, when refocusing the constituent SAR images to different heights, the scatterer’s response lays over to different ground-range positions, in accordance with (3.6). For these configurations, it is clear that the slant-height resolution after conventional beamforming is much coarser than the slant-range resolution. Nonetheless, the characteristics of the point-spread function are well-predicted by (3.44)–(3.49), although the indicated position of the grating lobe (the magenta cross) is not exact due to the small deviations of the true grazing angles from the nominal uniformly spaced values. For $\Delta\bar{\psi} \leq 0.10^\circ$, the grating lobe is well-above the maximum imaged height of 40 m, but for $\Delta\bar{\psi} = 0.15^\circ$, this alias may occur within the vertical extent of the canopy, depending on the height of the trees and the operating angle, which would render a 3D image of a forest ambiguous and ultimately useless. The grating lobe is distorted compared to the true response due to differential layover (recall Figure 3.7): at a focus height above or below the scatterer, its response lays over to slightly different ground-range positions in the constituent SAR images, since they were acquired at slightly different grazing angles, so the scattering contributions do not coherently add perfectly during beamforming. Layover is an important difference between standard beamforming in angle and 3D SAR beamforming to pixel positions.

The adaptive MVDR images in Figures 3.10–3.12 were formed by recomputing the weight vector given by (3.27) at each steer height. When the steer height matches the target height, the beamformer selects weights that pass the target’s scattering response undistorted (ideally), but for all other steer heights, the interferometric phases generated by the target and contained in the covariance matrix serve to indicate the presence of interference, which the beamformer attempts to null out by selecting weights that depend on the inverse of the covariance matrix—this is why there are (almost) no sidelobes. Thus, the displayed images are not precisely point-spread functions, but are the result of scanning beams in height.

Comparing Figure 3.11 for nine channels and Figure 3.12 for three channels, it is

apparent that the fine resolution of the adaptive beamformer is only weakly dependent on the grazing angle interval $N_{ch}\overline{\Delta\psi}$ covered by the acquired channels—at least in this case where the interference occurs at a single height—whereas the resolution of the conventional beamformer is inversely proportional to this angular aperture.

Importantly, the grating lobe occurs regardless of the type of beamforming that is applied, as is evident by its ubiquitous presence when the channel spacing is 0.15° . This ambiguity is a direct consequence of the sampling of the spatial-frequency domain by the collection apertures.

The performance of an adaptive beamformer is, obviously, dependent on the scene. The single point-target scenario indicates the maximum performance that can be expected. The next step is to observe the MVDR performance for a more realistic scene.

3.4.5 Imaging of a random volume over ground

For the intended application of foliage penetration, it is important to demonstrate that 3D SAR beamforming works for vertically extended clutter as well as coherent point targets. Figure 3.13 shows height-range images for a simulated random-volume-over-ground (RVOG) scene consisting of many, randomly positioned, point scatterers (several per resolution cell) that together generate homogeneous surface clutter at ground level and volume clutter which is vertically distributed above the ground; the RVOG model is detailed in Section 3.9. Each 2D SAR image (not shown) appeared to contain just speckle. Nonetheless, the height-range images (particularly (a), (b) and (f)) clearly show that 3D SAR beamforming permits the ground and volume layers to be distinguished from each other and from the surrounding empty space (i.e. the areas below the ground and above the top of the volume). This indicates that the phase of the laid-over contributions from the distributed scattering elements is retained by the clutter in 2D SAR images and can be exploited to reveal the height of those elements.

The accuracy of this representation of the scene's vertical structure depends on (i) the number of channels N_{ch} , (ii) the mean grazing angle spacing $\overline{\Delta\psi}$ between channels, and (iii) the beamforming technique. In addition, if the beamformer is adaptive, as is MVDR, then the accuracy depends on the scene itself via the sample covariance matrix. In (a), where conventional beamforming is applied to nine image channels spaced by 0.15° , the ground layer at 0 m and the volume layer between 0 m and 20 m are illustrated reasonably clearly, but there is a bright alias of the ground at about 30 m (and an alias of the volume above that), which accords with the grating lobe in the point-spread function in Figure 3.10(c). Observe that the ground alias exhibits a slight slope, which arises because the grazing angle varies appreciably over the swath for these close-range acquisitions: in going from mid-range ($y=0$) to far-range ($y=40$ m), $\overline{\psi}$ in (3.38) decreases from 45° to 43.9° and $\overline{\Delta\psi}$ in (3.39) decreases from 0.15° to 0.145° , so h_{amb} in (3.41) increases from 30.6 m to 32.5 m. In (b), the adaptive-MVDR technique achieves better vertical resolution, as shown by the narrower ground response, but suffers from the same aliasing effect, as expected from Figure 3.10(d).

In (c) and (d) of Figure 3.13, where the number of channels is reduced to three, neither of the beamforming techniques are able to resolve the two layers or even distinguish them from the surrounding empty space because the small angular aperture and consequently coarse resolution causes the mainlobe at the top of the volume to merge with the grating lobe from the ground. For the conventional beamformer, this outcome is expected given the poor point-spread function in Figure 3.12(a), but for the MVDR beamformer, the point-spread function in Figure 3.12(b) is misleadingly optimistic—in that case, when steering to heights above or below 0 m, the scatterer at 0 m is the only source of interference, and it can be suppressed strongly by one of the two ($N_{ch}-1$) available nulls, but in the RVOG case here, the two nulls are insufficient to cover the vertical extent of the volume interference, which can be suppressed only weakly. The interference suppression offered by the adaptive beamformer is primarily limited by the number of channels N_{ch} , which provide up to $N_{ch}-1$ nulls, and the actual performance depends on the spatial distribution of the interference field. In (e) and (f), the channel spacing is reduced to 0.05° , so the height-of-ambiguity h_{amb} (given by (3.41)) is increased by a factor of three, and the ground alias would occur at about 90 m. Now the MVDR beamformer is at least able to broadly delineate the vertical extent of the scene, although it is still difficult to distinguish the ground and volume layers visually.




Figure 3.10: Point-spread functions in the height-range plane, where the horizontal axis is ground-range (y), the vertical axis is height (z), and the intensity in decibels is displayed using grey-scale with 40 dB dynamic range and the peak value in each image mapped to white. Each image was obtained according to Algorithm 3.1, with fixed angular spacing $\overline{\Delta\psi} = 0.15^\circ$ achieved by using nine apertures spaced by 5.25 m in the across-track (y) direction, at altitudes 500 m, 1000 m and 1500 m for grazing angles ($\bar{\psi}$) 30° , 45° and 60° , respectively. The 3D processing was done using either conventional (left-column) or adaptive MVDR (right-hand column) beamforming. The blue cross is the first null along slant-range, given by (3.44); the cyan cross is the first null along ground-range, given by (3.48); the red cross is the first null along slant-height, given by (3.45); the magenta cross is the grating lobe along slant-height, given by (3.49).




Figure 3.11: Point-spread functions in the height-range plane, as described in Figure 3.10, except that now the nominal grazing angle is fixed to $\bar{\psi} = 45^\circ$, and the angular spacing ($\overline{\Delta\psi}$) is varied over 0.15° , 0.1° and 0.05° by varying the horizontal spacing between apertures over 5.25 m, 3.5 m and 1.75 m, respectively.




Figure 3.12: Point-spread functions in the height-range plane, as described in Figure 3.11, except that now the 3D processing is applied to only three channels.

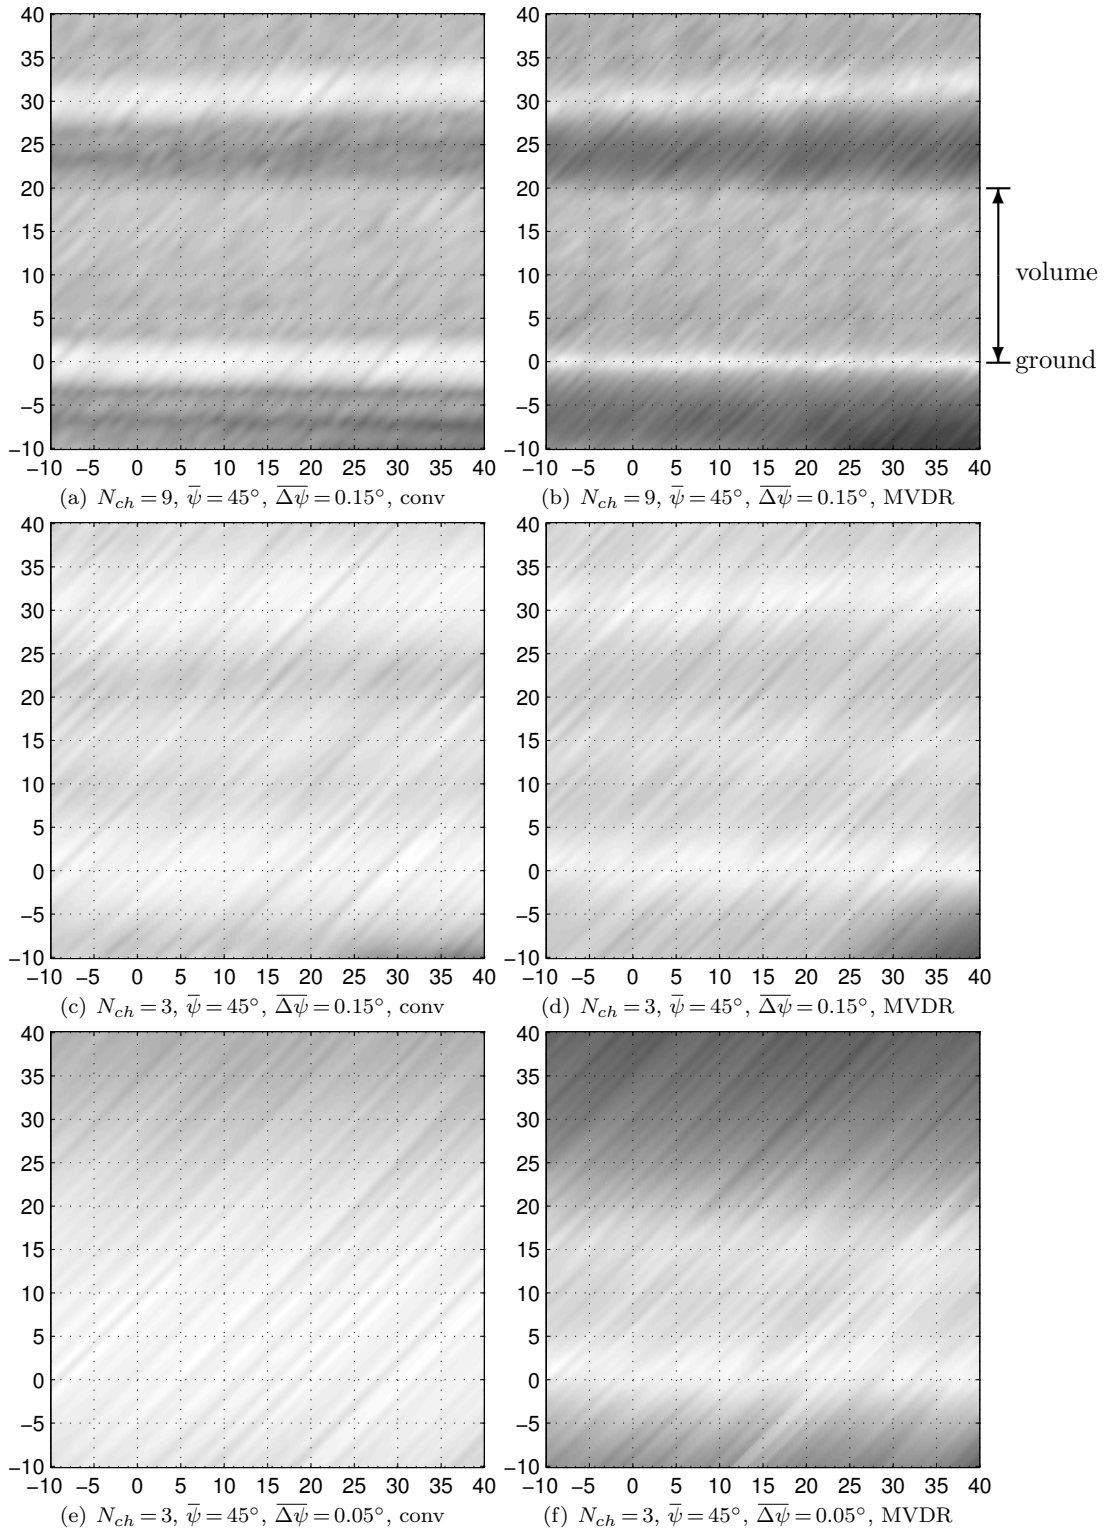


Figure 3.13: Height-range images for a simulated RVOG scene. The volume and ground clutter were generated by synthesising and coherently summing the responses of many point scattering elements: $10/\text{m}^2$ on the ground at 0 m and $5/\text{m}^3$ in the volume between 0 m and 20 m. The scattering intensity of the volume elements was designed to fall off with penetration through the volume at a rate of 0.1 dB/m. The volume power (the total scattering intensity of elements above 0 m) and the ground power (the scattering intensity of elements at 0 m) were set equal. The images were obtained following Algorithm 3.1, except that instead of taking a cut at one azimuth index, the intensity was averaged across azimuth, since the scene is horizontally uniform. The display parameters are the same as in Figures 3.10–3.12.

The performance of the adaptive MVDR beamformer is quantified in Figure 3.14 for different array configurations by averaging the intensity of the clutter (which is horizontally uniform) across azimuth and range to give plots of scattering intensity against steer height. The true vertical structure (the black dash-dotted line) was obtained by evaluating (3.112), (3.116) and (3.117) in Section 3.9.1 given the parameters for Figure 3.13. When the channel spacing ($\overline{\Delta\psi}$) is large (0.15° ; blue), the ground alias is visible at 31 m, and even for a medium spacing (0.10° ; green), the intensity at high steer heights rises due to the approaching ground alias, now at about 45 m. The vertical structure is clearly better approximated when the array contains nine channels (giving the dashed lines) instead of three channels (the solid lines). This contrasts with the point-scatterer case, where quite similar results were obtained after MVDR beamforming using either nine (Figure 3.11) or three (Figure 3.12) channels. To reiterate, the performance of the adaptive beamformer depends on the extent to which the available nulls can cover the spatial distribution of the interference. In this case, the performance is also better when the angular aperture $N_{ch}\overline{\Delta\psi}$ is larger, although this does not always hold for MVDR, as it depends on the precise geometry of the array and the structure of the scene. Overall, the three-channel configurations estimate the vertical structure of the RVOG scene poorly.




Figure 3.14: Variation of average intensity with steer height after adaptive MVDR beamforming of the RVOG scene simulated for Figure (3.13). The curves were obtained following Algorithm 3.1, except that instead of taking a cut at one azimuth index, the intensity was averaged across azimuth and range (and normalised for 0 m (ground) steer height). Six configurations are covered: either three (solid) or nine (dashed) channels, and a grazing angle spacing of either 0.15° (blue; $h_{amb} = 30.6$ m), 0.10° (green; $h_{amb} = 45.9$ m) or 0.05° (red; $h_{amb} = 91.8$ m). The dashed blue curve corresponds to Figure 3.13(b), the solid blue curve corresponds to Figure 3.13(d), and the solid red curve corresponds to Figure 3.13(f). The black dash-dotted lines indicate the true scattering intensity function.

However, the intended application of change detection on the ground does not require the vertical structure to be imaged. All that is required is that, when steered to the ground, the volume (i.e. canopy) interference is attenuated sufficiently to permit the ground coherence to be estimated with moderate accuracy. In the next section, this is shown to be achievable.

3.4.6 Coherent change detection of the ground under a random volume

Thus far, 3D processing has been discussed in terms of resolving vertical structure, which is the most common purpose of 3D SAR in the literature. However, the novel application considered here of coherent change detection on the ground under a forest canopy does not require the vertical structure to be imaged. Instead, the processing goal is to estimate the repeat-pass coherence of scattering at the ground level without significant bias from scattering above the ground. The essential requirements are that the canopy scattering response is attenuated sufficiently and that the ground scattering response is not distorted.

Figure 3.15 shows the results of applying coherent change detection at the ground height for a forest simulated as ground and volume clutter using the RVOG model employed in the previous section, although this time the total volume interference was

10 dB stronger than the ground signal. Raw scattering data for a variety of repeat-pass pairs of multichannel acquisitions were synthesised. The data from each channel was demodulated and focused to form a 2D SAR image, with the focus height matching the flat ground surface. The spatial-frequency apertures of the images from each pair of repeat passes were all trimmed to their common region, so as to avoid any surface decorrelation degrading the output coherence (3D SAR processing does not, in itself, require aperture trimming [Pincus *et al.* 2009]). Then the 3D SAR CCD process outlined in Figure 3.2 was undertaken: the multiple image channels from each pass were combined and steered to the ground height by beamforming, and then the repeat-pass coherence between the beamformer outputs was computed according to (2.1). The sliding windows for the spatial coherence average (for all CCDs) and the \hat{R} estimate (for MVDR) both covered fifty effective looks (i.e. independent spatial samples), which may be too demanding for real scenes, but is useful to show the best-case performance and avoid the misleading effect of estimation bias. Indeed, for the adaptive MVDR beamformer, it was found that using only eight or ten looks led to the ground response being distorted because of poor estimates of the covariance matrix.

To assess change detection performance, two extreme cases must be tested:

1. The ground is changed, but it is masked by a highly correlated volume. This occurs if the volume is unchanged and the two collection geometries are identical. This case is shown in the top row of Figure 3.15; here, the ground change was implemented by regenerating the random ground clutter using a different seed.
2. The ground is unchanged, but it is masked by a highly decorrelated volume. This occurs if the volume is changed and/or the repeat-pass collection geometries are different. This case is shown in the bottom row of Figure 3.15; here, the volume was constant but the two flight-tracks were offset by 0.25° in grazing angle, which at the close ranges considered here corresponds to realistic (though optimistic) offsets of four or five metres in stand-off and/or altitude.

Images (a) and (e) show the result of ordinary CCD processing of a single channel (i.e. using ordinary 2D SAR images): in (a), the correlated volume scattering masks the uncorrelated ground scattering, whereas in (e), the volume decorrelates, masking the correlated ground. Images (c) and (g) show the result of the proposed 3D SAR CCD processing using many (9) image channels per pass: in both cases, the true ground state is revealed i.e. changed in (c) and unchanged in (g). Clearly, even with the relatively strong volume interference, the proposed processing scheme works in these cases, permitting the ground coherence to be accurately estimated after the volume interference is attenuated.

The line plots in (d) and (h) broadly follow the expected trend that the accuracy of the ground coherence estimate improves as the number of image channels N_{ch} is increased, which accords with the results of the previous section. In addition, performance is usually better when the angular aperture $N_{ch}\overline{\Delta\psi}$ formed by the set of collected channels is larger. Indeed, for large apertures, the conventional and adaptive MVDR beamformers sometimes perform similarly. However, the difficulties and expense of building a multichannel airborne radar system drive the use of as few channels as possible. For the given scene and the tested scenarios, a minimal configuration would seem to require adaptive processing using five, widely spaced channels (the solid blue curve), for which complete change would be indicated by a coherence of 0.35 (from (d)), and no change would be indicated by a coherence of 0.87 (from (h)). Observe that the effect of the volume mask is to limit the dynamic range of the coherence estimator.

Two anomalous features stand out. Firstly, in (d), when the number of channels is small ($N_{ch} \leq 4$), the MVDR beamformer performs better when the angular spacing is smaller (i.e. the solid red curve for $\overline{\Delta\psi} = 0.05^\circ$ is below the solid blue curve for $\overline{\Delta\psi} = 0.15^\circ$), which is counter to the expected trend. An important property of this beamformer is that it performs better (steers a deeper null) when the matrix \hat{R} is more ill-conditioned, although this is accompanied by a greater sensitivity to array calibration errors [d'Assumpcao & Mountford 1984]. In this case, the small angular spacing and small number of channels leads to a highly ill-conditioned matrix that permits better performance. Secondly, in (h), when the number of channels is small ($N_{ch} \leq 4$), the conventional beamformer for the smaller angular spacing (the dashed red curve) actually performs worse than when no 3D processing is undertaken at all ($N_{ch} = 1$), again counter to the expected trend. In this case, the two complex outputs of

the beamformer for the two passes happen to be approximately out-of-phase, leading to a very small coherence. Both of these anomalous features will be studied further when the theoretical covariance matrix for the RVOG model is analysed in later sections.

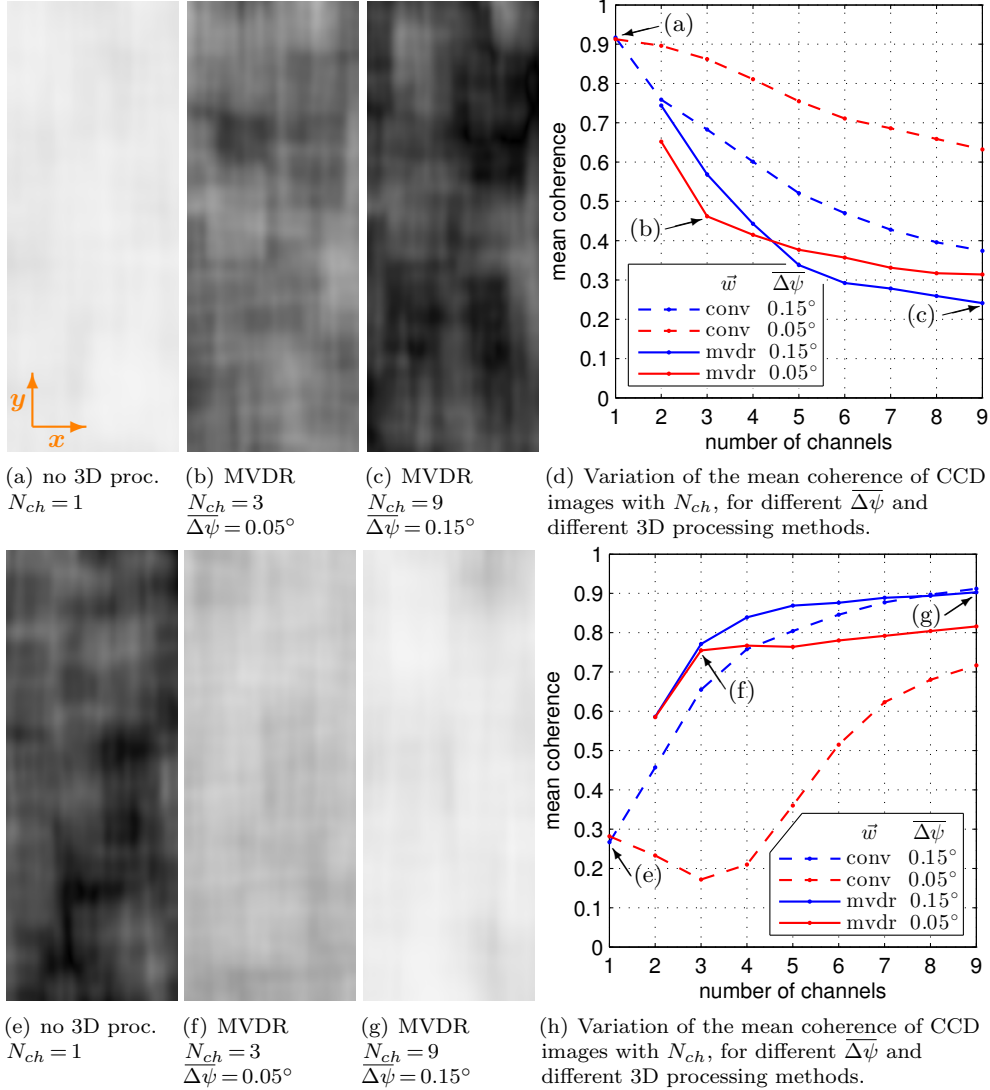


Figure 3.15: Coherent change detection (CCD) at ground height between repeat passes of the RVOG scene simulated for Figure (3.13), except that here the volume power was set 10 dB above the ground power. For the top row, the ground clutter was completely changed between passes (the random-number generator used to determine the position of scattering elements was seeded differently) but the volume was constant and the collection geometries were identical. For the bottom row, the ground and volume clutter were constant between passes, but the collection geometries were offset by 0.25° in grazing angle ($\bar{\psi}$). The example CCD images show the coherence magnitude (white is unity (perfect) correlation, black is zero correlation) in the azimuth and range dimensions at the ground height for different collection configurations (number of channels N_{ch} and grazing angle spacing $\Delta\psi$); (a) and (e) show the ordinary single-channel case, whereas (b), (c), (f) and (g) show the 3D CCDs after adaptive-MVDR SAR beamforming of the channels from each pass. The line plots show the mean coherence magnitude (averaged over azimuth and range) of many such CCD images, formed after either conventional or adaptive-MVDR SAR beamforming.

The key concept of 3D SAR CCD has now been demonstrated for one example scene. However, it is not clear what range of scenes this process will work on, nor how to go about designing a suitable multichannel radar system. How much volume attenuation is required? How should the radar system be configured to achieve this attenuation efficiently? In the following sections, a novel analysis framework is formulated which permits these questions to be answered.

3.5 Multichannel coherence

Repeat-pass coherence, as defined in (2.1), is now reformulated to allow for multiple channels per pass, combined as a weighted sum. For the application considered in this work, each ‘channel’ corresponds to a conventional 2D SAR image acquired at a particular grazing angle, and the weights are selected so as to attenuate scattering from above the ground. For now, the concept of coherence between two complex random scalars is simply extended to the vector case.

Consider a set of N_{ch} radar observations $\mathbf{x} = [x_1, x_2, \dots, x_{N_{ch}}]^T$ acquired at grazing angles $\boldsymbol{\psi} = [\psi_1, \psi_2, \dots, \psi_{N_{ch}}]^T$; treat the individual observations as zero-mean complex random variables and denote the power of the i^{th} observation x_i as σ_i^2 . The coherence γ_{ij} between any pair of observations x_i and x_j is

$$\gamma_{ij} = \frac{\mathbb{E}\{x_i x_j^*\}}{\sqrt{\mathbb{E}\{x_i x_i^*\} \mathbb{E}\{x_j x_j^*\}}} = \frac{\mathbb{E}\{x_i x_j^*\}}{\sqrt{\sigma_i^2 \sigma_j^2}}, \quad (3.50)$$

so the covariance matrix R can be expressed as

$$R = \mathbb{E}\{\mathbf{x}\mathbf{x}^H\} = [\mathbb{E}\{x_i x_j^*\}] = \left[\sqrt{\sigma_i^2 \sigma_j^2} \gamma_{ij} \right] \quad (i, j = 1, 2, \dots, N_{ch}) \quad (3.51)$$

where the last expression indicates that the $(i, j)^{\text{th}}$ element of R is equal to the coherence γ_{ij} between the i^{th} and j^{th} elements of \mathbf{x} scaled by the geometric mean of their respective powers σ_i^2 and σ_j^2 .

In the special case of constant channel power,

$$\sigma^2 = \sigma_i^2 \forall i \quad \Rightarrow \quad R = \sigma^2 \Gamma \quad (3.52)$$

where

$$\Gamma = [\gamma_{ij}] \quad (3.53)$$

is a matrix of pair-wise coherences. Note that $\gamma_{ij} = 1$ for $i = j$, so the leading diagonal of Γ will be all ones.

Weighting and summing the channels gives a new, combined, observation $y = \mathbf{w}^H \mathbf{x}$, as given in (3.23), for a selected weight vector $\mathbf{w} = [w_1, w_2, \dots, w_{N_{ch}}]^T$.

Now consider two multichannel radar observations \mathbf{x}_a and \mathbf{x}_b at grazing angles ψ_a and ψ_b with associated weight vectors \mathbf{w}_a and \mathbf{w}_b that generate combined observations $y_a = \mathbf{w}_a^H \mathbf{x}_a$ and $y_b = \mathbf{w}_b^H \mathbf{x}_b$. Using (3.24), the coherence $\gamma_{y_{ab}}$ between the two combined observations is

$$\gamma_{y_{ab}} = \frac{\mathbb{E}\{y_a y_b^*\}}{\sqrt{\mathbb{E}\{y_a y_a^*\} \mathbb{E}\{y_b y_b^*\}}} = \frac{\mathbf{w}_a^H R_{ab} \mathbf{w}_b}{\sqrt{\mathbf{w}_a^H R_a \mathbf{w}_a \mathbf{w}_b^H R_b \mathbf{w}_b}} \quad (3.54)$$

where $R_a = \mathbb{E}\{\mathbf{x}_a \mathbf{x}_a^H\}$ and $R_b = \mathbb{E}\{\mathbf{x}_b \mathbf{x}_b^H\}$ are the covariance matrices for all channels within passes a and b , respectively, and $R_{ab} = \mathbb{E}\{\mathbf{x}_a \mathbf{x}_b^H\}$ is the cross-covariance matrix for all channels across the two passes. Note that R_{ab} will not be Hermitian or indeed symmetric in any way. While the multichannel coherence $\gamma_{y_{ab}}$ will be analysed in this chapter using the covariance-matrix expression on the right-hand side of (3.54), data processing would involve forming the two combined observations separately and then computing their coherence, as in (3.23) and the left-hand side of (3.54).

For change detection, repeat passes would ideally observe the scene from the same angular viewpoint; in practice, the passes are likely to have some small angular offset. For multichannel acquisition, the relative collection geometries between channels for one pass will be approximately the same as the relative collection geometries between channels for a second pass. Hence, in typical acquisition scenarios, the grazing angles can be characterised by the following constraints, which are not onerous:

$$|\psi_{b_i} - \psi_{a_i}| \lesssim 1^\circ \forall i \quad \text{and} \quad \Delta\psi_{ij} = \Delta\psi_{a_{ij}} \approx \Delta\psi_{b_{ij}} \forall i, j \quad (3.55)$$

where $\Delta\psi_{a_{ij}}$ denotes the grazing angle separation between channels i and j on pass a .

Assume that the mean power received by all channels in a single multichannel acquisition is the same. This requires that the scene be observed from similar angular geometries, as indicated in (3.55), such that the dominant scattering behaviour observed

in all channels is the same, and in addition, that there be no miscalibration between the channels causing a power imbalance—or rather, any power imbalance in the radar system during acquisition has been corrected. Applying this mild assumption,

$$\left. \begin{array}{l} \sigma_a^2 = \sigma_{a_i}^2 \forall i \\ \sigma_b^2 = \sigma_{b_i}^2 \forall i \end{array} \right\} \Rightarrow \left. \begin{array}{l} R_a = \sigma_a^2 \Gamma_a \\ R_b = \sigma_b^2 \Gamma_b \\ R_{ab} = \sqrt{\sigma_a^2 \sigma_b^2} \Gamma_{ab} \end{array} \right\} \Rightarrow \gamma_{y_{ab}} = \frac{\mathbf{w}_a^H \Gamma_{ab} \mathbf{w}_b}{\sqrt{\mathbf{w}_a^H \Gamma_a \mathbf{w}_a \mathbf{w}_b^H \Gamma_b \mathbf{w}_b}}. \quad (3.56)$$

Note that there is no requirement that $\sigma_a^2 = \sigma_b^2$. Note too that because $\Gamma_{ab} = [\gamma_{a_i b_j}]$ contains the coherences across passes, its leading diagonal will not be all ones. If $N_{ch} = 1$, then $\Gamma_a = \Gamma_b = 1$ and $\Gamma_{ab} = \gamma_{ab}$, so for weights $w_a = w_b = 1$, the multichannel coherence $\gamma_{y_{ab}}$ would reduce to the usual single-channel coherence γ_{ab} .

In some scenarios, it may be reasonable to make the approximation that the two sets of coherences observed in the two multichannel acquisitions are the same. For example, if the scene itself is constant and the repeat-pass array geometries are similar, then the set of correlations between channels on pass a will be approximately equal to the set of correlations between channels on pass b . In this case, the weight vectors should sensibly be constrained to be equal. If applicable, this approximation permits a significant simplification of (3.56):

$$\Gamma_a \approx \Gamma_b \Rightarrow \left. \begin{array}{l} \Gamma = (\Gamma_a + \Gamma_b)/2 \\ \mathbf{w} = \mathbf{w}_a = \mathbf{w}_b \end{array} \right\} \Rightarrow \gamma_{y_{ab}} \approx \frac{\mathbf{w}^H \Gamma_{ab} \mathbf{w}}{\mathbf{w}^H \Gamma \mathbf{w}} \quad (3.57)$$

This approximation has been widely employed in the related field of polarimetric SAR interferometry (PolInSAR), where the channels correspond to different pairs of transmit-receive polarisations, and each weight vector is interpreted as a polarimetric scattering mechanism onto which the data is projected [Tabb *et al.* 2002, Flynn *et al.* 2002, Colin *et al.* 2005a, Cloude 2010 ch. 6.2.2]. Indeed, the multichannel radar acquisition modes and the coherence formulations presented in this chapter are directly analogous to those for PolInSAR, except that where PolInSAR exploits diversity in polarisation, the 3D SAR CCD processing here relies on diversity in grazing angle.

3.6 Dual-layer coherence

Section 3.1 discussed phenomenological issues relating to forest propagation and change detection. Now the forest is represented as a two-layer vertical structure consisting of a dense volume of scattering elements above a ground surface [Cloude 2010 ch. 5.2.4, 7.2–7.4]. It will be shown that the interferometric scattering response of the structure can be expressed in terms of the coherences of the two layers weighted by their relative power.

Model the forest canopy as a volume containing many scattering elements that together provide a macroscopic scattering response but also permit lossy propagation. Model the ground underneath as a propagation boundary that provides a surface scattering response. This dual-layer structure is depicted in Figure 3.16. The reference level $z = 0$ is the height at which the input SAR images are formed. The height of the ground above the reference level is denoted z_g ; ideally, $z_g = 0$, but achieving this requires accurate knowledge of the terrain height to be incorporated into the image formation.




Figure 3.16: The dual-layer forest model, consisting of a volume, which is a lossy propagation medium, above a ground surface at $z = z_g$, which is a hard propagation boundary. In the case illustrated here, the ground is below the reference level $z = 0$, so $z_g < 0$.

The models are assumed reasonable at L-band, but at lower frequencies (P-band and below) there is potential for ground penetration, which should be accounted for [Jao *et al.* 1999], and at much higher frequencies (X-band and above), the ground echo may be negligible due to limited propagation through the full vertical extent of the canopy.

Despite the simplicity of the physical model in Figure 3.16, treating the total scattering response of the structure as just the sum of independent ground and volume

contributions is not entirely satisfactory, because the energy incident on the ground clearly depends on propagation through the volume, and the energy may be scattered back-and-forth between the layers before returning to the radar receiver. Forest scattering is commonly considered in terms of independent scattering mechanisms, with three dominant mechanisms: volume scattering from the canopy elements, surface scattering from the ground, and double bounce (two successive specular reflections) between the surface and the volume (realistically, between the ground and tree trunks) [Durden *et al.* 1989, Freeman & Durden 1998, Cloude 2010 ch. 7.3]. Each of these independent scattering mechanisms manifest interferometrically as a phase centre at some height. The phase centres of the direct surface and double-bounce mechanisms will be at the same height z_g , because the total propagation distance in the two cases is the same [Cloude 2010 ch. 7.3].

For the purpose of 3D SAR beamforming, it is useful to group the independent scattering mechanisms into two categories according to phase centre height: scattering at the ground height z_g , both the direct-surface and double-bounce mechanisms, and scattering from above the ground [Cloude 2010 ch. 7.4.1]. For simplicity, the two categories will be referred to as ground g and volume v scattering sources, and they will be treated as independent, although the underlying separation into independent scattering mechanisms with different phase centre heights should be kept in mind. The ground component is modulated by a volume propagation factor p that accounts for two-way attenuation (due to the dissipation of energy through scattering) and possible delay of the propagating wave. (In addition, the polarisation of the wave may change as it propagates through the volume layer Cloude [2010 ch. 4.2.6].)

The SAR images are focused to $z = 0$. Ideally, $z_g = 0$ i.e. the focal surface matches the ground surface, as in Figure 3.1, so g is the desired scattering component at $z = 0$, and v models all the interfering scattering contributions at $z > 0$, which would lay over onto the ground in the conventional SAR images. If $z_g \neq 0$ ($z_g < 0$ implies that the ground is below the focal surface), then the ground will lay over in the images, and the 3D processing will attenuate the ground as if it were interference.

Given these definitions of the scattering sources, decompose a radar observation x of the dual-layer structure in Figure 3.16 as the coherent sum of independent ground g , volume v and noise n components, all of which are treated as zero-mean complex random variables. The ground component is scaled by the complex volume propagation factor p , where $0 \leq |p| \leq 1$. Allow for an unknown complex distortion d as well, in order to later assess the calibration requirements. In the case of airborne radar imaging, the observation would actually be a complex pixel value in a SAR image, with the distortion arising due to system and environmental effects not caused by the scene itself, principally signal modulation in the hardware, uncompensated platform motion, antenna beampattern variations over the aperture, and possible RF interference. Hence, the signal model is

$$x = d(pg + v) + n \quad (3.58)$$

and, since the components are mutually uncorrelated, the mean power is

$$\sigma_x^2 = E\{xx^*\} = |d|^2(|p|^2\sigma_g^2 + \sigma_v^2) + \sigma_n^2 \quad (3.59)$$

where σ_g^2 and σ_v^2 are the expected scattering powers from the ground and volume, and σ_n^2 is the expected noise power. Define the effective signal-to-noise ratio (SNR) and effective ground-volume power ratio μ as follows:

$$SNR = (|p|^2\sigma_g^2 + \sigma_v^2)/\sigma_n^2 \quad (3.60)$$

$$\mu = |p|^2\sigma_g^2/\sigma_v^2. \quad (3.61)$$

For convenience, also define the modified signal

$$\begin{aligned} \hat{x} &= x/d = pg + v + n/d \\ \Rightarrow \sigma_{\hat{x}}^2 &= \sigma_x^2/|d|^2 = |p|^2\sigma_g^2 + \sigma_v^2 + \sigma_n^2/|d|^2. \end{aligned} \quad (3.62)$$

Observe that the signal model in (3.58) presumes that the distortion is applied to the signal before the noise is added, which is standard in the SAR calibration literature [van Zyl 1990, Freeman 1991] and makes sense if the dominant source of distortion is

in the antenna and the dominant source of noise is in the receiver. A more elaborate system model is considered by Just & Bamler [1994], with separate transfer functions to account for the acquisition hardware (applied before noise) and the processing software (applied after noise).

Now consider two radar observations, x_a and x_b , at grazing angles ψ_a and ψ_b , with different propagation factors p_a and p_b , different distortions d_a and d_b , and uncorrelated noises n_a and n_b . Following Zebker & Villasenor [1992], split the contribution from each scattering source into a part that is correlated, c , across the observations, and parts that are uncorrelated, u_a and u_b :

$$\begin{aligned} g_a &= g_c + g_{u_a}, & v_a &= v_c + v_{u_a}, & n_a &= n_{u_a}, \\ g_b &= g_c + g_{u_b}, & v_b &= v_c + v_{u_b}, & n_b &= n_{u_b}. \end{aligned} \quad (3.63)$$

Thus, for the scattering components,

$$\begin{aligned} \sigma_{g_a}^2 &= \mathbb{E}\{g_a g_a^*\} = \sigma_{g_c}^2 + \sigma_{g_{u_a}}^2, & \sigma_{v_a}^2 &= \mathbb{E}\{v_a v_a^*\} = \sigma_{v_c}^2 + \sigma_{v_{u_a}}^2, \\ \sigma_{g_b}^2 &= \mathbb{E}\{g_b g_b^*\} = \sigma_{g_c}^2 + \sigma_{g_{u_b}}^2, & \sigma_{v_b}^2 &= \mathbb{E}\{v_b v_b^*\} = \sigma_{v_c}^2 + \sigma_{v_{u_b}}^2, \\ \mathbb{E}\{g_a g_b^*\} &= \sigma_{g_c}^2 e^{j\phi_g}, & \mathbb{E}\{v_a v_b^*\} &= \sigma_{v_c}^2 e^{j\phi_v}, \end{aligned} \quad (3.64)$$

and for the observed signals,

$$\begin{aligned} \sigma_{x_a}^2 &= \mathbb{E}\{x_a x_a^*\} = |d_a|^2 (|p_a|^2 \sigma_{g_a}^2 + \sigma_{v_a}^2) + \sigma_{n_a}^2, \\ \sigma_{x_b}^2 &= \mathbb{E}\{x_b x_b^*\} = |d_b|^2 (|p_b|^2 \sigma_{g_b}^2 + \sigma_{v_b}^2) + \sigma_{n_b}^2, \\ \mathbb{E}\{x_a x_b^*\} &= d_a d_b^* (p_a p_b^* \sigma_{g_c}^2 e^{j\phi_g} + \sigma_{v_c}^2 e^{j\phi_v}) = |d_a| |d_b| e^{j\phi_{gd}} (p_a p_b^* \sigma_{g_c}^2 + \sigma_{v_c}^2 e^{j\phi_{vg}}), \end{aligned} \quad (3.65)$$

where

$$\phi_{gd} = \phi_g - (\phi_{d_b} - \phi_{d_a}) \quad (3.66)$$

is the interferometric phase ϕ_g for the ground layer relative to the distortion phase difference $\phi_{d_b} - \phi_{d_a}$, and

$$\phi_{vg} = \phi_v - \phi_g \quad (3.67)$$

is the interferometric phase ϕ_v for the volume layer relative to the ground. Following the discussion in Section 3.3, the interferometric phases ϕ_g and ϕ_v each correspond physically to the difference in propagation distance for the two observations due to the vertical offset between the focal surface and the respective scattering source, either the ground surface for ϕ_g , or the effective phase centre of the volume for ϕ_v (see (3.22)). The interferometric phase difference ϕ_{vg} is not dependent on the choice of focal surface, but rather directly indicates the height

$$h_{v,\text{eff}} = \phi_{vg} / k_{z0} \quad (3.68)$$

of the effective phase centre of the volume above the ground [Papathanassiou & Cloude 2001]. The interferometric wavenumber k_{z0} , defined in (3.20), linearly transforms scattering height (above the focal surface) into interferometric phase.

Assume that, across observations, signal components from the same source have the same average power, and the propagation modulation is the same. This is reasonable if the set of objects in the scene does not change (though they may be rearranged) and if the scene is observed from similar angular geometries such that the overall scattering behaviour is effectively the same. The assumption of constant ground power formalises the notion of ‘subtlety’ in the context of change detection, as discussed in Section 3.1.5. Applying this assumption,

$$\left. \begin{aligned} \sigma_g^2 &= \sigma_{g_a}^2 = \sigma_{g_b}^2 \\ \sigma_v^2 &= \sigma_{v_a}^2 = \sigma_{v_b}^2 \\ \sigma_n^2 &= \sigma_{n_a}^2 = \sigma_{n_b}^2 \\ p &= p_a = p_b \end{aligned} \right\} \Rightarrow \begin{aligned} \sigma_{g_u}^2 &= \sigma_{g_{u_a}}^2 = \sigma_{g_{u_b}}^2 \\ \sigma_{v_u}^2 &= \sigma_{v_{u_a}}^2 = \sigma_{v_{u_b}}^2 \\ \sigma_{x_a}^2 &= |d_a|^2 (|p|^2 \sigma_g^2 + \sigma_v^2) + \sigma_n^2 = |d_a|^2 \sigma_{\tilde{x}_a}^2 \\ \sigma_{x_b}^2 &= |d_b|^2 (|p|^2 \sigma_g^2 + \sigma_v^2) + \sigma_n^2 = |d_b|^2 \sigma_{\tilde{x}_b}^2 \\ \mathbb{E}\{x_a x_b^*\} &= |d_a| |d_b| e^{j\phi_{gd}} (|p|^2 \sigma_{g_c}^2 + \sigma_{v_c}^2 e^{j\phi_{vg}}). \end{aligned} \quad (3.69)$$

The coherence γ_{ab} between the two observations can now be decomposed into terms γ_g and γ_v specifying the coherence for a single scattering source in isolation, the

power ratio μ between scattering sources, and the decorrelation γ_{SNR} due to noise. Starting from (2.1) and using (3.60), (3.61) and (3.69),

$$\gamma_{ab} = \frac{\mathbb{E}\{x_a x_b^*\}}{\sqrt{\mathbb{E}\{x_a x_a^*\}\mathbb{E}\{x_b x_b^*\}}} = e^{j\phi_{gd}} \frac{|p|^2 \sigma_{g_c}^2 + \sigma_{v_c}^2 e^{j\phi_{vg}}}{\sqrt{\sigma_{x_a}^2 \sigma_{x_b}^2}} \quad (3.70)$$

$$= e^{j\phi_{gd}} \frac{|p|^2 \sigma_{g_c}^2 + \sigma_{v_c}^2 e^{j\phi_{vg}}}{|p|^2 \sigma_g^2 + \sigma_v^2} \cdot \frac{|p|^2 \sigma_g^2 + \sigma_v^2}{\sqrt{\sigma_{x_a}^2 \sigma_{x_b}^2}} \quad (3.71)$$

$$= e^{j\phi_{gd}} \left(\frac{\mu}{1+\mu} \gamma_g + \frac{1}{1+\mu} \gamma_v \right) \gamma_{SNR} \quad (3.72)$$

$$= e^{j\phi_{gd}} \frac{\gamma_g + \gamma_v / \mu}{1 + 1/\mu} \gamma_{SNR} \quad (3.73)$$

where

$$\gamma_g = \sigma_{g_c}^2 / \sigma_g^2 \quad (\text{ground-only coherence; real}) \quad (3.74)$$

$$\gamma_v = (\sigma_{v_c}^2 / \sigma_v^2) e^{j\phi_{vg}} \quad (\text{volume-only coherence; complex}) \quad (3.75)$$

$$\gamma_{SNR} = \frac{|p|^2 \sigma_g^2 + \sigma_v^2}{\sqrt{\sigma_{x_a}^2 \sigma_{x_b}^2}} = \frac{1}{\sqrt{\left(1 + \frac{1}{|d_a|^2 SNR}\right) \left(1 + \frac{1}{|d_b|^2 SNR}\right)}}. \quad (3.76)$$

As expressed in (3.72), scattering from the ground and volume contribute to the total coherence as a weighted sum, where the weights depend on just the effective ground-volume power ratio. Observe that the coherence γ_{ab} depends on the magnitudes $|d_a|$ and $|d_b|$ of the distortions only via the decorrelation γ_{SNR} due to noise. In the traditional formulation of this problem it is assumed that $d_a = d_b = 1$, in which case γ_{SNR} simplifies to $(1 + SNR^{-1})^{-1}$ [Zebker & Villasenor 1992]. The expression in (3.73) takes the same form as (7.38) in the book by Cloude [2010], except that there it was assumed that $\gamma_g = 1$ i.e. there was no scope for the ground layer to change.

An alternative product decomposition is possible where each factor contains the decorrelation due to a single source. Extending (3.71),

$$\gamma_{ab} = e^{j\phi_{gd}} \underbrace{\frac{|p|^2 \sigma_{g_c}^2 + \sigma_{v_c}^2 e^{j\phi_{vg}}}{|p|^2 \sigma_g^2 + \sigma_v^2 e^{j\phi_{vg}}}}_{\tilde{\gamma}_g} \cdot \underbrace{\frac{|p|^2 \sigma_g^2 + \sigma_{v_c}^2 e^{j\phi_{vg}}}{|p|^2 \sigma_g^2 + \sigma_v^2}}_{\tilde{\gamma}_v} \cdot \underbrace{\frac{|p|^2 \sigma_g^2 + \sigma_v^2}{\sqrt{\sigma_{x_a}^2 \sigma_{x_b}^2}}}_{\gamma_{SNR}}. \quad (3.77)$$

For $\tilde{\gamma}_g$, the denominator is greater than the numerator by only the additional power $\sigma_{g_u}^2$ of the uncorrelated ground component. For $\tilde{\gamma}_v$, the denominator is greater than the numerator by the additional power $\sigma_{v_u}^2$ of the uncorrelated volume component (and the lack of the ϕ_{vg} phase term). For γ_{SNR} , if there is no distortion, the denominator is greater than the numerator by only the additional noise power σ_n^2 . Expressing the total coherence as a product of decorrelation factors in this way is standard in the SAR literature [Zebker & Villasenor 1992, Hagberg *et al.* 1995, Weber-Hoen & Zebker 2000, Cloude 2010 ch. 5.2.2, 5.2.5]. However, the weighted sum in (3.72) is more convenient for this work.

In the special case when $x_a = x_b$, each signal component, and the total signal, would obviously be perfectly correlated with itself ($\gamma_g = \gamma_v = \gamma_{SNR} = \gamma_{ab} = 1$), and there would be no interferometric phase ($\phi_g = \phi_v = 0$) since $\psi_a = \psi_b$. Note that even though noise may be present ($\sigma_n^2 > 0$; finite SNR), it would be perfectly correlated with itself, so it would not cause decorrelation. (Mathematically, $\mathbb{E}\{x_a x_b^*\}$ in (3.65) would need an additional noise power term, which would then flow through (3.70)–(3.72) to give unit coherence.)

For completeness, letting $\mathbf{x} = [x_a, x_b]^T$, observe that the covariance matrix R in (3.51) can be expressed as the sum of the covariance matrices R_g , R_v and R_n for the independent ground, volume and noise sources. Using (3.69), (3.74) and (3.75),

$$R = \begin{bmatrix} \mathbb{E}\{x_a x_a^*\} & \mathbb{E}\{x_a x_b^*\} \\ \mathbb{E}\{x_a^* x_b\} & \mathbb{E}\{x_b x_b^*\} \end{bmatrix} = R_g + R_v + R_n \quad (3.78)$$

where

$$R_g = |p|^2 \sigma_g^2 \begin{bmatrix} |d_a|^2 & |d_a||d_b|e^{j\phi_{gd}}\gamma_g \\ |d_a||d_b|e^{-j\phi_{gd}}\gamma_g & |d_b|^2 \end{bmatrix} \quad (3.79)$$

$$R_v = \sigma_v^2 \begin{bmatrix} |d_b|^2 & |d_a||d_b|e^{j\phi_{gd}}\gamma_v \\ |d_a||d_b|e^{-j\phi_{gd}}\gamma_v^* & |d_b|^2 \end{bmatrix} \quad (3.80)$$

$$R_n = \sigma_n^2 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \quad (3.81)$$

The goal of this chapter is to determine how best to estimate the ground-only coherence γ_g . This parameter is the measure of scene change that we seek: if $\gamma_g \sim 1$, then the ground has not changed between observations, but if $\gamma_g \sim 0$, then the ground has changed.

Consider using the magnitude of the observed coherence γ_{ab} as an estimator for the ground coherence γ_g . From (3.72) it is clear that $\hat{\gamma}_g = |\gamma_{ab}|$ would be biased by the volume, to the extent determined by the ground-volume power ratio μ : if the ground response was relatively strong ($\mu \gg 1$), then the volume bias would be small, but otherwise the volume may distort the estimate significantly. Furthermore, the nature of the error depends on the true ground coherence as well as the volume coherence: if $\gamma_g = 1$ (no change), then any decorrelation exhibited by the volume ($|\gamma_v| < 1$) will bias the estimate downwards i.e. make the ground appear less coherent than it really is, whereas, if $\gamma_g = 0$ (complete change), then any correlation exhibited by the volume ($|\gamma_v| > 0$) will bias the estimate upwards i.e. make the ground appear more coherent than it really is.

Some examples indicate the extent of the bias. Let $\mu = 1$ (equal ground and volume power) and $\gamma_{SNR} = 1$ (negligible noise). If $|\gamma_v| = 1$ (identical repeat-pass geometry and unchanged canopy, so no volume decorrelation) but $\gamma_g = 0$ (complete change on the ground), then $|\gamma_{ab}| = 0.5$. If $\gamma_v = 0$ (complete change in the canopy due, for example, to wind) but $\gamma_g = 1$ (no change on the ground), then $|\gamma_{ab}| = 0.5$ again. In the worst-case, $\gamma_v = -1$ (no decorrelation together with a 180° interferometric phase offset) and $\gamma_g = 1$ (no change on the ground), generating $\gamma_{ab} = 0$ i.e. zero observed coherence.

How can the estimator be made less sensitive to the effects of the volume? One possibility is to attenuate the volume response by vertically beamforming multiple channels acquired on each pass. To establish the basis for this approach, the dual-layer coherence is extended to the multichannel case.

3.7 Multichannel dual-layer coherence

The dual-layer forest coherence model in Section 3.6 and the multichannel coherence formulation in Section 3.5 will now be combined. This is the foundation of the analysis undertaken in this chapter, because it enables repeat-pass 3D SAR data to be modelled directly at the coherence level.

Consider repeat-pass acquisition over a forest using N_{ch} channels per pass at grazing angles ψ_a and ψ_b . In the most general case, the total coherence $\gamma_{y_{ab}}$ is given by (3.54). Say that the scene can be modelled as a dual-layer ground-volume structure, as depicted in Figure 3.16, and that the assumptions in (3.69) of constant power for signal components from the same scattering or noise source, and constant propagation, are satisfied for all $2N_{ch}$ channels.

Impose the calibration requirement that the channels within a pass are corrected for magnitude imbalance, such that the residual distortions at least have equal power, and the decorrelation due to noise, given by (3.76), is constant within a pass, as indicated below:

$$\left. \begin{array}{l} |d_a| = |d_{a_i}| \quad \forall i \\ |d_b| = |d_{b_i}| \quad \forall i \end{array} \right\} \Rightarrow \begin{array}{l} \gamma_{SNR_{a_i}} = (1 + (|d_a|^2 SNR)^{-1})^{-1} \\ \gamma_{SNR_{b_i}} = (1 + (|d_b|^2 SNR)^{-1})^{-1} \\ \gamma_{SNR_{ab}} = [(1 + (|d_a|^2 SNR)^{-1})(1 + (|d_b|^2 SNR)^{-1})]^{-\frac{1}{2}}. \end{array} \quad (3.82)$$

All channels $i = 1, 2, \dots, N_{ch}$ within a pass now have the same power, so the total coherence can be expressed according to (3.56).

Observe that even though the formulation in Section 3.6 presumed a repeat-pass channel pair, the results from that section can also be applied to pairs of channels acquired on the same pass. Hence, every channel-pair coherence in Γ_a , Γ_b and Γ_{ab} in

(3.56) can be modelled by (3.72), that is, as a weighted sum of the individual ground and volume coherences, giving

$$\Gamma_a = \left[\left(e^{j\phi_{gd}} \left(\frac{\mu}{1+\mu} \gamma_g + \frac{1}{1+\mu} \gamma_v \right) \gamma_{SNR} \right)_{a_{ij}} \right] = \Phi_{gd_a} \circ \left(\frac{\mu}{1+\mu} \Gamma_{g_a} + \frac{1}{1+\mu} \Gamma_{v_a} \right) \circ \Gamma_{SNR_a} \quad (3.83)$$

$$\Gamma_b = \left[\left(e^{j\phi_{gd}} \left(\frac{\mu}{1+\mu} \gamma_g + \frac{1}{1+\mu} \gamma_v \right) \gamma_{SNR} \right)_{b_{ij}} \right] = \Phi_{gd_b} \circ \left(\frac{\mu}{1+\mu} \Gamma_{g_b} + \frac{1}{1+\mu} \Gamma_{v_b} \right) \circ \Gamma_{SNR_b} \quad (3.84)$$

$$\Gamma_{ab} = \left[\left(e^{j\phi_{gd}} \left(\frac{\mu}{1+\mu} \gamma_g + \frac{1}{1+\mu} \gamma_v \right) \gamma_{SNR} \right)_{ab_{ij}} \right] = \Phi_{gd_{ab}} \circ \left(\frac{\mu}{1+\mu} \Gamma_{g_{ab}} + \frac{1}{1+\mu} \Gamma_{v_{ab}} \right) \gamma_{SNR_{ab}} \quad (3.85)$$

where

$$\Gamma_{SNR} = [\gamma_{SNR_{ij}}] \quad \text{with} \quad \gamma_{SNR_{ij}} = \begin{cases} 1 & i = j \\ \gamma_{SNR} & \text{otherwise} \end{cases} \quad (3.86)$$

$$\Phi_{gd} = [e^{j\phi_{gd_{ij}}}] \quad (\text{matrix of interferometric phasors}) \quad (3.87)$$

and \circ indicates the Hadamard product. By virtue of the aforementioned constant-power assumptions, the ground-volume power ratio μ is constant for all channels, and the SNR-dependent noise decorrelation γ_{SNR} is given by (3.82), except when the two constituent channels are in fact identical, as discussed in Section 3.6; Γ_{SNR} in (3.86) handles this special case. Similarly, for (3.83) and (3.84), $\phi_{gd_{ij}} = 0$ if $i = j$. Note that the matrices within a pass (Φ_{gd_a} , Γ_{g_a} , Γ_{v_a} and Γ_{SNR_a} for (3.83), and the equivalent for (3.84)) are Hermitian positive semi-definite, because they are components of a scaled covariance matrix (Γ_a or Γ_b), whereas the equivalent matrices across passes ($\Phi_{gd_{ab}}$, $\Gamma_{g_{ab}}$, $\Gamma_{v_{ab}}$) need not be Hermitian.

Assume that the complex scattering responses from the ground observed by all channels within a pass are perfectly correlated. This requires that accurate k -space trimming is applied during image formation, in order that the ground exhibit zero surface decorrelation (see the discussion in Section 2.1). Multichannel acquisition, as proposed in Figure 3.1, enables the observations to be made simultaneously, thus avoiding the possibility of temporal decorrelation; if the individual channels are instead acquired by successive passes of a single-channel radar, then it must simply be assumed that the ground scene has not changed. Between the two passes a and b , the ground may change, so the scattering responses observed across passes may be only partially correlated; given the perfect correlation assumed for all channel pairs within each pass, the level of correlation measured by all channel pairs across passes must be the same. Applying this mild assumption,

$$\gamma_{g_{a_{ij}}} = \gamma_{g_{b_{ij}}} = 1 \quad \forall i, j \Rightarrow \begin{cases} \Gamma_{g_a} = \Gamma_{g_b} = \mathbf{1} \\ \Gamma_{g_{ab}} = \gamma_g \mathbf{1} \end{cases} \quad (3.88)$$

where $\mathbf{1}$ is an $N_{ch} \times N_{ch}$ matrix of ones and γ_g is the ground coherence between passes a and b —this is the scene change parameter to be estimated. Rewriting (3.83)–(3.85) gives

$$\Gamma_a = \Phi_{gd_a} \circ \left(\frac{\mu}{1+\mu} \mathbf{1} + \frac{1}{1+\mu} \Gamma_{v_a} \right) \circ \Gamma_{SNR_a} \quad (3.89)$$

$$\Gamma_b = \Phi_{gd_b} \circ \left(\frac{\mu}{1+\mu} \mathbf{1} + \frac{1}{1+\mu} \Gamma_{v_b} \right) \circ \Gamma_{SNR_b} \quad (3.90)$$

$$\Gamma_{ab} = \Phi_{gd_{ab}} \circ \left(\frac{\mu}{1+\mu} \gamma_g \mathbf{1} + \frac{1}{1+\mu} \Gamma_{v_{ab}} \right) \gamma_{SNR_{ab}}. \quad (3.91)$$

From (3.19) and (3.20), the interferometric phase ϕ_g depends on both the height z_g of the ground above the focal surface, which is fixed for all channels, and the grazing angle separation $\Delta\psi$, which is different for different pairs of channels. Therefore, regardless of the distortion, each Φ_{gd} matrix cannot be reduced to $\mathbf{1}$ unless the focal surface matches the true ground surface (i.e. $z_g = 0$). However, observe that the

vertical wavenumber is only a weak function of the absolute geometry (i.e. $\bar{\psi}$). For example, for $\lambda = 23\text{ cm}$ (L-band), $\Delta\psi = 0.1^\circ$ (typical) and $z_g = 100\text{ m}$ (very large), the interferometric phases in the two cases $\bar{\psi} = 38^\circ$ and $\bar{\psi} = 39^\circ$ differ by only 9.7° . Therefore, given the constraints in (3.55) on the collection geometries, the two sets of ground interferometric phases will be approximately equal, that is,

$$\phi_{g_{ij}} = \phi_{g_{a_{ij}}} \approx \phi_{g_{b_{ij}}} \quad \forall i, j. \quad (3.92)$$

The complex coherence of the volume is strongly dependent on the relative geometry of the two observations. In particular, even small changes in the grazing angle separation $\Delta\psi$ can result in significant changes in the volume coherence, because the coherent superposition of the scattering responses from the constituent elements of the volume will change—this is simply the speckle effect again [Goodman 1975]. Given a set of observations at different grazing angles, as depicted in Figure 3.1, the matrix Γ_v of coherences for all possible pairs of observations will generally have full rank for a fully realised volume. A finite-bandwidth sensor can resolve a dense scene only in the sense of evaluating the net response in a small volume observed from a particular angle. This distributed-target characteristic of natural landscapes stands in contrast to radar observation of a single identifiable target that provides a strong coherent response over a spatial or temporal aperture.

Notwithstanding this strong dependence on the relative collection geometry, the volume coherence may be a weak function of the absolute geometry. For example, consider three pairs of channels at grazing angles $(44^\circ, 44.1^\circ)$, $(44^\circ, 45^\circ)$ and $(45^\circ, 45.1^\circ)$; the first and second pairs will probably exhibit very different volume coherences due to the large difference in angular separation, but the first and third pairs may exhibit very similar volume coherences because the average scattering characteristics may not change significantly over 1° . This argument is backed up by calculations using the random volume model of vegetation coherence, to be introduced in a later section, where the incident energy undergoes exponential decay [Cloude *et al.* 2000]. For example, for $\lambda = 23\text{ cm}$ (L-band), $\Delta\psi = 0.1^\circ$ (typical), volume height $h_v = 30\text{ m}$ (typical) and attenuation rate $\sigma_e^{\text{dB}} = 0.1\text{ dB/m}$ (typical), the volume coherences in the two cases $\bar{\psi} = 37^\circ$ and $\bar{\psi} = 38^\circ$ will be $\gamma_v = 0.655 \angle 145.6^\circ$ and $\gamma_v = 0.644 \angle 147.1^\circ$, which are very similar. Therefore, given the constraints in (3.55) on the collection geometries, the two sets of volume coherences will be approximately equal, that is,

$$\Gamma_v = \Gamma_{v_a} \approx \Gamma_{v_b}. \quad (3.93)$$

Importantly, this assumes that the volume itself appears unchanged, which is not unreasonable if it is observed using a long wavelength over a short time-scale.

At this point, it would be convenient to apply the approximation in (3.57). First, though, the distortions must be restricted such that each channel pair within one pass is subject to the same distortion as the corresponding pair within the other pass. When combined with the previous requirement in (3.82), that the distortion magnitudes within a pass be the same, the calibration requirement on the distortion magnitudes is strict:

$$\begin{aligned} |d| = |d_{a_i}| = |d_{b_j}| \quad \forall i, j &\Rightarrow \gamma_{SNR} = \gamma_{SNR_a} = \gamma_{SNR_b} = \gamma_{SNR_{ab}} = (1 + (|d|^2 SNR)^{-1})^{-1} \\ \Gamma_{SNR} = \Gamma_{SNR_a} = \Gamma_{SNR_b}. & \end{aligned} \quad (3.94)$$

The phase calibration requirement is less strict in that different channels within and across passes may still have different distortion phases (at this stage), as long as the phase difference for any channel pair within one pass is the same as that for the corresponding pair within the other pass. Together with the approximation in (3.92), that the ground interferometric phases for corresponding channel pairs are the same, this calibration requirement gives the following result:

$$\left. \begin{aligned} \phi_{g_{ij}} = \phi_{g_{a_{ij}}} \approx \phi_{g_{b_{ij}}} \quad \forall i, j \quad (\text{from (3.92)}) \\ \phi_{d_{a_j}} - \phi_{d_{a_i}} = \phi_{d_{b_j}} - \phi_{d_{b_i}} \quad \forall i, j \end{aligned} \right\} \Rightarrow \Phi_{gd} = \Phi_{gd_a} \approx \Phi_{gd_b}. \quad (3.95)$$

Putting together the assumptions in (3.69), (3.88), (3.92) and (3.93), and the calibration requirements in (3.82), (3.94) and (3.95), the coherence matrices Γ_a and

Γ_b in (3.89) and (3.90) will be approximately equal, and will take the following form:

$$\Gamma = \Gamma_a = \Phi_{gd} \circ \left(\frac{\mu}{1+\mu} \mathbf{1} + \frac{1}{1+\mu} \Gamma_v \right) \circ \Gamma_{SNR} \approx \Gamma_b. \quad (3.96)$$

Hence, applying the approximation in (3.57), the output coherence is

$$\begin{aligned} \gamma_{y_{ab}} &\approx \frac{\mathbf{w}^H \Gamma_{ab} \mathbf{w}}{\mathbf{w}^H \Gamma \mathbf{w}} \\ &= \frac{\mathbf{w}^H \left[\Phi_{gd_{ab}} \circ \left(\frac{\mu}{1+\mu} \gamma_g \mathbf{1} + \frac{1}{1+\mu} \Gamma_{v_{ab}} \right) \gamma_{SNR} \right] \mathbf{w}}{\mathbf{w}^H \left[\Phi_{gd} \circ \left(\frac{\mu}{1+\mu} \mathbf{1} + \frac{1}{1+\mu} \Gamma_v \right) \circ \Gamma_{SNR} \right] \mathbf{w}} \\ &= \frac{\mu \mathbf{w}^H \Phi_{gd_{ab}} \mathbf{w} \gamma_{SNR} \gamma_g + \mathbf{w}^H \Phi_{gd_{ab}} \circ \Gamma_{v_{ab}} \mathbf{w} \gamma_{SNR}}{\mu \mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \mathbf{w} + \mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \circ \Gamma_v \mathbf{w}} \\ &= \frac{\frac{\mu \mathbf{w}^H \Phi_{gd_{ab}} \mathbf{w} \gamma_{SNR}}{\mu \mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \mathbf{w}} \gamma_g + \frac{\mathbf{w}^H \Phi_{gd_{ab}} \circ \Gamma_{v_{ab}} \mathbf{w} \gamma_{SNR}}{\mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \circ \Gamma_v \mathbf{w}} \cdot \frac{\mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \circ \Gamma_v \mathbf{w}}{\mu \mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \mathbf{w}}}{1 + \frac{\mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \circ \Gamma_v \mathbf{w}}{\mu \mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \mathbf{w}}} \\ &= \frac{\delta \gamma_g + \beta_v \alpha_v / \mu}{1 + \alpha_v / \mu} \end{aligned} \quad (3.97)$$

where

$$\alpha_v = \frac{\mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \circ \Gamma_v \mathbf{w}}{\mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \mathbf{w}}, \quad \beta_v = \frac{\mathbf{w}^H \Phi_{gd_{ab}} \circ \Gamma_{v_{ab}} \mathbf{w} \gamma_{SNR}}{\mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \circ \Gamma_v \mathbf{w}}, \quad \delta = \frac{\mathbf{w}^H \Phi_{gd_{ab}} \mathbf{w} \gamma_{SNR}}{\mathbf{w}^H \Phi_{gd} \circ \Gamma_{SNR} \mathbf{w}}. \quad (3.98)$$

Note that α_v is a real non-negative number because its constituent matrices are Hermitian positive semi-definite and, by the Schur product theorem, their Hadamard product is too [Horn & Johnson 2013 ch. 7.5.3]. Moreover, $0 \leq \alpha_v, |\beta_v|, |\delta| \leq 1$, because the volume coherences in Γ_v and $\Gamma_{v_{ab}}$ and the noise decorrelation γ_{SNR} necessarily have magnitude less than or equal to one, and the majority of the volume coherences across passes in $\Gamma_{v_{ab}}$ will be less than or equal to the corresponding coherences within a pass in Γ_v , by virtue of the collection geometry.

α_v can be interpreted as a multichannel volume attenuation factor, acting to increase the effective ground-volume power ratio within a pass from μ (before beamforming) to μ' (after beamforming), where

$$\mu' = \mu / \alpha_v \quad (0 \leq \alpha_v \leq 1). \quad (3.99)$$

The expression for α_v in (3.98) takes its ideal (i.e. simplest) form when the focal surface matches the true ground surface under the canopy, so that the interferometric phases from the ground are zero. In addition, the differential phase distortions between all pairs of channels within a pass must be zero, and the SNR must be high enough that the decorrelation due to noise is negligible.

$$\left. \begin{aligned} z_g = 0 \Rightarrow \phi_{g_{ij}} = \phi_{g_{a_{ij}}} = \phi_{g_{b_{ij}}} = 0 \quad \forall i, j \\ \phi_{d_{a_j}} - \phi_{d_{a_i}} = \phi_{d_{b_j}} - \phi_{d_{b_i}} = 0 \quad \forall i, j \\ \gamma_{SNR} = 1 \Rightarrow \Gamma_{SNR} = \mathbf{1} \end{aligned} \right\} \Rightarrow \Phi_{gd} = \mathbf{1} \Rightarrow \alpha_v = \frac{\mathbf{w}^H \Gamma_v \mathbf{w}}{\mathbf{w}^H \mathbf{1} \mathbf{w}}. \quad (3.100)$$

To match the focal surface to the true ground surface, accurate knowledge of the topography is required. As discussed in Section 3.1, estimation of the topography of forested terrain is a SAR research topic in its own right, with both PolInSAR and 3D techniques able to provide rough measures of ground height [Reigber *et al.* 2005]. The condition in (3.100) of zero differential phase distortion within a pass is a very strict calibration requirement. If the channels were acquired simultaneously by a fixed and well-characterised antenna array, as proposed in Figure 3.1, then this requirement could be met by measuring and correcting all channel imbalances (which would usually be complex but static for the radar system) and by accurately tracking and compensating for the dynamic array geometry (including rotations induced by the platform attitude) during image focusing. On the other hand, if the channels were themselves acquired by repeat passes of a single antenna, then this calibration requirement would be much more difficult to achieve because the relative platform

motion between all passes must be mutually compensated; accurate phase calibration would likely require a data-driven autofocus across passes, as discussed in Section 3.1.4.

β_v can be interpreted as a ratio of the cross-pass ($\Gamma_{v_{ab}}$) and single-pass (Γ_v) volume coherences, after beamforming. The expression for β_v in (3.98) reduces to unity when the collection geometries for passes a and b are identical (the ideal for change detection), so that the volume coherence and the ground interferometric phase for a channel pair within a pass are the same for the corresponding pair across passes (recall that the volume itself was already assumed to be constant for (3.93)). In addition, the phase distortion for corresponding channels across passes must be the same, and the effect of noise must be negligible.

$$\left. \begin{aligned} \psi_{a_i} = \psi_{b_i} \forall i &\Rightarrow \Gamma_v = \Gamma_{v_{ab}} \\ \phi_{g_{ij}} = \phi_{g_{a_i b_j}} \forall i, j & \\ \phi_{d_{a_i}} = \phi_{d_{b_i}} \forall i & \\ \gamma_{SNR} = 1 &\Rightarrow \Gamma_{SNR} = \mathbf{1} \end{aligned} \right\} \Rightarrow \Phi_{gd} = \Phi_{gd_{ab}} \Rightarrow \beta_v = 1. \quad (3.101)$$

δ acts to degrade the ground coherence. It does not depend on the volume. The expression for δ in (3.98) reduces to unity when the ground interferometric phases, the distortion phases, and the noise satisfy the constraints in (3.101). Note that the condition on the ground interferometric phases requires either identical collection geometries (as in (3.101)) or matching focal and ground surfaces (as in (3.100)).

Consider using the magnitude of the multichannel coherence $\gamma_{y_{ab}}$ in (3.97) as an estimator for the ground coherence γ_g . It is clear that $\hat{\gamma}_g = |\gamma_{y_{ab}}|$ would be biased by the coherence of the volume as captured by β_v , to an extent determined by the effective ground-volume power ratio μ' in (3.99) after beamforming. In the ideal case when the repeat-pass collection geometries are identical, so that (3.101) is applicable and $\beta_v = \delta = 1$, the coherence estimation error ϵ_{γ_g} would be

$$\epsilon_{\gamma_g} = |\gamma_{y_{ab}}| - \gamma_g = \frac{1}{1 + \mu'} (1 - \gamma_g) = \begin{cases} \frac{1}{1 + \mu'} & \gamma_g = 0 \\ 0 & \gamma_g = 1. \end{cases} \quad (3.102)$$

Thus, for identical collection geometries, the worst-case error occurs when the changed ground ($\gamma_g = 0$) is masked by the coherent volume ($|\gamma_v| > 0$). Figure 3.17 quantifies the error in this case as a function of the volume attenuation α_v . Limiting the coherence error to, say, 0.1, sets the required α_v , as follows:

$$\epsilon_{\gamma_g} \leq 0.1 \Rightarrow \mu' \geq 9 \Rightarrow \alpha_v \leq \mu/9 \quad (\alpha_v^{(\text{dB})} \leq \mu^{(\text{dB})} - 9.5 \text{ dB}). \quad (3.103)$$

That is, the beamformer must attenuate the volume such that the effective ground-volume power ratio μ' is at least 9.5 dB. The smaller the initial ratio μ , the larger the required volume attenuation, quantified as a smaller value of the factor α_v . For example, if $\mu = 0$ dB (equal ground and volume power), then $\alpha_v < -9.5$ dB is needed, or if $\mu = -10$ dB, then $\alpha_v < -19.5$ dB is needed.




Figure 3.17: Variation of the ground coherence estimation error ϵ_{γ_g} in (3.102) with the multi-channel volume attenuation factor α_v for different ground-volume power ratios μ . Here, the true ground coherence is zero and the repeat-pass pair have the same collection geometry, so the high coherence of the volume biases the coherence estimate upwards. Note that this plot holds for any array geometry.

Clearly, the beamforming weight vector \mathbf{w} should be selected so as to minimise the multichannel volume attenuation factor α_v , in order that the effective ground-volume power ratio μ' is maximised and the ground coherence γ_g can be estimated with low error. Compared to the single-channel case, the multichannel data offers extra degrees of freedom in terms of choosing how to combine the channels; beamforming exploits this flexibility.

3.8 3D SAR beamforming for volume suppression

To recap, a signal model for the multichannel dual-layer coherence was derived in Section 3.7, culminating in an expression for the output coherence $\gamma_{y_{ab}}$ in (3.97). The magnitude of this coherence serves as an estimate of the desired coherence γ_g of the ground layer, with the accuracy of the estimate critically dependent on the multichannel volume attenuation factor α_v in (3.98), which acts to suppress the interference from the volume layer and thus raise the ground-volume power ratio, as indicated in (3.99). Under certain conditions, α_v takes the ideal form given in (3.100), dependent on just the volume coherences between channels within a pass via the matrix Γ_v , and the weight vector \mathbf{w} .

The question now is how best to select the weight vector \mathbf{w} for 3D SAR beamforming so as to suppress the volume interference and accurately estimate the ground coherence.

Recall from Section 3.6 that, for one interferometric pair of observations, the volume appears as an effective phase centre at one height $h_{v,\text{eff}} = \phi_{vg}/k_{z0}$ (as given by (3.68)), where ϕ_{vg} is the interferometric phase of the volume layer relative to the ground layer, and k_{z0} is the interferometric wavenumber that characterises the angular geometry of the pair (as given by (3.20)). Given multiple channels at slightly different grazing angles, different pairs of observations will have different interferometric wavenumbers and measure different interferometric phases. These differences may not quite cancel out, so the resulting apparent height of the volume phase centre would be slightly different for different pairs. In beamforming terms, the volume interference, as quantified by the matrix of volume coherences Γ_v , manifests across the array as slightly non-stationary.

As an example, phase centre heights can be computed using volume coherences predicted by (3.122) for the random volume model to be described in Section 3.9.1. Setting the scene parameters of volume height $h_v = 20\text{m}$ and attenuation rate $\sigma_e^{\text{dB}} = 0.1\text{dB/m}$, together with the collection parameters $\psi_1 = 35^\circ$, $\psi_2 = 35.05^\circ$ and $\psi_3 = 35.1^\circ$ for three channels, the apparent height $h_{v,\text{eff}}^{(i,j)}$ of the the volume phase centre for pair (i,j) will be $h_{v,\text{eff}}^{(1,2)} = 12.617\text{m}$, $h_{v,\text{eff}}^{(2,3)} = 12.614\text{m}$ and $h_{v,\text{eff}}^{(1,3)} = 12.785\text{m}$. Figure 3.20(d) shows that the apparent height is a weak function of the angular spacing between channels.

Nonetheless, a first attempt at 3D SAR beamforming might be to steer a single null onto the mean apparent height of the volume phase centre predicted by the random volume model, using the null-steer technique described in Section 3.4.2 and illustrated in Figure 3.9. However, this is not as effective as might be hoped, because the volume is truly spatially distributed, so nulling one height still leaves interference from other heights; the phase centre concept is just a geometric interpretation of the net interferometric phase, which in this case is the result of scattering contributions from many different heights.

More formally, the volume exhibits some vertical structure $f_v(z)$ (such as the exponential profile in (3.112) for the random volume model) through its scattering response, and the act of beamforming modulates this structure by the beampattern $b(z)$ (given by (3.36)) of the applied weight vector. The output 3D SAR image pixels contain the coherent sum of scattering contributions from the modulated vertical structure $b(z) \times f_v(z)$.

Multichannel data offers the opportunity to measure and exploit the correlation structure of distributed interference. Beamforming permits the steering of multiple nulls, as discussed in Section 3.4. Where should these nulls be steered? A beamforming approach is sought that optimally suppresses the volume interference.

3.8.1 Optimal volume suppression

The goal of this section is to determine the optimal weight vector \mathbf{w}_{opt} that minimises the multichannel volume attenuation factor α_v in (3.100), and thereby maximises the effective ground-volume power ratio $\mu' = \mu/\alpha_v$ in (3.99) after beamforming.

First note that Γ_v is Hermitian positive semi-definite, since it is a scaled covariance matrix. In fact, Γ_v will almost certainly be positive definite if (1) any canopy is present (i.e. $h_v \neq 0$) and (2) the channels have different grazing angles (i.e. $\Delta\psi_{ij} \neq 0 \forall i \neq j$). Therefore, in what follows, Γ_v is assumed to be invertible.

Observe from (3.100) that

$$\frac{1}{\alpha_v} = \frac{\mathbf{w}^H A \mathbf{w}}{\mathbf{w}^H B \mathbf{w}} = \frac{\mathbf{w}^H C C^{-1} A C^{-H} C^H \mathbf{w}}{\mathbf{w}^H C C^H \mathbf{w}} = \frac{\mathbf{w}'^H D \mathbf{w}'}{\mathbf{w}'^H \mathbf{w}'} \quad (3.104)$$

where $A = \mathbf{1}$ (an $N_{ch} \times N_{ch}$ matrix of ones) is Hermitian, $B = \Gamma_v$ is Hermitian positive definite, $B = C C^H$ is the Cholesky factorisation of B [Horn & Johnson 2013 pp. 90,441], $D = C^{-1} A C^{-H}$, and $\mathbf{w}' = C^H \mathbf{w}$ is the transformed weight vector. The right-hand side of (3.104) fits the general form of a Rayleigh quotient, whose bounds are the minimum and maximum eigenvalues of D , achieved when \mathbf{w}' equals the corresponding eigenvectors [Horn & Johnson 2013 pp. 234–235]. The eigenvalue problem can be expressed in either the transformed space $(\lambda, \mathbf{w}', D)$ or the original space (λ, \mathbf{w}, E) , as follows:

$$D \mathbf{w}' = \lambda \mathbf{w}' \quad (3.105)$$

$$\Rightarrow C^{-1} A C^{-H} C^H \mathbf{w} = \lambda C^H \mathbf{w} \quad (3.106)$$

where

$$E = C^{-H} C^{-1} A C^{-H} C^H = (C C^H)^{-1} A (C C^{-1})^H = B^{-1} A. \quad (3.107)$$

In this case, $E = B^{-1} A = \Gamma_v^{-1} \mathbf{1}$ has unit rank because the matrix of ones causes the columns of E to be all the same. The only non-trivial eigenvalue $\lambda_{\max} > 0$ and corresponding eigenvector \mathbf{w}_{\max} are

$$\left. \begin{array}{l} \mathbf{w}_{\max} = c \Gamma_v^{-1} \mathbf{v}_0 \\ \lambda_{\max} = \mathbf{v}_0^T \Gamma_v^{-1} \mathbf{v}_0 \end{array} \right\} \Rightarrow E \mathbf{w}_{\max} = \Gamma_v^{-1} \mathbf{1} c \Gamma_v^{-1} \mathbf{v}_0 = \underbrace{c \Gamma_v^{-1} \mathbf{v}_0}_{\mathbf{w}_{\max}} \underbrace{\mathbf{v}_0^T \Gamma_v^{-1} \mathbf{v}_0}_{\lambda_{\max}} = \lambda_{\max} \mathbf{w}_{\max} \quad (3.108)$$

where \mathbf{v}_0 is a vector of ones, defined previously in (3.25), $\mathbf{1} = \mathbf{v}_0 \mathbf{v}_0^T$, and c is any complex scalar—eigenvectors are defined up to a scale factor only. Since λ_{\max} is the upper bound on $1/\alpha_v$ (from (3.104)), the lower bound on α_v is

$$\alpha_{v,\min} = \frac{1}{\lambda_{\max}} = \frac{1}{\mathbf{v}_0^T \Gamma_v^{-1} \mathbf{v}_0}, \quad (3.109)$$

achieved by the optimal weight vector

$$\mathbf{w}_{\text{opt}} = \frac{1}{\mathbf{v}_0^T \Gamma_v^{-1} \mathbf{v}_0} \Gamma_v^{-1} \mathbf{v}_0. \quad (3.110)$$

where $c = 1/\lambda_{\max}$ is chosen so that the gain in the steer direction corresponding to \mathbf{v}_0 is unity i.e. from (3.24), $\mathbf{w}_{\text{opt}}^H \mathbf{v}_0 \mathbf{v}_0^T \mathbf{w}_{\text{opt}} = 1$. The normalisation does not change the optimisation of (3.104), since the scalar cancels out, but it does mean that beam patterns plotted for different \mathbf{w} are comparable.

The solution in (3.110) matches the form of the MVDR weight vector in (3.27), except that the latter uses the sample covariance matrix \hat{R} , which will contain the ground signal plus volume interference, whereas the former depends on the volume only via Γ_v , which is assumed known. However, it was shown in (3.32) that if the covariance matrix of the desired signal has unit rank, and if the beamformer is steered towards the desired signal, then the MVDR beamformer will depend on the interference only. From (3.89) and (3.90), the multichannel dual-layer model developed in Section 3.7 fits this scenario in the ideal case when the focal surface matches the ground surface (so the ground interferometric phase is zero for all channel pairs) and noise is negligible—the resulting ground covariance matrix $(\mu/(1+\mu)) \mathbf{v}_0 \mathbf{v}_0^T$ has unit rank. Hence, if the sample covariance matrix matches the ideal covariance matrix for the dual-layer decomposition, then the MVDR beamformer will match the optimal beamformer in (3.110). In essence, the derivation here shows nothing more than that the MVDR technique is optimal for volume suppression in this beamforming scenario, as expected.

In the general case of non-identical passes, there is no requirement that $\mathbf{w}_a = \mathbf{w}_b$; this restriction was imposed only to make the mathematical formulation simpler. It would be reasonable to optimise the volume suppression within each pass by recomputing (3.110) using Γ_{v_a} for pass a and Γ_{v_b} for pass b .

Note that \mathbf{w}_{opt} in (3.110) is optimal only in terms of volume suppression as quantified by the volume attenuation factor α_v in (3.100), but the ultimate measure of performance is the coherence estimation error $\epsilon_{\gamma_g} = |\gamma_{y_{ab}}| - \gamma_g$ over the whole range of γ_g , where $\gamma_{y_{ab}}$ is given by (3.54) (not some approximation to it). If the collection geometries are identical and the various power and noise assumptions hold, so that $\beta_v = \delta = 1$, then \mathbf{w}_{opt} will be optimal in terms of ϵ_{γ_g} too, as can be seen by inspection of (3.97). Otherwise, there is no guarantee that \mathbf{w}_{opt} provides the best possible performance, because the scattering responses obtained on two non-identical passes may give rise to a value of β_v (across passes) that undermines any method that optimises for α_v (within each pass). Nonetheless, given no extra information about the scene, including making no assumption about γ_g , volume suppression is assumed to be a reasonable general-purpose approach to take.

It is important to distinguish between the optimisation here and the optimisations commonly undertaken in PolInSAR [Flynn *et al.* 2002, Cloude 2010 ch. 6.2]. In PolInSAR, the aim is to maximise the total coherence magnitude across passes, in order to obtain the most precise estimate of interferometric phase [Just & Bamler 1994]. By contrast, the aim here is to estimate the magnitude of one component of the coherence across passes, in order to detect changes on the ground. Therefore it would not be appropriate here to just maximise the overall coherence; doing so may mask areas of genuine low ground coherence. When the weight vectors are constrained to be equal, as in (3.57), both optimisations can be formulated as an eigenvalue problem of the form given in (3.106) and (3.107). For PolInSAR, the coherence is maximised by setting $E = \Gamma^{-1}\Gamma_{ab}$.

The random volume model described in the next section provides a way to predict volume coherence. This will be used to populate Γ_v and then quantify and compare the attenuation performance of different beamformers $\{\mathbf{w}\}$ for different array geometries $\{(N_{ch}, \bar{\psi}, \Delta\psi)\}$ and different model canopies. Furthermore, it will be seen that the degradation caused by non-identical geometries is actually quite small, so \mathbf{w}_{opt} is generally a good choice.

3.9 Random-volume-over-ground model of a forest

It would be useful to be able to compare the performance of different radar designs for different model scenes. If the scattering response of the scenes were described by an analytical model that used only a small number of physical parameters, then a given radar design could be characterised by an operating envelope covering a limited range of suitable scenes. To that end, a very simple random-volume-over-ground (RVOG) model of a forest is now described [Treuhaft *et al.* 1996, Treuhaft & Siqueira 2000, Papathanassiou & Cloude 2001]. Where previous sections made no assumptions about the structure of the volume, in this section the intensity of scattering from the volume is assumed to decay exponentially with penetration. This model is commonly used in the field of PolInSAR, where the objective is to estimate the physical parameters for a given dataset via model inversion [Cloude & Papathanassiou 2003, Hajnsek *et al.* 2009b].

3.9.1 Random volume coherence

Model the volume as a homogenous medium consisting of many scattering elements that are randomly oriented and randomly distributed in the vertical dimension between the ground at $z = z_g$ and the top of the volume at $z = z_g + h_v$. Propagation through the volume is subject to exponential decay due to scattering and absorption losses, parameterised by an extinction coefficient σ_e (units: m^{-1}), where $1/\sigma_e$ is the penetration length (in the propagation direction) for which the (one-way) signal power is attenuated by the scale factor $1/e$ [Weber-Hoen & Zebker 2000, Dall 2007].

A point in the volume at height $z' = z - z_g$ ($0 \leq z' \leq h_v$) above the ground will be observed from above through a volume thickness

$$t_v(z') = (h_v - z') / \sin \psi. \quad (3.111)$$

by a radar tilted down at an angle ψ below the horizontal (this is the complement of the incidence angle). For backscatter, the signal propagates through this volume

thickness twice. Therefore, the signal power suffers two-way attenuation according to the scale factor

$$f_v(z') = e^{-2t_v\sigma_e} = e^{-p_1(h_v - z')} \quad (3.112)$$

where

$$p_1 = 2\sigma_e / \sin\psi \quad (3.113)$$

is the two-way extinction coefficient in the vertical direction. By this definition, the scale at the top of the volume is $f_v(h_v) = 1$, and it decreases with decreasing height z' , going through the point $f_v(h_v - \sin\psi/\sigma_e) = 1/e^2$, until it reaches $f_v(0) = e^{-p_1 h_v} = e^{-2t_{v_0}\sigma_e} > 0$ at ground level, where $t_{v_0} = h_v/\sin\psi$ is the total volume thickness. Thus, two-way loss in the propagation direction parameterised by $2\sigma_e$ can be equivalently expressed as a function of vertical height z' parameterised by p_1 in (3.113). The equation for $f_v(z')$ in (3.112) specifies the effective vertical structure of the volume in terms of the scene parameters h_v and σ_e .

Figure 3.18 sketches the thickness and intensity properties of a random volume. This extends a similar, simpler diagram in Cloude [2010 Fig. 5.16].




Figure 3.18: Effective vertical structure of a random volume. (a) shows the effective volume thickness $t_v(z')$ as a function of the height z' above the ground at $z = z_g$, given a planar wavefront incident at grazing angle ψ (the double-headed hook arrows indicate propagation of the wavefront into and out of the volume). (b) shows the exponential decay $f_v(z')$ of scattering intensity with penetration depth, referenced to the top of the volume at height h_v ; when plotted in decibels, the attenuation is linear.

The extinction coefficient σ_e can be linked to the more intuitive concept of an attenuation rate in units of decibels per metre. For propagation through a lossy medium, define the attenuation rate (AR) as

$$\text{AR} = -\frac{10 \log_{10} P_{\text{out}}/P_{\text{in}}}{l} \quad (3.114)$$

where P_{in} and P_{out} are the input and output powers on entry to and exit from the medium, respectively, and l is the propagation length through the medium. (The minus sign ensures that AR takes positive decibel values, even though $P_{\text{out}} < P_{\text{in}}$).

For the case at hand, $l = 2t_{v_0}$ i.e. twice the total volume thickness, and $P_{\text{out}}/P_{\text{in}} = f(0) = e^{-2t_{v_0}\sigma_e}$ i.e. the two-way attenuation scale factor for propagation through the full extent of the volume. Hence,


$$\text{AR} = -\frac{10 \log_{10} e^{-2t_{v_0}\sigma_e}}{2t_{v_0}} = (10 \log_{10} e)\sigma_e = \frac{10}{\ln 10}\sigma_e = \sigma_e^{\text{dB}}. \quad (3.115)$$

It is common to characterise a homogenous random volume by specifying σ_e^{dB} [Cloude *et al.* 2000]. Cloude [2010 p. 230] states the conversion to σ_e in (3.115), but it is not clear from most of the SAR literature that σ_e^{dB} can be interpreted as a simple attenuation rate, akin to cable loss. To reiterate, both the extinction coefficient σ_e (units: m^{-1}) and the attenuation rate σ_e^{dB} (units: dB/m) characterise the volume in terms of one-way power loss in the propagation direction.

Figure (3.19) shows different vertical structures for different values of the random volume parameters h_v and σ_e^{dB} .

The volume is modelled as homogeneous, so every scattering element in the volume would exhibit the same scattering intensity, σ^2 , if observed in isolation. Since the

Figure 3.19: Effective vertical structure $f_v(z')$ for different volumes, where each volume is specified by a height h_v and attenuation rate σ_e^{dB} . Note that the vertical axis is vertical height z' , but the attenuation rate is specified in the propagation direction, which in this case is tilted at $\psi = 45^\circ$. The slope of each line is given by $2\sigma_e^{\text{dB}}/\sin\psi$.



elements are observed as part of the volume, their intensity is modulated by the propagation loss specified by $f_v(z')$, giving an effective scattering intensity

$$\sigma_z^2(z') = f_v(z')\sigma^2, \quad (3.116)$$

which is a function of height. The random volume simulated for Figures 3.13 and 3.14 had a true scattering intensity plotted as the black dash-dotted line in Figure 3.14—this line corresponds to $\sigma_z^2(z')$.

For a SAR image focused to the ground surface, the scattering response of the volume will layover onto the ground. More specifically, at each horizontal position, the observed scattering response will contain the ground response plus the responses of the volume scattering elements at the same range and therefore illuminated by the same wavefront as the ground element, with all elements (from the ground and the volume) weighted by the propagation attenuation factor $f_v(z')$ in (3.112) according to their height.

The total scattering power σ_v^2 of the volume is the incoherent sum of the volume contributions along the wavefront. The length of the wavefront in the volume is $h_v/\cos\psi$ and the effective two-way extinction coefficient along the wavefront is $p_1 \cos\psi$, where p_1 in (3.113) is the two-way extinction coefficient in the vertical direction, so the attenuation at ground level is $f_v(0) = e^{-p_1 h_v}$, as expected from (3.112). Because the model volume is homogeneous, the summation can be equivalently expressed in the vertical direction. Treating the volume as a continuous medium and using (3.112) and (3.116) gives

$$\sigma_v^2 = \int_0^{h_v} \sigma_z^2(z') dz' = \sigma^2 \int_0^{h_v} e^{-p_1(h_v-z')} dz' = \frac{1 - e^{-p_1 h_v}}{p_1} \sigma^2 \quad (3.117)$$

where σ^2 and $\sigma_z^2(z')$ technically have units of power per unit height (so the summation is over elemental scattering power contributions $\sigma_z^2(z') dz'$).

To simulate a random volume using discrete scattering elements, it is necessary to assign scattering intensities to the elements according to (3.116) such that the volume exhibits both a desired vertical structure $f_v(z')$ and a desired total power σ_v^2 . If the simulation uses L scattering elements per metre in height, then set

$$\sigma^2 = \frac{p_1}{1 - e^{-p_1 h_v}} \cdot \frac{\sigma_v^2}{L}. \quad (3.118)$$

Now consider the coherence between two observations, v_a and v_b , of the volume. Assume that the volume is unchanged between observations, but that the two collections have a small difference in grazing angle $\Delta\psi$. Since the model volume consists of randomly oriented scattering elements, a small difference in angle does not cause a difference in expected power (neglecting the very small difference due to the differential volume thickness in (3.111)). Therefore, the two observations exhibit the same scattering intensity function $\sigma_z^2(z')$ in (3.116) and the same total volume power σ_v^2 in (3.117). This validates the assumption $\sigma_v^2 = \sigma_{v_a}^2 = \sigma_{v_b}^2$ made in (3.69) when deriving the generic dual-layer ground-volume coherence.

However, the difference in angle does give rise to an interferometric phase from each scattering element due to the difference in the residual propagation phase after

image formation, as discussed in Section 3.3. From (3.19), the interferometric phase depends on both the difference in grazing angle (via the interferometric wavenumber k_{z0}) and the height of the scattering element above the focal surface. To make this latter dependence explicit for the constituent scattering elements of the vertically extended volume, express the interferometric phase as

$$\phi(z) = k_{z0}z = k_{z0}z_g + k_{z0}z' = \phi(z_g + z') \quad (3.119)$$

where $\phi_g = \phi(z_g) = k_{z0}z_g$ is the interferometric phase of the ground relative to the focal surface (introduced in (3.65)). Note that the interferometric phase $\phi(z)$ in (3.119) varies with the vertical dimension z referenced to the focal surface, whereas the scattering intensity function $f_v(z')$ in (3.112) varies with the vertical dimension z' referenced to the ground surface. (Recall that $z_g = z - z'$ is the height of the ground above the focal surface.)

Putting these arguments together, the net correlation $E\{v_a v_b^*\}$ (see (3.64)) between the two observations over the vertical extent of the volume can be expressed as a summation of the scattering intensity (common to the two observations) weighted by the interferometric phase, where both terms vary with height. This follows the general approach of Papathanassiou & Cloude [2001] and Cloude [2010 eq. 5.46].

$$\begin{aligned} E\{v_a v_b^*\} &= \int_0^{h_v} \sigma_z^2(z') e^{j\phi(z_g+z')} dz' && \text{(analogous to (3.117))} \\ &= \sigma^2 e^{j\phi_g} \int_0^{h_v} e^{-p_1(h_v-z')} e^{jk_{z0}z'} dz' && \text{(using (3.112), (3.116) \& (3.119))} \\ &= \sigma^2 e^{j\phi_g} e^{-p_1 h_v} \int_0^{h_v} e^{p_2 z'} dz' && \text{(setting } p_2 = p_1 + jk_{z0}\text{)} \\ &= \sigma^2 e^{j\phi_g} \frac{e^{p_2 h_v} - 1}{p_2 e^{p_1 h_v}} && \end{aligned} \quad (3.120)$$

$$= \sigma^2 e^{j\phi_g} \frac{e^{p_2 h_v} - 1}{p_2 e^{p_1 h_v}} \quad (3.121)$$

Hence, the net volume coherence γ_v , as defined in (3.75), is the ratio of $E\{v_a v_b^*\}$ in (3.121) and $\sigma_v^2 = E\{v_a v_a^*\} = E\{v_b v_b^*\}$ in (3.117), with the ground interferometric phase ϕ_g removed. (Note that (3.118) is not used here.)

$$\begin{aligned} \gamma_v &= \frac{E\{v_a v_b^*\}}{\sqrt{E\{v_a v_a^*\} E\{v_b v_b^*\}}} e^{-j\phi_g} = \frac{\sigma^2}{\sigma_v^2} \frac{e^{p_2 h_v} - 1}{p_2 e^{p_1 h_v}} \\ &= \frac{p_1}{1 - e^{-p_1 h_v}} \cdot \frac{e^{p_2 h_v} - 1}{p_2 e^{p_1 h_v}} \\ &= \frac{p_1 (e^{p_2 h_v} - 1)}{p_2 (e^{p_1 h_v} - 1)} \quad \left(\begin{array}{l} p_1 = 2\sigma_\epsilon / \sin \psi \\ p_2 = p_1 + jk_{z0} \end{array} \right) \end{aligned} \quad (3.122)$$

This is what we seek: an explicit expression for the interferometric volume coherence in terms of physical parameters, in this case the volume height h_v and the extinction coefficient σ_ϵ . Indeed, the utility of the random volume model is in large part due to the fact that its coherence can be expressed analytically. As presented for a general volume in Section 3.6, the coherence γ_v is complex, with phase $\angle \gamma_v = \phi_{vg} = \phi_v - \phi_g$ (see (3.67) and (3.75)), so the volume will manifest interferometrically as a single effective phase centre at height $h_{v,\text{eff}} = \phi_{vg}/k_{z0}$ above the ground (see (3.68)).

A special case of the random volume model is when the extinction is zero, so the scattering intensity does not change with penetration. The resulting coherence can be found from (3.122) using L'Hôpital's rule, as follows:

$$\begin{aligned} f(p_1) &= p_1 (e^{p_2 h_v} - 1) \Rightarrow \frac{\partial f(p_1)}{\partial p_1} = (1 + p_1 h_v) e^{p_2 h_v} - 1 \\ g(p_1) &= p_2 (e^{p_1 h_v} - 1) \Rightarrow \frac{\partial g(p_1)}{\partial p_1} = (1 + p_2 h_v) e^{p_1 h_v} - 1 \\ \lim_{p_1 \rightarrow 0} \gamma_v &= \lim_{p_1 \rightarrow 0} \frac{f(p_1)}{g(p_1)} = \lim_{p_1 \rightarrow 0} \frac{\partial f(p_1)/\partial p_1}{\partial g(p_1)/\partial p_1} = \frac{e^{jk_{z0} h_v} - 1}{jk_{z0} h_v} = e^{jk_{z0} h_v/2} \text{sinc} \frac{k_{z0} h_v}{2\pi} \end{aligned}$$

$$\begin{aligned}
p_1 &= 0 \\
f_v(z') &= 1 \\
\sigma_e = 0 &\Rightarrow \gamma_v = e^{jk_{z0}h_v/2} \operatorname{sinc} \frac{k_{z0}h_v}{2\pi} \quad (h_{\text{amb}} = h_v \Rightarrow \gamma_v = 0) \\
h_{v,\text{eff}} &= h_v/2.
\end{aligned} \tag{3.123}$$

Figure 3.20 plots the random volume coherence γ_v in (3.122) for different scene and collection parameters. From (a), for larger volume heights or lower attenuation rates, the radar ‘sees’ a greater vertical extent of volume, so the resulting coherence magnitude is lower, with zero coherence when $\sigma_e = 0$ and $h_v = h_{\text{amb}}$. From (c), the effective phase centre of the volume always lies between $h_v/2$ and h_v , with higher heights for larger attenuation rates. Indeed, the larger the attenuation rate, the closer the volume approximates a surface, with higher coherence (approaching unity) and higher phase centre (approaching the top of the volume). From (b), the larger the angular separation between the two observations of the volume, the lower the coherence magnitude, with zero coherence when $\sigma_e = 0$ and $\Delta\psi = 2\pi \cos \bar{\psi} / k_{r0} h_v$, which corresponds to $h_{\text{amb}} = h_v$. Observe from (a) and (b) that the special case of zero extinction always generates the minimum volume coherence, given by (3.123). From (d), the interferometric phase and the associated phase centre height are weak functions of $\Delta\psi$ —unlike the coherence magnitude, the phase centre is an intrinsic property of the volume, so approximately the same value should be recoverable by a wide range of collection geometries, up to a limit imposed by h_{amb} .

3.9.2 Random-volume-over-ground coherence

Combining the general dual-layer coherence formulation presented in Section 3.6 with the specific random volume structure described in Section 3.9.1 gives the random-volume-over-ground (RVOG) model. The dual-layer interferometric coherence γ in (3.73) can now be evaluated using the random volume expression for γ_v in (3.122), along with values for the ground coherence γ_g (the indicator of ground change) and the ground-volume power ratio μ . Note that typical formulations of single-pass interferometry for PolInSAR implicitly assume that $\gamma_g = 1$ i.e. that the only source of decorrelation is the volume [Treuhaft *et al.* 1996]. Of course, the wider point of this chapter is to treat γ_g as an unknown and estimate it.

In the general dual-layer signal model in (3.58), developed at the start of Section 3.9, the complex propagation factor p is used to represent modulation of the scattering response from the ground layer due to propagation through the volume layer. In particular, the ground power σ_g^2 is attenuated by $|p|^2$, as indicated in (3.59). For the random volume model, attenuation due to propagation losses is captured by the exponential scale factor $f_v(z')$ in (3.112). Relating these formulations, the attenuation at the ground level due to the random volume is

$$|p|^2 = f_v(0) = e^{-p_1 h_v} = e^{-2\sigma_e h_v / \sin \psi}. \tag{3.124}$$

The ground-volume power ratio μ in (3.61) becomes

$$\mu = \frac{|p|^2 \sigma_g^2}{\sigma_v^2} = \frac{p_1}{e^{p_1 h_v} - 1} \cdot \frac{\sigma_g^2}{\sigma^2} \tag{3.125}$$

where the total power σ_v^2 of the random volume is given by (3.117), with σ^2 being the underlying power of each scattering element in the (homogeneous) volume, before accounting for propagation losses. Note that σ_v^2 in (3.117) incorporates the effect of propagation, whereas σ_g^2 and σ^2 do not. Thus, (3.125) makes it clear that μ depends on the inherent scattering properties of the two layers as well as the propagation losses through the volume, and the random volume parameters h_v and σ_e^{dB} determine those propagation losses. For typical radar and scene parameters, the scale factor $p_1 / (e^{p_1 h_v} - 1)$ is much smaller than unity. For example, given $\bar{\psi} = 35^\circ$, $h_v = 20\text{m}$ and $\sigma_e^{\text{dB}} = 0.1\text{dB/m}$, the scale factor is 0.016, and if the scattering intensity of individual ground and volume scene elements is the same i.e. $\sigma_g^2 = \sigma^2$ (assuming that comparison is even meaningful), then $\mu = -17.9\text{dB}$. In general, μ may take any value from a very wide range (e.g. $-30\text{dB} \leq \mu \leq 30\text{dB}$).

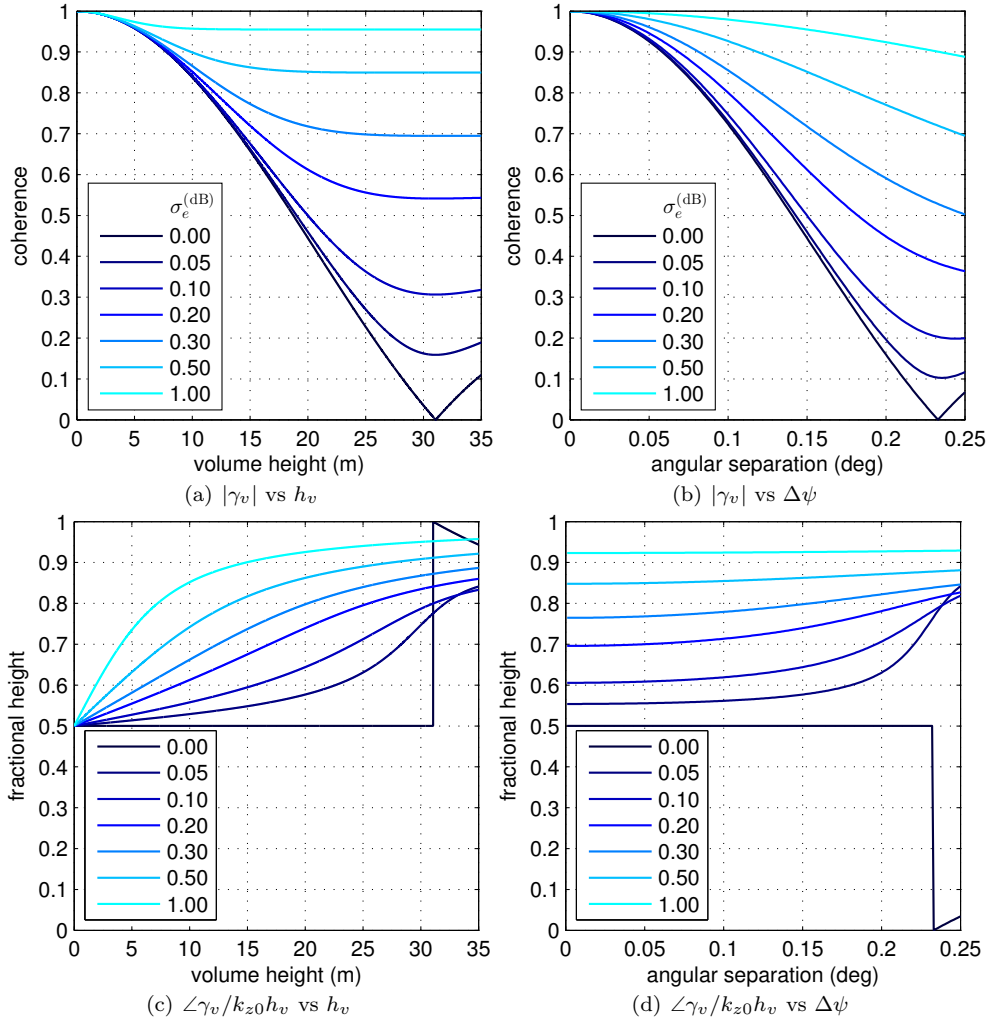


Figure 3.20: Random volume coherence γ_v in (3.122) for different scene and collection parameters. (a) and (b) show the magnitude of the coherence; (c) and (d) show a scaled version of the interferometric phase that corresponds to the height of the volume's effective phase centre as a fraction of the height of the volume itself. In (a) and (c), the volume height h_v is varied for different attenuation rates σ_e^{dB} , while the radar system and collection geometry are fixed, with angular separation $\Delta\psi = 0.15^\circ$, so $h_{\text{amb}} = 31$ m. In (b) and (d), $\Delta\psi$ is varied for different σ_e^{dB} and fixed $h_v = 20$ m; when $\Delta\psi = 0.233^\circ$, $h_{\text{amb}} = h_v$. In all cases, $\psi = 45^\circ$ and $\lambda = 23$ cm. Similar curves to (a) and (c) are plotted in Cloude [2010 Fig. 5.17 & 5.18], although their meaning is slightly obscured by unfortunate typographical errors.

Observe from (3.125) that μ and $|p|^2$ are directly proportional, so $\mu/\mu_{\text{ref}} = |p|^2/|p_{\text{ref}}|^2$ for some reference point $(\mu_{\text{ref}}, p_{\text{ref}})$. Figure 3.21 shows how the ground-volume power ratio μ varies with volume height h_v and attenuation rate σ_e^{dB} due to changing propagation attenuation $|p|^2$. As expected, an increase in both h_v and σ_e^{dB} causes a significant decline in μ .




Figure 3.21: Ground-volume power ratio μ in decibels (mapped to colour) as a function of volume height h_v and attenuation rate σ_e^{dB} , which cause propagation loss according to (3.124). The black cross marks the reference scene: $\mu = 0$ dB, $h_v = 20$ m and $\sigma_e^{\text{dB}} = 0.1$ dB/m. The black lines mark the ± 3 dB boundaries relative to the reference value of μ .

In order to appreciate the interaction between γ_v and μ , consider the particular case when $\gamma_g = \gamma_{SNR} = 1$, so the observed RVOG coherence in (3.73) reduces to

$$\gamma = e^{j\phi_{gd}} \frac{\gamma_v + \mu}{1 + \mu}. \quad (3.126)$$

This is the most common form of the RVOG model [Papathanassiou & Cloude 2001, Cloude 2010 eq. 7.38]. Recall that μ is real but γ_v is complex. When μ is small, the volume dominates, so $|\gamma| \approx |\gamma_v|$, and when μ is large, the ground dominates, so $|\gamma| \approx 1$. However, for moderate values of μ , say $|\mu| < 10$ dB (expected at L-band), the coherence magnitude $|\gamma|$ depends on how μ and γ_v interact, which is not immediately apparent. In particular, if $\angle\gamma_v \approx 180^\circ$ and $\mu \approx |\gamma_v|$, then $\gamma \approx 0$. Thus, even though the ground and volume may each be highly correlated in isolation, together they may generate a low net coherence.

Figure 3.22 illustrates this point by plotting the RVOG coherence magnitude $|\gamma|$ in (3.126) for different ground-volume power ratios μ and different angular separations $\Delta\psi$. Observe that, counter-intuitively, the decorrelation of the dual-layer structure does not vary monotonically with either μ (in (a)) or $\Delta\psi$ (in (b)), because γ_v (critically determined by $\Delta\psi$) is complex. For each value of $\Delta\psi$ in (a), there is a value of μ that minimises $|\gamma|$, and in (b), certain values of $\Delta\psi$ cause local minima in $|\gamma|$.

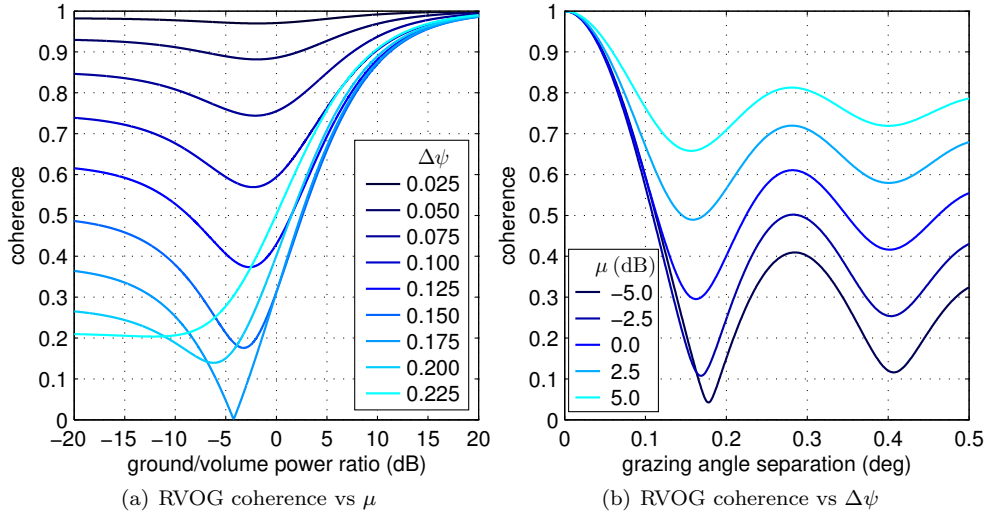


Figure 3.22: Magnitude of the dual-layer RVOG coherence γ in (3.126) as a function of (a) the ground-volume power ratio μ for different angular separations $\Delta\psi$, and (b) $\Delta\psi$ for different μ (same data as in (a), just displayed differently). Fixed parameters: $h_v = 20$ m, $\sigma_e^{\text{dB}} = 0.1$ dB/m, $\bar{\psi} = 45^\circ$, $\lambda = 23$ cm, $\gamma_g = 1$ (i.e. the ground is unchanged) and $\gamma_{SNR} = 1$ (i.e. neglecting noise). In (a), when $\mu \approx -20$ dB, the observed coherence is the volume coherence for that value of $\Delta\psi$, and when $\mu \approx 20$ dB, the observed coherence is unity, regardless of $\Delta\psi$. A related (but different) plot is shown in [Cloude 2010 Fig. 7.18].

3.10 RVOG beamformer performance

Define the RVOG beamformer \mathbf{w}_{rvog} as the optimal beamformer in (3.110) obtained when Γ_v is populated with volume coherences predicted by (3.122) according to the random volume model.

$$\mathbf{w}_{\text{rvog}} = \frac{\Gamma_v^{-1} \mathbf{v}_0}{\mathbf{v}_0^T \Gamma_v^{-1} \mathbf{v}_0} \quad \text{where } \Gamma_v = [\gamma_{v_{ij}}] \text{ and } \gamma_{v_{ij}} \text{ is given by (3.122) for pair } (i, j) \quad (3.127)$$

The RVOG beamformer serves two purposes. Firstly, it supports radar design by indicating the best-case performance of different multichannel systems operating over different scenes. In Section 3.10.1, performance is analysed in terms of volume suppression (α_v), and in Section 3.10.2, performance is analysed in terms of coherence estimation error (ϵ_{γ_g}). The performance is ‘best-case’ because the beamformer is matched to the (model) data.

Secondly, the RVOG beamformer could actually be applied to measured data. Implementation would require estimated or assumed values of the random volume parameters (volume height h_v and attenuation rate σ_e^{dB}) together with the known multichannel collection geometry (grazing angles $\{\psi_i\}$ for channels $i = 1, 2, \dots, N_{ch}$) and the radar wavelength λ . Performance would depend on how closely the modelled volume matches the actual canopy. Indeed, the assumptions of this approach are two-fold: firstly, that the canopy is well-modelled by a random volume in general, and secondly, that the canopy is well-modelled by the particular random volume specified by the selected parameters. The sensitivity to the second type of assumption will be tested in Section 3.10.3. The sensitivity to more general perturbations, particularly those that can be modelled as phase noise, will be tested in Section 3.10.4.

Compare the parametric RVOG beamformer $\mathbf{w}_{\text{rvog}} = \Gamma_v^{-1} \mathbf{v}_0$ in (3.127) with the adaptive MVDR beamformer $\mathbf{w}_{\text{mvdr}} = \hat{R}^{-1} \mathbf{v}_0$ in (3.27), neglecting scale factors. Where the former is clairvoyant, in that knowledge of the true volume interference Γ_v is required for maximal cancellation via Γ_v^{-1} , the latter instead requires an estimate of the total ground (signal) plus volume (interference) covariance matrix \hat{R} to adaptively cancel interference from heights not corresponding to the steering vector \mathbf{v}_0 . As discussed in Section 3.8.1, if the scene conforms to the ideal dual-layer model, and the sample covariance matrix accurately captures this, then the two beamformers will be equal. In cases where it is difficult to obtain a reliable estimate of \hat{R} , say due to a limited number of resolution cells to average over, it may be safer to instead estimate or assume values of the random volume parameters and employ \mathbf{w}_{rvog} . A mixed approach, whereby the adaptivity is constrained within reasonable RVOG bounds, may be possible. However, the RVOG beamformer may be comparatively poor if the canopy is not well-modelled by a random volume at all; tests for the validity of the RVOG model, given fully-polarimetric SAR data, have been proposed by López-Martínez & Alonso-González [2014].

3.10.1 Optimal RVOG volume suppression

First consider the best-case performance when the beamformer is matched to the (model) data. Figure 3.23(a) shows the volume attenuation achievable by the optimal RVOG beamformer for different collection parameters, given a typical RVOG scene. Note that, as the angular spacing $\Delta\psi$ decreases, the optimal performance initially improves (before plateauing) despite the resolution ρ_z in (3.42) of the conventional beamformer becoming coarser.

The key question is under what conditions the optimal RVOG beamformer provides sufficient volume attenuation to permit accurate ground coherence estimates. In the ideal change detection scenario where the repeat-pass collection geometries are identical, the required level of attenuation depends solely on the ground-volume power ratio μ according to (3.103); see Figure 3.17 in Section 3.7 for example values. Crucially, even with only three channels, judicious choice of the grazing angle spacing ($\Delta\psi \sim 0.05^\circ$) enables multichannel volume attenuation around -12 dB, which is sufficient to permit accurate ground coherence estimates when $\mu > -2.5$ dB. As the number of channels is increased (and the grazing angle spacing is adjusted to suit), the volume attenuation increases substantially, enabling greater sensitivity to ground changes even though the ground signal is weak; for example, nine channels (at $\Delta\psi \sim 0.15^\circ$) enables volume attenuation of -21 dB, thus permitting ground coherence estimates when $\mu > -11.5$ dB. (These results are for $\bar{\psi} = 35^\circ$; for $\bar{\psi} = 45^\circ$, the decibel attenuations are reduced by approximately six per cent.)

The high level of volume attenuation shown in Figure 3.23(a) for one scene is in fact achievable over a wide range of scenes. Figure 3.24(a) shows the performance of the optimal RVOG beamformer for different RVOG scene parameters, given a minimal three-channel radar system. Observe that for most scenes, the optimal multichannel volume attenuation is -10 dB or better. The performance is determined by both the height of the volume's effective phase centre and the effective width of the vertical structure: when there is little propagation loss (i.e. the attenuation rate σ_e^{dB} and the extinction coefficient σ_e are small), scattering contributions will be received from a wide range of heights, which the beamformer's ability to suppress is limited by its small number of channels, but when there is a large propagation loss, then the dominant scattering contributions must be generated from a relatively narrow vertical extent

Figure 3.23: Optimal RVOG beamformer performance for different collection parameters. (a) shows the minimum (best) achievable multichannel volume attenuation factor $\alpha_{v,\min}$ as a function of the grazing angle spacing $\Delta\psi$ for different numbers of channels N_{ch} forming an array centred at $\bar{\psi} = 35^\circ$ ($\lambda = 23\text{cm}$). The scene parameters are fixed: volume height $h_v = 20\text{m}$ and attenuation rate $\sigma_e^{\text{dB}} = 0.1\text{dB/m}$. Each point on each curve was obtained by populating the matrix Γ_v with volume coherences from (3.122) for all channel pairs in the given configuration, and then computing $\alpha_{v,\min}$ via (3.109). The curves in (a) have been cropped to exclude cases where the condition number of Γ_v , shown in (b), is greater than 10^6 . The role of the matrix condition number is discussed in Section 3.10.4.





Figure 3.24: Optimal RVOG beamformer performance for different scene parameters. (a) shows the minimum achievable multichannel volume attenuation factor $\alpha_{v,\min}$ as a function of the volume height h_v for different attenuation rates σ_e^{dB} . The collection parameters are fixed: $N_{ch} = 3$ channels spaced by $\Delta\psi = 0.05^\circ$ and centred at $\bar{\psi} = 35^\circ$ ($\lambda = 23\text{cm}$). Each point on each curve was obtained by generating the matrix Γ_v containing the volume coherences from (3.122) for all channel pairs observing the given scene, and then computing $\alpha_{v,\min}$ via (3.109). The curves in (a) have been cropped to exclude cases where the condition number of Γ_v , shown in (b), is greater than 10^6 . The role of the matrix condition number is discussed in Section 3.10.4.



near the top of the volume, which is just the kind of interference that the beamformer is best-suited to nulling i.e. interference that is (1) narrow and (2) far from the signal source. In the special case when there is no propagation loss and the vertical structure is uniform (from (3.123), $\sigma_e = 0 \Rightarrow f_v(z') = 1$), the effective width is always the full volume extent, and as the volume height increases, the same total interference power is just spread more thinly over a wider extent, so the beamformer performance is constant.

Comparison of the required volume attenuation in Figure 3.17 with the (best-case) achievable attenuation in Figures 3.23(a) and 3.24(a) shows that it is indeed feasible to do coherent change detection of the ground under a random volume-like forest canopy using only a handful of across-track channels.

The beam patterns in Figure 3.25 show how different beamformers suppress the volume. For the three-channel case in (a), the optimal RVOG beamformer (red), given by (3.127), simply steers two nulls inside the main lobe of the conventional beamformer (blue), just as was done explicitly to obtain the second null-steer beamformer shown in Figure 3.9. The difference is that the optimisation in (3.104)–(3.108) provides a recipe

for where to steer the nulls optimally, given knowledge of the RVOG parameters. For the nine-channel case in (b), the optimal beampattern (red) of the RVOG beamformer approximates the shape of the ideal notch filter (black), except that there is very large gain outside the scene extent. The volume attenuations achieved by the different beamformers are listed in Table 3.4.




Figure 3.25: Vertical beampatterns for conventional, null-steer and optimal (RVOG) beamformers. Collection parameters: (a) $N_{ch} = 3$, $\Delta\psi = 0.05^\circ$; (b) $N_{ch} = 9$, $\Delta\psi = 0.15^\circ$ ($\bar{\psi} = 35^\circ$ and $\lambda = 23$ cm for both). Scene parameters: $h_v = 20$ m, $\sigma_e^{\text{dB}} = 0.1$ dB/m. The null-steer beamformer (green) was constructed using control points $\{\mathbf{z}_c = [0, 13], \mathbf{a}_c = [1, 0]\}$, since 13 m is the approximate height of the volume's effective phase centre at these geometries. See Table 3.4 for values of the α_v metric. Compare the beampatterns to those in Figure 3.9.

Table 3.4: Multichannel volume attenuation factor α_v (in decibels), given by (3.100), for the beamformers shown in Figure 3.25.

beamformer	colour	$N_{ch}=3$	$N_{ch}=9$
conventional	blue	-1.8	-13.1
null-steer (single null at $h_{v,\text{eff}} \approx 13$ m from (3.68))	green	-7.4	-13.0
optimal RVOG ($\alpha_v = \alpha_{v,\text{min}}$ in (3.109))	red	-12.1	-21.2

3.10.2 Coherence estimation error

In the preceding section, beamformer performance was assessed in terms of volume attenuation within a pass. However, the ultimate test of performance is the error when estimating the ground coherence across passes, since this is the purpose of the whole exercise. From (3.102), if the repeat-pass collection geometries are identical i.e. $\psi_{a_i} = \psi_{b_i} \forall i = 1, 2, \dots, N_{ch}$, then the volume attenuation factor α_v is an appropriate proxy for coherence estimation error. In general though, it is necessary to explicitly evaluate the multichannel coherence $\gamma_{y_{ab}}$ and the consequent error $\epsilon_{\gamma_g} = |\gamma_{y_{ab}}| - \gamma_g$ over the full range $0 \leq \gamma_g \leq 1$ of possible ground coherence values.

In order to evaluate the multichannel coherence, multichannel RVOG data is generated. This data takes the multichannel dual-layer form given in (3.89)–(3.91), except that, for simplicity, it contains no distortion, no noise, and no ground interferometric phase (i.e. the focal surface matches the ground surface):

$$\Gamma_a = \frac{\mu}{1+\mu} \mathbf{1} + \frac{1}{1+\mu} \Gamma_{v_a} \quad (3.128)$$


$$\Gamma_b = \frac{\mu}{1+\mu} \mathbf{1} + \frac{1}{1+\mu} \Gamma_{v_b} \quad (3.129)$$

$$\Gamma_{ab} = \frac{\mu}{1+\mu} \gamma_g \mathbf{1} + \frac{1}{1+\mu} \Gamma_{v_{ab}}. \quad (3.130)$$

Note that it is not assumed that $\Gamma_{v_a} = \Gamma_{v_b}$ (see (3.93) for this approximation). The volume coherence matrices Γ_{v_a} , Γ_{v_b} and $\Gamma_{v_{ab}}$ are populated according to the random volume model by evaluating (3.122) for each pair of channels. For a given weight vector, the multichannel coherence $\gamma_{y_{ab}}$ is computed using (3.56) (not (3.57)).

Figure 3.26 shows how the optimal RVOG beamformer performs for different repeat-pass grazing-angle offsets. When the offset is zero (dark blue), the coherence estimation error is greatest at low ground coherences—the highly coherent volume masks changes on the ground. In this case, the error is given explicitly by (3.102). Given $\mu = -2.5$ dB, the maximum error (at $\gamma_g = 0$) is 0.1; this is the collection scenario discussed in the previous section. As the repeat-pass offset increases, the coherence error decreases for low γ_g but increases for high γ_g , quickly reaching the point where the error is greatest at high ground coherences—the low-coherence volume masks areas where the ground has not changed. Observe that the single-channel estimates (dash-dotted lines) that would be obtained by ordinary CCD are wildly in error over most of the range of γ_g . By suppressing the volume, the 3D CCD processing clearly provides reasonably accurate estimates of ground coherence over the full range of γ_g , at least for this level of ground-volume power balance μ . A repeat-pass offset of 1° is probably the worst-case upper limit for airborne change detection.

Figure 3.26: Ground coherence estimates $\hat{\gamma}_g$ as a function of the true ground coherence γ_g for different repeat-pass grazing angle offsets $\bar{\psi}_b - \bar{\psi}_a$. The black dashed line is the ideal. The solid lines show the multichannel estimates $\hat{\gamma}_g = |\gamma_{y_{ab}}|$ ($N_{ch} = 3$, $\Delta\psi = 0.05^\circ$) after determining the optimal RVOG beamformer \mathbf{w}_{rvog} in (3.127), generating multichannel RVOG data of the form in (3.128)–(3.130), and evaluating the multichannel coherence $\gamma_{y_{ab}}$ in (3.56). The dash-dotted lines show the single channel estimates $\hat{\gamma}_g = |\gamma_{ab}|$ given by (3.73) (with $\gamma_{SNR} = 1$, $\phi_{gd} = 0$). For the three repeat-pass pairs, the centre grazing angle of the first pass is fixed ($\psi_a = 35^\circ$), but ψ_b varies: 35° for dark blue, 35.4° for blue and 36° for cyan, giving rise to β_v values: $1 \angle 0^\circ$, $0.68 \angle 48.5^\circ$ and $0.24 \angle 83.5^\circ$, respectively. Scene parameters: $h_v = 20$ m, $\sigma_e^{\text{dB}} = 0.1$ dB/m and $\mu = -2.5$ dB. ($\lambda = 23$ cm, $z_g = 0$)



Hence, optimising for volume suppression (α_v) within each pass is a reasonable beamforming strategy, even when the passes are not identical. Having $\beta_v \neq 1$ (which, from (3.98), occurs when the volume coherences within a pass (Γ_v) are not equal to those across passes ($\Gamma_{v_{ab}}$)), is not detrimental to coherence estimation performance.

3.10.3 Sensitivity to assumed model parameters

The optimal RVOG beamformer in (3.127) depends on the parameters of the particular model random volume it is designed to suppress. In order to implement such a beamformer, knowledge of these parameters would be required, but realistically, this knowledge will always be limited. Indeed, for a real forest scene, the assumptions are two-fold: firstly, that the canopy is modelled reasonably well by a random volume, and secondly, that the canopy is modelled reasonably well by the particular random volume specified by the assumed parameter values. In this section, the beamformer's sensitivity to the second type of assumption is tested.

Consider a reference RVOG scene with volume height $h_v = 20$ m, attenuation rate $\sigma_e^{\text{dB}} = 0.1$ dB/m and ground-volume power ratio $\mu = 0$ dB (given three channels) or $\mu = -10$ dB (given nine channels). The optimal volume attenuation factor $\alpha_{v,\text{min}}$ given by (3.109) is -12 dB for three channels (at angular spacing $\Delta\psi = 0.05^\circ$) and -21 dB for nine channels ($\Delta\psi = 0.15^\circ$). This nominal operating point was highlighted in the discussion of Figures 3.23 and 3.24 in Section 3.10.1.

Somewhat arbitrarily, say that the coherence estimation error should be kept below 0.1. Given the reference values of μ stated above, volume attenuation of at least $\alpha_v = -9.5$ dB given three channels and $\alpha_v = -19.5$ dB given nine channels would

achieve this in the ideal change-detection case of identical repeat-pass geometries, according to (3.103). This sets the required performance.

Figure 3.27 shows the beamforming performance of the RVOG weight vector designed for the reference random volume when applied to different random volumes. Performance here is quantified by the effective volume attenuation factor α_v , given by (3.100). Broadly, red indicates poor performance, green indicates borderline acceptable performance, and blue indicates good performance. The black cross indicates the reference design point and the black loop marks the operating region for which the combination of volume height and attenuation rate permits acceptable performance.




Figure 3.27: Effective volume attenuation when the optimal RVOG weight vector in (3.127) designed for one model random volume is applied to scenes with different volume heights h_v and different attenuation rates σ_e^{dB} . The colours map to decibel values of the multichannel volume attenuation factor α_v , computed via (3.100), where the matrix Γ_v of volume coherences was generated for each $(h_v, \sigma_e^{\text{dB}})$ combination using (3.122). The black cross indicates the design point: 20 m, 0.1 dB/m. The black loop marks the region where α_v is below the (arbitrary) desirable threshold specified by $\alpha_{v,\text{accept}}$. Fixed parameters: $\bar{\psi} = 35^\circ$, $\lambda = 23$ cm.

It is clear that the RVOG beamformer is not very sensitive to attenuation rate. In fact, performance improves (i.e. α_v decreases) as the attenuation rate increases, going well beyond the optimal value $\alpha_{v,\text{min}}$, which is applicable only when the beamformer is matched to the scene. In beamforming terms, as the attenuation rate increases, the volume interference becomes more localised in height, so the nulls steered by the beamformer would more effectively cancel the volume response (see Figure 3.24(a)).

In the three-channel case shown in Figure 3.27(a), the RVOG beamformer is seen to be not overly sensitive to volume height. However, in the nine-channel case shown in Figure 3.27(b), performance drops rapidly for volume heights above the design point. This accords with the beamformer's beampattern, shown in Figure 3.25(b): the trade-off for the large attenuation achieved by steering $N_{ch} - 1$ nulls onto the volume (inside the mainlobe of the conventional beampattern) is large gain for any scene content outside the expected vertical extent. This suggests that, if many channels are available (more than three), it is safer to err on the high side when determining the volume height for which the beamformer is to be designed. For example, redesigning for $h_v = 24$ m in (b) would ensure acceptable performance up to 25 m with only a small performance penalty for heights below 20 m.

Importantly, the random volume parameters determine both the vertical structure for beamforming (relevant to α_v) and the propagation loss (captured in the ground-volume power ratio μ). Using the volume attenuation factor α_v as a performance indicator here is misleading, because it ignores the change in μ given by (3.125). Intuitively, a change in volume height or attenuation rate would change the energy level reaching the ground, thus changing the coherence estimation accuracy, but α_v does not depend on μ , so it is not a valid measure of overall performance unless μ is fixed. In this case, the performance is fully characterised by computing the actual coherence estimation error ϵ_{γ_g} given by (3.102) when $\gamma_g = 0$, thereby accounting for

both α_v and μ . Figure 3.28 shows the result of doing just this—it effectively combines Figure 3.27, which shows the beamformer response ignoring propagation, with Figure 3.21, which shows how the random volume parameters determine the propagation loss.




Figure 3.28: Coherence estimation error ϵ_{γ_g} when the optimal RVOG weight vector in (3.127) designed for one model random volume is applied to scenes with different volume heights h_v and different attenuation rates σ_e^{dB} . The colours map to ϵ_{γ_g} , computed via (3.102) for $\gamma_g = 0$, with the stated ground-volume power ratio μ at the design point modified for each $(h_v, \sigma_e^{\text{dB}})$ to suit the new propagation factor $|p|^2 = f_v(0)$ in (3.124) (recall that μ and $|p|^2$ are directly proportional). The black cross indicates the design point: 20 m, 0.1 dB/m and $\mu = 0$ dB for three channels or $\mu = -10$ dB for nine channels. The black loop marks the region where ϵ_{γ_g} is below the desirable threshold of 0.1. Fixed parameters: $\bar{\psi} = 35^\circ$, $\lambda = 23$ cm, $z_g = 0$.

It is now clear that as the attenuation rate increases from its assumed value at the design point, overall performance does not improve *ad infinitum*, because even though the volume becomes more localised in height, the ground signal also becomes weaker. On the other hand, as the attenuation rate decreases, μ increases, and this effect ensures that the coherence estimation error remains low, even though the beamformer is less effective at volume suppression.

For three channels, performance will be acceptable if the true volume height is within about three metres of the design height. In other words, the acceptable level of uncertainty in knowledge of the volume height is ± 3 m. The inversion methods of polarimetric SAR interferometry are able to estimate canopy height with this level of accuracy, as discussed in Section 3.1.2.

Overall, with appropriate precautions, the optimal RVOG beamformer is not particularly sensitive to the random volume parameter values for which it is designed. This makes it reasonable to apply this clairvoyant approach to real data, for which knowledge of the scene can only ever be approximate.

However, mismatch of the assumed and actual parameter values is not the only way in which the collected data could deviate from the selected RVOG model. Real hardware will impose noise, particularly phase noise, that limits the precision of the measured coherences. The next section analyses how sensitive different radar configurations are to noise-like perturbations.

3.10.4 Sensitivity to perturbations using the matrix condition number

Recall the processing diagram in Figure 3.2. As formalised in Section 3.5, the selected weight vectors \mathbf{w}_a and \mathbf{w}_b will be applied to the multichannel SAR image data vectors \mathbf{x}_a and \mathbf{x}_b , respectively, and the repeat-pass coherence computed according to (3.54). The derivation in Section 3.8.1 led to one possible choice of weight vector: one that is optimal within the confines of the dual-layer multichannel coherence model developed in Sections 3.5–3.7. The RVOG model described in Section 3.9 provides a way to predict what kind of coherence estimates might be observed.

Beyond the assumptions of these models, the measured data will be subject to perturbations from numerous sources, including noise from the RF receivers and processing errors due to uncertainties in the position sensors. Such perturbations

cause beamformer performance to degrade. A method is sought by which beamformers can be assessed and compared in terms of their sensitivity to perturbations. Note that this is not just a matter of choosing a particular weight vector during processing, but is essential to the task of radar design, since the sensitivity of the beamformer is critically dependent on the geometry of the array, particularly the number of channels and their angular spacing.

In this section, different array geometries are compared in terms of the volume attenuation achievable by the corresponding optimal RVOG beamformers, given volume coherences that nominally fit the RVOG model but are subject to phase noise, as a typical kind of perturbation. Before that though, the matrix condition number is introduced as a generic measure of sensitivity to perturbations. This parameter was plotted for different array geometries in Figure 3.23(b) and for different RVOG scene parameters in Figure 3.24(b).

The qualitative ‘condition’ of a functional mapping $f : x \rightarrow y = f(x)$ is a description of how sensitive f is to small perturbations in x : if y changes by a small amount, then the mapping is well-conditioned, but if y changes by a large amount, then the mapping is ill-conditioned [Trefethen & Bau 1997 p. 89].

This can be formalised for the ubiquitous problem of solving a system of linear equations $A\mathbf{x} = \mathbf{b}$, where A is invertible. First note that, for any vector norm $\|\cdot\|$ and compatible matrix norm $\|\|\cdot\|\|$ [Horn & Johnson 2013 pp. 343–344],

$$A\mathbf{x} = \mathbf{b} \Rightarrow \|\|A\|\|\|\mathbf{x}\| \geq \|A\mathbf{x}\| = \|\mathbf{b}\|. \quad (3.131)$$

Given a solution $\hat{\mathbf{x}}$, which has error $\mathbf{e} = \hat{\mathbf{x}} - \mathbf{x}$ and residual $\mathbf{r} = A\hat{\mathbf{x}} - \mathbf{b}$, an upper bound on the relative error of $\hat{\mathbf{x}}$, relative to the relative size of the residual, can be derived straightforwardly [Horn & Johnson 2013 p. 384]:

$$\begin{aligned} \|\mathbf{e}\| &= \|\hat{\mathbf{x}} - \mathbf{x}\| = \|A^{-1}A\hat{\mathbf{x}} - A^{-1}\mathbf{b}\| = \|A^{-1}\mathbf{r}\| \\ &\leq \|\|A^{-1}\|\|\|\mathbf{r}\| \\ &\leq \|\|A^{-1}\|\|\|\mathbf{r}\| \frac{\|\|A\|\|\|\mathbf{x}\|}{\|\mathbf{b}\|} \\ \therefore \frac{\|\mathbf{e}\|/\|\mathbf{x}\|}{\|\mathbf{r}\|/\|\mathbf{b}\|} &\leq \kappa(A) \end{aligned} \quad (3.132)$$

where the bound

$$\kappa(A) = \|\|A^{-1}\|\|\|\|A\|\| \geq 1 \quad (3.133)$$

is called the condition number of the matrix A . A well-conditioned matrix has a small condition number, so that both the relative error of the solution $\hat{\mathbf{x}}$ and the relative size of the residual \mathbf{r} will be similar, but an ill-conditioned matrix has a large condition number, so that $\hat{\mathbf{x}}$ may have significant error even though it was selected to make the residual small. Note the ‘may’ in the preceding sentence— $\kappa(A)$ is a conservative bound [Horn & Johnson 2013 p. 384].

The same condition number can be used to bound the error matrix $E = (A + \delta A)^{-1} - A^{-1}$ when directly computing the inverse of A suffering perturbation δA [Horn & Johnson 2013 p. 382]:

$$\frac{\|\|E\|\|\|\|A^{-1}\|\|}{\|\|\delta A\|\|\|\|A\|\|} \leq \frac{\kappa(A)}{1 - \kappa(A)\|\|\delta A\|\|\|\|A\|\|} \quad (3.134)$$

In this case, the inverse of an ill-conditioned matrix may have a large relative error even if the perturbation of the matrix is relatively small. A simple indicator that a matrix may be ill-conditioned is that some of the elements of its inverse are large compared to the elements of the matrix itself [Horn & Johnson 2013 p. 384]. If A is singular, then $\kappa(A)$ is infinity [Horn & Johnson 2013 p. 382].

Let the condition number be computed using the spectral norm, which is the matrix equivalent of the Euclidean two-norm for vectors [Horn & Johnson 2013 p. 346]. The condition number is now simply [Trefethen & Bau 1997 p. 94]


$$\kappa(A) = \frac{\sigma_{\max}}{\sigma_{\min}} \quad (3.135)$$

where σ_{\max} (σ_{\min}) is the largest (smallest) singular value of A . (If A is positive semi-definite, as covariances matrices are [Horn & Johnson 2013 p. 426], then its singular

values equal its eigenvalues [Horn & Johnson 2013 pp. 156,430].) This is what Matlab's `cond()` function returns. Interpreting A as a geometric transformation of the unit sphere (in \mathbb{C}^N) into an ellipse (in \mathbb{C}^N), $\kappa(A)$ in (3.135) is the ellipse's eccentricity [Trefethen & Bau 1997 pp. 94–95].

Trefethen & Bau [1997 p. 95] assert that solving a system of linear equations $A\mathbf{x} = \mathbf{b}$ entails losing $\log_{10} \kappa(A)$ digits of precision. The numerical experiment shown in Figure 3.29 supports this indirectly: the size of the residual error when computing the inverse of a matrix is linearly proportional to $\log_{10} \kappa(A)$. Hence, the condition number is indeed a good predictor of the precision loss that could arise when using a matrix operator.

Figure 3.29: Size of the residual error $R = AA^{-1} - I$ after computing the matrix inverse in Matlab compared with the condition number $\kappa(A)$. The size was computed using the spectral norm. Many realisations of a 5×5 matrix $A = \mathbf{1} + \sigma_n^2 Z$ were generated, where $\mathbf{1}$ was a fixed matrix of ones (singular $\Rightarrow \kappa(\mathbf{1}) = \infty$), Z was a complex random matrix whose elements were taken from a standard normal distribution ($\kappa(Z) \sim 10$) and σ_n^2 was a notional noise power (i.e. a scalar) that was varied between 1 and 10^{-13} in order to vary $\kappa(A)$. Each point is for one realisation of A .



In general, the loss of precision may be due to either numerical round-off or an accumulation of uncertainties in the data—both can be modelled as perturbations of the true input matrix [Horn & Johnson 2013 p. 381]. When solving a problem, one should consider how a potential algorithm degrades the available error budget, given the assumed precision of the input measurements and the desired precision of the output solution.

Reconsider $\mathbf{w}_{\text{rvog}} = \Gamma_v^{-1} \mathbf{v}_0$ in (3.127) and $\mathbf{w}_{\text{mvdR}} = \hat{R}^{-1} \mathbf{v}_0$ in (3.27) (neglecting scale factors): they can be rewritten in the general form $A\mathbf{x} = \mathbf{b}$ as $\Gamma_v \mathbf{w}_{\text{rvog}} = \mathbf{v}_0$ and $\hat{R} \mathbf{w}_{\text{mvdR}} = \mathbf{v}_0$. The latter fits the usual task of solving for \mathbf{x} given measured data in A , whereas the former instead involves determining the optimal solution \mathbf{x} using assumed knowledge of A before applying it to measured versions of A . In both cases, say that there is an ideal A that represents the truth, and a measured version A' that approximates A , subject to perturbations. For \mathbf{w}_{rvog} , the ideal is assumed, whereas for \mathbf{w}_{mvdR} , the ideal is what would be measured if there were no perturbations, including no estimation error. Mismatch between A and A' (i.e. between Γ_v and Γ'_v , or between R and \hat{R}) will cause a degradation in performance. For \mathbf{w}_{rvog} in particular, the degradation is quantified by a reduction in the achieved volume attenuation as given by (3.100) using Γ'_v in place of Γ_v . For a fixed Γ_v , how sensitive is the performance to the size of the perturbation? For a fixed level of perturbation, how sensitive is the performance given different Γ_v ?

Figure 3.30 provides answers to these questions for the case when the perturbation takes the form of phase noise, so the interferometric phases of the volume coherences in the off-diagonal elements of Γ_v are subject to random fluctuations. For the parameters under consideration (the same as those used in the three-channel case in Figure 3.25), the ideal volume attenuation factor (i.e. α_v when there is zero perturbation) is around -12 dB, as discussed in Section 3.8.1. As the level of perturbation increases (i.e. the phase noise applied to Γ_v has a higher variance), the average effective volume attenuation reduces (i.e. the factor $\bar{\alpha}_v$ moves closer to 0 dB). Most importantly, this reduction in attenuation is strongly dependent on the ideal Γ_v , as parameterised by the grazing angle spacing in (a) or the matrix condition number $\kappa(\Gamma_v)$ in (b): as the grazing angle spacing decreases, or more generally, the condition number increases, the average effective volume attenuation can reduce dramatically. Intuitively, as the array configuration becomes more extreme, the optimal beamformer becomes more sensitive to deviations from the assumed ideal. This sense of extremeness is quantified by the matrix condition number.

The same trend applies to the adaptive beamformer, but the interpretation is slightly different: since the matrix of interest (\hat{R}) must now be estimated, a significant contribution to the perturbation is likely to be estimation error. Thus, as the radar configuration becomes more extreme, the sensitivity to error increases, necessitating




Figure 3.30: Mean effective multichannel volume attenuation factor $\bar{\alpha}_v$ as a function of (a) grazing angle spacing $\Delta\psi$ and (b) condition number $\kappa(\Gamma_v)$ of the matrix Γ_v of volume coherences, for different levels of random phase-noise perturbation of Γ_v . Each point on each curve was obtained by first generating the set of model volume coherences in Γ_v (given $\Delta\psi$) via (3.122), along with the optimal RVOG weight vector \mathbf{w}_{rvoG} in (3.127), and then repeatedly computing the effective volume attenuation according to (3.100) when \mathbf{w}_{rvoG} is applied to many (~ 1000) realisations of the perturbed matrix $\Gamma'_v = \Gamma_v \circ \Gamma_{phn}$, where the perturbation $\Gamma_{phn} = [(e^{j\phi_{phn}})_{ij}]$ is a Hermitian matrix whose $(i, j)^{\text{th}}$ element models noise on the interferometric phase for channel pair (i, j) (if $i = j$, $\phi_{phn} = 0$; across pairs, the noises are independent). The phase noise samples are taken from a von Mises distribution (the circular analogue to the Gaussian distribution, suitable for modelling 2π -wrapped random phase noise [Berens 2009]), with standard deviation specified by σ_{phn} in degrees. Note that (a) and (b) plot the same data i.e. the same mean values of α_v obtained by averaging across the realised perturbations for each radar configuration $\Delta\psi$; the matrix condition number in (b) just provides a more general characterisation of Γ_v than the angular spacing in (a). Fixed parameters: $N_{ch} = 3$, $\bar{\psi} = 35^\circ$, $\lambda = 23$ cm, $h_v = 20$ m and $\sigma_e^{\text{dB}} = 0.1$ dB/m.

more accurate estimation.

Figures 3.23(b) and 3.24(b) show how the matrix condition number $\kappa(\Gamma_v)$ varies for different input parameters to Γ_v . Observe that the condition number increases rapidly as the grazing angle spacing decreases or, for a fixed spacing, as the number of channels increases. The variation with scene parameters is less dramatic, but it is clear that shorter volumes and higher attenuation rates give rise to larger condition numbers.

Comparison of Figures 3.23 and 3.30 reveals that there is a trade-off between performance and reliability when choosing the grazing angle spacing. As $\Delta\psi$ decreases and $\kappa(\Gamma_v)$ increases, the potential volume attenuation $\alpha_{v,\text{min}}$ in the ideal case improves (although the rate of change is small for a small number of channels), but the average effective volume attenuation $\bar{\alpha}_v$ in the perturbed case degrades. In general, larger condition numbers are associated with both higher performance [d'Assumpcao & Mountford 1984] and greater sensitivity to perturbations.

The radar designer must select a condition number that permits both useful and reliable performance. Consider a value of $\log_{10} \kappa(\Gamma_v) = 3.25$ as a potentially suitable compromise. From Figure 3.23(b), this is approximately achieved by the radar configurations ($N_{ch} = 3$, $\Delta\psi = 0.05^\circ$) and ($N_{ch} = 9$, $\Delta\psi = 0.18^\circ$), given typical parameters $\lambda = 23$ cm, $h_v = 20$ m and $\sigma_e^{\text{dB}} = 0.1$ dB/m. For the three-channel case shown in Figure 3.30(b), if phase noise was present with standard deviation $\sigma_{phn} = 5^\circ$, then the volume attenuation factor would degrade from -12 dB to about -9.5 dB, which according to (3.103) would limit accurate ground coherence estimation to cases where the ground-volume power ratio μ was at least 0 dB. This is chosen as an acceptable operating point in terms of both performance and sensitivity to perturbations.

3.10.5 Sensitivity to ground height

The analysis in Section 3.10 thus far has assumed that $z_g = 0$ i.e. that the focal surface at $z = 0$ onto which the SAR images are formed matches the ground scattering surface at $z = z_g$. In reality, the height of the ground in a forest would be known only roughly.

A vertical offset $z_g \neq 0$ would give rise to interferometric phases $\{\phi_{g_{ij}}\}$ given by (3.19) for the different channel pairs (i, j) . In the multichannel dual-layer coherence model in (3.89)–(3.91), these surface interferometric phases are contained in the Φ_{gd} matrices. How do these phases affect the output coherence estimate?

Figure 3.31 plots the variation of the multichannel coherence $|\gamma_{y_{ab}}|$ with the offset z_g for a particular collection scenario. This variation is periodic with period given by the unambiguous extent h_{amb} sampled by the array. In the ideal case when the offset is zero, $|\gamma_{y_{ab}}|$ is a reasonably good estimator for the true ground coherence γ_g , as discussed previously. However, even a very small non-zero offset ($|z_g| \gtrsim 2\text{m}$) induces interferometric phases that significantly distort the coherence estimate. Observe in particular that this distortion varies with γ_g : for example in (d), if $\gamma_g \sim 0$, $|\gamma_{y_{ab}}|$ can vary up to 0.84, whereas if $\gamma_g \sim 1$, $|\gamma_{y_{ab}}|$ can vary down to 0.46. The variation is not symmetric about $z = 0$. Clearly, a key requirement for 3D SAR CCD is accurate knowledge of the ground height in order to focus the SAR images onto the true ground surface and thereby avoid (or at least minimise) this distortion. Obtaining this terrain information is one of the most significant obstacles to 3D SAR CCD.

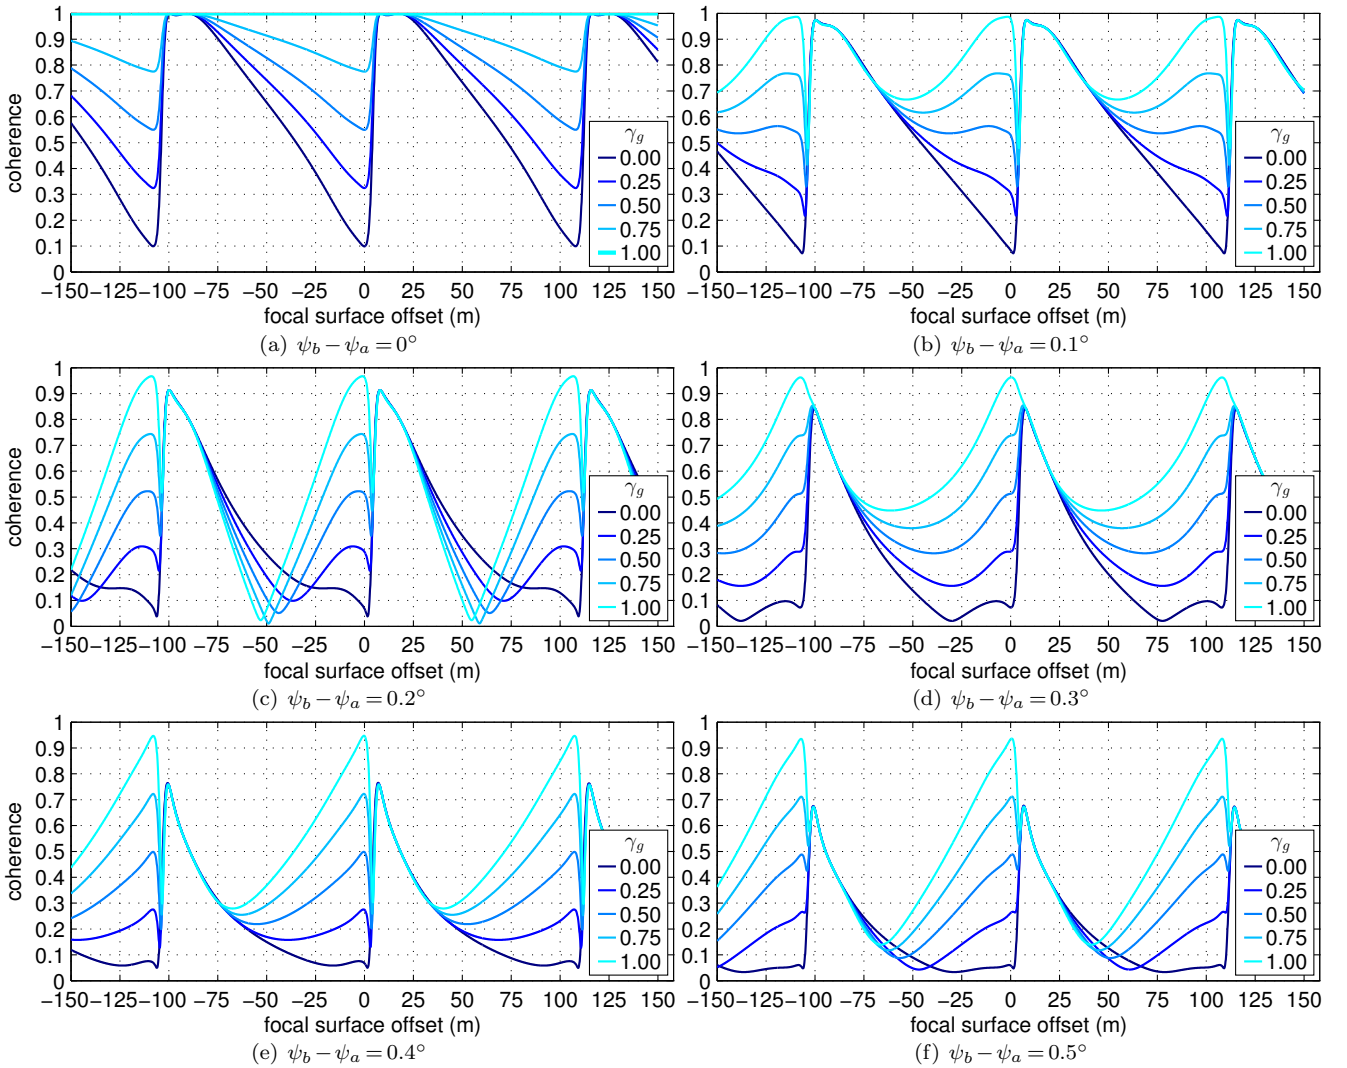


Figure 3.31: Output coherence magnitude $|\gamma_{y_{ab}}|$ as a function of the vertical offset between the SAR image focal surface (at $z=0$) and the ground surface (at $z=z_g$), for different values of the true ground coherence γ_g . Each point on each curve was obtained by generating multichannel RVOG data in the form given in (3.89)–(3.91), with the volume coherence matrices populated by evaluating (3.122) for each pair of channels, and then computing the multichannel coherence via (3.56) using the RVOG beamformer (fixed for all surface offsets). Scene parameters: $h_v=20\text{m}$, $\sigma_e^{\text{dB}}=0.1\text{dB/m}$ and $\mu=-2.5\text{dB}$. Collection parameters: $\lambda=23\text{cm}$, $N_{ch}=3$, $\Delta\psi=0.05^\circ$, $\psi_a=35^\circ$, ψ_b varied, and $\gamma_{SNR}=1$; the height of ambiguity given by (3.41) is $h_{amb}=108\text{m}$.

In beamforming terms, the beamformer is vertically steered to the focal surface by

the broadside steering vector \mathbf{v}_0 in (3.25), which is a vector of ones—this explicitly assumes that there are no phase differences between the channels. Any scattering not at the focal surface will be attenuated; indeed, the optimal beamformers (MVDR, RVOG etc) steer a null onto the off-height ‘interference’, which in this case is the desired ground signal. This ‘self-cancellation’ is a classic problem exhibited by optimal beamformers in general [Lombardini *et al.* 2009, Van Trees 2002 ch. 6.6]; variants of the adaptive MVDR technique have been proposed that selectively trade-off resolution for robustness [Lorenz & Boyd 2005, Gershman *et al.* 2006]. Note that if the true ground height was known, then the steering vector could be changed to account for the expected interferometric phases between channels, but this is irrelevant, because if the true ground height was known, then the SAR images could be focused onto the correct surface in the first place.

Numerous methods to estimate the ground height in a forest scene have been proposed in the literature [Homer & Longstaff 1995, Treuhaft & Siqueira 2004], typically using multiple interferometric and/or polarimetric channels, as exemplified by PolInSAR [Hajnsek *et al.* 2009b, Cloude 2010 ch. 8.1]. Given the data collection mode shown in Figure 3.1, an obvious approach is to use one multichannel dataset to image the vertical dimension via 3D SAR beamforming i.e. measure the vertical variation in scattering intensity by scanning beams in height [Reigber *et al.* 2005]. The ground height could then be set simply as the height of the bottom of the vertical band of scattering (assuming that the level of penetration to ground level is reasonable). For a typical RVOG scene illuminated with just three across-track channels, a height-range image generated by MVDR beamforming would look like the image in Figure 3.13(f), with a height resolution of the order of 5 m (see also the intensity-height plot in Figure 3.14). For ground-height estimation (as opposed to reflectivity estimation), it may be reasonable to apply subspace beamforming techniques that offer finer resolution [Reigber *et al.* 2015]. Analysis of these methods is beyond the scope of this project.

Nonetheless, it is worth noting that the extreme sensitivity of the multichannel coherence to the difference between the ground height and the focal height can be turned to our advantage, by using it to estimate the ground height, or at least refine a coarse initial estimate obtained in some other way. That is, instead of using the vertical variation in scattering intensity given one multichannel acquisition, the vertical variation in multichannel coherence between two multichannel acquisitions could be used. Observe in Figure 3.31 that the peak coherence (for $\gamma_g = 1$) is achieved when the surface height offset is zero or some multiple of the array height ambiguity (absolute height determination is not required for CCD). Therefore, given a repeat-pass pair of multichannel datasets, by repeatedly forming the SAR images at different focal heights, beamforming the images from each pass, and computing their coherence, the ground height can be found as the focal height for which the output coherence magnitude is maximised. This approach assumes that the ground is a source of high coherence (i.e. $\gamma_g = 1$), and that the volume is a source of low coherence (if the collection geometries were identical, as in Figure 3.31(a), then this method would not work). This would seem to go against the wider purpose of the CCD application, which is to estimate the ground coherence. However, for the purpose of ground-height estimation, the radar platform could make repeat passes over a scene in quick succession, such that it is reasonable to assume that the ground has not changed. For example, an operational system could initially make two passes over a scene, separated by, say, one hour, and then make subsequent passes over the scene once per day. The first two passes would be used to estimate the ground height by maximising the output coherence, and the subsequent passes would be used for the actual CCD application of finding areas of low ground coherence, where the ground scattering has changed.

3.11 3D SAR CCD simulation

The 3D SAR CCD concept is now demonstrated by simulation. This is the same simulation used throughout this thesis to demonstrate the proposed processing techniques. A key feature of the simulation is that the scattering response of clutter is synthesised at the raw pulse data level such that it matches the RVOG structural model and the resulting imagery matches the standard statistical model for SAR data.

Table 3.5 lists the main parameters of the 3D SAR CCD simulation. The grazing

Table 3.5: Parameters of the 3D SAR CCD simulation in this section. ('az' is the azimuth (along-track) direction and 'gr' is the ground-range (across-track) direction.)

operating frequency (wavelength λ)	1.32 GHz (22.7 cm)
chirp bandwidth	140 MHz
aperture	7.4°, 296 pulses
window	none (i.e. rect)
image resolution (az \times gr)	1 m \times 1.2 m
pixel spacing (az \times gr)	10 cm \times 10 cm (5 \times 5 pixel average for display)
swath (az \times gr)	180 m \times 120 m (cropped to 180 m \times 80 m for display)
volume height h_v	20 m
attenuation rate σ_e^{dB}	0.1 dB/m
ground-volume power ratio μ	0 dB
number of scatterers	216,000 (10/m ²) on ground 2,160,000 (5/m ³) in volume
ground change between passes	random shift of scatterers under text mask 12 cm standard deviation 2 m text line width
coherence average (az \times gr)	3.1 m \times 3.5 m (\sim 9 independent looks)
altitude (pass a , pass b)	6200 ft, 6210 ft
nominal range (pass a , pass b)	3.295 km, 3.276 km
nominal grazing angle (ψ_a, ψ_b)	35°, 35.3°
channel spacing	5 m across-track ($\Delta\psi \approx 0.05^\circ$)
number of channels per pass N_{ch}	3
raw data demodulation	matched filter
image formation	time-domain beamforming (see chapter 2)

angles of the individual channels are

$$\psi_a \approx \begin{bmatrix} 34.95^\circ \\ 35^\circ \\ 35.05^\circ \end{bmatrix}, \quad \psi_b \approx \begin{bmatrix} 35.25^\circ \\ 35.3^\circ \\ 35.35^\circ \end{bmatrix}. \quad (3.136)$$

Where noted, the array is modified by either doubling the channel spacing to 10 m, which doubles the angular spacing to approximately 0.1°, or increasing the number of channels to five.

Given the geometry and scene parameters listed in Table 3.5, the volume-only coherences and the total RVOG coherences can be computed via (3.122) and (3.72), respectively, for all pairs of channels within and across passes, assuming no scene change, no distortion, no noise, and no ground interferometric phase (i.e. the focal surface is matched to the ground surface). The resulting coherence matrices in (3.89)–(3.91) are

$$\begin{aligned} \Gamma_{v_a} &= \begin{bmatrix} 1\angle 0^\circ & 0.950\angle +42.6^\circ & 0.811\angle +86.2^\circ \\ 0.950\angle -42.6^\circ & 1\angle 0^\circ & 0.950\angle +42.4^\circ \\ 0.811\angle -86.2^\circ & 0.950\angle -42.4^\circ & 1\angle 0^\circ \end{bmatrix} & \Gamma_a &= \begin{bmatrix} 1\angle 0^\circ & 0.908\angle +20.7^\circ & 0.665\angle +37.5^\circ \\ 0.908\angle -20.7^\circ & 1\angle 0^\circ & 0.909\angle +20.7^\circ \\ 0.665\angle -37.5^\circ & 0.909\angle -20.7^\circ & 1\angle 0^\circ \end{bmatrix} \\ \Gamma_{v_b} &= \begin{bmatrix} 1\angle 0^\circ & 0.948\angle +43.2^\circ & 0.805\angle +87.6^\circ \\ 0.948\angle -43.2^\circ & 1\angle 0^\circ & 0.948\angle +43.1^\circ \\ 0.805\angle -87.6^\circ & 0.948\angle -43.1^\circ & 1\angle 0^\circ \end{bmatrix} & \Gamma_b &= \begin{bmatrix} 1\angle 0^\circ & 0.906\angle +21.0^\circ & 0.655\angle +37.9^\circ \\ 0.906\angle -21.0^\circ & 1\angle 0^\circ & 0.906\angle +21.0^\circ \\ 0.655\angle -37.9^\circ & 0.906\angle -21.0^\circ & 1\angle 0^\circ \end{bmatrix} \\ \Gamma_{v_{ab}} &= \begin{bmatrix} 0.241\angle -20.4^\circ & 0.261\angle +45.2^\circ & 0.251\angle +101.2^\circ \\ 0.267\angle -101.5^\circ & 0.241\angle -21.7^\circ & 0.261\angle +44.0^\circ \\ 0.411\angle -172.5^\circ & 0.268\angle -102.6^\circ & 0.241\angle -22.9^\circ \end{bmatrix} & \Gamma_{ab} &= \begin{bmatrix} 0.614\angle -3.9^\circ & 0.599\angle +8.9^\circ & 0.491\angle +14.5^\circ \\ 0.491\angle -15.4^\circ & 0.613\angle -4.2^\circ & 0.601\angle +8.7^\circ \\ 0.298\angle -5.1^\circ & 0.489\angle -15.5^\circ & 0.613\angle -4.4^\circ \end{bmatrix}. \end{aligned} \quad (3.137)$$

Observe that, in going from the volume-only matrices Γ_{v_a} and Γ_{v_b} to the total ground-plus-volume matrices Γ_a and Γ_b , the coherence magnitudes actually reduce. Across passes, in going from $\Gamma_{v_{ab}}$ to Γ_{ab} , the magnitude of element (3,1) decreases, but magnitude of element (1,3) increases. As discussed in Section 3.9.2, even though the ground is itself highly coherent, the net coherence of the full scene (ground and volume) can take counter-intuitive values, because the volume coherence is complex.

Given the perfect RVOG data in (3.137), the optimal clairvoyant beamformer \mathbf{w}_{RVOG} given by (3.127) would attenuate the volume interference by 12.1 dB (as listed in Table 3.4 for the same model parameters). Furthermore, the output coherence $\gamma_{y_{ab}}$ can be evaluated via (3.56) for different values of ground coherence γ_g , similar to the analysis of coherence estimation error in Section 3.10.2 (which used slightly different RVOG parameters). Of particular interest are the extreme values:

$$\begin{aligned}\gamma_g = 0 &\rightarrow |\gamma_{y_{ab}}| = 0.043 \\ \gamma_g = 1 &\rightarrow |\gamma_{y_{ab}}| = 0.977.\end{aligned}\tag{3.138}$$

This ostensibly provides a guide for how well the ground coherence can be recovered in the simulation. It will be seen that this guide is highly optimistic.

The key first step in the simulation is the synthesis of raw radar pulse echo data by accumulating the scattering responses from many point scatterers randomly positioned in 3D space. Each scatterer simply reflects the incident signal back to the receiver with a constant magnitude and phase (i.e. they have no angular or frequency dependence, and there are no multiple bounces). The magnitude is set according to the scatterer's height, in order to satisfy the chosen parameters of the RVOG model described in Section 3.9 (see (3.118) in particular); all scatterers on the ground have the same magnitude. The phase is random. Given the number of scatterers specified in Table 3.5 for the ground and volume layers, there will on average be several scatterers per horizontal resolution cell per metre of height. It is this high density of scattering elements that gives rise to somewhat realistic clutter: each SAR image will exhibit a speckle pattern whose statistical properties agree with the zero-mean complex Gaussian model for SAR data [Oliver & Quegan 1998 ch. 4.3], and the interferometric phase between pairs of images will exhibit the expected relationship to scattering height, as discussed in Section 3.3.

The data processing followed the steps shown in Figure 3.2. Six sets of raw radar data, three for pass *a* and three for pass *b*, were synthesised at the grazing angles in (3.136) and processed into six SAR images focused to ground level with a common spatial-frequency aperture. The three images from each pass were vertically beamformed to produce a new 3D SAR image whose complex pixel values are estimates of the reflectivity at ground level, with the volume suppressed. The complex coherence between the two 3D SAR images was computed according to (2.1) using a local spatial average. The resulting coherence magnitude provides a map of scene change, where low coherence (black) indicates substantial change and high coherence (white) indicates no change—this is the CCD image.

All six SAR images are shown in Figure 3.32 (both magnitude and phase). It is clear that the RVOG clutter scene contains no discernible structure; the images consist of just the complex speckle pattern that results from the coherent superposition of many scattering elements, including the layover of the elements above the ground. The small changes in geometry cause the images to appear to be different realisations of a random process, but the scattering responses are in fact entirely deterministic for the given scene. Moreover, there is no visible difference between the sets of images from the two passes, even though the ground was actually changed.

Figure 3.33 shows the output after adaptive MVDR beamforming of the SAR images from each pass in Figure 3.32. This output is 3D in the sense that it is steered to ground height, so the ground clutter is preserved but the above-ground scattering from the volume is suppressed. There is still no visible difference across passes.

The actual change in the ground scene between passes is illustrated in Figure 3.34. This shows the ‘ground-truth’ (literally), without the volume interference. The change was implemented by first constructing a text mask covering the scene and then randomly shifting the scattering elements under the letters in the mask. The mask spells out a hidden message—“HIDDEN MESSAGE”—with 2 m line width (compare to $1\text{ m} \times 1.2\text{ m}$ image resolution). The shift was normally distributed in distance, with zero mean and $12\text{ cm} (\approx \lambda/2)$ standard deviation (see Figure 3.4), and uniformly distributed in angle in the horizontal plane. The CCD in (b) is the best possible representation of the ground scene change that the CCD technique (as implemented in this thesis) can achieve, and all subsequent CCDs, formed for scenes containing volume as well as ground, should be interpreted as approximations to the ideal shown here. Therefore, the distortions seen in the letters in (b) should be expected in the later CCDs—they are a direct consequence of the different scattering responses due to the particular random shifts shown in (a),




Figure 3.32: Synthetic SAR images of an RVOG clutter scene (left: magnitude (30 dB dynamic range), right: phase ($(-180^\circ, 180^\circ)$)). See Table 3.5 for parameters.




Figure 3.33: 3D SAR images after adaptive MVDR beamforming of the SAR images from each pass in Figure 3.32. The sample covariance matrix was computed using a 19×23 m window (which covers 19×19 resolution cells). These images were used to form the CCD in Figure 3.35(f).

and are not a failing of the simulation. Note that the lines of low coherence marking out the letters in (b) are approximately 4 m wide, which is twice as wide as the actual scene change in (a), due to the spatial averaging involved in computing the coherence.

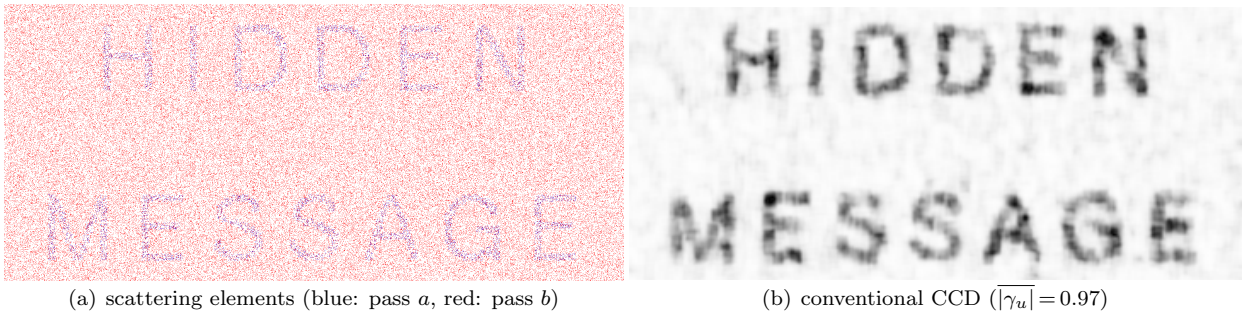


Figure 3.34: Ground-only scene components. (a) shows all the ground scattering elements as observed on the two passes, with the elements from pass b (red) overlaid on those from pass a (blue), so that where the elements have not been shifted, only the red version is visible. (On a computer, zoom in to see the individual elements.) (b) shows a conventional CCD between the two SAR images synthesised for pass a ($\psi_a = 35^\circ$) and pass b ($\psi_b = 35.3^\circ$); the images (not shown for brevity) look equivalent to those in Figure 3.32 i.e. just random clutter that is indistinguishable across passes.

The CCDs can be quantitatively compared by their average coherence magnitude $|\overline{\gamma_u}|$ in the unchanged areas, where the coherence should be unity. This value was computed in a window covering the middle area between the words, over the full azimuth extent of the image. The ground-only CCD in Figure 3.34(b) exhibits $|\overline{\gamma_u}| = 0.97$, possibly due to some minor imperfection in the aperture trim.

Figure 3.35 shows CCDs of the full RVOG scene. The ordinary single-channel CCD in (a), using just the middle image from each of the two passes (i.e. (c)/(d) and (i)/(j) in Figure 3.32), fails to recover the hidden message, because the volume causes decorrelation that contributes to an overall low coherence across the scene, including in areas where no actual ground change took place. Note that $|\overline{\gamma_u}| = 0.62$, which agrees with the RVOG coherence magnitude of 0.613 for the (2,2) element of Γ_{ab} in (3.137). The 3D CCD in (b), obtained using the conventional beamformer, whereby corresponding pixels in the three image channels from each pass are simply summed, is not much better, with $|\overline{\gamma_u}| = 0.65$. The array formed by the three closely

spaced channels offers very coarse vertical resolution ($\rho_z = 31.9\text{m}$) compared to the height of the volume ($h_v = 20\text{m}$), so little volume attenuation is achieved—this can be seen in the beampattern for the conventional beamformer in Figure 3.25(a). The RVOG beamformer, given by (3.127) for the RVOG model parameters in Table 3.5, does a better job of attenuating the volume component, leading to $|\overline{\gamma_u}| = 0.84$, which is sufficient to reveal the hidden message in (c). Its beampattern, also shown in Figure 3.25(a), shows that the beamformer steers two nulls onto the volume. The performance of the adaptive MVDR beamformer, given by (3.27), depends on the accuracy of the sample covariance matrix \hat{R} : in (d), the estimation window is small (covering 3×3 resolution cells), so the beamformer’s estimate of the interference field is partially in error, achieving only $|\overline{\gamma_u}| = 0.74$, whereas in (e), the estimation window is much larger (9×9 resolution cells), giving a better estimate of the covariance matrix, thus achieving better volume cancellation through the matrix inverse and a more accurate estimate of the ground coherence in the unchanged areas, with $|\overline{\gamma_u}| = 0.87$. Increasing the window size further provides diminishing returns, as shown in (f), where the window is very large (19×19 resolution cells), but $|\overline{\gamma_u}| = 0.88$, an increase of only 0.01. Nonetheless, comparison of (c), (e) and (f) shows that even a very small increase in $|\overline{\gamma_u}|$ (e.g. 0.03) can improve the visual separation of changed and unchanged areas.



Figure 3.35: CCDs of the RVOG scene using the SAR images from Figure 3.32. (a) shows the ordinary CCD between the images at $\psi_a = 35^\circ$ and $\psi_b = 35.3^\circ$. (b)–(f) show multichannel CCDs after beamforming the three channels on each pass. (b) used the conventional beamformer given by (3.26). (c) used the RVOG beamformer given by (3.127) (separately designed for pass a using Γ_{v_a} and pass b using Γ_{v_b} in (3.137)). (d)–(f) used the spatially adaptive MVDR beamformer given by (3.27), with the sample covariance matrix computed using different window sizes; the beamformed images used to form (f) are shown in Figure 3.33. Note that changing the window size does not change the horizontal resolution, because the weight vector is still separately applied at every pixel position (x, y) as $y(x, y) = \mathbf{w}_{\text{mvdr}}^H(x, y)\mathbf{x}(x, y)$ i.e. the output pixel value is not an average of the inputs pixels in the local neighbourhood around (x, y) .

Disconcertingly, the RVOG beamformer \mathbf{w}_{rvog} performs much worse than predicted

by (3.138), with the CCD in Figure 3.35(c) achieving only $|\overline{\gamma_u}| = 0.84$ instead of 0.977. This is because the synthetic data does not precisely match the RVOG model for which the beamformer was designed, and the closely spaced array is very sensitive to deviations from the model. To appreciate the problem, CCDs were generated for every pair of images in Figure 3.32, and their coherences $|\overline{\gamma_u}|$ measured, giving $\hat{\Gamma}_a$, $\hat{\Gamma}_b$ and $\hat{\Gamma}_{ab}$, which are measured (i.e. simulated) versions of the corresponding theoretical matrices in (3.137). The largest deviations between the simulated and theoretical matrices occur across passes, as follows:

$$|\hat{\Gamma}_{ab}| - |\Gamma_{ab}| = \begin{bmatrix} -0.0079 & 0.0045 & 0.0148 \\ -0.0075 & -0.0068 & 0.0026 \\ 0.0017 & -0.0083 & -0.0092 \end{bmatrix}, \quad \angle \hat{\Gamma}_{ab} - \angle \Gamma_{ab} = \begin{bmatrix} -1.46^\circ & -0.84^\circ & 0.90^\circ \\ -0.80^\circ & -0.77^\circ & -0.59^\circ \\ -0.18^\circ & 0.12^\circ & -0.41^\circ \end{bmatrix}. \quad (3.139)$$

Even though the simulated and theoretical matrices would seem to be quite similar, the small level of deviation (maximum magnitude: 0.0148, maximum phase: 1.46°) is enough to degrade the output coherence. Computing the multichannel coherence magnitude $|\hat{\gamma}_{y_{ab}}|$ via (3.56) by applying \mathbf{w}_{rvoG} to $\hat{\Gamma}_{ab}$ gives 0.836 (as observed in Figure 3.35(c)), whereas applying \mathbf{w}_{rvoG} to Γ_{ab} gives 0.977 (as listed in (3.138)).

The sensitivity of the RVOG beamformer to such small deviations is a result of the very small array spacing. Recall the analysis of the matrix condition number κ in Section 3.10.4. When the complex coherences between pairs of channels are subject to phase noise, the average volume attenuation achieved by the optimal beamformer is degraded by an amount that strongly depends on both the noise level and the array structure as quantified by $\kappa(\Gamma_v)$ —this was illustrated in Figure 3.30(b). The deviations here between the simulated and model data are not phase noise, but they are a similar small perturbation for which the same sensitivity is applicable.

To see this dependence on array structure, Figure 3.36 shows CCDs for the same scene using different arrays and beamformers. Table 3.6 lists the matrix condition numbers for the three different arrays under consideration and collates the values of $|\overline{\gamma_u}|$ measured in the CCDs in Figures 3.35 and 3.36. Figure 3.23(b) illustrates how the matrix condition number varies with array structure.

Table 3.6: Average coherence magnitude $|\overline{\gamma_u}|$ in the unchanged areas ($\gamma_g = 1$) for different arrays and beamformers, collated from Figures 3.35 and 3.36. The first column specifies the array and the second column gives the consequent matrix condition number. The third column lists the results for the RVOG beamformer, the fourth column lists the results for the MVDR beamformer implemented using a small window ($3 \times 3.6\text{m}$) for the estimation of the covariance matrix, and the fifth column lists the results for the MVDR beamformer implemented using a large window ($19 \times 23\text{m}$).

array	$\log_{10} \kappa(\Gamma_v)$	RVOG	MVDR (s)	MVDR (l)
$N_{ch} = 3, \Delta\psi = 0.1^\circ$	1.9	0.92	0.79	0.92
$N_{ch} = 3, \Delta\psi = 0.05^\circ$	3.2	0.84	0.74	0.88
$N_{ch} = 5, \Delta\psi = 0.05^\circ$	6.5	0.24	0.58	0.92

The pair-wise coherences $|\overline{\gamma_u}|$ for the two modified arrays in Figure 3.36 were checked in the same manner as for the reference array; the deviations between simulation and theory were similar to those shown in (3.139). Thus, in all cases, the data supplied to the beamformers were slightly perturbed versions of the RVOG ideal, and the level of perturbation was similar for the different arrays.

The wide variation in CCD quality is due to the differing levels of sensitivity of the arrays to perturbations. The matrix condition number captures this sensitivity. In Table 3.6, as the condition number increases, the RVOG beamformer becomes more sensitive to perturbations, or more precisely, it becomes more sensitive to mismatch between the data and the model for which it was designed, leading to a lower estimate of the coherence in the unchanged areas. Similarly, when the MVDR beamformer is implemented using a small window to compute the sample covariance matrix, the resulting inaccurate estimate of the true covariance matrix causes greater coherence loss when the condition number is larger. In both the RVOG and MVDR cases, the decorrelation is caused by the beamformer not only attenuating the volume suboptimally, but also mistakenly attenuating the ground component as if it were




Figure 3.36: CCDs of the RVOG scene using different arrays and beamformers. Compared to the array used in Figure 3.35 ($N_{ch} = 3$, $\Delta\psi = 0.05^\circ$), the channel spacing was doubled for the CCDs in the left-hand column (10 m across-track separation $\rightarrow \Delta\psi = 0.1^\circ$ grazing-angle spacing), and the number of channels was increased to five for the CCDs in the right-hand column. The first row was obtained using the RVOG beamformer (redesigned for the corresponding array), the second row was obtained using the MVDR beamformer with a small window for estimation of the covariance matrix, and the third row was obtained using the MVDR beamformer with a large window.

interference, leaving the low-coherence volume as the main component in the output. When MVDR is implemented using a large window, so that the sample covariance matrix is relatively accurate, then all of the arrays perform well, properly representing the unchanged areas with a high coherence. For the given arrays and the simulated data, $|\overline{\gamma_u}|$ seems limited to 0.92, not reaching the ground-only ideal of 0.97 in Figure 3.34(b). Nonetheless, it is clear that the 3D SAR CCD process using just three channels is able to reveal the subtle ground scene change under the volume.

For the results presented thus far, the repeat passes had a grazing angle difference of 0.3° , as given in (3.136). This was chosen as a typical separation given an airborne platform attempting to fly identical passes. Of course, the repeat-pass separation could reasonably take a wide range of values, although hopefully less than one degree (as assumed in (3.55)), because the underlying scattering behaviour of the objects in the scene should be essentially the same. Figure 3.37 shows the CCDs obtained when imaging the same RVOG scene using the parameters in Table 3.5 but shifting ψ_b so that the repeat-pass separation is varied from 0° to 0.5° in 0.1° increments. The main result is that the ordinary single-channel CCDs on the left fail to recover the hidden message, except in (a) when the repeat-passes are identical, but the multichannel CCDs on the right, obtained after MVDR beamforming, do recover the hidden message. Indeed, the beamformed coherence is quite resilient to the changing separation, with a small decline in $|\overline{\gamma_u}|$ because, at each pixel position, the two adaptive beamformers (for the two passes) will become increasingly different as the pass geometries become

more different, so the residual volume components in the two beamformed outputs will be increasingly different and therefore decreasingly correlated.

Two further features of Figure 3.37 deserve comment. Firstly, when the repeat passes are identical, the volume acts a source of high coherence, masking any areas of low coherence on the ground. The mean coherence of the text is 0.77 in (a) and 0.67 in (b), indicating that the beamformer has attenuated the volume component and improved the detectability of the ground change. This comes at the expense of a small (0.02) loss of coherence in the unchanged areas, likely due to imperfections in the simulation. Secondly, for the single-channel CCDs, the coherence $|\gamma_u|$ in the unchanged areas varies counter-intuitively with increasing angular separation: usually the coherence drops, as might be expected, but it jumps significantly between (e), where $\psi_{b_2} - \psi_{a_2} = 0.2^\circ$, and (g), where $\psi_{b_2} - \psi_{a_2} = 0.3^\circ$. In fact, this non-monotonic behaviour accords exactly with the RVOG model, as illustrated in Figure 3.22(b) (in that case, the nominal grazing angle was 45° , not 35° , but the shape of the curves is very similar).

Section 3.10.5 considered the effect of focusing to the wrong height i.e. when the focal surface is above or below the ground scattering surface. Figure 3.38 shows simulated results for this case. As the offset between the focal height (at $z = 0$) and the ground height (at $z = z_g$) increases, the coherence quickly degrades. When $z_g \neq 0$, the ground contributes non-zero interferometric phases between channels, so in beamforming terms, the ground is no longer at the ‘broadside’ steer direction, and the adaptive beamformer treats it as a source of interference, steering a null onto it. As predicted in Figure 3.31, the coherence variation is not symmetric about $z_g = 0$. However, note that the coherence in this simulation falls off even more steeply than that predicted in Figure 3.31(d), because the beamforming scenarios are different: in the latter case, the RVOG beamformer was kept fixed across all offsets, but here, the MVDR beamformer is allowed to adapt to the data at each offset, so as the magnitude $|z_g|$ of the offset increases, the null becomes deeper and more effective.

The fact that the multichannel coherence is so sensitive to the height offset between the focal and ground surfaces makes it ideally suited to the task of ground-height estimation, as discussed in Section 3.10.5. One drawback of this approach is that it would be computationally intensive, especially if the terrain is undulating, as the surface height search would need to be repeated at spatial intervals across the scene, within reasonable smoothness constraints. It is also worth reiterating that this approach assumes that maximising the output coherence is an appropriate objective, that is, it assumes that the ground is a source of high coherence and the volume is a source of low coherence, so maximising the output coherence can be interpreted as suppressing the volume response and preserving the ground response.

The results of the simulation indicate that the 3D SAR CCD concept, using only three channels, would be effective in detecting subtle changes on the ground under a forest canopy that approximates a random volume. However, if the channels are closely spaced ($\Delta\psi \sim 0.05^\circ$)—which is desirable in order to fit the multichannel system on an airborne platform—then the clairvoyant RVOG beamformer is not suitable, because it is excessively sensitive to deviations between the data and the model for which it was designed. Accurate coherence estimation requires the use of the adaptive MVDR beamformer, implemented with a large window (covering around 9×9 resolution cells, ideally more) for the computation of a sample covariance matrix that reliably estimates the true covariance matrix of the underlying scattering process. Such a large window restricts the technique to areas where the scene is spatially homogeneous.

3.11.1 Receiver operating characteristics

To complete the analysis of the proposed processing techniques, their receiver operating characteristics for change detection were evaluated. The approach taken here assumes that the CCD image serves as input to a simple detection algorithm that classifies each pixel position as either changed or unchanged by comparing the pixel coherence magnitude to some threshold. Given a CCD and ground-truth of a scene, the receiver operating characteristic (ROC) is the locus of the probability of detection (the proportion of changed pixel positions that were correctly classified) and the probability of false alarm (the proportion of unchanged pixel positions that were incorrectly classified) as the coherence threshold is varied between zero and one. Note that changes seen in CCDs are noisy (the coherence estimator has high variance where the true coherence is low, as




Figure 3.37: CCDs of the RVOG scene for different repeat-pass geometries. Pass a is fixed, with ψ_a given by (3.136), but pass b is varied for each row, as specified in the subcaptions. Left: single-channel CCDs; right: multichannel (MVDR) CCDs.



Figure 3.38: Multichannel CCDs of the RVOG scene after refocusing the constituent SAR images in Figure 3.32 to different heights. In each case, z_g specifies the height of the ground relative to the focal surface, with $z_g < 0$ implying that the ground is underneath. All CCDs were formed after MVDR beamforming using a $19 \times 23\text{m}$ window. Parameters as per Table 3.5 and (3.136).

shown in Appendix 2.A.1), and the changes often take the shape of extended, intermittent, curves, rather than blobs, so an automated detection algorithm that just thresholds each pixel independently is not optimal. However, it is adequate for the purpose of comparing the utility of the CCDs produced by the different processing techniques.

Figure 3.39 shows the receiver operating characteristics for selected CCDs in Figures

3.34 and 3.35. In the reference case (blue curve) when there is no intervening volume, the ROC is still not perfect (certain detection with zero false alarms) because the CCD resolution is coarser than the size of the true scene changes—compare the line widths of the text in the ground-truth (Figure 3.34(a)) and the ground-only CCD (Figure 3.34(b)). Recall that the CCD is formed using a spatial average, which blurs the boundary between changed and unchanged areas, and broadens the effective resolution compared to that of the input SAR images. If it were possible to use a finer SAR resolution, and concomitantly reduce the size of the averaging window, so it covered a smaller real-world area but the same number of resolution cells, then the blurred boundary region would be narrower, leading to a lower level of false alarms for a given detection probability. However, just reducing the size of the averaging window would lead to larger bias and greater variability of the coherence estimator (see Appendix 2.A.1), which may lead to overall worse detection performance.

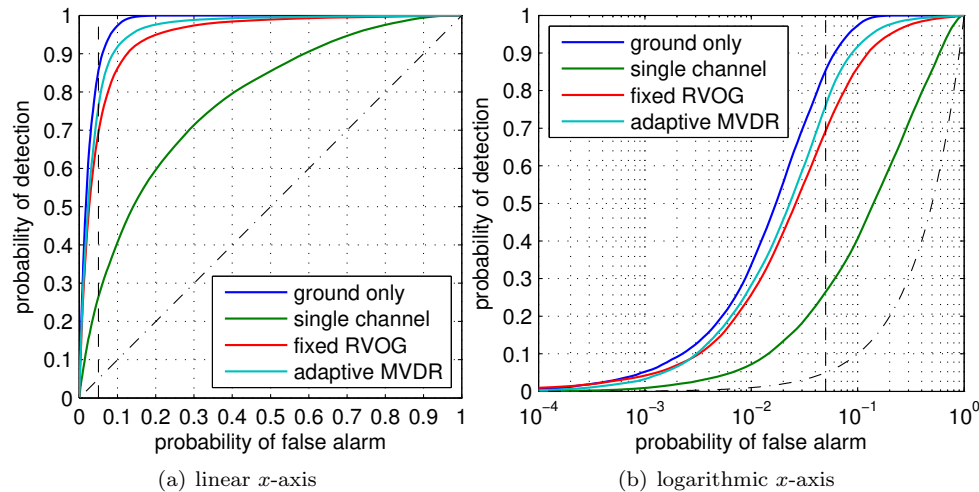


Figure 3.39: Receiver operating characteristics for CCDs generated by different processing techniques. Each curve shows the locus of the probability of detection and the probability of false alarm as the coherence threshold used to classify the pixels as either changed or unchanged is varied between zero and one. The ground-truth (the true change map) is given by Figure 3.34(a). The reference blue curve is for the CCD in Figure 3.34(b), when there is no intervening volume. The green curve is for the ordinary CCD in Figure 3.35(a), obtained without any 3D processing. The red curve is for the 3D CCD in Figure 3.35(c), obtained by applying the RVOG beamformer to three channels on each pass. The cyan curve is for the 3D CCD in Figure 3.35(f), obtained by applying the MVDR beamformer to three channels on each pass. For completeness, the same data are plotted in (a) and (b) using different scalings of the false-alarm axis.

The green curve in Figure 3.39, obtained using the ordinary single-channel CCD, exhibits poor detection performance compared to the red and cyan curves, obtained using the 3D CCDs. This is obvious from visual inspection of the corresponding CCDs in Figure 3.35. Best performance is exhibited by the cyan curve, obtained using the 3D CCD generated by applying MVDR beamforming (with a 19×23 m window for covariance matrix estimation) to three channels on each pass. To summarise detection performance at a false-alarm rate of, say, five per cent, the probability of detection is 0.86 for the ground-only reference case (blue curve), 0.26 for the single-channel CCD (green curve), 0.69 for the 3D CCD obtained using the fixed RVOG beamformer, and 0.76 for the 3D CCD obtained using the adaptive MVDR beamformer. Hence, the proposed 3D processing techniques significantly improve change detection performance for volume-obscured scenes.

3.12 3D SAR CCD simulation using PolSARproSim

The analysis and simulation of the 3D SAR CCD concept thus far has relied on the random-volume-over-ground (RVOG) model to represent forest scenes. Whilst this model is attractive for its small number of parameters and its approximate applicability

to a wide variety of forest types, it is obviously limited in its realism. In particular, the RVOG model ignores the scattering mechanism of ground-trunk double bounce, which can be much stronger than the desired ground backscatter response, especially at long wavelengths when the canopy penetration is good but the ground surface appears smooth, as discussed in Sections 3.1.2, 3.1.6 and 3.2.1. To account for this scattering mechanism, and independently verify the proposed processing techniques, the PolSARproSim program will now be employed.

PolSARproSim directly synthesises complex SAR images—not raw radar echo samples—of forest scenes. It is not a full-wave electromagnetic field solver, but it does follow a “physics-based approach” [Williams 2006 ch. 3.1.3] whereby each pixel value is the coherent superposition of the electromagnetic scattering responses of various canonical shapes (facets, cylinders and spheroids) that together represent a three-layer scene: a rough ground surface covered by a homogeneous understory of short vegetation and a stand of trees. Five scattering mechanisms are accounted for: direct backscattering from the constituent elements in each of the three layers, plus double-bounce scattering (a sequence of two forward reflections) between elements of the ground and elements in each of the other two layers. Note that the ground-tree interaction is dominated by ground-trunk double bounce. Broadly speaking, the direct backscattering generates the speckle pattern, which depends on the precise arrangement of scattering elements in the layers of clutter, whereas the double-bounce mechanisms are specular and give rise to isolated bright points in the image. This approach stands in contrast to that of the previous simulation in Section 3.11, where the scattering elements were all equivalent isotropic point scatterers with complex amplitudes but no shape, no electromagnetic properties, and no interaction (multiple bounces).

The resulting PolSARproSim imagery exhibits the statistical, textural and interferometric properties expected of real SAR imagery, as well as layover and shadowing effects. In addition, the simulation is fully polarimetric, accounting for horizontal and vertical polarisations on both transmit and receive. An ideal acquisition system is assumed, so real-world complications such as wavefront curvature, antenna beam-patterns and platform motion are ignored. PolSARproSim was developed by Mark Williams under contract to the University of Rennes, and the source code is publicly accessible from the website of the European Space Agency [Williams 2006].

The previous simulation in Section 3.11 assumed a homogeneous canopy and followed the RVOG model of extinction, whereby the propagating wave is subject to exponential decay parameterised by a single extinction coefficient σ_e . The resulting attenuation, given by (3.112), is a function just of depth. By contrast, the forest simulated by PolSARproSim consists of individual trees with different branch structures, and this spatial inhomogeneity is respected when accounting for extinction by evaluating the approximate attenuation at a 3D grid of points covering the scene [Williams 2006 ch. 8.3]. At each point, the precise path lengths through the living and dry parts of the crown of each tree, and through the short vegetation, are computed. These three types of vegetation are each characterised by a pair of extinction coefficients for horizontally and vertically polarised waves; these coefficients depend on the average permittivity across all elements of that type (as opposed to the individual permittivities of the actual intervening elements). When generating the output image, each scattering contribution is attenuated by the precomputed value at the nearest point in the grid.

The PolSARproSim process begins with the user specifying various scene and collection parameters. Scene parameters include the slope, surface roughness and moisture level of the ground, along with the species (pine/deciduous/hedge), mean height, and density of the canopy. Collection parameters include the centre frequency, the resolution, the number of passes, and for each pass, the slant range and incidence angle. The program then generates the three-layer scene. Each tree is individually realised, with a trunk, primary (radial) and secondary branches, and a crown volume filled with tertiary branches (twigs) and foliage (leaves or needles). Every scene element is assigned geometric and electromagnetic scattering properties through extensive use of the computer’s random number generator via the function `rand()`. The attenuation grid is evaluated. Finally, the program computes the total scattering response at each pixel position on each pass according to electromagnetic scattering theory, with various approximations to make it computationally practical [Williams 2006 ch. 8]. The output for each pass is a fully polarimetric set of four SAR images: HH, HV, VH (identical to HV) and VV, where each image pq is the response in polarisation q due to illumination

by polarisation p , and p and q are each either horizontal (H) or vertical (V).

Two modifications were made to the PolSARproSim source code in order to facilitate 3D SAR CCD processing. Firstly, to support 3D SAR beamforming, the program was extended to allow multiple interferometric channels (at different collection geometries) on each pass. Once the scene is generated, the SAR images for the different channels are computed in parallel. Secondly, to support CCD, a mechanism was added by which the scene could be changed in a controlled manner for the second pass. Two types of change were implemented:

- different realisations of the small-scale roughness scale (the complex Gaussian normalised scattering amplitude) of the facets making up the ground surface, and
- random perturbations of the orientation of the short vegetation, by up to 45° in azimuth and elevation.

These two types of change are intended to very roughly model the ground disturbance caused by one or more vehicles moving through a forest. They are conceptually different to the shifting of point scatterers implemented as part of the previous simulation in Section 3.11, but the consequent decorrelation is visually equivalent. In order to not disturb the sequence of random numbers generated by `rand()`, upon which subsequent calculations rely, these random changes were implemented using `rand_r()`, an equivalent alternative random number generator provided by the computer system. When synthesising the scene for the second pass, the changes were applied together along diagonal lines of some specified width. The trees were not changed.

Table 3.7 lists the parameters used to call PolSARproSim. Tables 3.8 and 3.9 list the scattering and propagation parameters, respectively, that characterise the layers generated by PolSARproSim; these values are the result of nominal scene constants hard-coded by the software’s author to “to ensure that predicted SAR backscattering coefficients are within observational limits” [Williams 2006 ch. 7.1.2]. Observe that direct backscatter from the ground surface is significantly weaker for HH compared to VV, whereas double bounce between the ground and the trees is significantly stronger for HH. The strongest mechanism is direct backscattering from the trees.

The kinds of scene changes that we seek to detect via CCD would likely manifest as changes in the direct backscattering from the ground or the short vegetation (or both); the other scattering mechanisms serve to obscure this change. In terms of the analysis in the previous sections, let the desired ‘ground’ signal be the direct backscattering from both the ground and the short vegetation (the latter is uniformly 30 cm tall, which effectively puts it at ground level given the limited vertical resolution of a realistic 3D SAR acquisition system), and let the undesired ‘volume’ interference be the direct backscattering from the trees. From Table 3.8, the ground-volume power ratio μ is therefore -4.41 dB for HH, -2.53 dB for HV and -0.13 dB for VV. Recall from Table 3.5 that the RVOG simulation in Section 3.11 used $\mu=0$ dB and $\sigma_e^{\text{dB}}=0.1$ dB/m, which are broadly consistent with PolSARproSim.

In order to verify that 3D SAR beamforming works correctly with the PolSARproSim data, 61 image channels were generated on each of the two passes, a and b . However, in order to show the performance of a minimal system for 3D SAR CCD, only three channels per pass will be used. The grazing angles of the individual channels are

$$\psi_a \approx \begin{bmatrix} 39.8^\circ \\ 39.85^\circ \\ 39.9^\circ \end{bmatrix}, \quad \psi_b \approx \begin{bmatrix} 40.1^\circ \\ 40.15^\circ \\ 40.2^\circ \end{bmatrix}. \quad (3.140)$$

Observe that the separation between passes is 0.3° , which is the same as that for the previous simulation in Section 3.11 (see (3.136)).

Two forest scenes were generated by PolSARproSim, one sparse and one dense. They are depicted in Figure 3.40. Most of the analysis will employ the sparse scene, in order to clearly demonstrate how an individual tree obscures the changes on the ground. As a final test of performance, the processing techniques will be applied to the more difficult dense scene.

Figure 3.41 shows the HH- and VV-polarised SAR images synthesised by PolSARproSim for the middle channel of the three channels from each pass listed in (3.140). The other channels look almost identical, except for very small differences around the trees due to the differences in layover at the slightly different grazing angles. The trees are brighter in the HH image than the VV image primarily due to the

Table 3.7: PolSARproSim call parameters. (‘az’ is the azimuth (along-track) direction and ‘gr’ is the ground-range direction. The parameters actually supplied to PolSARproSim were for slant-range instead of ground-range and incidence angle instead of grazing angle, but they have been converted here for easy comparison with Table 3.5. The roughness and moisture level of the ground are specified by the user on a nominal 0–10 scale; the program uses these to determine physical parameters of the model surface [Williams 2006 ch. 4.2.1, 5.4.8, 8.13.3]. Between passes, the ground was changed along diagonal lines of the specified width; the trees were not changed.)

operating frequency (wavelength λ)	1.32 GHz (22.7 cm)
window	2D Gaussian
image resolution (az \times gr)	0.8 m \times 1.6 m
pixel spacing (az \times gr)	0.53 m \times 1.06 m (2×1 pixel average for display)
swath (az \times gr)	127 m \times 168 m
ground roughness	5/10 (nominal simulation units)
ground moisture	3/10 (nominal simulation units)
mean short vegetation height	30 cm
tree species	pine (with mixed conical and spheroidal crowns)
mean tree height	20 m
number of trees	sparse: 72 (91 per hectare; 10.5 m spacing) dense: 207 (272 per hectare; 6.1 m spacing)
forest area	0.753 hectares (diameter 98 m)
type of ground change	different realisations of surface roughness scale and perturbed orientation of short vegetation
size of ground change	thick: 13 m line width thin: 3 m line width
coherence average (az \times gr)	3.17 m \times 3.17 m (~ 7.9 independent looks)
altitude	2109 ft
nominal range	1 km
nominal grazing angle ψ	40°
channel spacing	1.35 m across-track ($\Delta\psi \approx 0.05^\circ$)
number of channels per pass N_{ch}	up to 61

Table 3.8: PolSARproSim scattering parameters. (HV=VH)

scattering mechanism	backscattering coefficient (dB)			HH-VV correlation	
	HH	VV	HV	$ \cdot $ (dB)	\angle
direct ground	-30.0	-24.2	-41.1	-27.1	0.0°
direct short vegetation	-20.8	-20.9	-27.2	-23.6	0.2°
direct trees	-15.9	-19.1	-24.5	-20.0	0.0°
double bounce ground – short veg.	-24.7	-36.7	-41.6	-34.7	174.9°
double bounce ground – trees	-18.2	-27.6	-39.5	-24.5	176.7°

Table 3.9: PolSARproSim propagation parameters.

propagation medium	attenuation rate σ_e^{dB} (dB/m)	
	H	V
living crown	0.13	0.13
dry crown	0.02	0.03
short vegetation	0.38	0.38

ground-trunk double-bounce scattering mechanism being much stronger for horizontally polarised waves (at least according to the PolSARproSim scattering model). As was the case for the previous simulation in Section 3.11, there is no visible difference between the images across passes, even though the ground was actually changed.

Figure 3.42 shows the output after adaptive MVDR beamforming of the three PolSARproSim SAR images from each pass listed in (3.140). Compared to the ordinary

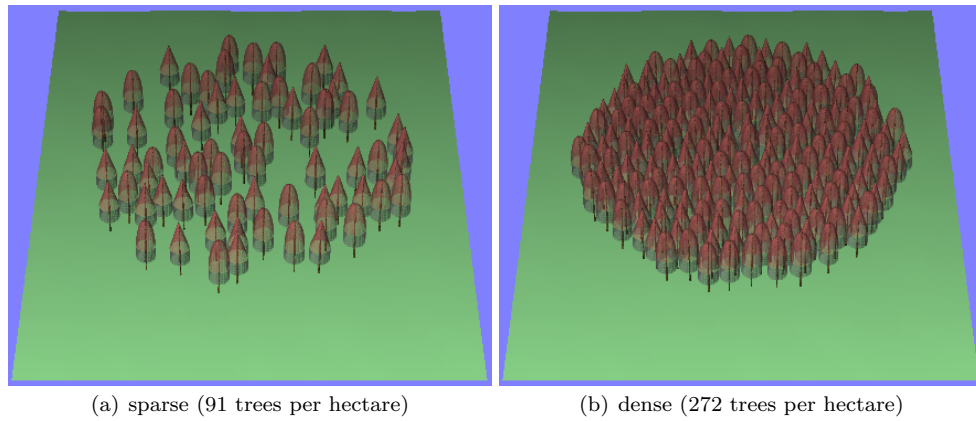


Figure 3.40: The sparse and dense forest stands generated by PolSARproSim using the parameters listed in Table 3.7, seen from the angular point-of-view of the nominal radar, with azimuth (x) across the page and near-range at the bottom edge. The ground surface (the green rectangle) is covered by short vegetation of mean height 30 cm. The forest consists of randomly positioned pine trees of mean height 20 m, each with a single straight trunk, branches, and a crown of twigs and needle-like foliage. The lower part of the crown (grey) is dry and cylindrical (mean diameter 5.65 m), whereas the upper part (brown) is living and either conical or spheroidal [Williams 2006].




Figure 3.41: PolSARproSim synthetic SAR images of the sparse forest stand (magnitude: 45 dB dynamic range, phase: $(-180^\circ, 180^\circ]$). x is the azimuth (along-track) direction and y is the ground-range (across-track) direction, with near-range at the top edge. The images in the upper row were collected at 39.85° on pass a ; the images in the lower row were collected at 40.15° on pass b . See Table 3.7 for parameters.

SAR images in Figure 3.41, the differing scattering responses of the trees are due to attenuation of the branch and crown responses, leaving just the ground-trunk double-bounce component, whose interferometric phase centre is located at the ground height. Despite this attenuation of the above-ground scattering response, there is still no visible difference across passes.

It is important to verify that 3D SAR beamforming of PolSARproSim images genuinely resolves scattering in the vertical dimension. To that end, height-range images were generated according to the method described in Algorithm 3.2, which is a modified version of Algorithm 3.1 for the case when the raw pulse data is not available, so the 3D SAR images must be formed using only the 2D SAR images. There is nothing




Figure 3.42: 3D SAR images after adaptive MVDR beamforming of the three PolSARproSim image channels from each pass specified in (3.140). The beamformer is steered to the ground, as usual. The sample covariance matrix \hat{R} from (3.28) was computed using a 3.17×3.17 m window (which covers approximately 4×2 resolution cells). Given the relatively small window (compared to the windows used in the previous simulation in Section 3.11), a very small amount of diagonal loading was added to stabilise the sample covariance matrix: $\hat{R}' = \hat{R} + \sigma_{dl}^2 I$ where σ_{dl}^2 was set equal to 0.1% of the mean power along the leading diagonal of \hat{R} . These images were used to form the CCDs in Figures 3.45(c) and 3.45(g).

special about this algorithm; the constituent steps are listed here just for clarity.

Figure 3.43 shows height-range images for different numbers of channels (61 and 3) and different beamforming methods (conventional and MVDR). From (3.42), assuming a uniform angular channel spacing of 0.05° , the nominal vertical half-power resolution of the conventional beamformer is 1.45 m given 61 channels and 29.44 m given 3 channels. These values agree with the results in (b) and (d): in (b), tree bases and tree crowns are successfully separated in height, but their respective scattering responses are represented by a sinc-like point-spread function with peak-to-null width approximately 1.5 m and sidelobes rotated in the height-range plane due to the oblique observation geometry ($\psi \sim 40^\circ$), as discussed in Section 3.4.4, whereas in (d), the mainlobe encompasses the whole vertical extent of the scene (mean tree height is 20 m), with the first null observed around $z = 30$ m. Note that the bright response at the base of each tree is due to ground-trunk double bounce. As shown in (c) for 61 channels, the adaptive MVDR beamformer is able to separate the tree components without generating sidelobes; also observe that the ground (which is covered in short vegetation with mean height 30 cm) is imaged with narrower vertical resolution compared to (b). Importantly, given only three channels, the MVDR beamformer in (e) still provides vertical resolution, although admittedly coarse. These results confirm that 3D SAR beamforming of PolSARproSim images works as expected.

Now return to the central problem of change detection. Two patterns of ground change were tested. The true patterns are shown in the ground change masks in Figure 3.44; the black lines indicate where the scene changes discussed previously were applied. Most of the analysis will employ the pattern formed by the thick lines of change, in order to clearly demonstrate how an individual tree obscures the changes on the ground; the width of the thick lines (13 m) is larger than the mean diameter of the tree crowns (5.65 m). As a final test of performance, the processing techniques will be applied to the more difficult, but more realistic, thin lines of change; the width of the thin lines (3 m) is the smallest width that was found to still give a visually discernible decorrelation, given

Algorithm 3.2 Height-range image formation using only SAR images

output data: A height-range image consisting of a stack of range-lines at one azimuth index, each steered to a different height.

input definitions:

- N_{ch} SAR images at slightly different grazing angles.
- A set of steer heights (relative to the original focus height at $z=0$).

- 1: **for all** steer heights **do** #Refocusing
- 2: **for all** SAR images **do**
- 3: Resample the image along ground-range y to account for the range layover shift $z \tan \psi$ [Jakowatz *et al.* 1996 ch.2.5] (assuming a flat focal plane), where z is the new steer height and ψ is the original grazing angle.
- 4: Correct the propagation range phase of each resampled pixel at (y, z) (relative to the origin in the scene) by adding the phase $4\pi r/\lambda$ where $r = \sqrt{(\text{standoff} + y)^2 + (\text{altitude} - z)^2}$ and the platform position at aperture centre is (standoff, altitude). (This works for PolSARproSim images, but SAR images formed differently may require a slightly different phase correction [Jakowatz & Wahl 2001].)
- 5: **end for**
- 6: Form a 3D SAR image steered to this height by coherently combining the N_{ch} refocused 2D SAR images using one of the beamforming methods described in Section 3.4.2.
- 7: Extract from the 3D SAR image a line along range at one azimuth index and insert it into the output height-range image.
- 8: **end for**

the limited image resolution and the implemented types of scene change. The CCDs should approximate these masks by exhibiting (black) lines of decorrelation surrounded by a (white) background of high coherence, with minimal distortion from the trees.

Figure 3.45 shows CCDs of the sparse forest. In all CCDs, the black lines of decorrelation match the thick lines of change in Figure 3.44. Thus, once again it can be seen that the magnitude of the complex coherence is a good detector of subtle scene changes that are not easily visible in either ordinary SAR images (Figure 3.41) or beamformed images steered to the height of interest i.e. the ground (Figure 3.42). In the ordinary single-channel CCDs in (a) and (d), formed using just the image channels at $\psi_a = 39.85^\circ$ and $\psi_b = 40.15^\circ$ (shown in Figure 3.41), the trees obscure the lines of decorrelation on the ground. The coherence at the location of each tree is the net result of (i) volume decorrelation of scattering from the branches and the crown due to the 0.3° repeat-pass separation, (ii) high coherence due to the ground-trunk double bounce, which is constant across passes, and (iii) the actual ground coherence. Furthermore, each tree is circumscribed by strong decorrelation because the types of scattering elements contributing to the net response at the tree edges changes considerably across passes due to the differential layover shift depicted in Figure 3.7. This effect was not observed in the previous simulation in Section 3.11 because in that case the above-ground scattering elements making up the volume were homogeneous and uniformly distributed, so any change in the particular scattering elements contributing to one pixel had little effect. In addition, PolSARproSim accounts for shadowing, so scattering from the ground on the far-range side of each tree is relatively weak, giving rise to an especially low net coherence on that side.

The multichannel CCDs in (b) and (f) in Figure 3.45, formed after conventional beamforming of the three channels from each pass listed in (3.140), suffer almost the same level of obscuration by the trees as the single-channel CCDs; this is expected, because the three-channel conventional beamformer offers little attenuation of scattering from above the ground. After MVDR beamforming in (c) and (g), the obscuration of the unchanged areas by the trees is significantly reduced, because the beamformer has attenuated scattering from the above-ground scattering elements, mostly removing the source of volume decorrelation. However, the trees still obscure the changed areas due to the high coherence of the ground-trunk double-bounce scattering component, which manifests in the CCDs as solid white circles of high coherence at the base of the trees, magnified by the spatial averaging involved in CCD processing. Importantly, adding extra channels, as shown in (d) and (h), has no effect on this high-coherence bias,

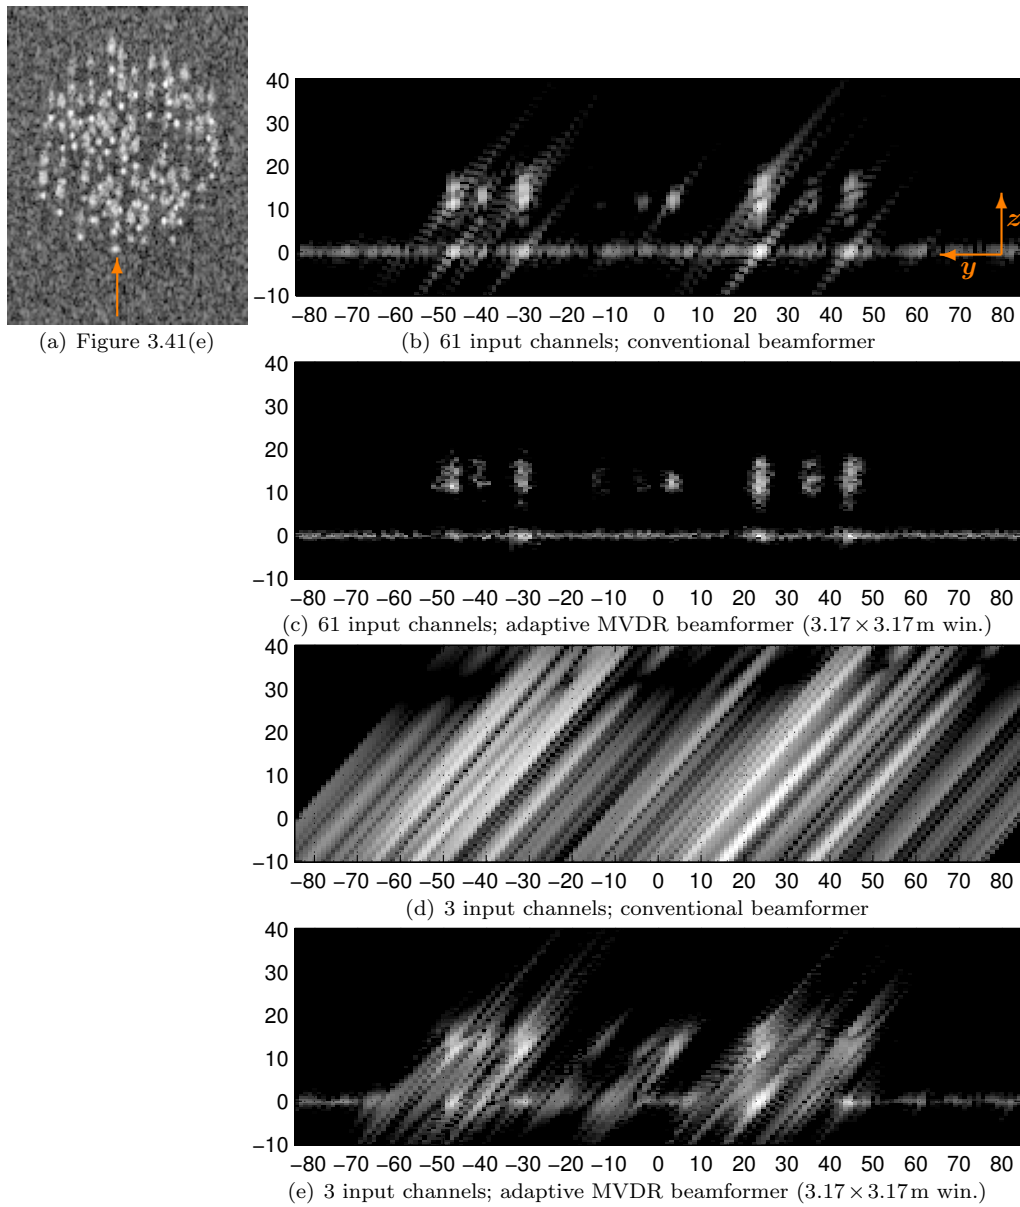


Figure 3.43: Height-range images ((b)–(e)) after beamforming PolSARproSim HH images from pass b . The horizontal axis is ground-range (y), the vertical axis is height (z), and the intensity in decibels is displayed using grey-scale with 40 dB dynamic range and the peak value in each image mapped to white. Each image was obtained according to Algorithm 3.2, with the 3D processing implemented as specified in the subcaption, and the cut taken along range at the azimuth column indicated by the orange arrow in (a). For (b) and (c), the 61 input image channels were centred on $\psi = 40^\circ$ and spaced by $\Delta\psi \approx 0.05^\circ$, as listed in Table 3.7. For (d) and (e), the three input image channels were those at ψ_b in (3.140).




Figure 3.44: Two change masks that specify where the ground scene changes between passes a and b were applied; white indicates no change and black indicates change. The lines of change are 13 m wide for the ‘thick’ change on the left and 3 m wide for the ‘thin’ change on the right. The actual changes implemented in the black lines were described previously and are listed in Table 3.7.




Figure 3.45: CCDs of the sparse forest scene with thick lines of change between passes. The upper row is for HH and the lower row is for VV. (a) and (d) show the ordinary CCDs between the PolSARproSim SAR images in Figure 3.41 ($\psi_a = 39.85^\circ$ and $\psi_b = 40.15^\circ$). (b) and (f) show the CCDs obtained after conventional beamforming of the three channels from each pass listed in (3.140). (c) and (g) show the CCDs obtained after adaptive MVDR beamforming of the same three channels from each pass; the beamformer output is shown in Figure 3.42, along with details of the processing. (d) and (h) show the CCDs obtained after MVDR beamforming of 61 channels.

because the double-bounce component is located interferometrically at the ground height—no 3D processing can remove this bias.

Figure 3.46 shows the CCDs obtained for the same scene when the repeat-pass separation is varied from 0° to 0.4° in 0.1° increments. For ordinary single-channel CCDs, the net coherence of the trees varies with the repeat-pass separation due to the strong geometry dependence of the volume decorrelation of the branches and crown; this variation is not monotonic with increasing angular separation, which agrees with Figure 3.22(b) for a random volume and Figure 3.37 for the previous simulation. When the passes are identical, the trees act solely as a source of high coherence, obscuring the areas of change on the ground; when the repeat-pass separation is non-zero, the trees obscure both the changed and unchanged areas on the ground. By contrast, after adaptive MVDR beamforming using three channels per pass, the multichannel CCDs vary little as the repeat-pass separation is increased; the above-ground scattering contributions that cause volume decorrelation have been attenuated in all cases, leaving a small residual decorrelation over the unchanged areas as well as patches of high coherence due to ground-trunk double bounce that obscure the areas of ground change. This relative resilience to the changing repeat-pass separation was also observed in Figure 3.37 for the previous simulation.

3.12.1 Polarimetric filtering of ground-trunk double bounce

The PolSARproSim simulation demonstrates the coherence bias caused by the ground-trunk double-bounce scattering mechanism, which is likely to be a relatively strong component of the total scattering response at L-band and below. Vertical resolution through grazing angle diversity is unable to remove the bias, because the phase centre of this scattering mechanism is located at the same height as the ground. However, the double-bounce mechanism is distinguished by its polarimetric signature [Durden *et al.* 1989, Freeman & Durden 1998]. Polarimetric diversity offers the possibility of a polarimetric filter that removes, or at least attenuates, the double-bounce component




Figure 3.46: CCDs of the sparse forest scene for different repeat-pass geometries. The separation between passes increases by 0.1° down the rows: $\psi_b - \psi_a = 0^\circ$ for the first row, $\psi_b - \psi_a = 0.1^\circ$ for the second row, etc. For all rows, $(\psi_a + \psi_b)/2 = 40^\circ$ and the spacing between channels is 0.05° . The single-channel CCDs use only the middle channel from each pass; the multichannel CCDs were formed after adaptive MVDR beamforming of three channels using the same processing parameters as specified for Figure 3.42.

and the associated coherence bias. This section considers how such a filter could be constructed and how effective it would be.

Reconsider the dual-antenna across-track interferometer depicted in Figure 3.1, which generates an across-track array of three SAR image channels at different phase centres and therefore different grazing angles. Let both antennas be dual-linear-polarised, so each can transmit and receive on H and V, where H and V denote the horizontal and vertical components of the standard linear polarisation basis. (Any other basis, such as the circular basis, could be used, but the linear basis is assumed here just for ease of understanding.) If the radar collects data in a fully polarimetric mode, whereby the system cycles pulse-to-pulse through the transmitters at the different polarisations (H and V) and positions (1 and 3), and samples the echoes using all the receivers, then the sixteen transmit-receive pairs can be grouped into nine distinct polarimetric-spatial channels, as shown in Table 3.10, such that at each of the three effective phase centre positions (1, 2 and 3) there are three transmit-receive polarimetric pairs (HH, HV and VV). The Intermap radar system acquires data in this way.

Table 3.10: Channels generated by a pair of dual-polarised antennas operating in a fully-polarimetric radar mode. The antenna positions are denoted 1 and 3 and the polarisations are denoted H and V. Assuming that the system is calibrated and the scattering is reciprocal, the sixteen transmit-receive pairs can be grouped into the nine channels listed in the top row, corresponding to three polarimetric channels at each of the three effective phase centre positions.

Pulse	Tx	Rx				HH1	HV1	VV1	HH2	HV2	VV2	HH3	HV3	VV3
1	H1	H1	V1	H3	V3	H1H1	H1V1		H1H3	H1V3				
2	V1	H1	V1	H3	V3		V1H1	V1V1		V1H3	V1V3			
3	H3	H1	V1	H3	V3				H3H1	H3V1		H3H3	H3V3	
4	V3	H1	V1	H3	V3					V3H1	V3V1		V3H3	V3V3

The grouping of transmit-receive pairs into the polarimetric-spatial channels shown in Table 3.10 requires that the radar system is calibrated and assumes that scattering from the scene is both spatially and polarimetrically reciprocal. That is, the complex scattering coefficient $s_{PQ}(\hat{\mathbf{n}}_i, \hat{\mathbf{n}}_s)$, which connects polarisation P incident from direction $\hat{\mathbf{n}}_i$ with polarisation Q scattered in direction $\hat{\mathbf{n}}_s$, is equal to coefficient $s_{QP}(-\hat{\mathbf{n}}_s, -\hat{\mathbf{n}}_i)$, which connects polarisation Q incident from direction $-\hat{\mathbf{n}}_s$ with polarisation P scattered in direction $-\hat{\mathbf{n}}_i$ [Cloude 2010 ch. 1.4.1], [Ulaby & Long 2014 p.168]. Note that scattering reciprocity does not technically permit equating all four pairs grouped under the HV2 banner, but if the geometry is pseudo-monostatic, as considered here, so that the directions $\hat{\mathbf{n}}_i$ and $-\hat{\mathbf{n}}_s$ are nearly parallel, then this approximation would seem reasonable.

Let $\check{\mathbf{x}} = [x_{\text{HH}}, \sqrt{2}x_{\text{HV}}, x_{\text{VV}}]^T$ denote the HH, HV and VV scattering responses at one pixel position measured at one phase centre ($\check{\mathbf{x}}$ is used to distinguish the polarimetric vector from \mathbf{x} , which was defined in Section 3.5 for measurements made at multiple across-track phase centres using one polarisation). $\check{\mathbf{x}}$ can be seen as one possible vectorisation of the traditional scattering matrix representation, assuming that reciprocity is applicable, so $x_{\text{HV}} = x_{\text{VH}}$ (the $\sqrt{2}$ factor is necessary to conserve the total power) [Cloude 1985, Cloude & Pottier 1996]. Furthermore, let $\check{\mathbf{w}} = [w_{\text{HH}}, w_{\text{HV}}, w_{\text{VV}}]^T$ denote a polarimetric scattering mechanism, with $\|\check{\mathbf{w}}\| = \sqrt{\check{\mathbf{w}}^H \check{\mathbf{w}}} = 1$. The weighted sum $\check{\mathbf{y}} = \check{\mathbf{w}}^H \check{\mathbf{x}}$ can be interpreted as the projection of the data in $\check{\mathbf{x}}$ onto the scattering mechanism represented by $\check{\mathbf{w}}$; this formulation is central to polarimetric SAR interferometry [Cloude & Papathanassiou 1998].

The dual-layer model introduced in Section 3.6 could be extended from a scalar scattering model to allow for a polarimetric vector response from each layer. The multichannel dual-layer coherence formulation in Section 3.7 could then be modified in some way to incorporate polarimetry into a single coherence-based change metric. (Polarimetric extensions to conventional coherent and incoherent change detection were mentioned in Section 3.1.5.) However, to limit the scope of this chapter, such extensions will not be pursued.

Instead, the intent here is to first determine the scattering mechanism $\check{\mathbf{w}}_{\text{gtdb}}$ corresponding to ground-trunk double bounce in the SAR imagery, and then project the polarimetric data onto an orthogonal mechanism $\check{\mathbf{w}}_{\text{ogtdb}}$, in order to filter out

the ground-trunk double-bounce component. Given the small swath relative to the nominal range (see Table 3.7), and the small variation in the grazing angles of the across-track channels, these scattering mechanisms are assumed to have no spatial or grazing-angle dependencies, so the same polarimetric projection will be applied at all pixels in all SAR images.

Double bounce between two perpendicular perfectly conductive surfaces is characterised polarimetrically by the scattering mechanism $\check{\mathbf{w}}_{\text{edb}} = [1, 0, -1]/\sqrt{2}$, so there is zero cross-polarisation and the co-polarisations have equal magnitude (hence the ‘e’ in the subscript) and opposite sign [Cloude 2010 ch. 1.4.2.2]. A suitable orthogonal scattering mechanism in this case would be $\check{\mathbf{w}}_{\text{esb}} = [1, 0, 1]/\sqrt{2}$, corresponding to equal-magnitude single bounce. However, if the surfaces are instead made of different dielectric materials, neither of which are perfect conductors, as in the case of ground-trunk double bounce, then the co-polarisations will no longer have equal magnitude, and their phase offset may be large (perhaps close to 180°) for only a limited range of grazing angles centred on 45° [Cloude & Pottier 1996, Cloude 2010 ch. 3.1.2]. Therefore, $\check{\mathbf{w}}_{\text{esb}}$ is unlikely to be orthogonal to this dielectric double bounce. Note that, in general, the scattering mechanisms of surface backscatter from rough ground and ground-trunk double bounce are approximately orthogonal [Cloude & Pottier 1996].

Since cross-polarisation is not relevant to double bounce (ignoring any rotations induced by terrain slope [Cloude 2010 ch. 3.1.5]), the HV components of $\check{\mathbf{w}}_{\text{gtdb}}$ and $\check{\mathbf{w}}_{\text{ogtdb}}$ can be set to zero. For orthogonality,

$$\check{\mathbf{w}}_{\text{ogtdb}}^H \check{\mathbf{w}}_{\text{gtdb}} = w_{\text{ogtdbHH}}^* w_{\text{gtdbHH}} + w_{\text{ogtdbVV}}^* w_{\text{gtdbVV}} = 0 \Rightarrow \frac{w_{\text{ogtdbHH}}}{w_{\text{ogtdbVV}}} = \left(-\frac{w_{\text{gtdbVV}}}{w_{\text{gtdbHH}}} \right)^* \quad (3.141)$$

The co-polarisation ratio for $\check{\mathbf{w}}_{\text{gtdb}}$ on the right-hand side of (3.141) can be estimated from the data simply by averaging the co-polarisation ratio $x_{\text{VV}}/x_{\text{HH}}$ for pixels dominated by ground-trunk double bounce. Hence, the main task is to find such pixels.

The alpha-entropy-span polarimetric decomposition provides a way to quantify and interpret the type and strength of scattering mechanisms contributing to each pixel over a wide, continuous range of possible mechanisms [Cloude & Pottier 1997, Pottier & Cloude 1997]. This decomposition is described in detail in Section 4.A.1. At each pixel position, the local coherency matrix (a transformed version of the sample covariance matrix, with the same eigenvalues) is estimated [Cloude 1985], and then the three decomposition metrics are evaluated using the eigenvalues and eigenvectors of this matrix [Cloude & Pottier 1997]. The alpha angle indicates the mean scattering mechanism; the entropy indicates the relative purity of this scattering mechanism; the span indicates the intensity (brightness) of the pixel relative to the rest of the scene. Mapping alpha, entropy and span to the hue, saturation and value (HSV) components of a colour image generates a convenient false-colour representation of the decomposed polarimetric scattering response across the scene. Figure 3.47 illustrates this for PolSARproSim polarimetric SAR imagery of the sparse and dense forest scenes. Note that the window used for coherency matrix estimation is quite small in order to avoid diluting the polarimetric signature at the base of each tree with the responses from the surrounding area; as a result, the metrics suffer from significant finite-sample bias and variance, but not enough to prevent pixel comparison within the image [López-Martínez *et al.* 2005, Lee *et al.* 2008].

If a pixel contained pure equal-magnitude double bounce (i.e. $\check{\mathbf{w}}_{\text{edb}}$), then the alpha angle would be maximum and the entropy zero; such a pixel would appear pure red according to the false-colour mapping described in Figure 3.47. Pixels dominated by ground-trunk double bounce will roughly approximate this, so the alpha angle should be high and the entropy low; such pixels would appear yellow or orange. Observe in Figure 3.47 that there is no red, but there are small patches of yellow and yellowish-green, which can be interpreted as locations where there is a strong ground-trunk double-bounce scattering mechanism, which is also slightly dipole-like in that one polarisation, in this case HH, is significantly stronger than the other. Furthermore, it is apparent in (b) that the decomposition effectively separates the scattering responses from different tree components, not by their vertical position, but by their polarimetric signature: the yellowish ground-trunk double bounce is locked at the base of each tree, but the blue-green single bounce from branches in the canopy lays over towards the platform (at the top of the page), leaving the tree bases bare at the far range edge




Figure 3.47: Alpha-entropy-span polarimetric decomposition of the set of HH, HV and VV PolSARproSim SAR images of the sparse and dense forest scenes at $\psi_b = 40^\circ$. The hue is set by the alpha angle, varying between blue for equal-magnitude ($|\text{HH}| = |\text{VV}|$) single bounce, through green for a dipole-like response ($|\text{HH}| \gg |\text{VV}|$ or $|\text{HH}| \ll |\text{VV}|$), to red for equal-magnitude double bounce. The saturation (i.e. the colour purity or white-level) is set by the entropy. The intensity (i.e. brightness) relative to the other pixels is set by the span. The local coherency matrix at each pixel was estimated by averaging the Hermitian outer products of the polarimetric vectors at the surrounding pixels in a $3.17 \times 3.17\text{m}$ window, giving ~ 8 independent looks.

and the canopy alone at the near-range edge.

Making use of the ability of this polarimetric decomposition to separate scattering mechanisms, Algorithm 3.3 describes how pixels dominated by ground-trunk double bounce can be found in polarimetric SAR imagery.

Algorithm 3.3 Ground-trunk double bounce pixel search

output data: Positions of pixels with a dominant ground-trunk double-bounce scattering response.

input data:

- HH, HV and VV SAR images acquired at one phase centre.
- Number of pixels N_{gtdb} to return.

- 1: Compute the alpha-entropy-span decomposition of the input polarimetric dataset [Cloude & Pottier 1997, Pottier & Cloude 1997]. This involves doing an eigendecomposition of the local coherency matrix at each pixel position.
 - 2: Rank all pixels by their alpha angle (in descending order), entropy (ascending) and span (descending). For each pixel, denote these ranks as r_a , r_e and r_s .
 - 3: Rank all pixels by the single metric $r = w_a r_a + w_e r_e + w_s r_s$ (ascending), where the weights w_a , w_e and w_s quantify the relative importance of the alpha angle (the type of scattering), the entropy (the purity of the scattering), and the span (the brightness of the pixel in the scene). Setting $w_a = 3$, $w_e = 2$ and $w_s = 1$ was found to work well for the identification of tree trunks in PolSARproSim images, but these values are in no sense optimal. Using the alpha-entropy-span ranks instead of their actual values avoids the problem of how to combine the three quantities into a single metric.
 - 4: Select the top N_{gtdb} pixels according to their rank r (lower is better), enforcing a minimum spatial separation δ_{gtdb} between selected pixels in order to ensure that they are at different trees. For the PolSARproSim images here, $N_{\text{gtdb}} = 20$ and $\delta_{\text{gtdb}} = 6\text{m}$ were found to be suitable.
-

The ground-trunk double-bounce scattering mechanisms obtained by averaging over the pixels found by applying Algorithm 3.3 to the PolSARproSim imagery for the sparse and dense scenes shown in Figure 3.47 are

$$\text{sparse: } \check{\mathbf{w}}_{\text{gtdb}} = \begin{bmatrix} 0.9198 \angle +179.553^\circ \\ 0 \\ 0.3924 \angle +0.447^\circ \end{bmatrix}, \quad \text{dense: } \check{\mathbf{w}}_{\text{gtdb}} = \begin{bmatrix} 0.9205 \angle +176.056^\circ \\ 0 \\ 0.3908 \angle +3.944^\circ \end{bmatrix}. \quad (3.142)$$

The associated orthogonal mechanism $\check{\mathbf{w}}_{\text{ogtdb}}$, given by (3.141), serves as a polarimetric rejection filter for ground-trunk double bounce in each scene. Applying the filter is

just a matter of computing the projection $\tilde{\mathbf{w}}_{\text{ogtdb}}^H \tilde{\mathbf{x}}$ for the polarimetric data vector $\tilde{\mathbf{x}}$ at every pixel position.

Figure 3.48 compares the filtered SAR image to the ordinary HH, VV and HH+VV polarimetric channels for the sparse scene. Observe that the bright points of ground-trunk double bounce at the base of the trees in the HH and VV images are still present in the HH+VV images but are significantly attenuated in the filtered images. This is particularly striking in the 3D image in (h), the result of the polarimetric filter first attenuating the ground-trunk double bounce, and then the MVDR beamformer attenuating the canopy response, together effectively removing the trees from the scene.




Figure 3.48: 2D (upper row) and 3D (lower row) SAR images of the sparse forest stand on pass b for different polarimetric combinations. The 3D images are the output of the three-channel adaptive MVDR beamforming process used in Figure 3.42. (c) and (g) show the Pauli sum HH+VV (scaled by $1/\sqrt{2}$ to conserve power). (d) and (h) show the weighted sum of HH and VV where the weight vector $\tilde{\mathbf{w}}_{\text{ogtdb}}$ filters out the ground-trunk double-bounce scattering component. Note that for (g) and (h), first the particular polarimetric combination was formed at each grazing angle using the ordinary 2D SAR images, and then these combined images were beamformed to produce the 3D image.

Figure 3.49 compares 3D SAR CCDs of the sparse scene obtained for different polarimetric channels. Visually, the points of high coherence, caused by ground-trunk double bounce, which distort the lines of decorrelation through the forest stand, are much reduced in the polarimetrically filtered CCD in (d) compared to the other CCDs. However, there is clearly still some distortion of the ground coherence due to residual scattering from the trees. In fact, there will always be some high-coherence memory of the trees in a CCD, because no ground change can occur at the base where the trunk meets the ground (unless the tree is chopped down).

In the previous simulation in Section 3.11, the CCDs were compared in terms of their average coherence $|\overline{\gamma_u}|$ in the unchanged areas, because the usual effect of the volume interference is to lower the coherence. Here, the CCDs will be compared primarily in terms of their average coherence $|\overline{\gamma_c}|$ in the changed areas, because the effect of the ground-trunk double-bounce interference is to raise the coherence. This average coherence was computed using pixels within the circular forest stand that are located under the appropriate change mask in Figure 3.44. Change detectability ultimately depends on the coherence contrast $\Delta\overline{\gamma_{uc}} = |\overline{\gamma_u}| - |\overline{\gamma_c}|$.

Somewhat surprisingly, $|\overline{\gamma_c}|$ shows only a small decrease (0.005) within the forest stand for the polarimetrically filtered CCD in (d) compared to the VV CCD in (b). In fact, although it is difficult to see visually, the average coherence over the lines of




Figure 3.49: CCDs of the sparse forest scene for different polarimetric combinations. All CCDs were obtained after adaptive MVDR beamforming of the three channels from each pass listed in (3.140), according to the processing detailed for Figure 3.42. The constituent images from pass *b* are shown in (e)–(h) in Figure 3.48.

change outside the forest stand is actually higher by 0.022 in (d) compared to (b). To understand why, note that PolSARprosim accounts for double bounce between not only the ground and the trunks but also the ground and the short vegetation, and since the orientation of the short vegetation is changed between passes, this double-bounce scattering component would change, contributing to the observed decorrelation. Although designed to suppress ground-trunk double bounce, the polarimetric filter will likely suppress all double bounce to a significant extent, so applying this filter effectively suppresses the decorrelation that would have been observed due to any change in double-bounce scattering. Thus, in seeking to suppress an undesirable source of high coherence (ground-trunk double bounce), the filter also suppresses an important source of low coherence (changes in the short vegetation).

Now consider the dense scene with thin lines of change between passes. This is a much more difficult change detection scenario, because the changes are much smaller, and the obscuration by the canopy is more severe. Nonetheless, as shown in Figure 3.50, the polarimetric and 3D processing again successfully remove most of the scattering from the trees in the forest. In particular, comparison of the 3D image in (b) with the 2D image in (a) demonstrates the reduction in the canopy response achieved by weighting and summing just three across-track channels on each pass using \mathbf{w}_{mvdr} , and comparison of the polarimetrically filtered image in (d) with the VV image in (c) demonstrates the attenuation of the bright points of ground-trunk double bounce achieved by weighting and summing the polarimetric channels using $\check{\mathbf{w}}_{\text{ogtdb}}$.




Figure 3.50: 2D ((a)) and 3D ((b)–(d)) SAR images of the dense forest stand on pass *b* for different polarimetric combinations. The MVDR beamforming used to produce (b)–(d) was the same as that employed in Figure 3.42 and Figure 3.48.

The CCDs in Figure 3.51 show that the proposed processing does not just reduce the intensity of scattering from the trees, but actually reveals the scattering from the ground underneath. In the conventional single-channel CCD in (a), the ground coherence is obscured by volume decorrelation of the canopy, and in the HH 3D CCD in (b), the ground coherence is obscured by strong ground-trunk double bounce. The double-bounce scattering mechanism is weaker for vertical polarisation, so the thin

lines of change on the ground are visible in the VV 3D CCD in (c). However, the best result is seen in the 3D CCD in (d) using the polarimetric combination that is orthogonal to the observed ground-trunk double bounce: the average coherence $|\overline{\gamma_c}|$ of the changed areas within the forest stand is lower in (d) than (c) by 0.07. Admittedly, the coherence $|\overline{\gamma_u}|$ of the unchanged areas is slightly lower too (by 0.022), because the high-coherence double bounce is not present, leaving just the residual volume decorrelation after beamforming, but on balance, the polarimetric filtering increases the coherence contrast $\overline{\Delta\gamma_{uc}}$ by 0.048, making it easier to detect the changes.




Figure 3.51: CCDs of the dense forest scene for different polarimetric combinations. The CCDs in (b)–(d) were obtained after adaptive MVDR beamforming of the three channels from each pass listed in (3.140), according to the processing detailed for Figure 3.42. The constituent images from pass *b* are shown in Figure 3.50.

Despite the difficult challenge posed by the dense forest, the proposed polarimetric and spatial array processing clearly reveals the thin lines of change on the ground underneath.

3.12.2 Receiver operating characteristics for PolSARproSim

The detection performance permitted by the various processing techniques is now compared in terms of receiver operating characteristics. This follows the approach in Section 3.11.1. The ground-truth change maps are shown in Figure 3.44. In what follows, the detection and false alarm statistics that form the ROC curves were computed using only the pixels within the circular forested region in the middle of the scene (diameter 98 m; see Table 3.7 and Figure 3.40), because the aim here is to assess the performance specifically in forested regions.

Figure 3.52 shows the ROC curves for selected CCDs in Figures 3.45 and 3.49 of the sparsely forested scene subject to thick lines of change. As in Figure 3.39 for the previous simulation, the detection performance of the single-channel CCD (blue) is comparatively poor, with a detection probability of only 0.16 at a false-alarm rate of 5 per cent. After MVDR beamforming of three channels on each pass, the HH (green) and VV (red) 3D CCDs achieve detection probabilities of 0.60 and 0.76 at the same false-alarm rate. Clearly, detection is strongly dependent on polarisation. The polarimetric combination that best rejects ground-trunk double bounce (cyan) achieves the highest detection probability of 0.83. Hence, the combination of (i) 3D SAR beamforming to suppress the canopy interference and (ii) polarimetric filtering to suppress the ground-trunk double bounce permits the best detection performance.

Figure 3.53 shows the ROC curves for the CCDs in Figure 3.51 of the densely forested scene subject to thin lines of change. Compared to the thick lines of change in the sparsely forested scene, this is a much harder detection problem. Change detection using single-channel HH imagery (blue) is actually slightly worse than just tossing a coin (e.g. heads is change, tails is no-change), because the probability of detection is slightly less than the probability of false alarm; for example, the detection probability is 0.04 when the false-alarm rate is 5 per cent. Again, the best detection performance (cyan) is achieved by applying MVDR beamforming to the three-channel data after polarimetric filtering of the ground-trunk double bounce, giving a detection probability




Figure 3.52: Receiver operating characteristics for CCDs generated by different processing techniques. The ground-truth (the true change map) is given by Figure 3.44 (thick lines of change). The blue curve is for the ordinary HH CCD in Figure 3.45(a), obtained without any 3D processing. The green, red and cyan curves are for the HH, VV and polarimetrically filtered 3D CCDs in Figures 3.45(c), 3.45(g) and 3.49(d), respectively, obtained by applying the MVDR beamformer to three channels of the specified polarisation on each pass. For completeness, the same data are plotted in (a) and (b) using different scalings of the false-alarm axis.




Figure 3.53: Receiver operating characteristics for CCDs generated by different processing techniques. The ground-truth (the true change map) is given by Figure 3.44 (thin lines of change). The blue curve is for the ordinary HH CCD in Figure 3.51(a), obtained without any 3D processing. The green, red and cyan curves are for the HH, VV and polarimetrically filtered 3D CCDs in Figures 3.51(b), 3.51(c) and 3.51(d), respectively, obtained by applying the MVDR beamformer to three channels of the specified polarisation on each pass. For completeness, the same data are plotted in (a) and (b) using different scalings of the false-alarm axis.

of 0.5 at the aforementioned false-alarm rate. Using VV polarisation directly (red) gives a detection probability of 0.44.

It is worth reiterating that better detection performance could be achieved by a more sophisticated detection procedure tailored to the expected shape of the ground disturbance (e.g. lines). However, developing different detection methods is beyond the scope of this chapter. The approach taken here of just thresholding the coherence of each pixel independently is sufficient for comparing the essential utility of the CCD images produced by the different processing techniques. The ROC curves in Figures 3.52 and 3.53 quantify the ease by which the lines of change can be visually distinguished in the CCDs in Figures 3.45, 3.49 and 3.51. In a forest-obscured environment, the

techniques of adaptive SAR beamforming and polarimetric filtering together generate CCDs that permit much better detection of changes at ground level, compared to ordinary single-channel CCDs or beamformed CCDs generated using one standard polarisation.

3.13 Conclusion

Collection and processing methods have been described that would facilitate the novel concept of repeat-pass coherent change detection (CCD) on the ground under a forest canopy. The key requirement is that multiple SAR image channels at different grazing angles are acquired on each pass and coherently combined to obtain coarse vertical resolution that suppresses the scattering response of the canopy whilst preserving the complex scattering response of the ground. A minimal, but realistic, three-channel system would be a dual-antenna across-track interferometer operating in an alternating-transmit, simultaneous-receive mode to synthesise three effective phase centres at slightly different grazing angles. Additional requirements were analysed, notably that the operating wavelength must be both long enough to penetrate the canopy and short enough to be sensitive to changes on the ground, and that the radar system must achieve a high signal-to-noise performance (noise-equivalent sigma-nought: -36 dB) due to the weak ground backscatter response which is further attenuated during propagation through the canopy.

The forest scene was modelled simply as two layers: a volume of scattering elements, permitting lossy propagation, above a ground surface. Measurement of ground change was then formulated as a multichannel coherence estimation problem, where the ground coherence is taken as the desired indicator of change, and the volume coherence biases the output coherence estimate in a manner that depends on the repeat-pass collection geometry, the relative scattering intensity of the ground and volume layers, and the true ground coherence. The effect of the beamformer (the complex weight vector used to combine the multiple channels) was captured by a multichannel volume attenuation factor, which raises the effective ground-volume power ratio, making the estimation of ground coherence less sensitive to volume interference. The weight vector that optimises this factor was derived, and it was found to be equivalent to the usual MVDR (i.e. Capon) beamformer.

The random-volume-over-ground (RVOG) model, commonly used in the radar remote-sensing literature for forest parameter estimation, was repurposed to provide a framework for multichannel radar design, performance assessment and simulation. In particular, this model provides a simple way to generate values of volume coherence, parameterised by just the volume height and the attenuation rate of propagation through the volume. A key result is that, for a typical RVOG scene, the optimal three-channel beamformer can provide up to 12 dB of volume attenuation; if the ground scattering intensity is no more than 2.5 dB below that of the volume (after accounting for propagation loss), then this level of volume attenuation permits ground coherence estimation with error no greater than 0.1 (i.e. 10%). The sensitivity of beamformer performance to perturbations of the RVOG model was investigated in terms of the condition number of the matrix of volume coherences between channels; it was shown that this sensitivity is critically dependent on the design of the radar, that is, the number of channels and their spacing; as the spacing decreases, the condition number, and the sensitivity to perturbations, increase rapidly. Beamformer performance was also shown to be extremely sensitive to knowledge of the ground height—an offset of just a few metres between the focal surface used for SAR image formation and the true ground scattering surface leads to low estimates of ground coherence, because the adaptive beamformer steers a null onto the ground. However, this sensitivity can be exploited for ground height estimation by maximising the repeat-pass multichannel coherence of unchanged areas of the scene.

The proposed concept was demonstrated in simulation. The scattering response of RVOG clutter was synthesised as raw pulse echo samples using many randomly positioned point scatterers. SAR images of this clutter were formed at different grazing angles and combined using the conventional, adaptive-MVDR and fixed-RVOG beamformers, where the latter is clairvoyant, using assumed or estimated RVOG model parameters. The accuracy of the final repeat-pass coherence estimate was

found to depend on the accuracy of the model used to design the beamformer, to an extent determined by the radar design. The RVOG beamformer is highly sensitive to deviations from the RVOG model, making it impractical. The MVDR beamformer requires an accurate estimate of the covariance matrix, which implies a large number of independent looks; for the example here, an estimation window covering 3×3 resolution cells worked poorly, whereas a window covering 9×9 resolution cells worked well.

The proposed concept was also demonstrated by adapting PolSARproSim, an open-source program that synthesises polarimetric SAR images according to electromagnetic theory. The relatively strong scattering mechanism of ground-trunk double bounce was found to act as a source of high coherence, obscuring potential changes on the ground near trees. Since the phase centre of this scattering mechanism is located at ground height, vertical beamforming is not helpful. Instead, it was shown that this scattering mechanism could be suppressed by estimating its average polarimetric signature using a polarimetric decomposition of fully polarimetric data, and then projecting the data onto an orthogonal scattering mechanism, thereby filtering-out the ground-trunk double bounce.

Whilst implementation of the novel concept of CCD under a forest canopy faces numerous practical obstacles, this chapter has sought to show that reasonable performance is feasible.

Chapter 4

Polarimetric calibration of circularly polarised SAR data

4.1 Preface

This chapter is an extended version of the following published article:

Pincus, P., Preiss, M., Goh, A. S. & Gray, D., “Polarimetric calibration of circularly polarized synthetic aperture radar data”, *IEEE Transactions on Geoscience & Remote Sensing*, 55(12):6824-6839, Dec. 2017.

The extra material in this chapter is listed below.

- The calibration algorithm of Michelson *et al.* [1997] is cited and compared to the proposed algorithm. (The author became aware of their paper only recently.) See Section 4.2.
- The particular difficulty of achieving good angular alignment between a dihedral corner reflector deployed in the field and an airborne radar subject to wind is highlighted. See Section 4.2.
- Exact expressions for estimating the cross-talk parameters using a trihedral corner reflector in the circular basis are derived. The approximate expressions given in the published paper are then obtained from the exact expressions. See (4.50) and (4.51) in Section 4.5.
- Reflection symmetry is described in greater detail, including a brief summary of the experimental verification provided by Antar & Hendry [1985]. See before and after (4.53) in Section 4.6.
- The first-order distortion caused by cross-talk of the correlation between the co-polarised and cross-polarised responses from reflection-symmetric clutter is expressed in the linear basis as well as the circular basis. See (4.67) and (4.68) in Section 4.6.
- The eigendecomposition-based colourmap used to display polarimetric SAR images is explained in slightly more detail. See Appendix 4.A.1.

Minor changes to the literature review and the language have also been made.

4.2 Introduction

Since the first (ill-fated) NASA/JPL imaging radar polarimeter was flown in 1985 [Zebker *et al.* 1987], airborne and spaceborne synthetic aperture radar systems have been built that measure the whole scattering matrix in order to fully characterise the electromagnetic scattering response at every point in an illuminated landscape. These ‘fully polarimetric’ radars typically use a dual-polarised antenna system operated in an alternating-transmit, simultaneous-receive mode, which generates four polarimetric channels corresponding to the elements of the scattering matrix [Christensen *et al.* 1994, Stacy *et al.* 2003, Touzi & Shimada 2009]. These systems require calibration to ensure that the scattering behaviour is faithfully represented. Polarimetric calibration is the task of estimating and correcting for imbalances and impurities (cross-talk) in

the channels delivering the different polarimetric responses [Freeman 1992]. Absolute (radiometric) calibration is not considered here.

The first step in developing a radar calibration process is to select a suitable distortion model consisting of specific imbalance and cross-talk parameters that can be estimated from the data. Early work established the general matrix representation of a quad-channel radar system (shown in (4.8)) where the diagonal elements model four complex gains (two desired paths on transmit and two on receive) and the off-diagonal elements model four complex cross-talks (two leakage paths on transmit and two on receive) [Sheen *et al.* 1989]. Factoring-out two gains (for one reference transmit-receive path with some net absolute gain and phase) leaves two imbalance parameters and four cross-talk ratios as the standard polarimetric distortion parameters [Freeman 1991, Quegan 1994, Goh *et al.* 2007]. More elaborate models may be needed if the radar system is dynamic, such as when the receive gains are adjusted pulse-to-pulse and not properly compensated after sampling [Freeman 1991]. Simpler models may be justified for a monostatic radar if, for example, channel isolation is guaranteed everywhere except at the antennas; since the antennas are passive and behave reciprocally on transmit and receive, the system cross-talk will be reciprocal and can be represented by just two parameters, one for each antenna [Sarabandi & Ulaby 1990, Gau & Burnside 1995].

The second step is to estimate the distortion parameters. Initial approaches compared the observed data to the known scattering matrices of specially constructed corner reflectors placed in the scene [Barnes 1986, Yueh *et al.* 1990, Freeman 1992]. Soon after, the set of reference responses was augmented with the known structure of the covariance matrix for natural clutter that exhibits ‘reflection symmetry’, whereby the average scattering response is mirrored on either side of the vertical plane containing the line-of-sight axis [van Zyl 1990, Nghiem *et al.* 1992]. A single corner reflector gives a parameter estimate at a single position in the scene, whereas clutter permits parameter estimates over a wide area, which is useful given that the polarimetric distortion may vary with the orientations of the antennas relative to the ground, and therefore will vary over an image [Touzi *et al.* 1993]. The formulation of the clutter approach leads to non-linear equations for the distortion parameters, which have been solved most commonly by iterative estimation and correction using linearised approximations [Klein 1992, Ainsworth *et al.* 2006, López-Martínez *et al.* 2007, Goh *et al.* 2007]. In almost all cases, the theoretical scattering and covariance matrices, against which the observed data is compared, are specified in the linear basis i.e. the elements of the matrices are the scattering coefficients and correlations for horizontally and vertically polarised waves. This suits the standard case when the data is acquired using linearly polarised antenna elements.

This study extends the literature in two ways. Firstly, the standard system model of Freeman [1991] and Quegan [1994] is re-established in a unified framework for two generic radar designs which either support or do not support the assumption of cross-talk reciprocity, depending on whether the switch that toggles the transmit polarisation is placed before or after the main amplifier. Freeman [2009] informally outlined the simplification of the cross-talk calibration problem permitted by using a low-power switch design, but here both distortion models are formally derived in a consistent way, thus bringing together disparate works in the literature. A key step in the derivation is a new factorisation of the polarimetric distortion matrix, which enables the distortion expressions to be conveniently simplified when the cross-talk can be assumed reciprocal.

Secondly, a calibration process is developed for data acquired in the circular basis i.e. when the antenna elements are left and right circularly polarised and the reference matrices are for left and right circularly polarised waves. Although the utility of circular polarisation, and more broadly, basis diversity, has been widely studied in the imaging radar literature [Lee & Pottier 2009], the calibration of non-linear-basis data has seen little attention. Raney [2007] promoted the calibration advantages of a hybrid system where the transmitted signal is circularly polarised but the echo is received by two linearly polarised antennas. This study, however, is motivated by the *Ingara* airborne radar (see Section 1.3), which was modified to operate at L-band using two helical antennas that transmit and receive the two circular polarisations, as shown in Figure 4.1. The author knows of only one article that directly addresses the problem of calibrating an airborne imaging radar system that is circularly polarised; in that article, an estimation technique is proposed that uses corner reflectors, including two dihedrals oriented at different angles around the radar line-of-sight direction [Michelson

et al. 1997]. However, the response of a dihedral is particularly sensitive to angular misalignment: for the strong specular double bounce from a dihedral to return in the backscatter direction, and therefore be received by a monostatic radar, the dihedral seam must be perpendicular to the radar line-of-sight. Achieving good alignment between a corner reflector deployed in the field and an airborne radar subject to wind is difficult.




Figure 4.1: The helical antenna pair of the Defence Science and Technology Group's *Ingara* L-band airborne radar. The helices turn in opposite directions, giving orthogonal left and right circular polarisations.

In this chapter, both corner reflectors and clutter are considered for distortion estimation. In contrast to the linear-basis case, it is found here that the cross-talk parameters in the circular basis, whether assumed reciprocal or not, cannot be estimated from reflection-symmetric clutter using linearised approximations. An alternative approach using only corner reflectors is proposed, although it is limited to cases when cross-talk reciprocity is applicable. Gridded trihedrals [Sheen *et al.* 1992] can be used in place of dihedrals, thus alleviating the difficulty of alignment.

The proposed calibration steps are verified using Monte Carlo simulations and demonstrated using *Ingara* L-band radar data. The simulations enable the accuracy of the distortion estimation to be assessed over a wide, multidimensional parameter space [Goh *et al.* 2007], whereas tests on real data, whilst important, are inevitably limited by the number of available test scenarios and only approximate or assumed knowledge of the true (undistorted) scattering response. Moreover, since the proposed estimation method is iterative, it is important to show that it consistently converges to the true distortion values [Ainsworth *et al.* 2006]. It is not enough to show that some distortion parameters can be found which, when inverted, seem to approximately calibrate one or two datasets, because the underlying non-linear estimation problem may not have a unique solution, and so the potential solution may not correspond to the true system distortion. In common with the literature [Klein 1992, Ainsworth *et al.* 2006, López-Martínez *et al.* 2007], no proof is offered here that the distortion estimation problem has a unique solution, nor that the proposed technique is guaranteed to find it. Nonetheless, the success of the simulations over a wide parameter space indicates that the proposed technique reliably converges to the true distortion, if the cross-talk is small (< -10 dB total cross-talk power) and reciprocal.

To establish the context and explain the notation in this chapter, consider an ideal, fully polarimetric imaging radar system, which generates four image channels such that each pixel position is characterised by four complex values corresponding to the elements of the scattering matrix S . For arbitrary orthogonal polarisations p and q , the scattering matrix can be expressed in the p - q basis as [Graves 1956]

$$S^{p-q} = \begin{bmatrix} s_{pp} & s_{pq} \\ s_{qp} & s_{qq} \end{bmatrix} \quad (4.1)$$

where the two subscripts specify the polarisations of a response (first subscript) due to a stimulus (second subscript). For example, s_{pq} represents the complex coefficient for scattering of polarisation p arising from incident electromagnetic waves of polarisation q . In the linear basis, p - q is h - v (horizontal and vertical), and in the circular basis, p - q is l - r (left and right). Given a measurement of S in one polarimetric basis, the scattering response due to any other polarisation can be synthesised [van Zyl *et al.* 1987]. The observed scattering matrix \hat{S} delivered by a real radar system will suffer distortions due to the transmitters, receivers and antennas, which polarimetric calibration attempts to correct. These distortions will vary with direction and frequency; the latter dependence is not considered in this work. The distortions are expressed here in the framework established by Freeman [1991] and Quegan [1994] and refined by Goh *et al.* [2007].

In Section 4.3 the polarimetric distortion models are derived. In Section 4.4 the basic mathematical tools associated with polarimetric scattering and a polarimetric basis are summarised, and the effect of transforming distorted data between the linear and circular bases is analysed. Calibration techniques using corner reflectors and clutter are developed in Sections 4.5 and 4.6, respectively, with a particular emphasis on the circular basis. In Section 4.7 some of these calibration techniques are demonstrated using data from the *Ingara* L-band radar.

4.3 Polarimetric distortion models

Real radar systems are complex and idiosyncratic (i.e. somewhat unique) in their construction, but from the point of view of polarimetric calibration, the key issue is the set of different RF signal paths that different polarisations will take through the system. A common way to acquire fully polarimetric data with a dual-antenna pulsed radar is to alternate the transmit polarisation from pulse to pulse, but receive both polarisations for every pulse. Both the X-band and L-band variants of the *Ingara* radar system operate in this way. The system design can then be reduced, conceptually, to a simple set of amplifiers m , antennas a and one toggle switch s , together providing two paths through the system on transmit and two on receive. These paths would ideally be balanced (equal amplification and delay) and pure (allow only one polarisation); the distortion models will account for imbalance and impurity.

Figure 4.2 shows two possible models of the transmit path and one model of the receive path. These models of generic designs allow for independent scaling factors t and r on transmit and receive for each device (specified by a superscript m , s or a) along each polarimetric path (specified by a subscript p or q). The amplifiers scale the input signal i by factors t^m on transmit and the input (echo) signal $[i_p, i_q]^T$ by factors r^m on receive, with the two polarisations p and q potentially scaled differently, leading to channel imbalance. The antennas scale the signals by factors t^a and r^a , with different polarisations scaled differently, potentially causing both channel imbalance and cross-talk. For example, antenna P will transmit polarisation p with scaling factor t_{pp}^a and polarisation q with scaling factor t_{qp}^a (ideally zero). The switch alternates pulse-to-pulse between the two polarimetric paths on transmit (indicated by even pulse index τ and odd pulse index $\tau+1$).

In transmit model m - s - a shown in Figure 4.2(a), the switch (s) goes after the common amplifier (m), operating on the high-power signal. Typically this is some kind of PIN diode switch with imperfect isolation, so that when it is ‘open’ there may be non-negligible leakage down the polarimetric path which should not be transmitting. Assuming it toggles symmetrically, the high-power switch can be modelled by non-ideal scaling factors t_c^s and t_o^s when closed and open, respectively. In the alternative transmit model s - m - a , shown in Figure 4.2(b), the switch operates on the low-power signal, and there are separate amplifiers for the two polarimetric paths. In this case it may be assumed that the isolation is perfect and the switch has an ideal one-zero characteristic (to achieve this in practice, each amplifier should be unpowered when its associated signal path is switched off). The two designs demonstrate the trade-off between minimising the potential for channel imbalance by using common components for all channels, as in m - s - a , and minimising the potential for cross-talk by ensuring complete isolation between channels, as in s - m - a .

The models in Figure 4.2 and the analysis below assume that the radar system parameters are constant during acquisition. However, it is common practice to adjust the receive gains pulse-to-pulse, in order to make best use of the available dynamic range of the receiver. A simple example of this for linearly polarised antennas is where the gain for the cross-polarised channels is selected to be larger than that for the co-polarised channels, in order to approximately accommodate the expected difference (~ 6 dB) in scattering power. This differential gain is implemented by toggling the gains r_p^m and r_q^m pulse-to-pulse, in synchrony with the switching transmit polarisation. The *Ingara* radar system does this, as did the NASA/JPL DC-8 system [Freeman 1991]. Any deliberate dynamic gain control such as this must be separately compensated, prior to polarimetric calibration, because the receive model in Figure 4.2(c) assumes fixed scaling factors. If the gain adjustments have to be treated as unknowns, then they can potentially be estimated as part of a more elaborate distortion model [Freeman 1991].

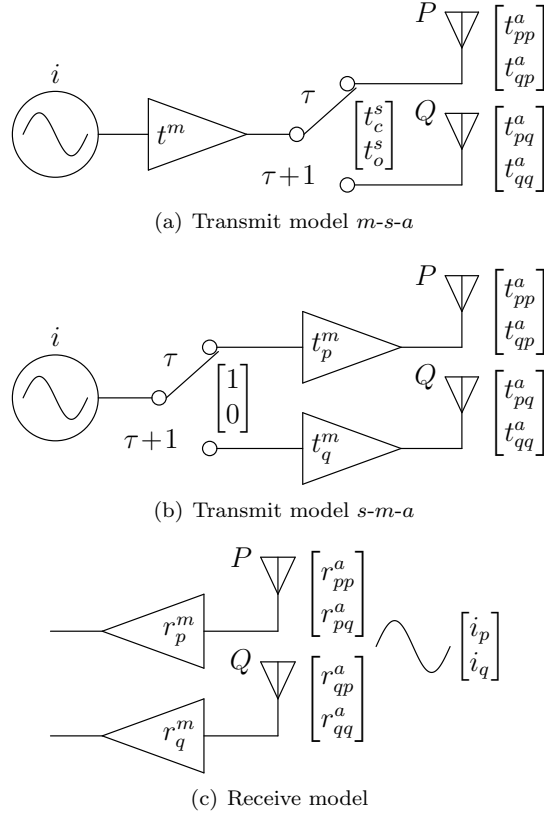


Figure 4.2: General models of the transmit and receive paths in typical fully polarimetric pulsed radar systems, accounting for polarimetric distortion.

Using transmit model m - s - a , the polarimetric vector signal generated on pulse τ is

$$\begin{bmatrix} t_c^m t_{pp}^a + t_o^m t_{pq}^a \\ t_c^m t_{qp}^a + t_o^m t_{qq}^a \end{bmatrix} i = \underbrace{\begin{bmatrix} t_{pp}^a & t_{pq}^a \\ t_{qp}^a & t_{qq}^a \end{bmatrix}}_{T_a} \begin{bmatrix} t_c^s \\ t_o^s \end{bmatrix} t^m i \quad (4.2)$$

and on pulse $\tau + 1$ is

$$\begin{bmatrix} t_o^m t_{pp}^a + t_c^m t_{pq}^a \\ t_o^m t_{qp}^a + t_c^m t_{qq}^a \end{bmatrix} i = \begin{bmatrix} t_{pp}^a & t_{pq}^a \\ t_{qp}^a & t_{qq}^a \end{bmatrix} \begin{bmatrix} t_o^s \\ t_c^s \end{bmatrix} t^m i. \quad (4.3)$$

For transmit model s - m - a , the signal generated on pulse τ is

$$\begin{bmatrix} t_p^m t_{pp}^a \\ t_p^m t_{qp}^a \end{bmatrix} i = \begin{bmatrix} t_{pp}^a & t_{pq}^a \\ t_{qp}^a & t_{qq}^a \end{bmatrix} \begin{bmatrix} t_p^m & 0 \\ 0 & t_q^m \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} i \quad (4.4)$$

and on pulse $\tau + 1$ is

$$\begin{bmatrix} t_q^m t_{pq}^a \\ t_q^m t_{qq}^a \end{bmatrix} i = \begin{bmatrix} t_{pp}^a & t_{pq}^a \\ t_{qp}^a & t_{qq}^a \end{bmatrix} \begin{bmatrix} t_p^m & 0 \\ 0 & t_q^m \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} i. \quad (4.5)$$

The received signal for all pulses is

$$\begin{bmatrix} r_p^m r_{pp}^a i_p + r_p^m r_{pq}^a i_q + n_p \\ r_q^m r_{qp}^a i_p + r_q^m r_{qq}^a i_q + n_q \end{bmatrix} = \begin{bmatrix} r_p^m & 0 \\ 0 & r_q^m \end{bmatrix} \underbrace{\begin{bmatrix} r_{pp}^a & r_{pq}^a \\ r_{qp}^a & r_{qq}^a \end{bmatrix}}_{R_a} \begin{bmatrix} i_p \\ i_q \end{bmatrix} + \begin{bmatrix} n_p \\ n_q \end{bmatrix} \quad (4.6)$$

where n_p and n_q allow for noise in each receive channel.

The antennas are passive and therefore can be assumed to behave reciprocally. Using antenna P as an example, the transmit gain t_{pp}^a and receive gain r_{pp}^a will be equal, and the transmit cross-talk t_{qp}^a and receive cross-talk r_{pq}^a will also be equal. Hence, the antenna distortion matrices T_a and R_a are related as follows [Sarabandi & Ulaby 1990, Touzi *et al.* 1993]:

$$R_a = T_a^T. \quad (4.7)$$

The antennas' responses will vary as a function of their orientation in azimuth and elevation (i.e. they will each have some beampattern), and therefore T_a and R_a are likely to vary with both range over the illuminated swath and azimuthal squint off the antenna boresight [Touzi *et al.* 1993]. Nonetheless, when forming the synthetic aperture image, the beampattern is usually assumed constant over the integration angle. (Similarly, the antennas' responses will vary with frequency, but they are typically assumed constant over the operating bandwidth.)

The transmitted and received signals are connected by scattering matrices of the form in (4.1) for all points in the illuminated scene. Standard SAR processing will split the pulse data into four channels covering the possible (transmit, receive) polarimetric combinations and produce four images in which each complex pixel value ostensibly corresponds to an element of the scattering matrix at the pixel position. Considering one position, it can be seen from (4.2)–(4.6) that a general model for the observed scattering matrix \tilde{S} is

$$\underbrace{\begin{bmatrix} \tilde{s}_{pp} & \tilde{s}_{pq} \\ \tilde{s}_{qp} & \tilde{s}_{qq} \end{bmatrix}}_{\tilde{S}} = \underbrace{\begin{bmatrix} r_{pp} & r_{pq} \\ r_{qp} & r_{qq} \end{bmatrix}}_R \underbrace{\begin{bmatrix} s_{pp} & s_{pq} \\ s_{qp} & s_{qq} \end{bmatrix}}_S \underbrace{\begin{bmatrix} t_{pp} & t_{pq} \\ t_{qp} & t_{qq} \end{bmatrix}}_T + \underbrace{\begin{bmatrix} n_{pp} & n_{pq} \\ n_{qp} & n_{qq} \end{bmatrix}}_N \quad (4.8)$$

where

$$\begin{aligned} r_{pp} &= r_p^m r_{pp}^a, & r_{pq} &= r_p^m r_{pq}^a \\ r_{qp} &= r_q^m r_{qp}^a, & r_{qq} &= r_q^m r_{qq}^a \end{aligned} \quad (4.9)$$

and, using transmit model *m-s-a*,

$$\begin{aligned} t_{pp} &= t^m (t_c^s t_{pp}^a + t_o^s t_{pq}^a), & t_{pq} &= t^m (t_c^s t_{pq}^a + t_o^s t_{pp}^a) \\ t_{qp} &= t^m (t_c^s t_{qp}^a + t_o^s t_{qq}^a), & t_{qq} &= t^m (t_c^s t_{qq}^a + t_o^s t_{qp}^a) \end{aligned} \quad (4.10)$$

or, using transmit model *s-m-a*,

$$\begin{aligned} t_{pp} &= t_p^m t_{pp}^a, & t_{pq} &= t_q^m t_{pq}^a \\ t_{qp} &= t_p^m t_{qp}^a, & t_{qq} &= t_q^m t_{qq}^a. \end{aligned} \quad (4.11)$$

The form of (4.8) matches the standard polarimetric data model used for calibration [van Zyl 1990, Freeman 1991, Quegan 1994]. Observe that the four output channels have separate realisations of additive receiver noise because the two receivers operate at two different times in order to measure the full scattering matrix: say, n_{pp} arises at (receive channel, time index) (P, τ) , n_{pq} at $(P, \tau+1)$, n_{qp} at (Q, τ) and n_{qq} at $(Q, \tau+1)$. For image post-processing, the speckle observed in radar images may be modelled statistically as multiplicative noise [Oliver & Quegan 1998 ch. 4.5, Lee & Pottier 2009 ch. 5.1.1], but for calibration, only the physical system should be modelled, and the dominant noise contributions are added in the receiver [Gray *et al.* 1990].

The observed scattering matrix in (4.8) can be vectorised using the Kronecker product \otimes , giving a single polarimetric distortion matrix $D = R \otimes T^T$.

$$\underbrace{\begin{bmatrix} \tilde{s}_{pp} \\ \tilde{s}_{pq} \\ \tilde{s}_{qp} \\ \tilde{s}_{qq} \end{bmatrix}}_{\tilde{\mathbf{s}}} = \underbrace{\begin{bmatrix} r_{pp} t_{pp} & r_{pp} t_{qp} & r_{pq} t_{pp} & r_{pq} t_{qp} \\ r_{pp} t_{pq} & r_{pp} t_{qq} & r_{pq} t_{pq} & r_{pq} t_{qq} \\ r_{qp} t_{pp} & r_{qp} t_{qp} & r_{qq} t_{pp} & r_{qq} t_{qp} \\ r_{qp} t_{pq} & r_{qp} t_{qq} & r_{qq} t_{pq} & r_{qq} t_{qq} \end{bmatrix}}_{D = R \otimes T^T} \underbrace{\begin{bmatrix} s_{pp} \\ s_{pq} \\ s_{qp} \\ s_{qq} \end{bmatrix}}_{\mathbf{s}} + \underbrace{\begin{bmatrix} n_{pp} \\ n_{pq} \\ n_{qp} \\ n_{qq} \end{bmatrix}}_{\mathbf{n}} \quad (4.12)$$

The scattering coefficients for natural landscapes may be treated as zero-mean complex normal random variables [Oliver & Quegan 1998 ch. 4.3], which can then be characterised by their variances and complex correlations, collectively captured by the covariance matrix $C = E\{\mathbf{ss}^H\}$. Assuming that the scatter and noise are uncorrelated, and that the noise channel components are uncorrelated with each other [van Zyl 1990], the observed covariance matrix is

$$\tilde{C} = E\{\tilde{\mathbf{s}}\tilde{\mathbf{s}}^H\} = E\{(D\mathbf{s} + \mathbf{n})(D\mathbf{s} + \mathbf{n})^H\} = DCD^H + C_n \quad (4.13)$$

where the noise covariance matrix $C_n = E\{\mathbf{nn}^H\}$ is diagonal with $E\{n_{pp}n_{pp}^*\} = E\{n_{pq}n_{pq}^*\}$ for receiver P and $E\{n_{qq}n_{qq}^*\} = E\{n_{qp}n_{qp}^*\}$ for receiver Q because noise realisations from the same receiver have the same average power.

The general linear distortion model D in (4.12) can be factorised to separate the different types of distortion: absolute gain and phase, channel imbalance, and cross-talk. This can be done in two ways, either as a column-based factorisation

$$D = yMAK \quad (4.14)$$

or as a row-based factorisation

$$D = yKAM \quad (4.15)$$

where

$$K = \begin{bmatrix} k^2 & 0 & 0 & 0 \\ 0 & k & 0 & 0 \\ 0 & 0 & k & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad A = \begin{bmatrix} \alpha & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \alpha & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad M = \begin{bmatrix} 1 & v & w & vw \\ z & 1 & wz & w \\ u & uv & 1 & v \\ uz & u & z & 1 \end{bmatrix} \quad (4.16)$$

and the absolute gain and phase is

$$y = D(4,4) = r_{qq}t_{qq}, \quad (4.17)$$

the channel imbalance parameters (ideally 1) are

$$k = D(2,2)/y = r_{pp}/r_{qq} \quad (4.18)$$

$$\alpha = D(3,3)/yk = \frac{t_{pp}/t_{qq}}{r_{pp}/r_{qq}} \quad (4.19)$$

and the cross-talk parameters (ideally 0) are, for factorisation $D = yMAK$ (4.14),

$$u = D(3,1)/yk^2\alpha = r_{qp}/r_{pp} \quad (4.20)$$

$$v = D(3,4)/y = t_{qp}/t_{qq} \quad (4.21)$$

$$w = D(2,4)/y = r_{pq}/r_{qq} \quad (4.22)$$

$$z = D(2,1)/yk^2\alpha = t_{pq}/t_{pp}, \quad (4.23)$$

or, for factorisation $D = yKAM$ (4.15),

$$u = D(3,1)/yk\alpha = r_{qp}/r_{qq} = r_{qp}^a/r_{qq}^a \quad (4.24)$$

$$v = D(3,4)/yk\alpha = t_{qp}/t_{pp} (= t_{qp}^a/t_{pp}^a \text{ for } s\text{-}m\text{-}a \text{ model only}) \quad (4.25)$$

$$w = D(2,4)/yk = r_{pq}/r_{pp} = r_{pq}^a/r_{pp}^a \quad (4.26)$$

$$z = D(2,1)/yk = t_{pq}/t_{qq} (= t_{pq}^a/t_{qq}^a \text{ for } s\text{-}m\text{-}a \text{ model only}). \quad (4.27)$$

Note that although the structure of the cross-talk matrix M in the two factorisations is the same, as specified in (4.16), the constituent cross-talk ratios are defined differently, as listed in (4.20)–(4.23) and (4.24)–(4.27).

The $D = yMAK$ (4.14) factorisation was first formulated by Quegan [1994] and was the basis of his own and others' calibration algorithms [Ainsworth *et al.* 2006, López-Martínez *et al.* 2007]. Goh *et al.* [2007] developed the formulation into a unified framework for the comparison of different polarimetric calibration algorithms, and showed that, given data in the linear basis exhibiting reflection symmetry, the task of calibrating for y and k can be completely separated from the estimation and correction of the other parameters. This formulation was subsequently used to calibrate data from the *Ingara* X-band system [Goh 2012], for which transmit model *m-s-a* is applicable. The *Ingara* L-band variant, however, uses separate transmit amplifiers for the two polarimetric channels, so transmit model *s-m-a* is applicable.

The $D = yKAM$ (4.15) factorisation is advantageous when transmit model *s-m-a* in Figure 4.2(b) is applicable because the amplifier factors cancel out from the cross-talk expressions, leaving just the antenna factors, as shown in (4.24)–(4.27). The total system, represented by T and R , is in general not reciprocal, but the antenna, represented by T_a and R_a , is reciprocal [Sarabandi & Ulaby 1990, Gau & Burnside 1995]. Combining (4.7) with (4.24)–(4.27), the number of cross-talk parameters reduces from four to two, with each one characterising the cross-talk of one antenna, as shown below.

$$\left. \begin{array}{l} \text{Transmit model } s\text{-}m\text{-}a \\ \text{Factorisation } D = yKAM \\ \text{Antenna reciprocity } R_a = T_a^T \end{array} \right\} \Rightarrow \begin{cases} v = w & \text{cross-talk parameter for antenna P} \\ u = z & \text{cross-talk parameter for antenna Q} \end{cases} \quad (4.28)$$

This result follows naturally from consideration of the transmit and receive structures in Figures 4.2(b) and 4.2(c), where the polarimetric paths are separate and isolated everywhere, except at the antennas. Since each antenna is reciprocal, its polarimetric leakage is reciprocal, and is characterised by just one cross-talk parameter. The $D = yKAM$ (4.15) factorisation simply enables this physical result to be represented mathematically.

In summary, two polarimetric distortion models have been derived. Firstly, the general model, consisting of six complex parameters: two channel imbalances, k and α , and four cross-talks, u , v , w and z . Secondly, the specific model, for when the conditions in (4.28) are met, consisting of four complex parameters: two channel imbalances as above, and two cross-talks, u and v . The absolute gain y is the radiometric distortion common to all channels, and although always present, its correction is not considered part of the polarimetric calibration problem here. Note that polarimetric calibration does not require the original eight complex perturbation parameters in (4.8) to be recovered; instead, only an estimate of D in (4.12), up to a complex scaling factor, is needed.

Selecting the $D = yKAM$ (4.15) factorisation, the observed data would be polarimetrically calibrated by first correcting for channel imbalance and then correcting for cross-talk:

$$\hat{\mathbf{s}} = \hat{M}^{-1} \hat{A}^{-1} \hat{K}^{-1} \hat{\mathbf{s}}. \quad (4.29)$$

Later sections will consider how the channel imbalance and cross-talk parameters can be estimated and how the estimation algorithms perform in the presence of each other and in the presence of noise. First, the mathematical framework for the circular basis will be established.

4.4 Polarimetric basis change

In this work, scattering matrices are defined using the backscatter alignment (BSA) convention, whereby the same Cartesian axes are used to define both the incident and scattered waves, along with the associated scattering coefficients [Lee & Pottier 2009 pp. 62-63]. With this convention, if the scattering is reciprocal then the cross-polarised responses will be equal i.e. $s_{pq} = s_{qp}$.

Given a scattering matrix S^{x-y} defined for Cartesian polarimetric components x and y , the polarimetric basis of S^{x-y} can be changed using a unitary basis-change matrix U_{pq} via a unitary consimilarity transformation [Lüneburg 1995], or equivalently, a unitary congruent transformation [Cloude 2010 p. 55], as follows:

$$S^{p-q} = U_{pq}^T S^{x-y} U_{pq}. \quad (4.30)$$

The unexpected, non-Hermitian form of (4.30) is due to the coordinate manipulations involved in constructing the BSA reference frame [Lüneburg 1995]. Importantly, if S^{x-y} is symmetric, then S^{p-q} will also be symmetric. Vectorising (4.30) using the Kronecker product \otimes , as in (4.12), gives the scattering vector and the covariance matrix in the new basis.

$$\mathbf{s}^{p-q} = U_{4pq} \mathbf{s}^{x-y} \quad \text{where } U_{4pq} = (U_{pq} \otimes U_{pq})^T \quad (4.31)$$

$$\therefore C^{p-q} = \mathbb{E}\{\mathbf{s}^{p-q}(\mathbf{s}^{p-q})^H\} = U_{4pq} C^{x-y} U_{4pq}^H \quad (4.32)$$

The columns of U_{pq} are the orthogonal polarisations p and q forming the new basis expressed in terms of the x and y components of the Cartesian basis. When U_{pq} is the identity matrix, then the polarimetric basis is a pair of linear polarisations, nominally horizontal ($p = h$) and vertical ($q = v$), aligned with the Cartesian axes x and y . When U_{pq} is

$$U_{lr} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & j \\ j & 1 \end{bmatrix} \quad (4.33)$$

then the polarimetric basis is a pair of circular polarisations, nominally left ($p = l$) and right ($q = r$), both of which have Cartesian components that are always equal in magnitude and 90° out-of-phase; for l , y leads x , and for r , x leads y . Note that Michelson *et al.* [1997] used an older definition of r as $[1, -j]^T$, leading to a basis-change matrix whose determinant is $-j$. The definition of the circular basis in (4.33) ensures

that polarisations l and r have a mutually consistent absolute phase, and that the basis-change matrix U_{lr} has unit determinant [Cloude 1994, Lee & Pottier 2009 pp. 41-43].

Converting distorted data acquired in a general p - q basis to the linear h - v basis (i.e. obtaining the polarimetric components along Cartesian axes x and y) leads to a modified distortion matrix \check{D} . To see this, combine the change of basis in (4.31) with the distortion model in (4.12), ignoring noise.

$$\begin{aligned} \check{\mathbf{s}}^{p-q} &= D\mathbf{s}^{p-q} \\ \rightarrow U_{4pq}^{-1}\check{\mathbf{s}}^{p-q} &= U_{4pq}^{-1}DU_{4pq}U_{4pq}^{-1}\mathbf{s}^{p-q} \\ \therefore \check{\mathbf{s}}^{x-y} &= \check{D}\mathbf{s}^{x-y} \quad \text{where } \check{D} = U_{4pq}^{-1}DU_{4pq} \end{aligned} \quad (4.34)$$

$$\therefore \check{C}^{x-y} = \mathbb{E}\{\check{\mathbf{s}}^{x-y}(\check{\mathbf{s}}^{x-y})^H\} = \check{D}C^{x-y}\check{D}^H \quad (4.35)$$

\check{D} in (4.34) can be interpreted in two ways. On the one hand, \check{D} can be treated as containing the distortion parameters $\{y, k, \alpha, u, v, w, z\}$ from the original p - q basis, but in a modified structure which incorporates the conversion to the linear basis. Using the row-based factorisation in (4.15),

$$\check{D} = U_{4pq}^{-1}DU_{4pq} = y \underbrace{U_{4pq}^{-1}KU_{4pq}}_{\check{K}} \underbrace{U_{4pq}^{-1}AU_{4pq}}_{\check{A}} \underbrace{U_{4pq}^{-1}MU_{4pq}}_{\check{M}} = y\check{K}\check{A}\check{M}. \quad (4.36)$$

On the other hand, \check{D} can itself be factorised via (4.17)–(4.19) and (4.24)–(4.27) to represent the distortion using parameters $\{y_{xy}, k_{xy}, \alpha_{xy}, u_{xy}, v_{xy}, w_{xy}, z_{xy}\}$ in the new linear basis (specified explicitly), but arranged in the standard structure of (4.16).

$$\check{D} = y_{xy}K_{xy}A_{xy}M_{xy} \quad (4.37)$$

Both interpretations of \check{D} can be equivalently expressed using the column-based factorisation in (4.14).

The particular focus of this study is when the acquired data is in the circular basis. In this case, U_{4lr} , obtained from (4.31) and (4.33), substantially modifies the distortion matrix via (4.34), such that the channel imbalance and cross-talk parameters constituting \check{D} in (4.37) (in the output linear basis) are each a complicated mixture of the channel imbalance and cross-talk parameters constituting D (in the input circular basis). In particular, what are usually small cross-talk values in the measurement basis could become large values in the new basis. To see this, consider distortion due to circular-basis channel imbalance k and α only: after applying U_{4lr} via (4.34) and factorising the resultant linear-basis distortion \check{D} via (4.14) or (4.15), there is no linear-basis channel imbalance ($k_{xy} = \alpha_{xy} = 1$) and the linear-basis cross-talk parameters are

$$u_{xy} = -w_{xy} = j \frac{1-k}{1+k} \quad \text{and} \quad z_{xy} = -v_{xy} = j \frac{1-k\alpha}{1+k\alpha}. \quad (4.38)$$

Observe that the magnitudes of the cross-talk parameters in the linear basis are strongly dependent on the phases of the channel imbalance parameters in the circular basis. When $|1+k|$ or $|1+k\alpha|$ are small, then $|u_{xy}|$ or $|v_{xy}|$ may be large.

Unfortunately, standard polarimetric calibration algorithms have been developed for the linear basis only, and they assume small cross-talk values to justify first-order approximations [Klein 1992, Quegan 1994, Ainsworth *et al.* 2006]. Simply converting the acquired circular-basis data to the linear basis and then applying a standard algorithm could be problematic if the effective cross-talk values after basis conversion are large. Instead, the proposed approach is to estimate and apply a calibration solution in the circular basis; at the very least, channel imbalance should be corrected so that $|1+k|$ and $|1+k\alpha|$ in (4.38) are not small.

After circular-basis calibration and conversion to the linear basis, a standard algorithm could still be applied to verify or refine the calibration solution. Any residual distortion $D_{xy}^{(r)}$ in the linear basis (typically estimated in terms of the column-based factorisation in (4.14)) can be converted back to the circular basis via

$$D^{(r)} = U_{4lr}D_{xy}^{(r)}U_{4lr}^{-1} \quad (4.39)$$

and refactorised to give parameter estimates which can be combined with the initial circular-basis estimates. Expanding (4.39) to express the desired circular-basis parameters in terms of the estimated linear-basis parameters shows that the transformation

for each circular-basis parameter requires all of the linear-basis parameters except for the absolute gain and phase y_{xy} ; it is not possible to transform just a subset of the relative distortion parameters.

4.5 Calibration via corner reflectors

A corner reflector is a man-made metallic structure which should provide a scattering response in accordance with a known theoretical scattering matrix [Duboc 1943]. The signal-to-noise ratio (SNR) for these targets is usually very large. By measuring the deviation of the acquired data from the theoretical scattering matrix, one or more polarimetric distortion parameters can be estimated. Each parameter estimate is valid for the particular antenna orientation used to illuminate the target.

Considering both the linear h - v and circular l - r polarimetric bases, the scattering matrices for a trihedral are

$$S_{th}^{h-v} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad S_{th}^{l-r} = j \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}; \quad (4.40)$$

for a dihedral are

$$S_{dh}^{h-v} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad S_{dh}^{l-r} = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}; \quad (4.41)$$

for a dihedral rotated by θ_{los} around the line-of-sight axis are

$$S_{rd}^{h-v} = \begin{bmatrix} \cos 2\theta_{los} & \sin 2\theta_{los} \\ \sin 2\theta_{los} & -\cos 2\theta_{los} \end{bmatrix}, \quad S_{rd}^{l-r} = \begin{bmatrix} e^{j2\theta_{los}} & 0 \\ 0 & -e^{-j2\theta_{los}} \end{bmatrix}; \quad (4.42)$$

and for a gridded trihedral [Sheen *et al.* 1992] (see photograph in Figure 4.8) observed at azimuth angle θ and grazing angle ψ are

$$S_{gt}^{h-v} = \frac{1}{m_{gt}^2} \begin{bmatrix} r_{gt}^2 & p_{gt} \\ p_{gt} & i_{gt}^2 \end{bmatrix}, \quad S_{gt}^{l-r} = \frac{1}{2} \begin{bmatrix} c_{gt}^2/m_{gt}^2 & j \\ j & -c_{gt}^{*2}/m_{gt}^2 \end{bmatrix} \quad (4.43)$$

where $r_{gt} = \sin \theta$, $i_{gt} = -\cos \theta \sin \psi$, $p_{gt} = r_{gt} i_{gt}$, $c_{gt} = r_{gt} + j i_{gt}$ and $m_{gt}^2 = r_{gt}^2 + i_{gt}^2$. Note that the circular-basis expressions for the dihedral and rotated dihedral given by Lee & Pottier [2009 p. 95] have missing negative signs; see Cloude [2010 p. 56] for the correct expressions. Observe from S_{rd}^{l-r} in (4.42) that the magnitude of the dihedral response in the circular basis does not depend on the line-of-sight rotation θ_{los} : the co-polarised components are always equal in magnitude and there is never any cross-polarised component. For the gridded trihedral response S_{gt}^{l-r} in the circular basis, all polarimetric components have equal magnitude ($1/2$), and the phases of the co-polarised elements sum to 180° , regardless of the angular collection geometry (θ, ψ) .

From (4.29), the data should be corrected for channel imbalance first, and since in practice this is usually the more significant of the two types of distortion, it is proposed to first estimate and correct for channel imbalance, and then estimate and correct for cross-talk.

Reconsider the observed scattering vector in (4.12), assuming that the cross-talk and noise are negligible; later the estimation process will be tested in the presence of each type of perturbation. Selecting the $D = yKAM$ (4.15) factorisation, the observed scattering vector is

$$\underbrace{\begin{bmatrix} \tilde{s}_{pp} \\ \tilde{s}_{pq} \\ \tilde{s}_{qp} \\ \tilde{s}_{qq} \end{bmatrix}}_{\tilde{\mathbf{s}}} = \underbrace{\begin{bmatrix} yk^2 \alpha s_{pp} \\ yk s_{pq} \\ yk \alpha s_{qp} \\ y s_{qq} \end{bmatrix}}_{yKAs}. \quad (4.44)$$

The channel imbalance parameters can then be estimated using ratios of the responses from appropriately selected corner reflectors, as listed below for both linear and circular

bases.

$$\hat{\alpha} = \frac{\tilde{s}_{qp}/\tilde{s}_{pq}}{s_{qp}/s_{pq}} \quad (\text{use } S_{rd}^{h-v} \text{ or } S_{gt}^{h-v}; \quad S_{th}^{l-r} \text{ or } S_{gt}^{l-r}) \quad (4.45)$$

$$\hat{k} = \pm \sqrt{\frac{\tilde{s}_{pp}/\tilde{s}_{qq}}{\hat{\alpha}s_{pp}/s_{qq}}} \quad (\text{use any except for } S_{th}^{l-r}) \quad (4.46)$$

$$\hat{k} = \frac{\tilde{s}_{pp}/\tilde{s}_{qp}}{s_{pp}/s_{qp}} \quad (\text{use } S_{rd}^{h-v} \text{ or } S_{gt}^{h-v}; \quad S_{gt}^{l-r}) \quad (4.47)$$

$$\hat{k} = \frac{\tilde{s}_{pq}/\tilde{s}_{qq}}{s_{pq}/s_{qq}} \quad (\text{use } S_{rd}^{h-v} \text{ or } S_{gt}^{h-v}; \quad S_{gt}^{l-r}) \quad (4.48)$$

The sign ambiguity in the estimate of \hat{k} obtained via (4.46) can be resolved by comparison with the estimates from (4.47) and (4.48) [Goh 2012 p. 150].

Trihedral corner reflectors are to be preferred, when appropriate, because their response is approximately constant over a wide angular range (see the theoretical and measured plots by Sarabandi & Chiu [1996]), whereas the responses of gridded trihedrals and dihedrals are in some way sensitive to the orientation $(\theta, \psi, \theta_{los})$ of the antenna relative to the target, and this relative orientation can be difficult to control or accurately measure for an airborne radar platform subject to wind. In particular, for the backscattering response of a dihedral to be in accordance with (4.41) or (4.42), which express the polarimetric components that arise due to specular double bounce, the dihedral seam must be perpendicular to the radar line-of-sight.

Correcting the observed data for channel imbalance gives an intermediate estimate $\hat{\mathbf{s}}'$, which still suffers from cross-talk.

$$\underbrace{\begin{bmatrix} \hat{s}'_{pp} \\ \hat{s}'_{pq} \\ \hat{s}'_{qp} \\ \hat{s}'_{qq} \end{bmatrix}}_{\hat{\mathbf{s}}'} = \underbrace{\begin{bmatrix} \tilde{s}_{pp}/\hat{k}^2\hat{\alpha} \\ \tilde{s}_{pq}/\hat{k} \\ \tilde{s}_{qp}/\hat{k}\hat{\alpha} \\ \tilde{s}_{qq} \end{bmatrix}}_{\hat{A}^{-1}\hat{K}^{-1}\hat{\mathbf{s}}} \approx y \underbrace{\begin{bmatrix} s_{pp} + vs_{pq} + ws_{qp} + vws_{qq} \\ s_{pq} + zs_{pp} + ws_{qq} + wzs_{qp} \\ s_{qp} + us_{pp} + vs_{qq} + uvs_{pq} \\ s_{qq} + us_{pq} + zs_{qp} + uzs_{pp} \end{bmatrix}}_{yMs} \quad (4.49)$$

Now consider a corner reflector whose scattering matrix is symmetric with zero co-polarisation. In the linear basis, a dihedral rotated by 45° would satisfy this requirement, and in the circular basis, a trihedral would be suitable. The ratio of the co-polarised and cross-polarised responses will also be zero; non-zero values arise due to cross-talk, specifically, leakage of the co-polarised component. If transmit model $s\text{-}m\text{-}a$ is applicable, then (4.28) can be invoked, leading to estimates of the two reciprocal cross-talk parameters:

$$\left. \begin{array}{l} s_{pp} = s_{qq} = 0 \\ s_{pq} = s_{qp} \\ v = w \\ u = z \end{array} \right\} \Rightarrow \left. \begin{array}{l} r_1 = \frac{\hat{s}'_{pp}}{\hat{s}'_{pq}} = \frac{v+w}{1+wz} = \frac{2v}{1+uv} \\ r_2 = \frac{\hat{s}'_{qq}}{\hat{s}'_{qp}} = \frac{u+z}{1+uv} = \frac{2u}{1+uv} \end{array} \right\} \Rightarrow \begin{array}{l} r_1 u^2 - 2u + r_2 = 0 \Rightarrow \hat{u} = \frac{\kappa}{r_1} \approx \frac{r_2}{2} \\ r_2 v^2 - 2v + r_1 = 0 \Rightarrow \hat{v} = \frac{\kappa}{r_2} \approx \frac{r_1}{2} \end{array} \quad (4.50)$$

where $\kappa = 1 - \sqrt{1 - r_1 r_2} \approx r_1 r_2 / 2$. The approximation is justified by the assumption that the co-polarised leakage is small compared to the desired cross-polarised signal i.e. $\hat{s}'_{pp}\hat{s}'_{qq} \ll \hat{s}'_{pq}\hat{s}'_{qp}$. (For the same reason, the smaller of the two solutions to each of the quadratic equations is selected.) Although not strictly required, the approximation is attractive because it ties each parameter to one physical type of cross-talk: \hat{u} measures leakage of polarisation q and \hat{v} measures leakage of polarisation p . Note that the leakage must still be large compared to the noise, so that it can actually be measured from the imagery.

Correcting the intermediate data for cross-talk gives a vectorised estimate $\hat{\mathbf{s}} = \hat{M}^{-1}\hat{\mathbf{s}}'$ of the true scattering matrix.

The estimates of the channel imbalance and cross-talk parameters can be refined by repeating the total estimation-correction process until the values have stabilised. This is useful because the channel imbalance is being estimated in the presence of cross-talk, and the cross-talk is being estimated using first-order approximations. Note that standard polarimetric calibration algorithms, although motivated by different physical

concepts, similarly involve first-order linearisations of cross-talk expressions [Klein 1992, Quegan 1994, Ainsworth *et al.* 2006], and are usually also iterative; none of the published iterative algorithms is accompanied by proofs of uniqueness or convergence, although usually they do converge to the correct values [Goh *et al.* 2007].

The intermediate parameter estimates from iterations $n = 1, 2, \dots, N$ can be accumulated using the formulae below to give estimates of the distortion parameters for the original data. The expressions were derived by linearising the expansion of multiple distortion matrices, and are equivalent to those in Goh *et al.* [2007] when $\hat{k}_n = 1 \forall n$.

$$\begin{aligned} \hat{k} &= \prod_{n=1}^N \hat{k}_n, & \hat{u} &= \hat{k} \sum_{n=1}^N \frac{1}{\prod_{m=1}^n \hat{k}_m} \hat{u}_n, & \hat{v} &= \frac{1}{\hat{\alpha} \hat{k}} \sum_{n=1}^N \prod_{m=1}^n \hat{\alpha}_m \hat{k}_m \hat{v}_n, \\ \hat{\alpha} &= \prod_{n=1}^N \hat{\alpha}_n, & \hat{w} &= \frac{1}{\hat{k}} \sum_{n=1}^N \prod_{m=1}^n \hat{k}_m \hat{w}_n, & \hat{z} &= \hat{\alpha} \hat{k} \sum_{n=1}^N \frac{1}{\prod_{m=1}^n \hat{\alpha}_m \hat{k}_m} \hat{z}_n \end{aligned} \quad (4.52)$$

The preceding equations suggest that the four cross-talk parameters will be distinct. However, when the cross-talk satisfies (4.28) i.e. $u = z$ and $v = w$, then the estimate $\hat{\alpha}$ obtained via (4.45) using any symmetric cross-polarising corner reflector will in fact not be perturbed by this reciprocal cross-talk (this can be shown algebraically), so the first iteration should give the exact value of α , and subsequent iterations $n > 1$ should give $\hat{\alpha}_n = 1$ (ignoring noise), leading to accumulated estimates $\hat{u} = \hat{z}$ and $\hat{v} = \hat{w}$ in (4.52) as desired.

Monte Carlo simulations were conducted to verify the proposed estimation techniques. For each trial, the ideal, circular-basis scattering responses for a trihedral, dihedral and gridded trihedral ($\theta = 45^\circ, \psi = 30^\circ$) were distorted by random values of y, k, α, u, v, w, z and, when specified, noise N . In the figures below, total cross-talk power is $|u|^2 + |v|^2 + |w|^2 + |z|^2$ and total noise power is $|n_{pp}|^2 + |n_{pq}|^2 + |n_{qp}|^2 + |n_{qq}|^2$, both relative to unity signal power from the ideal trihedral.

Figure 4.3 shows how non-iterative estimation of channel imbalance in the circular basis performs in the presence of cross-talk. The estimation accuracy using trihedral and dihedral corner reflectors is high: the relative error is less than 5% for total cross-talk powers up to -10 dB. The gridded trihedral is much more sensitive to cross-talk, and estimation of α in particular is prone to suffer error: for the relative error in $\hat{\alpha}$ to be less than 5%, the total cross-talk power must be less than -31 dB, whereas for \hat{k} , the total cross-talk power must be less than -25 dB. Hence, when estimating α , a trihedral should be used, and when estimating k , there is a trade-off between the better cross-talk performance of the dihedral and the wider beamwidth (enabling easier acquisition) of the gridded trihedral.

Figure 4.4 shows how iterative estimation of channel imbalance and cross-talk in the circular basis performs when the cross-talk is reciprocal i.e. $u = z$ and $v = w$. Importantly, even though the initial estimates of the distortions may have significant error (because they mutually perturb each other and the cross-talk estimates depend on first-order approximations), the subsequent estimates reliably converge to give the true parameter values, with the remaining error after just two iterations less than 1% for total cross-talk powers up to -10 dB. As mentioned earlier, $\hat{\alpha}$ is not perturbed by reciprocal cross-talk, so the initial estimate should match the true value.

Figure 4.5 shows how the iterative estimation process performs when complex noise N is added to the distorted corner reflector responses. The iterations reduce the estimation error to a floor set by the noise level, and this holds regardless of the choice of corner reflector for estimating k , even though the sizes of the errors on the initial distortion estimates may be quite different. When the total noise power is less than -45 dB and the total cross-talk leakage is -20 dB, then the cross-talk estimation error after two iterations is less than 5%.

4.6 Calibration via clutter

Corner reflectors provide estimates of the distortion parameters which are valid only for the particular antenna orientations used to illuminate the targets, but, as noted in Section 4.3, the system distortion may vary with antenna orientation. This is




Figure 4.3: Simulation results using corner reflectors for circular-basis channel imbalance estimation in the presence of cross-talk, without added noise. 1000 Monte Carlo trials were conducted with uniform random realisations of the magnitude and phase of all distortion parameters in the limited ranges $0.5 \leq |y|, |k|, |\alpha| \leq 1.5$ and $|u|, |v|, |w|, |z| \leq 0.3$ (four independent cross-talk distortions were applied), with all phases unrestricted. In each trial, α was estimated via (4.45) using the trihedral (blue cross) and gridded trihedral (red dot), whereas k was estimated via (4.46) (along with (4.47)–(4.48) for sign resolution) using the dihedral (blue cross) and a gridded trihedral (red dot). The estimates were made without iteration i.e. without cross-talk correction.




Figure 4.4: Simulation results using corner reflectors for circular-basis channel imbalance and cross-talk estimation in the presence of reciprocal cross-talk, without added noise. 1000 Monte Carlo trials were conducted as in Fig. 4.3 except that $u = z$ and $v = w$. In each trial, first the channel imbalance was estimated (α using the trihedral, k using the gridded trihedral) and corrected, and then the cross-talk was estimated (via (4.50)–(4.51) using the trihedral, assuming cross-talk reciprocity) and corrected. The estimation-correction process was repeated twice more, with successive intermediate parameter estimates accumulated via (4.52). The estimation errors for \hat{k} and \hat{u} are plotted in (a) and (b), respectively, for the first (red dot), second (blue cross) and third (green diagonal cross) iterations (the error for \hat{k} on the third iteration was well-below $10^{-3}\%$; the error for \hat{v} had the same shape as that for \hat{u}).

problematic when the antennas are fixed to the platform, which may yaw or roll during flight. Furthermore, corner reflectors may not respond according to their theoretical ideal due to mutual coupling with the ground, misalignment with the radar antenna, imperfect construction or simply insufficient size relative to the operating wavelength [Sarabandi *et al.* 1994, 1995]. These limitations can be avoided by using a scattering model for natural clutter observed over a large area, enabling calibration for a range




Figure 4.5: Simulation results using corner reflectors for circular-basis cross-talk estimation in the presence of reciprocal cross-talk and noise. 1000 Monte Carlo trials were conducted with different realisations of additive complex noise N and fixed distortion parameters: $y = 1\angle 0^\circ$, $k = 1.35\angle 72^\circ$, $\alpha = 0.6\angle -105^\circ$, $u = z = 0.05\angle 37^\circ$, $v = w = 0.05\angle 61^\circ$ (total cross-talk power -20 dB). In each trial, the iterative process of estimation and correction for channel imbalance and cross-talk was undertaken, as in Fig. 4.4, except that k was estimated in one of two ways: using a dihedral (a) or a gridded trihedral (b). The estimation error for \hat{u} is plotted for the first (red dot) and second (blue cross) iterations (the third showed negligible change).

of azimuth and depression angles off the antenna boresight.

Some natural landscapes exhibit reflection symmetry, whereby the average polarimetric scattering response is mirrored on either side of the vertical plane containing the line-of-sight axis [Nghiem *et al.* 1992]. (Azimuthal symmetry extends reflection symmetry by requiring that every plane containing the line-of-sight axis serves as a mirror plane.) Examples include ploughed fields viewed perpendicular to the furrows and homogeneous volumes such as agricultural crops, forest and snow, viewed from any direction [Nghiem *et al.* 1992]. In addition, natural landscapes usually scatter reciprocally. Summing the independent contributions from both sides of the mirror plane leads to a simple structure for the covariance matrix C_{rs}^{h-v} of reciprocal and reflection-symmetric clutter in the linear basis whose vertical (horizontal) component is parallel (perpendicular) to the plane of symmetry [Nghiem *et al.* 1992, Cloude & Pottier 1996]:

$$C_{rs}^{h-v} = \begin{bmatrix} \sigma_{hh}^2 & 0 & 0 & \sigma_{hhvv} \\ 0 & \sigma_{hv}^2 & \sigma_{hv}^2 & 0 \\ 0 & \sigma_{hv}^2 & \sigma_{hv}^2 & 0 \\ \sigma_{hhvv}^* & 0 & 0 & \sigma_{vv}^2 \end{bmatrix}. \quad (4.53)$$

The fact that the co-polarised and cross-polarised responses are uncorrelated for a reflection-symmetric scene was verified by Antar & Hendry [1985], who measured the polarimetric scattering response of rain at X-band using a dual-linearly-polarised radar whose basis could be rotated with a phase-shifter, and found that the co-cross-correlation exhibited a deep null when the polarisations were parallel and perpendicular to the mean orientation of the rain drops. Their observation was noted by [van Zyl 1990], who first proposed exploiting this characteristic for polarimetric calibration. Note that if a radar illuminates terrain with some slope along azimuth (i.e. in the along-track direction), then the mirror plane will be rotated off the vertical, and the clutter will not exhibit reflection symmetry. To accommodate this in their polarimetric calibration algorithm, Ainsworth *et al.* [2006] allowed the co-cross-correlations to take non-zero values; this extension is not considered here.

Changing the polarimetric basis of C_{rs}^{h-v} in (4.53) via (4.32) using U_{4lr} obtained from

(4.31) and (4.33), reciprocity and reflection symmetry in the circular basis take the form

$$\begin{aligned}
C_{rs}^{l-r} &= \begin{bmatrix} \sigma_{ll}^2 & \sigma_{llr} & \sigma_{lrl} & \sigma_{llrr} \\ \sigma_{lrll} & \sigma_{lr}^2 & \sigma_{lrll} & \sigma_{lrrr} \\ \sigma_{rlll} & \sigma_{rllr} & \sigma_{rl}^2 & \sigma_{rlrr} \\ \sigma_{rrll} & \sigma_{rrlr} & \sigma_{rrrl} & \sigma_{rr}^2 \end{bmatrix} \quad (\text{Hermitian}) \\
&= \begin{bmatrix} \varsigma_s - \xi_s + \chi & \xi_d - j\varsigma_d & \xi_d - j\varsigma_d & -\varsigma_s + \xi_s + \chi \\ \xi_d + j\varsigma_d & \varsigma_s + \xi_s & \varsigma_s + \xi_s & -\xi_d - j\varsigma_d \\ \xi_d + j\varsigma_d & \varsigma_s + \xi_s & \varsigma_s + \xi_s & -\xi_d - j\varsigma_d \\ -\varsigma_s + \xi_s + \chi & -\xi_d + j\varsigma_d & -\xi_d + j\varsigma_d & \varsigma_s - \xi_s + \chi \end{bmatrix} \quad (4.54)
\end{aligned}$$

where

$$\begin{aligned}
\varsigma_s &= (\sigma_{hh}^2 + \sigma_{vv}^2)/4, & \xi_s &= (\sigma_{hhvv} + \sigma_{hhvv}^*)/4, & \chi &= \sigma_{hv}^2 \\
\varsigma_d &= (\sigma_{hh}^2 - \sigma_{vv}^2)/4, & \xi_d &= -j(\sigma_{hhvv} - \sigma_{hhvv}^*)/4
\end{aligned} \quad (4.55)$$

are all real. Now instead of zero co-cross-correlation, reflection symmetry leads to

$$\sigma_{ll}^2 = \sigma_{rr}^2, \quad \Im\{\sigma_{llrr}\} = 0 \quad \text{and} \quad \sigma_{llr} = \sigma_{lrl} = -\sigma_{lrr}^* = -\sigma_{rllr}^*. \quad (4.56)$$

Note that for a covariance matrix to be feasible, the leading diagonal elements must be non-negative, and that places additional restrictions on the linear-basis terms constituting C_{rs}^{l-r} in (4.54):

$$-\frac{1}{2}(\sigma_{hh}^2 + \sigma_{vv}^2) \leq \Re\{\sigma_{hhvv}\} \leq \frac{1}{2}(\sigma_{hh}^2 + \sigma_{vv}^2) + 2\sigma_{hv}^2. \quad (4.57)$$

Reconsider the observed covariance matrix in (4.13) for reciprocal and reflection-symmetric clutter, assuming that the cross-talk and noise are negligible; later the estimation process will be tested in the presence of each type of perturbation. Selecting factorisation $D = yKAM$ (4.15), the observed covariance matrix in the circular-basis is

$$\begin{aligned}
\tilde{C}_{rs}^{l-r} &= |y|^2 \underbrace{\begin{bmatrix} |k|^4 |\alpha|^2 \sigma_{ll}^2 & |k|^2 k \alpha \sigma_{llr} & |k|^2 k |\alpha|^2 \sigma_{llr} & k^2 \alpha \sigma_{llrr} \\ |k|^2 k^* \alpha^* \sigma_{llr}^* & |k|^2 \sigma_{lr}^2 & |k|^2 \alpha^* \sigma_{lr}^2 & -k \sigma_{llr}^* \\ |k|^2 k^* |\alpha|^2 \sigma_{llr}^* & |k|^2 \alpha \sigma_{lr}^2 & |k|^2 |\alpha|^2 \sigma_{lr}^2 & -k \alpha \sigma_{llr}^* \\ (k^2)^* \alpha^* \sigma_{llrr} & -k^* \sigma_{llr} & -k^* \alpha^* \sigma_{llr} & \sigma_{ll}^2 \end{bmatrix}}_{(yKA)C_{rs}^{l-r}(yKA)^H}. \quad (4.58)
\end{aligned}$$

The channel imbalance parameters can then be estimated by matching the observed response to the theoretical covariance matrix, as listed below.

$$|\hat{\alpha}| = \sqrt{\tilde{\sigma}_{rl}^2 / \tilde{\sigma}_{lr}^2} \quad (\text{balance the cross-pol. powers for reciprocity}) \quad (4.59)$$

$$\arg\{\hat{\alpha}\} = \arg\{\tilde{\sigma}_{llr}\} \quad (\text{enforce reciprocity}) \quad (4.60)$$

$$|\hat{k}| = \sqrt[4]{\frac{\tilde{\sigma}_{ll}^2 / \tilde{\sigma}_{rr}^2}{|\hat{\alpha}|^2}} \quad (\text{balance the co-pol. powers for reflection symmetry}) \quad (4.61)$$

$$\arg\{\hat{k}\} = \frac{1}{2} \arg\{-\tilde{\sigma}_{llrl} / \tilde{\sigma}_{rrlr}\} + n\pi \quad (\text{where } n = 0, \pm 1, \pm 2, \dots) \quad (4.62)$$

$$\arg\{\hat{k}\} = \frac{1}{2} \arg\left\{-\frac{\tilde{\sigma}_{llr} / \tilde{\sigma}_{rrrl}}{\hat{\alpha} / \hat{\alpha}^*}\right\} + n\pi \quad (4.63)$$

$$\arg\{\hat{k}\} = \frac{1}{2} \arg\left\{\frac{\tilde{\sigma}_{llrr}}{\hat{\alpha}}\right\} + \frac{1 - \text{sgn}\sigma_{llrr}}{2} \cdot \frac{\pi}{2} + n\pi \quad (\text{sgn}\sigma_{llrr} = \pm 1 \text{ is unknown}) \quad (4.64)$$

Unfortunately, the phases of \hat{k} in (4.62) and (4.63) have 180° ambiguities due to the k^2 dependence, and (4.64) has a 90° ambiguity due to the additional unknown sign of σ_{llrr} . Furthermore, (4.64) depends on the value of the co-polarisation correlation $\sigma_{llrr} = -(\sigma_{hh}^2 + \sigma_{vv}^2)/4 + (\sigma_{hhvv} + \sigma_{hhvv}^*)/4 + \sigma_{hv}^2$ (see (4.54)–(4.55)), which may have a small magnitude and a noisy phase. A reference value of k , previously obtained from corner reflectors, is needed to resolve the sign.

Note that trying the same approach in the linear basis, with $\tilde{C}_{rs}^{h-v} = (yKA)C_{rs}^{h-v}(yKA)^H$, leads to the well-known estimate of α [Ainsworth *et al.* 2006], but no way of estimating k .

Figure 4.6 shows how the estimation of channel imbalance performs in the presence of cross-talk and noise. For small perturbations, the clutter-based estimation process is accurate, but Figure 4.6(a) shows that it is highly sensitive to cross-talk: the relative error is less than 5% only for total cross-talk powers up to about -32 dB. Compared to corner reflector-based estimation using trihedrals and dihedrals, shown in Figure 4.3, clutter-based estimation of the circular-basis channel imbalance parameters is significantly more sensitive to cross-talk.

Figure 4.6(b) depicts the case where the cross-talk is reciprocal and fixed to be very small (total leakage power -60 dB), but the noise is allowed to vary up to the nominal clutter power. The estimation of α is seen to not be affected by reciprocal cross-talk (although it is affected by non-reciprocal cross-talk, as shown in 4.6(a)), so is always noise-limited, but the estimation of k is limited by the cross-talk when the noise is low, with the error floor indicated by 4.6(a).

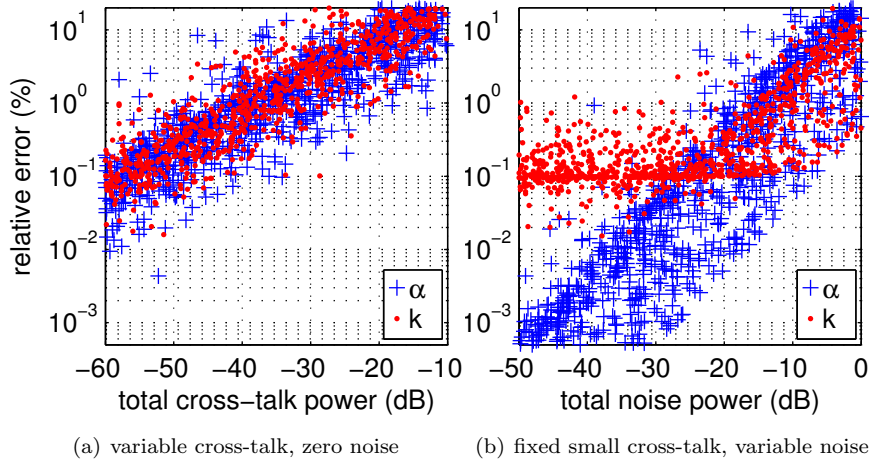


Figure 4.6: Simulation results using clutter for channel imbalance estimation in the presence of cross-talk and noise. 1000 Monte Carlo trials were conducted, beginning with uniform random realisations of the linear-basis, reflection-symmetric covariance matrix C_{rs}^{h-v} in (4.53), where $-20 \leq \sigma_{hh}^2 \leq 20$ dB, $-20 \leq \sigma_{hv}^2 \leq 0$ dB, $-20 \leq |\sigma_{hhvv}|^2 \leq 0$ dB ($0 \leq \angle \sigma_{hhvv} < 360^\circ$), all relative to $\sigma_{vv}^2 = 1$. In each trial, the covariance matrix was converted to the circular basis and perturbed according to (4.13): in (a), the polarimetric distortion parameters had magnitude and phase selected from uniform random distributions over the limited ranges $0.5 \leq |y|, |k|, |\alpha| \leq 1.5$ and $|u|, |v|, |w|, |z| \leq 0.3$ (four independent cross-talk distortions were applied), with all phases unrestricted, and no noise was added, whereas in (b), the distortion parameters were fixed to $y = 1 \angle 0^\circ$, $k = 1.35 \angle 72^\circ$, $\alpha = 0.6 \angle -105^\circ$, $u = z = 0.0005 \angle 37^\circ$ and $v = w = 0.0005 \angle 61^\circ$ (i.e. reciprocal cross-talk causing total leakage power -60 dB), and the noise powers in C_n were varied. Channel imbalance parameters α (blue cross) and k (red dot) were then estimated via (4.59)–(4.62), with the sign of k selected using the truth. The plots show the relative errors $|(\hat{\alpha} - \alpha)/\alpha|$ and $|(\hat{k} - k)/k|$.

Assuming the channel imbalance estimates are reasonable, the observed circular-basis data can be partially corrected, giving an intermediate estimate $\hat{C}'_{rs}{}^{l-r}$, which still suffers from cross-talk.

$$\hat{C}'_{rs}{}^{l-r} = \hat{A}^{-1} \hat{K}^{-1} \tilde{C}'_{rs}{}^{l-r} \hat{K}^{-H} \hat{A}^{-H} \approx |y|^2 M C_{rs}^{l-r} M^H \quad (4.65)$$

After converting to the linear basis, the intermediate data can be expressed in terms of the modified distortion \check{M} , as in (4.36), where the circular-basis cross-talk is incorporated with the change-of-basis matrices, giving

$$\hat{C}'_{rs}{}^{h-v} = U_{4lr}^{-1} \hat{C}'_{rs}{}^{l-r} U_{4lr}^{-H} = |y|^2 \check{M} C_{rs}^{h-v} \check{M}^H \quad \text{where } \check{M} = U_{4lr}^{-1} M U_{4lr}. \quad (4.66)$$

The task now is to find values of the cross-talk parameters for which the conditions of reflection symmetry are satisfied, either via (4.65), for the conditions listed in (4.56), or via (4.66), for the structure shown in (4.53) where the co-cross-correlations are all zero. The latter form of the problem is similar to the standard cross-talk and α

estimation problem addressed in the literature (see Goh *et al.* [2007] and references therein), except that α is already accounted for, and more importantly, where the standard linear-basis problem uses M , the circular-basis formulation here uses \tilde{M} .

A solution to this problem has not been found. Standard approaches to cross-talk estimation involve linearising expressions for the co-cross-correlation elements of a reflection-symmetric covariance matrix—compare the algorithms presented in a unified framework by Goh *et al.* [2007]. Following in this tradition, several attempts were made to modify the linearised, iterative algorithms of Klein [1992] and Ainsworth *et al.* [2006] for the circular basis, and various other linear and non-linear equation-solving techniques were tried, but Monte Carlo simulations of each proposed method consistently converged to the wrong set of cross-talk parameter values.

In fact, evidence is now presented which indicates that no unique solution exists for the linearised version of the estimation problem. The equation for the cross-talk-distorted data $\hat{C}'_{rs}{}^{h-v}$ in (4.66) can be expanded and linearised in the actual circular-basis cross-talk parameters u, v, w and z , whereby all terms containing products or powers of the parameters greater than first-order are ignored. The expressions for the observed non-zero co-cross-correlation elements can then be written in a convenient matrix form as

$$\begin{aligned} \begin{bmatrix} \tilde{\sigma}_{hhhv} \\ \tilde{\sigma}_{hhvh} \\ \tilde{\sigma}_{hvvv} \\ \tilde{\sigma}_{vhvv} \end{bmatrix} &= \frac{|y|^2}{2} \begin{bmatrix} \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hv}^2 \\ \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hv}^2 \\ \sigma_{vv}^2 & \sigma_{hhvv} & \sigma_{vv}^2 & \sigma_{hhvv} \\ \sigma_{hhvv} & \sigma_{vv}^2 & \sigma_{hhvv} & \sigma_{vv}^2 \end{bmatrix} \begin{bmatrix} u \\ v \\ w \\ z \end{bmatrix} + \begin{bmatrix} \sigma_{hhvv} & \sigma_{hh}^2 & \sigma_{hhvv} & \sigma_{hh}^2 \\ \sigma_{hh}^2 & \sigma_{hhvv} & \sigma_{hh}^2 & \sigma_{hhvv} \\ \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hv}^2 \\ \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hv}^2 \end{bmatrix} \begin{bmatrix} u \\ v \\ w \\ z \end{bmatrix}^* \\ &= \frac{|y|^2}{2} \begin{bmatrix} \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hhvv} & \sigma_{hh}^2 \\ \sigma_{hv}^2 & \sigma_{hv}^2 & \sigma_{hh}^2 & \sigma_{hhvv} \\ \sigma_{vv}^2 & \sigma_{hhvv} & \sigma_{hv}^2 & \sigma_{hv}^2 \\ \sigma_{hhvv} & \sigma_{vv}^2 & \sigma_{hv}^2 & \sigma_{hv}^2 \end{bmatrix} \begin{bmatrix} u+w \\ v+z \\ (u+w)^* \\ (v+z)^* \end{bmatrix}. \end{aligned} \quad (4.67)$$

Compare this to the equivalent expressions in terms of the linear-basis cross-talk parameters (subscript xy), when the distortion takes the form of M instead of \tilde{M} .

$$\begin{aligned} \begin{bmatrix} \tilde{\sigma}_{hhhv} \\ \tilde{\sigma}_{hhvh} \\ \tilde{\sigma}_{hvvv} \\ \tilde{\sigma}_{vhvv} \end{bmatrix} &= |y|^2 \begin{bmatrix} 0 & \sigma_{hv}^2 & \sigma_{hv}^2 & 0 \\ 0 & \sigma_{hv}^2 & \sigma_{hv}^2 & 0 \\ 0 & 0 & \sigma_{vv}^2 & \sigma_{hhvv} \\ \sigma_{hhvv} & \sigma_{vv}^2 & 0 & 0 \end{bmatrix} \begin{bmatrix} u_{xy} \\ v_{xy} \\ w_{xy} \\ z_{xy} \end{bmatrix} + \begin{bmatrix} 0 & 0 & \sigma_{hhvv} & \sigma_{hh}^2 \\ \sigma_{hh}^2 & \sigma_{hhvv} & 0 & 0 \\ \sigma_{hv}^2 & 0 & 0 & \sigma_{hv}^2 \\ \sigma_{hv}^2 & 0 & 0 & \sigma_{hv}^2 \end{bmatrix} \begin{bmatrix} u_{xy} \\ v_{xy} \\ w_{xy} \\ z_{xy} \end{bmatrix}^* \end{aligned} \quad (4.68)$$

Crucially, (4.67) implies that the four non-zero co-cross-correlations observed in the linear basis are, to first order, entirely due to the two coherent sums $u+w$ and $v+z$ of the cross-talk parameters from the circular-basis, so that although in general there will be four complex circular-basis cross-talk parameters, their first-order distortion of linear-basis reflection symmetry does not permit the individual parameters to be identified or estimated. Even if cross-talk reciprocity can be assumed, whereby $u=z$ and $v=w$, the parameters remain unidentifiable. The linearised expressions obtained for the standard linear-basis cross-talk problem, as shown in (4.68), do not suffer from this ambiguity. (Observe that (4.68) is closely related to (17)–(18) in [Ainsworth *et al.* 2006].) It would therefore seem that the circular-basis cross-talk parameters cannot be estimated using linearised expressions constrained solely by reflection symmetry. It is possible that other approaches, which avoid linearisation or add more constraints, may yield reliable estimation algorithms, but none have been found thus far.

Given that the circular-basis parameters k and α and the linear-basis parameters $\alpha_{xy}, u_{xy}, v_{xy}, w_{xy}$ and z_{xy} can be estimated from clutter [Goh *et al.* 2007], it may be tempting to extend the idea from Section 4.4 of correcting residual distortion in the linear basis by iteratively estimating and correcting for the two aforementioned sets of parameters using alternating basis conversions of the clutter covariance matrix. However, this will not fully calibrate the data, because a portion of the distortion will remain unaccounted for: expansion of (4.39) and its inverse reveals that this remnant distortion resides in the circular basis solely as a cross-talk component $u = -v = -w = z$ and resides in the linear basis solely as a channel-imbalance component $k_{xy} = (1+u^2)/(1-u^2+j2u)$. Since neither the circular-basis cross-talk nor the linear-

basis channel imbalance k_{xy} can be estimated using only reflection-symmetric clutter, full calibration requires additional measurements, typically from corner reflectors.

To summarise, given reflection-symmetric clutter in the circular basis, no cross-talk estimation is proposed here, and the channel imbalance parameters must be estimated non-iteratively in the presence of cross-talk. In the case of reciprocal cross-talk and relatively low noise, the cross-talk will limit the accuracy of \hat{k} but not $\hat{\alpha}$, as shown in Figure 4.6(b).

4.7 Data analysis

Given the analysis in Sections 4.5 and 4.6, the following polarimetric calibration algorithm is proposed for radar data acquired in the circular basis. The design of the radar system must permit the assumption of cross-talk reciprocity, as discussed in Section 4.3.

1. Iteratively estimate the polarimetric distortion parameters using corner reflectors, as detailed in Section 4.5.
 - (a) Estimate the channel imbalance parameter α via (4.45) using responses from trihedrals and the imbalance parameter k via (4.46) using responses from either dihedrals (at any rotation) or gridded trihedrals, with the sign of \hat{k} resolved by comparison with (4.47)–(4.48) evaluated for gridded trihedrals.
 - (b) Apply the partial channel imbalance correction.
 - (c) Estimate cross-talk parameters $u = z$ and $v = w$ (assuming cross-talk reciprocity) via (4.50)–(4.51) using responses from trihedrals.
 - (d) Apply the partial cross-talk correction.
 - (e) Repeat (a)–(d) until the parameter estimates converge to their ideal (unity imbalance, zero cross-talk), within some tolerance.
 - (f) Accumulate the parameter estimates via (4.52).
2. For verification, estimate the channel imbalance parameters α and k via (4.59)–(4.64), as described in Section 4.6, using responses from reflection-symmetric clutter e.g. forest, snow, crops or ploughed fields acquired at a similar orientation as that for the corner reflectors.
3. Correct the observed data $\tilde{\mathbf{s}} = D\mathbf{s} + \mathbf{n}$ (4.12) by constructing the polarimetric distortion matrix $\hat{D} = \hat{K}\hat{A}\hat{M}$ (4.15) and applying its inverse.
4. Convert the corrected data to the linear basis via (4.31) using U_{lr} in (4.33).
5. Estimate the residual distortion via standard linear-basis calibration techniques (see Goh [2012] and references therein). If all of the linear-basis parameters are available, they can be converted back to the circular basis via (4.39) and combined with the initial parameter estimates.

The estimated distortion parameters can then be used to calibrate other data acquired using the same radar system at a similar angular orientation.

This procedure is illustrated here using three polarimetric SAR images, shown in Figures 4.8–4.10, of two sites, acquired on the 8th and 11th of July during the initial trial of the *Ingara* L-band radar in 2013. As mentioned previously, this system used the helical antennas shown in Figure 4.1 to collect fully-polarimetric data in the circular basis, and followed the *s-m-a* design shown in Figure 4.2(b), permitting the assumption of cross-talk reciprocity.

For each polarimetric channel, the raw pulse echo samples were match-filtered and processed via a time-domain beamforming algorithm [Pincus *et al.* 2013] to form a SAR image focused to the terrain surface obtained from a smoothed digital elevation model [Gallant *et al.* 2011]. No autofocus was applied. Before image formation, the pulse samples were scaled to compensate for pulse-to-pulse adjustments in the receiver gains during acquisition. After image formation, the four images corresponding to the four polarimetric channels *ll*, *lr*, *rl* and *rr* were translated to align them in range and azimuth (this was facilitated by a fine pixel grid which significantly oversampled the resolved scene); this translation was necessary because the phase centres of the two antennas are not co-located, as is obvious from Figure 4.1, and their offset was not accounted for during image formation.

The images presented here were formed using an aperture (set of pulses) covering a slant-plane integration angle of 5.9° centred on the antenna boresight, accounting

for flight-track squint and platform yaw—the same part of the antenna beam pattern should be used when generating, comparing and applying calibration solutions. (The one-way 3 dB beamwidth is 30° .) Given the transmitted chirp bandwidth of 140 MHz, the nominal grazing angle of 32° and the application of a Hamming window, the effective azimuth and ground-range resolutions are 1.50 m and 1.72 m respectively.

The images are displayed via a false-colour mapping of the $\bar{\alpha}$ -angle, entropy and span scattering metrics [Cloude & Pottier 1997] into the hue, saturation and value (HSV) image components, respectively (see Appendix 4.A.1 for more details). The metrics are obtained from an eigendecomposition of the local polarimetric coherency matrix in the linear basis [Cloude 1985], which was estimated here using a simple rectangular window covering 3×3 resolution cells. This small sample size was selected to minimise mixing of the point-target and clutter signatures for validation of the polarimetric calibration. Homogeneous regions of clutter could be averaged much more to lower the finite-sample bias and variance [López-Martínez *et al.* 2005, Lee *et al.* 2008].

For pure (rank-one; zero-entropy) scattering, the $\bar{\alpha}$ -angle serves as a smooth measure of the co-polarisation ratio s_{vv}/s_{hh} , which is mapped to colour according to Figure 4.12(b). Ideal corner reflector responses, given as scattering matrices in (4.40)–(4.43), can be expressed as simple $\bar{\alpha}$ -angles via (4.73) and (4.77), as listed in Table 4.1. Note that the $\bar{\alpha}$ -angle is independent of target orientation (assuming the target response conforms to its characteristic scattering matrix).

Table 4.1: Ideal corner reflector responses in terms of $\bar{\alpha}$ -angle and displayed colour.

Corner reflector	s_{vv}/s_{hh}	$\bar{\alpha}$	Colour
trihedral (TH)	1	0°	blue
upright/rotated dihedral (DH/RD)	-1	90°	red
left/right gridded trihedral (GT _l /GT _r)	orientation dependent	45°	green

The corner reflector site, shown in Figures 4.8 and 4.9, contained two 2.45 m trihedrals (TH), four 2.45 m gridded trihedrals (two with fins parallel to the left face (GT_l; $+\theta$) and two with fins parallel to the right face (GT_r; $-\theta$)), and two 0.54 m dihedrals (one upright (DH) and one rotated (RD) by 22.5° around the line-of-sight)—see the photos in Figure 4.8. The site was imaged four times during the trial, but two of the acquisitions were unsuitable for calibration because of either receiver saturation or excessive yaw (13°). In the latter case, for an aperture centred on the antenna boresight, buildings located outside the intended swath were strongly illuminated and their bright responses aliased into the imaged scene due to an insufficiently high PRF, corrupting the desired response. This degradation can be seen to a small extent in the displayed figures as differing ‘smudges’ of unexpected colour.

The iterative estimation-correction process for polarimetric distortion, summarised in step 1 of the above calibration procedure, was applied separately to the images acquired on the 8th (Figure 4.8) and the 11th (Figure 4.9). Three iterations were sufficient to achieve convergence. In each iteration, corresponding parameter estimates from the same type of corner reflector (trihedral or dihedral) were averaged, with the gridded trihedrals used only to resolve the sign of \hat{k} . Table 4.2 lists the initial individual estimates of the polarimetric distortion parameters in the circular basis for every calibration source, along with the combined estimates obtained on successive iterations. Figure 4.7 shows how the individual corner reflector responses vary as successive calibration solutions are applied.

The individual parameter estimates from corner reflectors of the same type within an image show good consistency. The estimates across days are broadly consistent, although the magnitude of $\hat{\alpha}$ differs moderately (1.67 vs 1.78), and the phase of \hat{u} is not consistent at all. The iterations clearly converge very quickly, which agrees with the simulations in Section 4.5. Figure 4.7 indicates that the calibration does a good job of driving the polarimetric responses of the corner reflectors towards their ideal $\bar{\alpha}$ -angle, as listed in Table 4.1. The trihedral responses are changed most substantially, with the calibration seeking to balance their cross-polarisations via $\hat{\alpha}$ and null both co-polarisations via \hat{u} and \hat{v} in order to match (4.40) in the circular basis. Observe in particular that the additional effect of cross-talk correction, apparent when going from 0.5 iterations (i.e. channel imbalance correction only) to 1 iteration, is appreciable, so the cross-talk for this radar system is clearly not insignificant and should not be

Table 4.2: Estimated magnitude and phase of the circular-basis channel imbalance parameters α and k and reciprocal cross-talk parameters u and v for the *Ingara* L-band radar.

Day	Source	$\hat{\alpha}$	\hat{k}	\hat{u}	\hat{v}
8	TH (tl)	$1.67\angle-63^\circ$		$0.093\angle-155^\circ$	$0.100\angle-163^\circ$
	TH (br)	$1.67\angle-63^\circ$		$0.092\angle-133^\circ$	$0.073\angle-151^\circ$
	DH		$0.64\angle-50^\circ$		
	RD		$0.68\angle-56^\circ$		
	GT _r (tl)	$1.67\angle-64^\circ$	$0.70\angle-63^\circ$		
	GT _r (br)	$1.67\angle-63^\circ$	$0.80\angle-48^\circ$		
	GT _l (tl)	$1.67\angle-64^\circ$	$0.70\angle-44^\circ$		
	GT _l (br)	$1.67\angle-63^\circ$	$0.71\angle-28^\circ$		
	iter. 1	$1.67\angle-63^\circ$	$0.66\angle-53^\circ$	$0.091\angle-144^\circ$	$0.086\angle-158^\circ$
	iter. 2	$1.67\angle-63^\circ$	$0.65\angle-53^\circ$	$0.089\angle-144^\circ$	$0.088\angle-158^\circ$
	iter. 3	$1.67\angle-63^\circ$	$0.65\angle-53^\circ$	$0.089\angle-144^\circ$	$0.088\angle-158^\circ$
11	TH (tl)	$1.77\angle-62^\circ$		$0.100\angle-53^\circ$	$0.080\angle-179^\circ$
	TH (br)	$1.79\angle-62^\circ$		$0.076\angle-69^\circ$	$0.091\angle-124^\circ$
	DH		$0.65\angle-58^\circ$		
	RD		$0.62\angle-68^\circ$		
	GT _r (tl)	$1.77\angle-62^\circ$	$0.83\angle-68^\circ$		
	GT _r (br)	$1.80\angle-62^\circ$	$0.84\angle-26^\circ$		
	GT _l (tl)	$1.77\angle-62^\circ$	$0.62\angle-49^\circ$		
	GT _l (br)	$1.81\angle-62^\circ$	$0.64\angle+1^\circ$		
	iter. 1	$1.78\angle-62^\circ$	$0.63\angle-63^\circ$	$0.087\angle-60^\circ$	$0.076\angle-150^\circ$
	iter. 2	$1.78\angle-62^\circ$	$0.63\angle-63^\circ$	$0.086\angle-59^\circ$	$0.076\angle-150^\circ$
iter. 3	$1.78\angle-62^\circ$	$0.63\angle-63^\circ$	$0.086\angle-59^\circ$	$0.076\angle-150^\circ$	
11	clutter	$1.68\angle-58^\circ$	$0.54\angle-77^\circ$	N/A	N/A

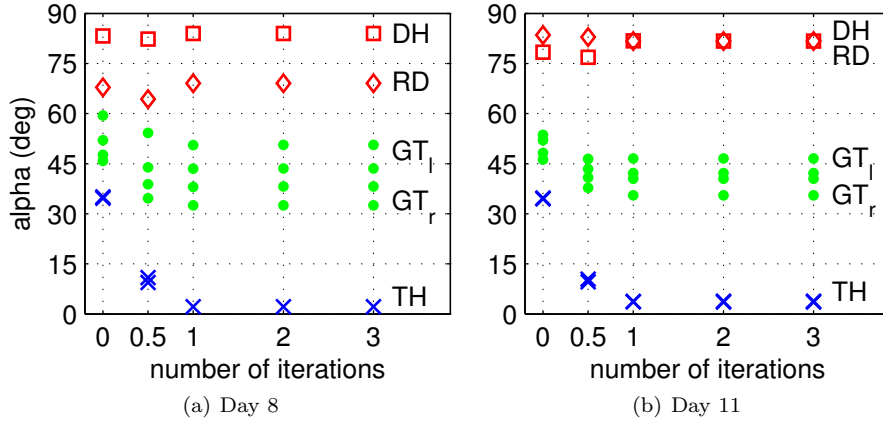


Figure 4.7: Variation of the individual corner reflector responses, in terms of $\bar{\alpha}$ -angle, as successive iterations of the calibration solution are applied. Zero iterations corresponds to the initial, uncalibrated data, and 0.5 iterations corresponds to correcting for the initial estimates of channel imbalance only. The two trihedrals (TH) are marked as blue crosses (on top of each other), the four gridded trihedrals (GT_l (upper two) and GT_r (lower two)) as green dots, the upright dihedral (DH) as a red square and the rotated dihedral (RD) as a red diamond.

neglected, as it has the effect of reducing the observable polarimetric dynamic range. Also note that calibration improves the responses of the gridded trihedrals (i.e. pushes their $\bar{\alpha}$ -angle closer to 45° , on average), even though they were not used to estimate the distortion parameters (except for a sign disambiguation), thus indicating in a small way the validity of the calibration. Anomalous behaviour is exhibited by the rotated dihedral on the 8th, likely due to the -5.5° azimuthal misalignment between the target boresight and the selected aperture causing a non-ideal response; the dihedral sensitivity to angular misalignment may be exacerbated here by its small size ($\sim 2\lambda$).

These calibration effects are visible in the SAR images in Figures 4.8 and 4.9, particularly by the changes in colour of the trihedrals, from green-cyan to pure blue, and the overall clutter background, from green to light blue. The latter change is principally due to correction for the large channel imbalance, which is attributable mainly to the receivers because, from Table 4.2 and the parameter definitions in (4.18)–(4.19), $|\hat{k}| = |r_{ll}/r_{rr}| = 0.65$ but $|\hat{k}\hat{\alpha}| = |t_{ll}/t_{rr}| = 1.09$. From Fig. 4.12(b), this change corresponds to making the co-polarisation ratio in the linear basis closer to (but not equal to) unity, which is arguably more reasonable for a natural scene that provides moderate surface scatter from rough ground (at L-band) and is not covered by polarimetrically selective, dipole-like scattering elements. The yellow points behind the gridded trihedrals indicate a second phase centre with a longer path delay; the cause of this scattering response is not clear.

The clutter site, shown in Figure 4.10, contains a stand of plantation pine trees on flat ground, which is assumed to give rise to reflection symmetry. Unfortunately, it was imaged only on the 10th and 11th, not the 8th, and the former acquisition was at a significantly steeper grazing angle, so only the latter image is presented. Step 2 of the calibration procedure was applied using the covariance matrix computed for a homogeneous region of forest. The phase estimate $\angle\hat{k}$ via (4.64) was discounted because the observed correlation magnitude $|\hat{\sigma}_{llrr}|$ was low and the associated phase $\angle\hat{\sigma}_{llrr}$ very noisy. The resulting channel imbalance estimates, listed in Table 4.2, show relatively good agreement with the estimates obtained from corner reflectors, although the magnitude of k differs moderately (0.54 vs 0.63). This mutual agreement provides some confidence that the circular-basis methods outlined in this work lead to accurate calibration for channel imbalance. Repeated estimates using different clutter sites would permit an increase in confidence, but unfortunately, most of the other vegetated landscapes imaged during the trial were undulating, which would complicate clutter-based calibration [Ainsworth *et al.* 2006].

Overall, the polarimetric distortion parameter estimates listed in Table 4.2 are broadly consistent across days and types of sites (except for the phase of \hat{u}), allowing for some variability in the true scattering responses and the acquisition geometries. Repeated estimates using sites containing both corner reflectors and suitable clutter would permit a fuller validation, particularly if the residual linear-basis distortion could be estimated and transformed back to the circular basis, as outlined in step 5 of the proposed calibration procedure. Unfortunately, such data were not collected during this trial.

Finally, Figure 4.11 shows an undulating forest scene imaged on the 11th containing a 1.2 m trihedral (slightly obscured by foliage) and a vehicle parked between trees on the side of a road. After applying the calibration solution from the 8th (see Table 4.2), the $\bar{\alpha}$ -angle of the trihedral response changed from 33° to 15°, roughly as expected, although it may indicate that some distortion remains. The vehicle was illuminated approximately perpendicular to its side, giving rise to a ground-vehicle double bounce which, when mixed with the response of the surrounding vegetation, generates a stand-out yellow ($\bar{\alpha} = 65^\circ$) point response expected of partial dihedral scattering. Overall, the measured and calibrated radar data is sufficiently accurate to permit physical interpretation; this provides a limited but encouraging validation of the calibration solution and, moreover, the procedure used to obtain it.

4.8 Conclusion

Two novel aspects of polarimetric calibration for imaging radar systems have been addressed. Firstly, the radar system model has been formulated in the context of two generic transmitter designs, either a single amplifier followed by a high-power switch or a low-power switch followed by two amplifiers. It has been shown that the associated polarimetric distortion matrix can be factorised in two different ways to give, in the general case, the standard distortion model of two channel imbalance parameters and four cross-talk parameters. Moreover, when the low-power-switch model is applicable, one of the factorisations permits a significant simplification where the cross-talk can be represented using just two reciprocal parameters, one for each antenna.

Secondly, calibration techniques for circularly polarised antennas have been derived.

It was shown that the circular-basis responses from a set of corner reflectors could (continues on p. 199)




Figure 4.8: Polarimetric SAR image from July 8 of the corner reflector calibration site (500×500 m) before (upper) and after (lower) calibration, depicted using an $\bar{\alpha}$ -entropy-span decomposition mapped to HSV. Acquisition geometry: synthetic aperture centred on antenna boresight, grazing angle 33.4° , flight-track squint 4.1° , platform yaw 4.0° , azimuthal illumination direction -5.5° off the targets' boresight. See Table 4.1 for the expected corner reflector responses and Figure 4.7(a) for quantitative comparison of the actual responses before and after calibration.




Figure 4.9: Polarimetric SAR image from July 11 of the corner reflector calibration site, displayed as for Figure 4.8. Acquisition geometry: synthetic aperture centred on antenna boresight, grazing angle 30.6° , flight-track squint 1.2° , platform yaw 0.9° , azimuthal illumination direction $+3.6^\circ$ off the targets' boresight. See Table 4.1 for the expected corner reflector responses and Figure 4.7(b) for quantitative comparison of the actual responses before and after calibration.




Figure 4.10: Polarimetric SAR image from July 11 of the clutter calibration site (700×500 m), displayed as for Figure 4.8 using the circular-basis calibration solution in Table 4.2 from July 8. Acquisition geometry: synthetic aperture centred on antenna boresight, grazing angle 35.5° , flight-track squint 0.0° , platform yaw 0.6° . The red arrow indicates where the photo (bottom-right) was taken. The red box indicates the flat, forested region used to compute the covariance matrix for clutter-based estimation of the channel imbalance parameters. The buildings in the top-left area give rise to a partial (yellow) dihedral-like response due to ground-wall double bounce. It may be expected that the tree trunks would lead to a ground-trunk double-bounce scattering mechanism in the forested region; the fact that this is not observed may be due to the trunks being obscured by low attached branches (see photo) and the ground being obscured by a rough, uneven layer of fallen branches and foliage. Indeed, this absence agrees with the multifrequency observations reported by Fleischman *et al.* [1996 Fig. 6], where a significant phase difference between s_{hh} and s_{vv} , indicative of double bounce, occurred only at P-band.




Figure 4.11: Polarimetric SAR image from July 11 of a forest site (500×500 m) containing a vehicle and a 1.2 m trihedral (TH), after circular-basis calibration using the solution in Table 4.2 from July 8. Acquisition geometry: synthetic aperture centred on antenna boresight, grazing angle 30.2° , flight-track squint 1.6° , platform yaw 1.0° . The red arrow indicates where the photo (bottom-left) was taken. The $\bar{\alpha}$ -entropy 2D histogram of all pixels (bottom-right; red (blue) is more (less) popular) shows that the main forest response is a mixture of rough surface scattering and canopy propagation effects ($28^\circ \leq \bar{\alpha} \leq 37^\circ$ and $0.67 \leq H \leq 0.80$) [Cloude & Pottier 1997]. As in Figure 4.10, there is little ground-trunk double bounce from the forest, possibly due to the undulating terrain as well as the rough ground and obscured trunks [Cloude 2010 ch. 3.1.3–5]. (For the histogram, the $\bar{\alpha}$ -entropy decomposition was recomputed using a rectangular window covering 9×9 resolution cells to lower the finite-sample bias and variance [López-Martínez *et al.* 2005].)

(continues from p. 195) be used to iteratively estimate the channel imbalance and cross-talk parameters, if the cross-talk can be assumed reciprocal. Reflection-symmetric clutter can provide an alternative estimate of channel imbalance. However, in contrast to the linear basis, no algorithm could be found to estimate the cross-talk from clutter; indeed, it was shown that the distortion of reflection-symmetric clutter caused, to first-order, by cross-talk in the circular basis leads to the cross-talk parameters being unidentifiable.

The calibration techniques were applied to data acquired by the *Ingara* L-band radar using left- and right-polarised helical antennas. Sets of parameter estimates obtained across two days and two sites were broadly consistent. A wider range of data would enable the *Ingara* radar to be better characterised and further validate the proposed circular-basis calibration techniques.

4.A Appendix

4.A.1 Eigendecomposition of polarimetric SAR imagery

The eigendecomposition metrics $\bar{\alpha}$ -angle, entropy H and span P_s are discussed here, with explanations of their computation, interpretation and mapping into colour for the display of polarimetric SAR images. (The $\bar{\alpha}$ -angle is unrelated to the channel imbalance parameter α .)

Let s_{hh} , s_{hv} , s_{vh} and s_{vv} denote the complex polarimetric values at one pixel position after calibration in the collection basis and conversion to the linear basis. Instead of using the standard scattering vector \mathbf{s} in (4.12), Cloude [1985] advocates using the more physically meaningful Pauli scattering vector \mathbf{s}_P , defined as

$$\mathbf{s}_P = \frac{1}{\sqrt{2}} \begin{bmatrix} s_{hh} + s_{vv} \\ s_{hh} - s_{vv} \\ s_{hv} + s_{vh} \\ j(s_{hv} - s_{vh}) \end{bmatrix}. \quad (4.69)$$

The Hermitian inner product of \mathbf{s}_P (and of \mathbf{s}) is the total scattering power P_s , sometimes called the span [Cloude 1994]. The mean Hermitian outer product of \mathbf{s}_P is the coherency matrix C_P , which is the Pauli-analogue of the covariance matrix C [Cloude 1985]. In practice, C_P is estimated by a spatial average $\langle \dots \rangle$ of the outer products in the neighbourhood. (In the polarimetric radar literature, \mathbf{s}_P and C_P are often denoted \mathbf{k} and T [Cloude 1985, Lee & Pottier 2009 ch. 3.2.2, 3.3.2].)

$$P_s = \mathbf{s}_P^H \mathbf{s}_P = \mathbf{s}^H \mathbf{s} = |s_{hh}|^2 + |s_{hv}|^2 + |s_{vh}|^2 + |s_{vv}|^2 \quad (4.70)$$

$$C_P = \mathbf{E}\{\mathbf{s}_P \mathbf{s}_P^H\} \quad (4.71)$$

Let $\{\lambda_n\}$ and $\{\mathbf{v}_n\}$ (where $n = 1, 2, \dots, N$ and $N = 4$) denote the ordered, non-increasing eigenvalues and the corresponding unit-length eigenvectors of C_P . Note that C and C_P are both Hermitian positive semidefinite and are unitarily similar, with equal, nonnegative eigenvalues and related sets of eigenvectors [Cloude 1994, Lee & Pottier 2009, Horn & Johnson 2013 pp. 94,95,228,229,426,429-431,438]. A standard relation for any square matrix A is that the sum of its eigenvalues is equal to the sum of its diagonal elements i.e. the trace $\text{Tr}(A)$. Hence,

$$P_\lambda = \sum_{n=1}^N \lambda_n = \text{Tr}(C_P) = \text{Tr}(C) = \mathbf{E}\{|s_{hh}|^2\} + \mathbf{E}\{|s_{hv}|^2\} + \mathbf{E}\{|s_{vh}|^2\} + \mathbf{E}\{|s_{vv}|^2\} = \mathbf{E}\{P_s\}. \quad (4.72)$$

For the usual case of reciprocal scattering, the cross-polarisations will be equal, except for system noise, so no more than $N - 1$ independent sources of scattering can be identified via an eigendecomposition. It is assumed here that the smallest eigenvalue (λ_N) is due to the noise; it can be subtracted from the other eigenvalues as a noise reduction measure [Hajnsek *et al.* 2001]. The remaining $N - 1$ eigenvectors (having dropped \mathbf{v}_N) can be interpreted as independent scattering mechanisms, each occurring with relative strength or ‘probability’ p_n according to the relative size of its associated eigenvalue. Using a simple parameterisation of each eigenvector [Cloude 1994, Cloude & Pottier 1995], the n^{th} mechanism can be represented by a single angle α_n , obtained from the first element $(\mathbf{v}_n)_1$ of the n^{th} eigenvector as shown in (4.73) below.

$$\alpha_n = \arccos |(\mathbf{v}_n)_1| \quad (0 \leq \alpha_n \leq 90^\circ) \quad (4.73)$$

$$p_n = \frac{\lambda_n - \lambda_N}{P_\lambda - \lambda_N} \quad (4.74)$$

The total scattering response at the pixel position can then be expressed in terms of the mean scattering mechanism $\bar{\alpha}$ (how the scattering object polarises incident waves) and the entropy H (the extent to which the scattering object depolarises incident waves by inducing random polarimetric disorder or noise).

$$\bar{\alpha} = \sum_{n=1}^{N-1} p_n \alpha_n \quad (0 \leq \bar{\alpha} \leq 90^\circ) \quad (4.75)$$

$$H = - \sum_{n=1}^{N-1} p_n \log_{N-1} p_n \quad (0 \leq H \leq 1) \quad (4.76)$$

Zero entropy occurs when there is only one non-zero eigenvalue, indicating just one (deterministic) mechanism corresponding to a single scattering matrix [Cloude & Pottier 1995]. As the entropy increases from zero, it modulates the interpretation of the $\bar{\alpha}$ -angle by indicating increased (slight to medium) roughness for surfaces or a wider distribution of orientations for anisotropic particles in a random volume [Cloude & Pottier 1997]. Medium or high entropy may indicate multiple scattering contributions such as both ground and canopy responses when a forest is illuminated with a foliage-penetrating wavelength. Maximum entropy occurs when all eigenvalues are equal, indicating the generation of white polarimetric noise [Cloude & Pottier 1995].

Fig. 4.12(a) specifies how the $\bar{\alpha}$ -angle, entropy and span metrics have been mapped to the hue, saturation and value (HSV) image components, respectively. A standard transformation then gives the red, green and blue (RGB) colour intensities [Smith 1978]. Observe that the basic colour is specified solely by the $\bar{\alpha}$ -angle; the other metrics modulate how pure (saturated) and how bright the selected colour will be when displayed.




Figure 4.12: Colourmap (a) for the depiction of polarimetric SAR images via an eigendecomposition. The $\bar{\alpha}$ -angle specifies the hue, from blue ($\bar{\alpha} = 0^\circ$), through green ($\bar{\alpha} = 45^\circ$), to red ($\bar{\alpha} = 90^\circ$), without wrap. The entropy H specifies the saturation, where lower entropy implies more pure, less random, scattering, depicted as less white. The span P_s specifies the value of brightness i.e. the pixel intensity relative to the rest of the image. (In this work, the span is converted to decibels before being linearly mapped to brightness.) In the special case of a pure (rank-one; zero-entropy) scattering response, with negligible cross-polarisation, (b) shows how the $\bar{\alpha}$ -angle, given by (4.73) and (4.77), varies with the magnitude (radial scale from -18 dB at centre to $+18$ dB at edge) and phase (polar angle in degrees) of the co-polarisation ratio $d = s_{vv}/s_{hh}$.

To see how the $\bar{\alpha}$ -angle represents a continuous range of physical scattering mechanisms, consider the eigendecomposition for a rank-one (pure) scattering object S_{bd}^{h-v} , parameterised by co-polarisation ratio $d = s_{vv}/s_{hh}$ and cross-polarisation ratio $b = s_{hv}/s_{hh} \ll 1$ [Cloude 2010 ch. 4.1.3].

$$S_{bd}^{h-v} = \begin{bmatrix} 1 & b \\ b & d \end{bmatrix} \xrightarrow[\text{of } \mathbf{s}_P \mathbf{s}_P^H]{\text{eigen d.}} \begin{matrix} \lambda_1 = 1 + |d|^2 + 2|b|^2, & \lambda_2 = \lambda_3 = \lambda_4 = 0 \\ \mathbf{v}_1 = \frac{1}{\sqrt{2\lambda_1}} [1 + d, 1 - d, 2b, 0]^T \end{matrix} \quad (4.77)$$

When $d = 1$ (and b is negligibly small), $\bar{\alpha} = 0^\circ$, depicted as pure blue; this occurs for odd-bounce reflection from flat plates and trihedral corner reflectors and for scattering from a random volume of spheres [Cloude & Pottier 1996]. When $d = -1$, $\bar{\alpha} = 90^\circ$, depicted as pure red; this occurs for even-bounce reflection from dihedral corner reflectors, regardless of whether they are upright (4.41) or rotated (4.42); the eigendecomposition metrics used here are invariant to rotations around the line-of-sight. Double bounce from ground-wall and ground-trunk pairs may approximate a dihedral-like response ($60^\circ \lesssim \bar{\alpha} \lesssim 90^\circ$; yellow/orange), depending on the relative scattering properties of the two right-angled surfaces [Cloude 2010 ch. 3.1.4]. When $|d| \ll 1$ or $|d| \gg 1$, $\bar{\alpha} \approx 45^\circ$, depicted as pure green; this occurs for dipole-like scattering from wires and random volumes of anisotropic particles at similar orientations (e.g. crops) [Cloude & Pottier 1997]. Interestingly, gridded trihedrals (4.43) have $\bar{\alpha} = 45^\circ$ regardless of orientation

(θ, ψ) (accounting for non-zero b in this case). As illustrated in Fig. 4.12(b), the $\bar{\alpha}$ -angle smoothly captures the full range of the co-polarisation ratio d , and through this provides an indication of the scattering mechanism.

Chapter 5

Conclusion

5.1 Summary

This thesis has investigated the potential use of airborne or spaceborne imaging radar to detect disturbances of the ground in a forest. The established technique of coherent change detection, typically applied to a repeat-pass pair of SAR images of open ground, is extended to tree-covered scenes. The coherence metric is sensitive to subtle changes in the clutter because it measures the magnitude of the complex correlation between the complex speckle patterns in the two SAR images. Interference from the canopy, undulations in the terrain, and polarimetric distortion in the radar system, all pose significant obstacles. Novel solutions to these problems have been developed.

Chapter 2 considered how the radar images should be formed in order to avoid the spurious coherence loss that arises when the images are focused to a flat plane but the scattering surface is undulating. Following the approach of Blacknell and Andre (and co-authors) [Blacknell *et al.* 2010, 2011, André *et al.* 2013], an image-formation algorithm was detailed that exploited the flexibility offered by back-projection not only to focus the image onto a surface matched to the scene topography but also to do spatially adaptive trimming of the spatial-frequency supports from the two passes. In order to choose this trim properly, the dependence of the supports on local slope and collection geometry was derived—this mathematical formulation is the key novel contribution in this chapter. The properties of the resulting point-spread function and image spectrum were studied. Improved coherence was demonstrated using two airborne radar datasets.

The new idea of 3D SAR CCD for monitoring the ground under a forest canopy was proposed in Chapter 3. It was formulated as a multichannel dual-layer coherence estimation problem, where the coherence of scattering from the ground surface is estimated after suppressing interference from the canopy layer by beamforming multiple image channels acquired at slightly different grazing angles on each pass. The choice of operating wavelength was considered in terms of the trade-off between foliage penetration and change sensitivity. A framework for comparing radar designs and beamforming algorithms, predicting coherence estimation performance, and assessing the sensitivity to error, was built around the random-volume-over-ground (RVOG) model of forest scattering. In particular, it was shown that applying optimal (MVDR/Capon) beamforming to an L-band array of just three channels, spaced by 0.05° in grazing angle, can attenuate the interference due to a 20 m tall volume by 12 dB, and if the ground scattering intensity is no more than 2.5 dB below that of the volume, then the ground coherence can be estimated with an error of no more than 0.1 i.e. 10%. Such an array can be formed by a dual-antenna radar system operating in an alternating-transmit, simultaneous-receive mode to synthesise three effective phase centres at different across-track positions. The 3D SAR CCD concept was demonstrated using an RVOG clutter simulation developed by the author and a modified version of the physics-based SAR image simulator PolSARproSim [Williams 2006]. Receiver operating characteristics computed using the RVOG simulation, with ground and volume scattering intensities set equal, showed that whilst ordinary (single-channel) CCD achieved a detection probability of only 26% at a false-alarm rate of 5%, 3D SAR CCD using a three-channel system achieved a detection probability of 76%. Given fully polarimetric radar data, performance was further improved using a novel polarimetric filtering algorithm to remove contributions from the ground-trunk double-bounce scattering mechanism, which may mask changes on the ground near trees.

Polarimetric calibration is a prerequisite to making physically meaningful interpretations of fully polarimetric data. In Chapter 4, polarimetric distortion due to channel imbalance and cross-talk was formulated using a flexible radar system model that unified disparate works in the literature and allowed a simplification of the cross-talk representation from four to two parameters if the transmitter followed a low-power switch design. Distortion estimation techniques were then investigated for circularly polarised data acquired using left- and right-polarised antennas—this problem has seen little attention in the literature. An algorithm was proposed that used the scattering responses from ordinary and gridded trihedral corner reflectors, and optionally dihedrals, to iteratively estimate the two channel imbalance parameters and two cross-talk parameters. Monte Carlo simulations showed that the method reliably converged to the true parameter values. Airborne radar data was calibrated using the method, with broadly consistent parameter estimates obtained across flights. Arguably the most interesting result, however, was that the distortion of reflection-symmetric clutter caused, to first-order, by cross-talk in the circular basis does not permit the individual cross-talk parameters to be identified. This precludes the clutter-based approaches commonly used to estimate cross-talk in the linear basis.

5.2 Future Work

The obvious next step is to demonstrate 3D SAR CCD using real forest data. A likely source for this data would be the newly upgraded *Ingara* airborne imaging radar, which offers good SNR, but has only one phase centre in the across-track direction [Preiss *et al.* 2018]. Therefore, the set of image channels required for 3D SAR beamforming could not be acquired in a single pass, and instead must be acquired by making repeat passes in quick succession at slightly different altitudes or standoffs, which is the typical collection method for 3D SAR. This leads to a significant additional challenge: residual motion-compensation errors affecting the input images (particularly the pixel phase) must be estimated and removed to permit coherent addition across passes. Techniques for this array calibration or ‘autofocus’ step have been proposed in the 3D SAR literature, as discussed in Section 3.1.4. Note that the SAR imagery would be formed and polarimetrically calibrated using the methods described in Chapters 2 and 4.

To provide confidence in the real-data validation, it would be useful to develop a measure of the purity of the ground signal extracted by 3D SAR beamforming. Using the RVOG model for forest scattering, it has been shown that the task of volume attenuation is highly sensitive to system errors, with this sensitivity quantified by the condition number of the matrix of volume coherences. In addition, when knowledge of the scene topography is inaccurate, the extracted signal may not even contain the ground response. If some feature of ground backscatter that distinguished it from other scattering mechanisms could be identified, then perhaps some metric could be developed that indicated the extent to which the extracted signal exhibited this feature. This metric would provide a model-based verification that the extracted signal had indeed arisen from the ground surface and was relatively free of contamination from the canopy. Polarimetric decomposition quantities such as $\bar{\alpha}$ -angle and entropy, evaluated after beamforming fully polarimetric data [Ferro-Famil *et al.* 2012], have been used informally for this purpose [Pardini & Papathanassiou 2017].

Going one step further, alternative methods of extracting the ground backscatter component from the forest scattering response should be investigated. Canopy suppression via 3D SAR beamforming, specifically the MVDR (i.e. Capon) method, has been shown to work well given only three image channels when the ground and canopy scattering intensities are at similar levels. However, if the ground response was significantly weaker, then the required level of canopy attenuation would be concomitantly greater, necessitating a much larger number of channels for beamforming—probably too many to acquire in a single pass of an airborne platform. Recent analysis of L-band SAR imagery of a tropical forest showed that scattering from the ground is often 15 dB below that from the canopy [Pardini *et al.* 2018]. In such cases, instead of traditional beamforming, the polarimetric-interferometric methods of Tebaldini [2009] or Pardini & Papathanassiou [2017], which are partially model-based, may point the way to more effective retrieval of a weak ground signal. It is worth reiterating that the processing must preserve phase, and that the output image must accurately represent the complex speckle pattern due to scattering at ground level.

Bibliography

Abbreviations of journal and conference names

Abbreviation	Full name
IEEE Proc.	Proceedings of the IEEE
IEEE TAES	IEEE Transactions on Aerospace and Electronic Systems
IEEE TAP	IEEE Transactions on Antennas and Propagation
IEEE TGRS	IEEE Transactions on Geoscience and Remote Sensing
IEEE TIP	IEEE Transactions on Image Processing
IEEE TSP	IEEE Transactions on Signal Processing
IEEE GRS Letters	IEEE Geoscience and Remote Sensing Letters
Asia-Pacific SAR Conf.	IEEE Asia-Pacific Conference on Synthetic Aperture Radar
Euro. Radar Conf.	European Radar Conference
EuSAR Conf.	European Conference on Synthetic Aperture Radar
ICASSP Conf.	IEEE International Conference on Acoustics, Speech and Signal Processing
IEEE Radar Conf.	IEEE Radar Conference
IGARSS	IEEE International Geoscience and Remote Sensing Symposium
Int. Radar Conf.	International Conference on Radar
Sig., Sys. & Comp. Asilomar Conf.	Asilomar Conference on Signals, Systems and Computers
SPIE Alg. SAR Im. Conf.	SPIE Algorithms for Synthetic Aperture Radar Imagery

Ainsworth, T. L. & Williams, M. L., “A simulation study of topographic effects on PolSAR classification of forests and crops”, in *Proc. IGARSS*, Jul. 2011. [78]

Ainsworth, T. L., Ferro-Famil, L. & Lee, J.-S., “Orientation angle preserving *a posteriori* polarimetric SAR calibration”, *IEEE TGRS*, 44(4):994–1003, Apr. 2006. [3, 7, 176, 177, 181, 183, 186, 188, 189, 191, 195]

Albinet, C., Borderies, P., Koleck, T., Rocca, F., Tebaldini, S., Villard, L., Toan, T. L., Hamadi, A. & Minh, D. H. T., “TropiSCAT: A ground based polarimetric scatterometer experiment in tropical forests”, *J. Selected Topics in Applied Earth Observations and Remote Sensing*, 5(3):1060–1066, Jun. 2012. [86]

Andre, D., Morrison, K., Blacknell, D., Muff, D. & Nottingham, M., “Very high resolution coherent change detection”, in *Proc. IEEE Radar Conf.*, May 2015. [75]

André, D. B., Blacknell, D. & Morrison, K., “Spatially variant incoherence trimming for improved SAR CCD”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 8746, May 2013. [4, 15, 39, 203]

Antar, Y. M. M. & Hendry, A., “Correlation measurements in precipitation at linear polarisation using dual-channel radar”, *Electronics Letters*, 21(22):1052–1054, Oct. 1985. [3, 175, 188]

Askne, J., Dammert, P. B. G., Ulander, L. M. H. & Smith, G., “C-band repeat-pass interferometric SAR observations of the forest”, *IEEE TGRS*, 35(1):25–35, Jan. 1997. [81]

- Askne, J., Santoro, M., Smith, G. & Fransson, J. E. S., “Multitemporal repeat-pass SAR interferometry of boreal forests”, *IEEE TGRS*, 41(7):1540–1550, Jul. 2003. [81, 82]
- Ausherman, D. A., Kozma, A., Walker, J. L., Jones, H. M. & Poggio, E. C., “Developments in radar imaging”, *IEEE TAES*, 20(4):363–400, Jul. 1984. [36, 40, 41]
- Ausherman, D. A., “SAR digital image-formation processing”, in *Proc. SPIE Digital Image Processing Conf.*, vol. 528, Jul. 1985. [41]
- Ballester-Berman, J. D., Lopez-Sanchez, J. M. & Fortuny-Guasch, J., “Retrieval of biophysical parameters of agricultural crops using polarimetric SAR interferometry”, *IEEE TGRS*, 43(4):683–694, Apr. 2005. [81]
- Ballester-Berman, J. D. & Lopez-Sanchez, J. M., “Coherence loci for a homogeneous volume over a double-bounce ground return”, *IEEE GRS Letters*, 4(2):317–321, Apr. 2007. [81]
- Ballester-Berman, J. D. & Lopez-Sanchez, J. M., “Applying the Freeman-Durden decomposition concept to polarimetric SAR interferometry”, *IEEE TGRS*, 48(1):466–479, Jan. 2010. [78]
- Ballester-Berman, J. D. & Lopez-Sanchez, J. M., “Combination of direct and double-bounce ground responses in the homogeneous oriented volume over ground model”, *IEEE GRS Letters*, 8(1):54–58, Jan. 2011. [81]
- Bamler, R. & Hartl, P., “Synthetic aperture radar interferometry”, *Inverse Problems*, 14(4):1–54, Aug. 1998. [13, 30, 60, 81, 95]
- Barber, J. & Kogon, S., “Probabilistic three-pass SAR coherent change detection”, in *Proc. Sig., Sys. & Comp. Asilomar Conf.*, Nov. 2012. [87]
- Barber, J., “A generalized likelihood ratio test for coherent change detection in polarimetric SAR”, *IEEE GRS Letters*, 12(9):1873–1877, Sep. 2015. [87]
- Barnes, R. M., “Antenna polarization calibration using in-scene reflectors”, Tech. Rep. TT-65, Lincoln Laboratory, Sep. 1986. [3, 176]
- Benitz, G. R., “High-definition vector imaging”, *Lincoln Laboratory Journal*, 10(2):147–170, 1997. [22]
- Berens, P., “CircStat: A MATLAB toolbox for circular statistics”, *J. Statistical Software*, 31(10), Sep. 2009. [143]
- Bessette, L. A. & Ayasli, S., “Ultra-wideband P-3 and CARABAS II foliage attenuation and backscatter analysis”, in *Proc. IEEE Radar Conf.*, May 2001. [1, 79, 80, 82]
- Binder, B. T., Toups, M. F., Ayasli, S. & Adams, E. M., “SAR foliage penetration phenomenology of tropical rain forest and northern U.S. forest”, in *Proc. IEEE Radar Conf.*, May 1995. [80, 82]
- Blacknell, D., Andre, D. B. & Finch, C. M., “SAR CCD over mountainous regions”, in *Proc. Intl conf. SAS/SAR*, Sep. 2010. [4, 15, 39, 44, 203]
- Blacknell, D., Andre, D., Finch, C., Nottingham, M. & Muff, D., “Life in the frequency domain: SAR CCD over mountains and the physics of vibrating targets”, in *Proc. DASP*, Jul. 2011. [4, 15, 39, 44, 203]
- Borden, B. & Cheney, M., “Radar imaging”, *Inverse Problems*, 29(5), May 2013. [11, 23]
- Born, M. & Wolf, E., *Principles of Optics*, 7th edn, CUP, 1999. [61]
- Boss, N., Ertin, E. & Moses, R., “Autofocus for 3D imaging with multipass SAR”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 7699, Apr. 2010. [84]
- Bracewell, R. N., “Strip integration in radio astronomy”, *Aust. J. Physics*, 9(2):198–217, 1956. [65]

- Bracewell, R. N., *Fourier analysis and imaging*, Kluwer Academic, 2003. [61, 64, 65]
- Buckingham, W. A., *Operation Ranch Hand: The Air Force and herbicides in Southeast Asia 1961-1971*, Office of Air Force History, 1982. [1]
- Bureau of Meteorology, “Severe Tropical Cyclone Yasi”, 2011, (also see Queensland Government National Parks, “Cyclone Yasi”). [32]
- Cafforio, C., Prati, C. & Rocca, F., “SAR data focusing using seismic migration techniques”, *IEEE TAES*, 27(2):194–207, Mar. 1991. [11]
- Cantalloube, H. M. J., Combernoux, A. & Oriot, H., “Repeat-pass autofocus for airborne polarimetric SAR tomography”, in *Proc. Intl Radar Symposium*, Jun. 2017. [84]
- Capon, J., Greenfield, R. J. & Kolker, R. J., “Multidimensional maximum-likelihood processing of a large aperture seismic array”, *Proc. IEEE*, 55(2):192–211, Feb. 1967. [103]
- Capon, J., “High-resolution frequency-wavenumber spectrum analysis”, *Proc. IEEE*, 57(8):1408–1418, Aug. 1969. [2, 85, 101]
- Capon, J. & Goodman, N. R., “Probability distributions for estimators of the frequency-wavenumber spectrum”, *Proc. IEEE*, 58(10):1785–1786, Oct. 1970, (minor corrections Jan. 1971). [85, 101]
- Carlson, G. E., *Signal and linear system analysis*, 2nd edn, John Wiley & Sons, 1998. [24, 28, 29, 61, 62, 70, 73]
- Carrara, W. G., Goodman, R. S. & Majewski, R. M., *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*, Artech House, Boston, 1995. [11, 30, 40, 63, 69, 83]
- Carter, G. C., Knapp, C. H. & Nuttall, A. H., “Statistics of the estimate of the magnitude-coherence function”, *IEEE Trans. Audio & Electroacoustics*, 21(4):388–389, Aug. 1973. [60]
- Chan, C. K. & Farhat, N. H., “Frequency swept tomographic imaging of three-dimensional perfectly conducting objects”, *IEEE TAP*, AP-29(2):312–319, Mar. 1981. [22, 84]
- Christensen, E. L., Madsen, S. N., Dall, J. & Skou, N., “The Danish polarimetric SAR for remote sensing applications”, in *Proc. IGARSS*, Aug. 1994. [175]
- Cloude, S. R., “Target decomposition theorems in radar scattering”, *Electronics Letters*, 21(1):22–24, Jan. 1985. [166, 167, 193, 200]
- Cloude, S. R., “Special unitary groups in polarimetry theory”, in *Proc. SPIE Polarization Analysis & Measurement II Conf.*, vol. 2265, Sep. 1994. [183, 200]
- Cloude, S. R. & Papathanassiou, K. P., “Polarimetric optimisation in radar interferometry”, *Electronics Letters*, 33(13):1176–1178, Jun. 1997. [3, 82]
- Cloude, S. R. & Papathanassiou, K. P., “Polarimetric SAR interferometry”, *IEEE TGRS*, 36(5):1551–1565, Sep. 1998. [3, 82, 166]
- Cloude, S. R., Papathanassiou, K. P., Reigber, A. & Boerner, W. M., “Multi-frequency polarimetric SAR interferometry for vegetation structure extraction”, in *Proc. IGARSS*, Jul. 2000. [53, 80, 82, 123, 129]
- Cloude, S. R. & Papathanassiou, K. P., “Three-stage inversion process for polarimetric SAR interferometry”, *IEE Proc. - Radar, Sonar & Navig.*, 150(3):125–134, Jun. 2003. [3, 82, 86, 128]
- Cloude, S. R., Corr, D. G. & Williams, M. L., “Target detection beneath foliage using polarimetric synthetic aperture radar interferometry”, *Waves in Random Media*, 14(2):S393–S414, 2004. [53, 83]

- Cloude, S. R., “Polarization coherence tomography”, *Radio Science*, 41(4):1–27, Aug. 2006. [81, 100]
- Cloude, S. R., *Polarisation: Applications in Remote Sensing*, OUP, 2010. [3, 5, 13, 32, 61, 64, 81, 82, 97, 117, 118, 120, 128, 129, 131, 133, 134, 145, 166, 167, 182, 184, 199, 201]
- Cloude, S. R. & Pottier, E., “Concept of polarization entropy in optical scattering”, *Optical Engineering*, 34(6):1599–1610, Jun. 1995. [200, 201]
- Cloude, S. R. & Pottier, E., “A review of target decomposition theorems in radar polarimetry”, *IEEE TGRS*, 34(2):498–518, Mar. 1996. [3, 166, 167, 188, 201]
- Cloude, S. R. & Pottier, E., “An entropy based classification scheme for land applications of polarimetric SAR”, *IEEE TGRS*, 35(1):68–78, Jan. 1997. [6, 86, 167, 168, 193, 199, 201]
- Colin, E., Titin-Schnaider, C. & Tabbara, W., “Coherence optimization methods for scattering centers separation in polarimetric interferometry”, *J. Electromagn. Waves and Appl.*, 19(9):1237–1250, 2005a. [82, 117]
- Colin, E., Titin-Schnaidér, C. & Tabbara, W., “FOPEN with polarimetric interferometry: Validations with experimental data at P-band”, in *Proc. PolInSAR Workshop*, Jan. 2005b. [83]
- Conradsen, K., Nielsen, A. A., Schou, J. & Skriver, H., “A test statistic in the complex Wishart distribution and its application to change detection in polarimetric SAR data”, *IEEE TGRS*, 41(1):4–19, Jan. 2003. [88]
- Couch, L. W., II, “Communication principles – intelligence coding”, in Whitaker, J. C., ed., *The electronics handbook*, 2nd edn, CRC Press, 2005. [66]
- Cumming, I. G. & Wong, F. H., *Digital processing of synthetic aperture radar data: Algorithms and implementation*, Artech House, Boston, 2005. [11, 26, 29]
- Curlander, J. C. & McDonough, R. N., *Synthetic aperture radar: Systems and signal processing*, John Wiley & Sons, 1991. [26, 29, 72]
- Cutrona, L. J., Vivian, W. E., Leith, E. N. & Hall, G. O., “A high-resolution radar combat-surveillance system”, *IRE Transactions on Military Electronics*, 5(2):127–131, Apr. 1961. [1, 11]
- Cutrona, L. J., “Synthetic aperture radar”, in Skolnik, M. I., ed., *Radar Handbook*, 2nd edn, McGraw-Hill, 1990. [89, 91]
- d’Alessandro, M. M. & Tebaldini, S., “Phenomenology of P-band scattering from a tropical forest through three-dimensional SAR tomography”, *IEEE GRS Letters*, 9(3):442–446, May 2012. [86, 100]
- d’Alessandro, M. M., Tebaldini, S. & Rocca, F., “Phenomenology of ground scattering in a tropical forest through polarimetric SAR tomography”, *IEEE TGRS*, 51(8):4430–4437, Aug. 2013. [86, 100]
- Dall, J., “InSAR elevation bias caused by penetration into uniform volumes”, *IEEE TGRS*, 45(7):2319–2324, Jul. 2007. [80, 128]
- d’Assumpcao, H. A. & Mountford, G. C., “An overview of signal processing for arrays of receivers”, *J. EEE, Australia*, 4(1):6–18, Mar. 1984, URL <http://www.dst.defence.gov.au/publication/overview-signal-processing-arrays-receivers>, accessed: 26/02/2019. [103, 114, 143]
- d’Assumpcao, H. & Gray, D., *Beamforming and Array Processing*, p. 218, University of Adelaide, 2007. [101, 103]
- Davidson, G. W. & Bamler, R., “Multiresolution phase unwrapping for SAR interferometry”, *IEEE TGRS*, 37(1):163–174, Jan. 1999. [14]
- Davis, M. E., *Foliage Penetration Radar*, SciTech, 2011. [1, 78, 84]

- Davis, M. E. & Reis, J. J., “Operation of GeoSAR foliage penetration modes in spectrum regulated environment”, in *Proc. IEEE Radar Conf.*, Oct. 2015. [89]
- DeGraaf, S. R., “SAR imaging via modern 2D spectral estimation”, *IEEE TIP*, 7(5):729–761, May 1998. [22]
- DeGraaf, S. R., “3-d fully polarimetric wide-angle superresolution-based SAR imaging for ATR”, in *Proc. Adaptive Sensor Array Processing Workshop*, Jun. 2005. [83]
- Deledalle, C.-A., Denis, L., Tupin, F., Reigber, A. & Jäger, M., “NL-SAR: A unified nonlocal framework for resolution-preserving (Pol)(In)SAR denoising”, *IEEE TGRS*, 53(4):2021–2038, Apr. 2015. [13, 101]
- Desai, M. D. & Jenkins, W. K., “Convolution backprojection image reconstruction for spotlight mode synthetic aperture radar”, *IEEE TIP*, 1(4):505–517, Oct. 1992. [4, 12]
- Dhar, T., Gray, D. & Menges, C., “FDTD based vegetation simulation – application to multiple wavelengths and vegetation species”, in *Proc. IGARSS*, Jul. 2012. [82]
- Doren, N. E., Jakowatz, C. V., Wahl, D. E. & Thompson, P. A., “General formulation for wavefront curvature correction in polar-formatted spotlight-mode sar images using space-variant post-filtering”, in *Proc. Im. Proc. Conf.*, vol. 3070, pp. 861–864, Oct. 1997. [12]
- Duboc, C. A., “Corner reflectors”, Tech. rep., MIT, Jan. 1943, <http://dspace.mit.edu/handle/1721.1/68311>. [184]
- Duque, S., Rossi, C. & Fritz, T., “Single-pass tomography with alternating bistatic TanDEM-X data”, *IEEE GRS Letters*, 12(2):409–413, Feb. 2015. [87]
- Durden, S. L., van Zyl, J. J. & Zebker, H. A., “Modeling and observation of the radar polarization signature of forested areas”, *IEEE TGRS*, 27(3):290–301, May 1989. [3, 81, 118, 164]
- Durden, S. L., Klein, J. D. & Zebker, H. A., “Polarimetric radar measurements of a forested area near Mt Shasta”, *IEEE TGRS*, 29(3):444–450, May 1991. [79, 81]
- Durden, S. L., Klein, J. D. & Zebker, H. A., “Measurement and simulation of signal fluctuations caused by propagation through trees”, *Radio Science*, 28(6):1049–1051, Nov. 1993. [80, 81, 82]
- Erten, E., Reigber, A., Ferro-Famil, L. & Hellwich, O., “A new coherent similarity measure for temporal multichannel scene characterization”, *IEEE TGRS*, 50(7):2839–2851, Jul. 2012. [88]
- Erten, E., Lopez-Sanchez, J. M., Yuzugullu, O. & Hajnsek, I., “Retrieval of agricultural crop height from space: A comparison of SAR techniques”, *Remote Sensing of Environment*, 187:130–144, Dec. 2016. [81]
- ESA Report, “Mission selection: Biomass (ESA SP-1324/1)”, May 2012. [82]
- ESA Sentinel-1, “S-1A & S-1B annual performance report for 2016”, Apr. 2017. [92]
- Farr, T. G., Rosen, P. A., Caro, E., Crippen, R., Duren, R., Hensley, S., Kobrick, M., Paller, M., Rodriguez, E., Roth, L., Seal, D., Shaffer, S., Shimada, J., Umland, J., Werner, M., Oskin, M., Burbank, D. & Alsdorf, D., “The shuttle radar topography mission”, *Reviews of Geophysics*, 45(2):1–33, Jun. 2007. [18, 52]
- Ferro-Famil, L., Huang, Y. & Reigber, A., “High resolution SAR tomography using full rank polarimetric spectral estimators”, in *Proc. IGARSS*, Jul. 2012. [86, 100, 204]
- Fleischman, J. G., Ayasli, S., Adams, E. M. & Gosselin, D. R., “Part I: Foliage attenuation and backscatter analysis of SAR imagery”, *IEEE TAES*, 32(1):135–144, Jan. 1996. [80, 82, 198]
- Flynn, T., Tabb, M. & Carande, R., “Coherence region shape extraction for vegetation parameter estimation in polarimetric SAR interferometry”, in *Proc. IGARSS*, Jun. 2002. [82, 117, 128]

- Fornaro, G., Lombardini, F. & Serafino, F., “Three-dimensional multipass SAR focusing: Experiments with long-term spaceborne data”, *IEEE TGRS*, 43(4):702–714, Apr. 2005. [2]
- Freeman, A., “A new system model for radar polarimeters”, *IEEE TGRS*, 29(5):761–767, Sep. 1991. [3, 118, 176, 177, 178, 180]
- Freeman, A., “SAR calibration: An overview”, *IEEE TGRS*, 30(6):1107–1121, Nov. 1992. [3, 176]
- Freeman, A. & Durden, S. L., “A three-component scattering model for polarimetric SAR data”, *IEEE TGRS*, 36(3):963–973, May 1998. [3, 80, 81, 118, 164]
- Freeman, A., “On the design of spaceborne polarimetric SARs”, in *Proc. IEEE Radar Conf.*, May 2009. [176]
- Frey, O., Magnard, C., Ruegg, M. & Meier, E., “Focusing of airborne synthetic aperture radar data from highly nonlinear flight tracks”, *IEEE TGRS*, 47(6):1844–1858, Jun. 2009. [11]
- Frey, O. & Meier, E., “3-d time-domain SAR imaging of a forest using airborne multibaseline data at L- and P-bands”, *IEEE TGRS*, 49(10):3660–3664, Oct. 2011a. [2, 86]
- Frey, O. & Meier, E., “Analyzing tomographic SAR data of a forest with respect to frequency, polarization, and focusing technique”, *IEEE TGRS*, 49(10):3648–3659, Oct. 2011b. [2, 86, 100]
- Fröling, P.-O., Gustavsson, A., Lundberg, M. & Ulander, L. M. H., “Circular-aperture VHF-band SAR for detection of vehicles in forest concealment”, *IEEE TGRS*, 50(4):1329–1339, Apr. 2012. [83]
- Gallant, J. C., Dowling, T. I., Read, A. M., Wilson, N., Tickle, P. & Inskeep, C., “1 second SRTM derived digital elevation models: user guide”, Geoscience Australia, 2011, URL www.ga.gov.au/topographic-mapping/digital-elevation-data.html, accessed: 18/12/2014. [18, 52, 192]
- Garestier, F., Dubois-Fernandez, P. C. & Champion, I., “Forest height inversion using high-resolution P-band Pol-InSAR data”, *IEEE TGRS*, 46(11):3544–3559, Nov. 2008. [81]
- Gatelli, F., Guarnieri, A. M., Parizzi, F., Pasquali, P., Prati, C. & Rocca, F., “The wavenumber shift in SAR interferometry”, *IEEE TGRS*, 32(4):855–865, Jul. 1994. [2, 4, 14, 39, 47, 97]
- Gau, J.-R. J. & Burnside, W. D., “New polarimetric calibration technique using a single calibration dihedral”, *IEE Proc. - Microwaves, Antennas & Propagation*, 142(1):19–25, Feb. 1995. [7, 176, 181]
- Gershman, A. B., Luo, Z.-Q. & Shahbazpanahi, S., “Robust adaptive beamforming based on worst-case performance optimization”, in Li, J. & Stoica, P., eds., *Robust Adaptive Beamforming*, Wiley, 2006. [85, 102, 145]
- Ghiglia, D. C., “IFSAR correlation improvement through local slope correction”, in *Proc. IGARSS*, Jul. 1998. [15]
- Giglio, D. A., “Overview of foliage/ground penetration and interferometric SAR experiments”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 2230, Apr. 1994. [79]
- Goh, A. S., Preiss, M., Gray, D. A. & Stacy, N. J. S., “Comparison of parameter estimation accuracy of distributed-target polarimetric calibration techniques”, in *Proc. IGARSS*, Jul. 2007. [3, 176, 177, 181, 186, 191]
- Goh, A. S., *Bistatic Synthetic Aperture Radar Data Processing and Analysis*, Ph.D. thesis, University of Adelaide, Adelaide, Australia, Feb. 2012, URL <https://digital.library.adelaide.edu.au/dspace/handle/2440/71871>. [181, 185, 192]

- Goodman, J. W., “Statistical properties of laser speckle patterns”, in Dainty, J., ed., *Laser Speckle and Related Phenomena*, vol. 9, Springer-Verlag, 1975. [1, 13, 59, 75, 123]
- Goodman, N. R., “Statistical analysis based on a certain multivariate complex Gaussian distribution (an introduction)”, *Annals of Mathematical Statistics*, 34(1):152–177, Mar. 1963. [101]
- Graham, L. C., “Synthetic interferometer radar for topographic mapping”, *IEEE Proc.*, 62(6):763–768, Jun. 1974. [13, 95]
- Graves, C. D., “Radar polarization power scattering matrix”, *IRE Proc.*, 44(2):248–252, Feb. 1956. [177]
- Gray, A. L., Vachon, P. W., Livingstone, C. E. & Lukowski, T. I., “Synthetic aperture radar calibration using reference reflectors”, *IEEE TGRS*, 28(3):374–383, May 1990. [180]
- Griffiths, H., Cohen, L., Watts, S., Mokole, E. & Baker, C., “Radar spectrum engineering and management: Technical and regulatory issues”, *IEEE Proc.*, 103(1):85–102, Jan. 2015. [73, 89]
- Guerci, J. R., “Principal components, covariance matrix tapers, and the subspace leakage problem”, *IEEE TAES*, 38(1):152–162, Jan. 2002. [85]
- Guinvarc’h, R. & Thirion-Lefevre, L., “Propagation in urban areas: Orientation, permittivity and entropy”, in *Proc. IGARSS*, Jul. 2015. [81]
- Hagberg, J. O., Ulander, L. M. H. & Askne, J., “Repeat-pass SAR interferometry over forested terrain”, *IEEE TGRS*, 33(2):331–340, Mar. 1995. [81, 120]
- Hajnsek, I., Papathanassiou, K. P. & Cloude, S. R., “Removal of additive noise in polarimetric eigenvalue processing”, in *Proc. IGARSS*, Jul. 2001. [200]
- Hajnsek, I., Jagdhuber, T., Schän, H. & Papathanassiou, K. P., “Potential of estimating soil moisture under vegetation cover by means of PolSAR”, *IEEE TGRS*, 47(2):442–454, Feb. 2009a. [1]
- Hajnsek, I., Kugler, F., Lee, S.-K. & Papathanassiou, K. P., “Tropical-forest-parameter estimation by means of Pol-InSAR: The INDREX-II campaign”, *IEEE TGRS*, 47(2):481–493, Feb. 2009b. [3, 82, 128, 145]
- Halversen, S. D., Nanis, J. G., Owirka, G. J. & Novak, L. M., “Comparison of ultrawideband SAR target detection algorithms”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 2230, Apr. 1994. [83, 87]
- Harcke, L., Weintraub, L., Yun, S.-H., Dickinson, R., Gurrola, E., Hensley, S. & Marechal, N., “Spotlight-mode synthetic aperture radar processing for high-resolution lunar mapping”, in *Proc. IEEE Radar Conf.*, May 2010. [12]
- Harris, F. J., “On the use of windows for harmonic analysis with the discrete fourier transform”, *IEEE Proc.*, 66(1):51–83, Jan. 1978. [63]
- Hayward, S. D., Blake, A., Denton, J. & Coe, D., “Coherent sub-band combination and polarimetric calibration in very wideband low-frequency SAR”, in *Proc. IEEE Radar Conf.*, May 2015. [8]
- Held, D. N., Brown, W. E., Freeman, A., Klein, J. D., Zebker, H., Sato, T., Miller, T., Nguyen, Q. & Lou, Y., “The NASA-JPL multifrequency, multipolarisation airborne SAR system”, in *Proc. IGARSS*, Sep. 1988. [79]
- Held, D. N. & O’Brien, J. D., “Norden systems’ 2-dimensional interferometric SAR flying testbed”, in *Proc. IGARSS*, May 1992. [87]
- Hellsten, H., “CARABAS – an UWB low frequency SAR”, in *Proc. IEEE MTT-S Microwave Symposium*, Jun. 1992. [79]

- Hellsten, H., Ulander, L. M. H., Gustavsson, A. & Larsson, B., “Development of VHF CARABAS II SAR”, in *Proc. SPIE Radar Sensor Tech. Conf.*, vol. 2747, pp. 48–60, Apr. 1996. [1, 12, 79, 83]
- Hensley, S., Michel, T., Neumann, M., Lavalley, M., Ahmed, R., Muellerschoen, R. & Chapman, B., “A comparison of multi-baseline polarimetric interferometry at La Amistad and La Selva, Costa Rica with a modified PolSARproSim scattering tool”, in *Proc. EuSAR Conf.*, Jun. 2014. [78]
- Hensley, S., Moller, D., Oveisgharan, S., Michel, T. & Wu, X., “Ka-band mapping and measurements of interferometric penetration of the Greenland ice sheets by the GLISTIN radar”, *J. Selected Topics in Applied Earth Observations and Remote Sensing*, 9(6):2436–2450, Jun. 2016. [81]
- Hector, R. T. & Kassam, S. A., “The unifying role of the coarray in aperture synthesis for coherent and incoherent imaging”, *IEEE Proc.*, 78(4):735–752, Apr. 1990. [22, 40, 84]
- Homer, J. & Longstaff, I. D., “Improving height resolution of interferometric SAR”, in *Proc. IGARSS*, Jul. 1995. [84, 145]
- Homer, J., Longstaff, I. D. & Callaghan, G., “High resolution 3-D SAR via multi-baseline interferometry”, in *Proc. IGARSS*, May 1996. [2, 84]
- Horn, R., “E-SAR – The experimental airborne L/C-band SAR system of DFVLR”, in *Proc. IGARSS*, Sep. 1988. [79]
- Horn, R., “The DLR airborne SAR project E-SAR”, in *Proc. IGARSS*, May 1996. [79]
- Horn, R. A. & Johnson, C. R., *Matrix Analysis*, 2nd edn, CUP, 2013. [102, 124, 127, 141, 142, 200]
- Huang, Y., Ferro-Famil, L. & Reigber, A., “Under-foilage object imaging using SAR tomography and polarimetric spectral estimators”, *IEEE TGRS*, 50(6):2213–2225, Jun. 2012. [2, 86, 100]
- Huang, Y., Zhang, Q., Schwäbisch, M., Wei, M. & Mercer, B., “Tomographic analysis for boreal forests using single-pass L-band PolInSAR data”, in *Proc. IGARSS*, Jul. 2013. [106]
- Huber, P. J. & Ronchetti, E. M., *Robust statistics*, 2nd edn, John Wiley & Sons, 2009. [101]
- IEEE Std 521-2002, “Standard letter designations for radar-frequency bands”, . [79]
- Institute of Medicine, *Veterans and Agent Orange: Health Effects of Herbicides Used in Vietnam*, The National Academies Press, 1994. [1]
- Jakowatz, C. V. & Wahl, D. E., “Three-dimensional tomographic imaging for foliage penetration using multiple-pass spotlight-mode SAR”, in *Proc. Sig., Sys. & Comp. Asilomar Conf.*, Nov. 2001. [2, 22, 84, 100, 162]
- Jakowatz, C. V. & Wahl, D. E., “Eigenvector method for maximum-likelihood estimation of phase errors in SAR imagery”, *J. Opt. Soc. Am. A*, 10(12):2539–2546, Dec. 1993. [12]
- Jakowatz, C. V., Wahl, D. E., Eichel, P. H., Ghiglia, D. C. & Thompson, P. A., *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*, Springer, Sandia National Laboratories Albuquerque, 1996. [1, 4, 11, 12, 13, 14, 15, 18, 19, 20, 21, 24, 29, 30, 32, 36, 40, 41, 59, 64, 65, 68, 69, 71, 75, 77, 84, 87, 94, 98, 162]
- Jakowatz, C. V., Wahl, D. E. & Yocky, D. A., “Beamforming as a foundation for spotlight-mode SAR image formation by backprojection”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 6970, Apr. 2008. [4, 12]
- Jakowatz, C. V. & Wahl, D. E., “Considerations for autofocus of spotlight-mode SAR imagery created using a beamforming algorithm”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 7337, Apr. 2009. [12, 44]

- James, R. R. & Hendrickson, C. R., “Efficacy of frequency on detecting targets in foliage using incoherent change detection”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 2230, Apr. 1994. [83]
- Jao, J. K., Lee, C. F., Ayasli, S. & Haywiser, E. J., “Analysis and results of the 1995 Yuma ground penetration SAR measurements”, *IEEE A & E Systems Magazine*, 14(6), Jun. 1999. [117]
- Joerg, H., Pardini, M., Hajnsek, I. & Papathanassiou, K. P., “On the separation of ground and volume scattering using multibaseline SAR data”, *IEEE GRS Letters*, 14(9):1570–1574, Sep. 2017. [103]
- Joughin, I. R., Winebrenner, D. P. & Percival, D. B., “Probability density functions for multilook polarimetric signatures”, *IEEE TGRS*, 32(3):562–574, May 1994. [60]
- Just, D. & Bamler, R., “Phase statistics of interferograms with applications to SAR”, *Applied Optics*, 33(20):4361–4368, Jul. 1994. [95, 98, 119, 128]
- Kak, A. C. & Slaney, M., *Principles of Computerized Tomographic Imaging*, IEEE Press, 1988. [12]
- Kankaku, Y., Suzuki, S. & Osawa, Y., “ALOS-2 mission and development status”, in *Proc. IGARSS*, Jul. 2013. [92]
- Kay, S. M., *Fundamentals of statistical signal processing: Estimation theory*, Prentice Hall, 1993. [59, 60, 101]
- Khatri, U., Singh, G. & Kumar, S., “Potential of space-borne PolInSAR for forest canopy height estimation over India – A case study using fully polarimetric L- C- and X-band SAR data”, *J. Selected Topics in Applied Earth Observations and Remote Sensing*, 11(7):2406–2416, Jul. 2018. [82]
- Klauder, J. R., Price, A. C., Darlington, S. & Albersheim, W. J., “The theory and design of chirp radars”, *Bell System Technical Journal*, 39(4):745–808, Jul. 1960. [1, 11, 22, 23]
- Klein, J. D., “Calibration of complex polarimetric SAR imagery using backscatter correlations”, *IEEE TAES*, 28(1):183–194, Jan. 1992. [3, 176, 177, 183, 186, 191]
- Knaell, K. K. & Cardillo, G. P., “Radar tomography for the generation of three-dimensional images”, *IEE Proc. - Radar, Sonar & Navig.*, 142(2):54–60, Apr. 1995. [22, 84]
- Kragh, T. J., “Minimum-entropy autofocus for three-dimensional SAR imaging”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 7337, Apr. 2009. [84]
- Krieger, G., Hajnsek, I., Papathanassiou, K., Younis, M. & Moreira, A., “Interferometric synthetic aperture radar missions employing formation flying”, *IEEE Proc.*, 98(5):816–843, May 2010. [87]
- Kugler, F., Schulze, D., Hajnsek, I., Pretzsch, H. & Papathanassiou, K. P., “TanDEM-X PolInSAR performance for forest height estimation”, *IEEE TGRS*, 52(10):6404–6422, Oct. 2014. [82]
- Kugler, F., Lee, S.-K., Hajnsek, I. & Papathanassiou, K. P., “Forest height estimation by means of PolInSAR data inversion: The role of the vertical wavenumber”, *IEEE TGRS*, 53(10):5294–5311, Oct. 2015. [82]
- Lamont-Smith, T., Hill, R. D., Hayward, S. D., Yates, G. & Blake, A., “Filtering approaches for interference suppression in low-frequency SAR”, *IEE Proc. - Radar, Sonar & Navig.*, 153(4):338–344, Aug. 2006. [73]
- Lavalle, M., Solimini, D., Pottier, E. & Desnos, Y.-L., “Forest parameters inversion using polarimetric and interferometric SAR data”, in *Proc. IGARSS*, Jul. 2009. [78]
- Lavalle, M., Simard, M. & Hensley, S., “A temporal decorrelation model for polarimetric radar interferometers”, *IEEE TGRS*, 50(7):2880–2888, Jul. 2012. [82]

- Lavalle, M. & Hensley, S., “Extraction of structural and dynamic properties of forests from polinsar data affected by temporal decorrelation”, *IEEE TGRS*, 53(9):4752–4767, Sep. 2015. [82]
- Lee, H. & Liu, J. G., “Analysis of topographic decorrelation in SAR interferometry using ratio coherence imagery”, *IEEE TGRS*, 39(2):223–232, Feb. 2001. [14]
- Lee, J.-S., “Speckle analysis and smoothing of synthetic aperture radar images”, *Computer Graphics & Image Processing*, 17(1):24–32, Sep. 1981. [2, 13]
- Lee, J.-S., Hoppel, K. W., Mango, S. A. & Miller, A. R., “Intensity and phase statistics of multilook polarimetric and interferometric SAR imagery”, *IEEE TGRS*, 32(5):1017–1028, Sep. 1994. [101]
- Lee, J.-S., Ainsworth, T. L., Kelly, J. P. & López-Martínez, C., “Evaluation and bias removal of multilook effect on entropy/alpha/anisotropy in polarimetric SAR decomposition”, *IEEE TGRS*, 46(10):3039–3052, Oct. 2008. [167, 193]
- Lee, J.-S. & Pottier, E., *Polarimetric Radar Imaging: From Basics to Applications*, CRC Press, 2009. [176, 180, 182, 183, 184, 200]
- Lee, S.-K., Kugler, F., Papathanassiou, K. P. & Hajnsek, I., “Quantification of temporal decorrelation effects at L-band for polarimetric SAR interferometry applications”, *J. Selected Topics in Applied Earth Observations and Remote Sensing*, 6(3):1351–1367, Jun. 2013. [82]
- Lee, S.-K. & Fatoyinbo, T. E., “TanDEM-X PolInSAR inversion for mangrove canopy height estimation”, *J. Selected Topics in Applied Earth Observations and Remote Sensing*, 8(7):3608–3618, Jul. 2015. [82]
- Li, F. K. & Goldstein, R. M., “Studies of multibaseline spaceborne interferometric synthetic aperture radars”, *IEEE TGRS*, 28(1):88–97, Jan. 1990. [2, 14]
- Lin, Y.-C. & Sarabandi, K., “A Monte Carlo coherent scattering model for forest canopies using fractal-generated trees”, *IEEE TGRS*, 37(1):440–451, Jan. 1999. [82]
- Lombardini, F., Ender, J., Röbbing, L., Galletto, M. & Verrazzani, L., “Experiments of interferometric layover solution with the three-antenna airborne AER-II SAR system”, in *Proc. IGARSS*, Sep. 2004. [87]
- Lombardini, F., Pardini, M., Fornaro, G., Serafino, F., Verrazzani, L. & Costantini, M., “Linear and adaptive spaceborne three-dimensional SAR tomography: a comparison on real data”, *IEEE TGRS*, 3(4):424–436, Aug. 2009. [85, 102, 145]
- Lombardini, F. & Lombardo, P., “Maximum likelihood array SAR interferometry”, in *Proc. IEEE Dig. Sig. Pro. Workshop*, Sep. 1996. [84]
- Lombardini, F. & Reigber, A., “Adaptive spectral estimation for multibaseline SAR tomography with airborne L-band data”, in *Proc. IGARSS*, Jul. 2003. [85]
- Lombardini, F., Montanari, M. & Gini, F., “Reflectivity estimation for multibaseline interferometric radar imaging of layover extended sources”, *IEEE TSP*, 51(6):1508–1519, Jun. 2003. [2, 85, 99, 101]
- Lombardini, F. & Pardini, M., “3D SAR tomography: The multibaseline sector interpolation approach”, *IEEE GRS Letters*, 5(4):630–634, Oct. 2008. [85, 104]
- Long, M. W., *Radar reflectivity of land and sea*, 3rd edn, Artech House, 2001. [91]
- López-Martínez, C., Pottier, E. & Cloude, S. R., “Statistical assessment of eigenvector-based target decomposition theorems in radar polarimetry”, *IEEE TGRS*, 43(9):2058–2074, Sep. 2005. [167, 193, 199]
- López-Martínez, C., Cortés, A. & Fàbregas, X., “Analysis and improvement of polarimetric calibration techniques”, in *Proc. IGARSS*, Jul. 2007. [3, 176, 177, 181]

- López-Martínez, C. & Alonso-González, A., “Assessment and estimation of the RVoG model in polarimetric SAR interferometry”, *IEEE TGRS*, 52(6):3091–3106, Jun. 2014. [135]
- Lorenz, R. G. & Boyd, S. P., “Robust minimum variance beamforming”, *IEEE TSP*, 53(5):1684–1696, May 2005. [85, 102, 145]
- Lundberg, M., Ulander, L. M. H., Pierson, W. E. & Gustavsson, A., “A challenge problem for detection of targets in foliage”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 6237, Apr. 2006. [83]
- Lüneburg, E., “Principles of radar polarimetry”, *IEICE Trans. Electron.*, E78-C(10):1339–1345, Oct. 1995. [182]
- Madsen, S. N. & Zebker, H. A., “Automated absolute phase retrieval in across-track interferometry”, in *Proc. IGARSS*, May 1992. [95]
- Marechal, N., “Tomographic formulation of interferometric SAR for terrain elevation mapping”, *IEEE TGRS*, 33(3):726–739, May 1995. [2, 4, 14]
- Marino, A., Cloude, S. R. & Woodhouse, I. H., “A polarimetric target detector using the Huynen fork”, *IEEE TGRS*, 48(5):2357–2366, May 2010. [83]
- Marino, A., Cloude, S. R. & Lopez-Sanchez, J. M., “A new polarimetric change detector in radar imagery”, *IEEE TGRS*, 51(5):2986–3000, May 2013. [84, 88]
- Marino, R. M. & Davis, W. R., “Jigsaw: A foliage-penetrating 3d imaging laser radar system”, *Lincoln Laboratory Journal*, 15(1):23–36, 2005. [81]
- Marzetta, T. L., “A new interpretation for Capon’s maximum likelihood method of frequency-wavenumber spectral estimation”, *IEEE Transactions on Acoustics, Speech and Signal Processing*, 31(2):445–449, 1983. [103]
- McCorkle, J. & Rofheart, M., “An order $N^2 \log N$ backprojector algorithm for focusing wide-angle wide-bandwidth arbitrary-motion synthetic aperture radar”, in *Proc. SPIE Radar Sensor Tech. Conf.*, vol. 2747, pp. 25–36, Apr. 1996. [12]
- McNairn, H., Jackson, T. J., Wiseman, G., Bélair, S., Berg, A., Bullock, P., Colliander, A., Cosh, M. H., Kim, S.-B., Magagi, R., Moghaddam, M., Njoku, E. G., Adams, J. R., Homayouni, S., Ojo, E. R., Rowlandson, T. L., Shang, J., Goäta, K. & Hosseini, M., “The soil moisture active passive validation experiment 2012 (SMAPVEX12): Prelaunch calibration and validation of the SMAP soil moisture algorithms”, *IEEE TGRS*, 53(5):2784–2801, May 2015. [1]
- MDA, “RADARSAT-2 product description”, Mar. 2016. [92]
- Mercer, B., Zhang, Q., Schwäbisch, M., Denbina, M. & Cloude, S., “Forest height and ground topography at L-band from an experimental single-pass airborne PolInSAR system”, in *Proc. PolInSAR Workshop*, Jan. 2009, (also see presentation). [82, 92, 107]
- Meyer, F. J., Nicoll, J. B. & Doulgeris, A. P., “Correction and characterization of radio frequency interference signatures in L-band synthetic aperture radar data”, *IEEE TGRS*, 51(10):4961–4972, Oct. 2013. [73]
- Michelson, D. G., Cumming, I. G. & Livingstone, C. E., “A calibration algorithm for circular polarimetric radars”, *J. Electromagnetic Waves & Applications*, 11(5):659–674, May 1997. [3, 7, 175, 176, 182]
- Migliaccio, M., Nunziata, F. & Buono, A., “SAR polarimetry for sea oil slick observation”, *Intl J. Remote Sensing*, 36(12):3243–3273, Jun. 2015. [88]
- Milman, A. S., “SAR imaging by $\omega - \kappa$ migration”, *International Journal of Remote Sensing*, 14(10):1965–1975, 1993. [12]
- Minh, D. H. T., Tebaldini, S., Rocca, F., Koleck, T., Borderies, P., Albinet, C., Villard, L., Hamadi, A. & Toan, T. L., “Ground-based array for tomographic imaging of the tropical forest in P-band”, *IEEE TGRS*, 51(8):4460–4472, Aug. 2013. [86]

- Minh, D. H. T., Tebaldini, S., Rocca, F., Toan, T. L., Borderies, P., Koleček, T., Albinet, C., Hamadi, A. & Villard, L., “Vertical structure of P-band temporal decorrelation at the Paracou forest: Results from TropiScat”, *IEEE GRS Letters*, 11(8):1438–1442, Aug. 2014. [86]
- Moghaddam, M. & Saatchi, S., “Analysis of scattering mechanisms in SAR imagery over boreal forest: Results from BOREAS ’93”, *IEEE TGRS*, 33(5):1290–1296, Sep. 1995. [81]
- Munson, D. C., O’Brien, J. D. & Jenkins, W. K., “A tomographic formulation of spotlight-mode SAR”, *IEEE Proc.*, 71(8):917–925, Aug. 1983. [1, 11, 22, 29, 32]
- Munson, D. C. & Sanz, J. L. C., “Image reconstruction from frequency-offset Fourier data”, *IEEE Proc.*, 72(6):661–669, Jun. 1984. [22]
- Murray, R. M., Li, Z. & Sastry, S. S., *A mathematical introduction to robotic manipulation*, CRC Press, 1994. [64]
- Nannini, M., Scheiber, R. & Moreira, A., “Estimation of the minimum number of tracks for SAR tomography”, *IEEE TGRS*, 47(2):531–543, Feb. 2009. [85]
- Nannini, M., Scheiber, R., Horn, R. & Moreira, A., “First 3D reconstructions of targets hidden beneath foliage by means of polarimetric SAR tomography”, *IEEE GRS Letters*, 9(1):60–64, Jan. 2012. [2, 86, 100]
- Neumann, M., Ferro-Famil, L. & Reigber, A., “Estimation of forest structure, ground, and canopy layer characteristics from multibaseline polarimetric interferometric SAR data”, *IEEE TGRS*, 48(3):1086–1104, Mar. 2010. [81]
- Newey, M., Benitz, G. & Kogon, S., “A generalized likelihood ratio test for SAR CCD”, in *Proc. Sig., Sys. & Comp. Asilomar Conf.*, Nov. 2012. [87]
- Nghiem, S. V., Yueh, S. H., Kwok, R. & Li, F. K., “Symmetry properties in polarimetric remote sensing”, *Radio Science*, 27(5):693–711, Sep. 1992. [3, 176, 188]
- Nguyen, L., Kapoor, R. & Sichina, J., “Detection algorithms for ultrawideband foliage-penetration radar”, in *Proc. SPIE Radar Sensor Tech. Conf.*, vol. 3066, Apr. 1997. [1, 83]
- North, D. O., “An analysis of the factors which determine signal/noise discrimination in pulsed-carrier systems”, *IEEE Proc.*, 51(7):1016–1027, Jul. 1963, (reprint of RCA Report PTR-6C, Jun. 1943). [26]
- Novak, L. M., “Change detection for multi-polarization, multi-pass SAR”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 5808, Apr. 2005a. [87]
- Novak, L. M., “Coherent change detection for multi-polarization SAR”, in *Proc. Sig., Sys. & Comp. Asilomar Conf.*, Oct. 2005b. [87]
- Oliver, C. & Quegan, S., *Understanding Synthetic Aperture Radar Images*, Artech House, 1998, (reprinted by SciTech Publishing, 2004). [13, 60, 101, 147, 180]
- Oriot, H., “Activity monitoring with airborne SAR imagery”, *NATO STO SET Panel 191*, Nov. 2014. [1, 87]
- Papathanassiou, K. P. & Cloude, S. R., “Single-baseline polarimetric SAR interferometry”, *IEEE TGRS*, 39(11):2352–2363, Nov. 2001. [3, 82, 119, 128, 131, 134]
- Papoulis, A., *Probability, random variables and stochastic processes*, 3rd edn, McGraw-Hill, 1991. [13, 59]
- Pardini, M., Cantini, A., Lombardini, F. & Papathanassiou, K., “3-D structure of forests: First analysis of tomogram changes due to weather and seasonal effects at L-band”, in *Proc. EuSAR Conf.*, Jun. 2014. [85, 101]
- Pardini, M. & Papathanassiou, K., “On the estimation of ground and volume polarimetric covariances in forest scenarios with SAR tomography”, *IEEE GRS Letters*, 14(10):1860–1864, Oct. 2017. [86, 103, 204]

- Pardini, M., Tello, M., Cazcarra-Bes, V., Papathanassiou, K. P. & Hajnsek, I., “L- and P-band 3D SAR Reflectivity Profiles Versus Lidar Waveforms: The AfriSAR Case”, *J. Selected Topics in Applied Earth Observations and Remote Sensing*, 11(10):3386–3401, Oct. 2018. [86, 204]
- Pincus, P. B., Preiss, M. & Gray, D. A., “Aspects of 3D tomography for multiple-pass spotlight-mode airborne SAR”, in *Proc. IGARSS*, Jul. 2009. [2, 84, 114]
- Pincus, P. B., Preiss, M., Goh, A., Stacy, N. J. S. & Gray, D., “Low frequency high resolution SAR imaging and polarimetric analysis of a Queensland tropical forest”, in *Proc. IGARSS*, Jul. 2013. [8, 13, 32, 83, 90, 192]
- Pincus, P. B., Preiss, M. & Gray, D. A., “3D SAR beamforming under a foliage canopy from a single pass”, in *Proc. ICASSP Conf.*, Apr. 2015a. [75]
- Pincus, P. B., Preiss, M., Stacy, N. J. S. & Gray, D. A., “Image formation on undulating terrain using the upgraded *Ingara* L-band radar system”, in *Proc. Asia-Pacific SAR Conf.*, Sep. 2015b. [3, 8, 19, 30, 39, 52, 92]
- Ponce, O., Prats-Iraola, P., Scheiber, R., Reigber, A. & Moreira, A., “First airborne demonstration of holographic SAR tomography with fully polarimetric multicircular acquisitions at L-band”, *IEEE TGRS*, 54(10):6170–6196, Oct. 2016. [84, 86]
- Pottier, E. & Cloude, S. R., “Application of the $H/A/\bar{\alpha}$ polarimetric decomposition theorem for land classification”, in *Proc. SPIE Wideband Interferometric Sensing and Imaging Polarimetry Conf.*, vol. 3120, Jul. 1997. [6, 167, 168]
- Praks, J., Kugler, F., Papathanassiou, K. P., Hajnsek, I. & Hallikainen, M., “Height estimation of boreal forest: Interferometric model-based inversion at L- and X-band versus HUTSCAT profiling scatterometer”, *IEEE GRS Letters*, 4(3):466–470, Jul. 2007. [3, 82]
- Praks, J., Antropov, O. & Hallikainen, M., “LIDAR-aided SAR interferometry studies in boreal forest: Scattering phase center and extinction coefficient at X- and L-band”, *IEEE TGRS*, 50(10):3831–3843, Oct. 2012. [82]
- Prats, P., Reigber, A. & Mallorqui, J. J., “Interpolation-free coregistration and phase-correction of airborne SAR interferograms”, *IEEE GRS Letters*, 1(3):188–191, Jul. 2004. [98]
- Prats, P., Reigber, A. & Mallorqui, J. J., “Topography-dependent motion compensation for repeat-pass interferometric SAR systems”, *IEEE GRS Letters*, 2(2):206–210, Apr. 2005. [98]
- Prats-Iraola, P., Scheiber, R., Marotti, L., Wollstadt, S. & Reigber, A., “TOPS interferometry with TerraSAR-X”, *IEEE TGRS*, 50(8):3179–3188, Aug. 2012. [14]
- Preiss, M., Gray, D. & Stacy, N., “Space variant filtering of polar format spotlight SAR images for wavefront curvature correction and interferometric processing”, in *Proc. IGARSS*, pp. 179–181, Jun. 2002. [12]
- Preiss, M. & Stacy, N. J. S., “Scene coherency at X-band from repeat pass polarimetric interferometry”, in *Proc. IGARSS*, pp. 1081–1084, IEEE, Jul. 2005. [81]
- Preiss, M., Goh, A. S. & Pincus, P. B., “First results from the *ingara* l-band sar”, in *Proc. Intl Conf. on Radar*, Aug. 2018. [8, 204]
- Preiss, M., Gray, D. A. & Stacy, N. J. S., “Detecting scene changes using synthetic aperture radar interferometry”, *IEEE TGRS*, 44(8):2041–2054, Sep. 2006. [1, 13, 87]
- Preiss, M. & Stacy, N. J. S., “Polarimetric SAR coherent change detection”, in *Proc. EuSAR Conf.*, Jun. 2008. [87]
- Q-PAR, “0.2-2 GHz dual linearly polarised sinuous antenna”, (the actual model was QSPDP0.2-2S from Q-PAR Angus; this datasheet (2015) is for the updated model from new parent Steatite). [9]

- Quegan, S., “A unified algorithm for phase and cross-talk calibration of polarimetric data—Theory and observations”, *IEEE TGRS*, 32(1):89–99, Jan. 1994. [3, 7, 176, 177, 180, 181, 183, 186]
- Raney, R. K., “Hybrid-polarity SAR architecture”, *IEEE TGRS*, 45(11):3397–3404, Nov. 2007. [176]
- Raney, K. I. & Khatri, H. C., “Modified difference change detector for small targets in SAR imagery”, *IEEE TAES*, 44(1):57–76, Jan. 2008. [87]
- Reale, D., Fornaro, G., Pauciuolo, A., Zhu, X. & Bamler, R., “Tomographic imaging and monitoring of buildings with very high resolution SAR data”, *IEEE GRS Letters*, 8(4):661–665, Jul. 2011. [85]
- Redding, N. J. & Payne, T. M., “Inverting the spherical radon transform for 3D SAR image formation”, in *Proc. Int. Radar Conf.*, Sep. 2003. [24]
- Reigber, A., Neumann, M., Guillaso, S., Sauer, S. & Ferro-Famil, L., “Evaluating PolInSAR parameter estimation using tomographic imaging results”, in *Proc. Euro. Radar Conf.*, Oct. 2005. [2, 85, 100, 124, 145]
- Reigber, A., Jäger, M., Fischer, J., Horn, R., Scheiber, R., Prats, P. & Nottensteiner, A., “Performance of the L- and P-band subsystems of the F-SAR airborne SAR instrument”, in *Proc. EuSAR Conf.*, Apr. 2012. [92]
- Reigber, A., Nannini, M., del Hoyo, A. M., del Campo Becerra, G. M. & Shkvarko, Y., “A comparative study of tomographic SAR focusing methods”, in *Proc. Euro. Radar Conf.*, Sep. 2015. [86, 145]
- Reigber, A., “Range dependent spectral filtering to minimize the baseline decorrelation in airborne SAR interferometry”, in *Proc. IGARSS*, Jun. 1999. [14]
- Reigber, A. & Moreira, A., “First demonstration of airborne SAR tomography using multibaseline L-band data”, *IEEE TGRS*, 38(5):2142–2152, Sep. 2000. [2, 22, 84]
- Rignot, E. J. M. & van Zyl, J. J., “Change detection techniques for ERS-1 SAR data”, *IEEE TGRS*, 31(4):896–906, Jul. 1993. [2, 13, 75, 87]
- Robinson, C. A., “Radar counters camouflage”, *Signal Magazine*, Jun. 2007. [84]
- Rodriguez, E. & Martin, J. M., “Theory and design of interferometric synthetic aperture radars”, *IEE Proceedings F – Radar and Signal Processing*, 139(2):147–159, Apr. 1992. [4, 14]
- Rogan, A. & Carande, R., “Improving the fast back projection algorithm through massive parallelizations”, in *Proc. SPIE Radar Sensor Tech. Conf.*, vol. 7669, Apr. 2010. [12]
- Rosa, R. A. S., Fernandes, D., Barreto, T. L. M., Wimmer, C. & Nogueira, J. B., “Change detection under the forest in multitemporal full-polarimetric P-band SAR images using Pauli decomposition”, in *Proc. IGARSS*, Jul. 2016. [84]
- Rosen, P. A., Hensley, S., Joughin, I. R., Li, F. K., Madsen, S. N., Rodríguez, E. & Goldstein, R. M., “Synthetic aperture radar interferometry”, *IEEE Proc.*, 88(3):333–382, Mar. 2000. [13, 95]
- Sabry, R., “A new coherency formalism for change detection and phenomenology in SAR imagery: A field approach”, *IEEE GRS Letters*, 6(3):458–462, Jul. 2009. [88]
- Sandia National Laboratories, “IED detector developed by Sandia Labs being transferred to Army”, Jun. 2014, URL https://share-ng.sandia.gov/news/resources/news_releases/copperhead_army. [75]
- Sarabandi, K. & Ulaby, F. T., “A convenient technique for polarimetric calibration of single-antenna radar systems”, *IEEE TGRS*, 28(6):1022–1033, Nov. 1990. [7, 176, 179, 181]

- Sarabandi, K., Pierce, L. E., Oh, Y., Dobson, M. C., Ulaby, F. T., Freeman, A. & Dubois, P., “Cross-calibration experiment of JPL AIRSAR and truck-mounted polarimetric scatterometer”, *IEEE TGRS*, 32(5):975–985, Sep. 1994. [187]
- Sarabandi, K., Pierce, L. E., Dobson, M. C., Ulaby, F. T., Stiles, J. M., Chiu, T., Roo, R. D., Hartikka, R., Zambetti, A. & Freeman, A., “Polarimetric calibration of SIR-C using point and distributed targets”, *IEEE TGRS*, 33(4):858–866, Jul. 1995. [187]
- Sarabandi, K. & Chiu, T.-C., “Optimum corner reflectors for calibration of imaging radars”, *IEEE TAP*, 44(10):1348–1361, Oct. 1996. [185]
- Sarabandi, K. & Lin, Y.-C., “Simulation of interferometric SAR response for characterizing the scattering phase center statistics of forest canopies”, *IEEE TGRS*, 38(1):115–125, Jan. 2000. [82]
- Schmitt, M. & Zhu, X. X., “Demonstration of single-pass millimeterwave sar tomography for forest volumes”, *IEEE GRS Letters*, 13(2):202–206, Feb. 2016. [87]
- Schwäbisch, M., Gopal, S., Mercer, B., Zhang, Q. & Wei, M., “Early results using single-pass L-band PolInSAR”, in *Proc. IGARSS*, Jul. 2008. [77, 87, 105]
- Schwäbisch, M., Mercer, B., Zhang, Q. & Huang, W., “Accurate focusing of single-pass airborne InSAR data at L-band”, in *Proc. IGARSS*, Jul. 2010. [11]
- Seipel, M. S., “Long wavelength radar; its purpose and applications”, in *Proc. SPIE Aerial Reconnaissance Systems Conf.*, vol. 0079, Sep. 1976. [78]
- Seymour, M. S. & Cumming, I. G., “Maximum likelihood estimation for SAR interferometry”, in *Proc. IGARSS*, Aug. 1994, (corrected version). [13, 59, 60]
- Shapiro, I. I., Zisk, S. H., Rogers, A. E. E., Slade, M. A. & Thompson, T. W., “Lunar topography: Global determination by radar”, *Science*, 178(4064):939–948, Dec. 1972. [13, 95]
- Sheen, D. R., Malinas, N. P., Kletzli, D. W., Lewis, T. B. & Roman, J. F., “Foliage transmission measurements using a ground-based ultrawide band 300-1300 mhz SAR system”, *IEEE TGRS*, 32(1):118–130, Jan. 1994. [79, 82]
- Sheen, D. R., VandenBerg, N. L., Shackman, S. J., Wiseman, D. L., Elenbogen, L. P. & Rawson, R. F., “P-3 Ultra-Wideband SAR: Description and examples”, *IEEE A & E Systems Magazine*, 11(11):25–30, Nov. 1996. [79]
- Sheen, D. R., Freeman, A. & Kasischke, E. S., “Phase calibration of polarimetric radar images”, *IEEE TGRS*, 27(6):719–731, Nov. 1989. [7, 176]
- Sheen, D. R., Johansen, E. L., Elenbogen, L. P. & Kasischke, E. S., “The gridded trihedral: A new polarimetric SAR calibration reflector”, *IEEE TGRS*, 30(6):1149–1153, Nov. 1992. [32, 177, 184]
- Silk, J. D. & Hunter, W., “Pixel level change detection for foliage penetrating SAR: Registration studies and sensitivity analyses”, in *Proc. IEEE Radar Conf.*, Apr. 1999. [83]
- Sletten, M. A. & Brozena, J., “Fopen target detection via joint space/angle variation”, *IEEE GRS Letters*, 13(6):762–766, Jun. 2016. [83]
- Smith, A. R., “Color gamut transform pairs”, in *Proc. SIGGRAPH*, 1978. [201]
- Soumekh, M., “Moving target detection in foliage using along track monopulse synthetic aperture radar imaging”, *IEEE TIP*, 6(8):1148–1163, Aug. 1997. [1, 84]
- Stacy, N. J. S. & Burgess, M. P., “*Ingara*: the Australian airborne imaging radar system”, in *Proc. IGARSS*, Aug. 1994. [8]
- Stacy, N. J. S., Badger, D. P., Goh, A. S., Preiss, M. & Williams, M. L., “The DSTO *Ingara* airborne X-band SAR polarimetric upgrade: First results”, in *Proc. IGARSS*, Jul. 2003. [8, 175]

- Sullivan, R., Nichols, A., Rawson, R., Haney, C., Darreff, F. & Schanne, J., “Polarimetric X/L/C-band SAR”, in *Proc. IEEE Radar Conf.*, Apr. 1988. [79]
- Sullivan, R., “Synthetic aperture radar”, in Skolnik, M. I., ed., *Radar Handbook*, 3rd edn, McGraw-Hill, 2008. [21, 66]
- Tabb, M., Orrey, J., Flynn, T. & Carande, R., “Phase diversity: A decomposition for vegetation parameter estimation using polarimetric SAR interferometry”, in *Proc. EuSAR Conf.*, 2002. [117]
- Tebaldini, S., “Algebraic synthesis of forest scenarios from multibaseline PolInSAR data”, *IEEE TGRS*, 47(12):4132–4142, Dec. 2009. [86, 204]
- Tebaldini, S. & Guarnieri, A. M., “On the role of phase stability in SAR multibaseline applications”, *IEEE TGRS*, 48(7):2953–2966, Jul. 2010. [84]
- Tebaldini, S., “Single and multipolarimetric SAR tomography of forested areas: A parametric approach”, *IEEE TGRS*, 48(5):2375–2387, May 2010. [86]
- Tebaldini, S. & Rocca, F., “Multibaseline polarimetric SAR tomography of a boreal forest at P- and L-bands”, *IEEE TGRS*, 50(1):232–246, Jan. 2012. [2, 86, 100]
- Tebaldini, S., Rocca, F., d’Alessandro, M. M. & Ferro-Famil, L., “Phase calibration of airborne tomographic SAR data via phase center double localization”, *IEEE TGRS*, 54(3):1775–1792, Mar. 2016a. [84]
- Tebaldini, S., Rocca, F., d’Alessandro, M. M. & Ferro-Famil, L., “Phase calibration of tomographic SAR data using volumetric natural targets”, in *Proc. EuSAR Conf.*, Jun. 2016b. [84]
- Thirion, L., Colin, E. & Dahon, C., “Capabilities of a forest coherent scattering model applied to radiometry, interferometry, and polarimetry at P- and L-band”, *IEEE TGRS*, 44(4):849–862, Apr. 2006. [82]
- Thirion-Lefevre, L. & Colin-Koeniguer, E., “Investigating attenuation, scattering phase center, and total height using simulated interferometric SAR images of forested areas”, *IEEE TGRS*, 45(10):3172–3179, Oct. 2007. [82]
- Torgrimsson, J., Dammert, P., Hellsten, H. & Ulander, L. M. H., “Factorized geometrical autofocus for synthetic aperture radar processing”, *IEEE TGRS*, 52(10):6674–6687, Oct. 2014. [12]
- Toups, M. F., Ayasli, S. & Fleischman, J. G., “Part II: Analysis of foliage-induced synthetic pattern distortions”, *IEEE TAES*, 32(1):145–155, Jan. 1996. [82, 83]
- Touzi, R., Livingstone, C. E., Lafontaine, J. & Lukowski, T. I., “Consideration of antenna gain and phase patterns for calibration of polarimetric SAR data”, *IEEE TGRS*, 31(6):1132–1145, Nov. 1993. [176, 179, 180]
- Touzi, R. & Lopes, A., “Statistics of the stokes parameters and of the complex coherence parameters in one-look and multilook speckle fields”, *IEEE TGRS*, 34(2):519–531, Mar. 1996. [1, 13, 59, 60]
- Touzi, R., Lopes, A., Bruniquel, J. & Vachon, P. W., “Coherence estimation for SAR imagery”, *IEEE TGRS*, 37(1):135–149, Jan. 1999. [13, 55, 60]
- Touzi, R. & Shimada, M., “Polarimetric PALSAR calibration”, *IEEE TGRS*, 47(12):3951–3959, Dec. 2009. [175]
- Trefethen, L. N. & Bau, D., III, *Numerical linear algebra*, SIAM, 1997. [141, 142]
- Treuhaft, R. N., Madsen, S. N., Moghaddam, M. & van Zyl, J. J., “Vegetation characteristics and underlying topography from interferometric radar”, *Radio Science*, 31(6):1449–1485, Nov. 1996. [3, 53, 81, 82, 128, 132]
- Treuhaft, R. N. & Cloude, S. R., “The structure of oriented vegetation from polarimetric interferometry”, *IEEE TGRS*, 37(5):2620–2624, Sep. 1999. [81]

- Treuhaft, R. N. & Siqueira, P. R., “Vertical structure of vegetated land surfaces from interferometric and polarimetric radar”, *Radio Science*, 35(1):141–177, Jan. 2000. [3, 53, 81, 128]
- Treuhaft, R. N. & Siqueira, P. R., “The calculated performance of forest structure and biomass estimates from interferometric radar”, *Waves in Random Media*, 14(2):345–358, 2004. [145]
- Tsang, L., Kong, J. A. & Ding, K.-H., *Scattering of electromagnetic waves: Theories and applications*, John Wiley & Sons, 2000. [24]
- Tsunoda, S. I., Pace, F., Stence, J., Woodring, M., Hensley, W. H., Doerry, A. W. & Walker, B. C., “Lynx: A high-resolution synthetic aperture radar”, in *Proc. IEEE Aerospace Conf.*, Mar. 2000. [2, 75, 87]
- Ulaby, F. T., Held, D., Dobson, M. C., McDonald, K. C. & Senior, T. B. A., “Relating polarization phase difference of SAR signals to scene properties”, *IEEE TGRS*, 25(1):83–92, Jan. 1987. [81]
- Ulaby, F. T. & Dobson, M. C., *Handbook of radar scattering statistics for terrain*, Artech House, 1989. [92]
- Ulaby, F. T., Whitt, M. W. & Dobson, M. C., “Measuring the propagation properties of a forest canopy using a polarimetric scatterometer”, *IEEE TAP*, 38(2):251–258, Feb. 1990a. [79]
- Ulaby, F. T., Sarabandi, K., McDonald, K., Whitt, M. & Dobson, M. C., “Michigan microwave canopy scattering model”, *Intl J. Remote Sensing*, 11(7):1223–1253, Jul. 1990b. [82]
- Ulaby, F. T. & Long, D. G., *Microwave radar and radiometric remote sensing*, University of Michigan Press, 2014. [89, 91, 166]
- Ulander, L. M. H., Frörlind, P.-O., Gustavsson, A., Hellsten, H. & Larsson, B., “Detection of concealed ground targets in CARABAS SAR images using change detection”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 3721, Apr. 1999. [83]
- Ulander, L. M. H., Hellsten, H. & Stenström, G., “Synthetic-aperture radar processing using fast factorized back-projection”, *IEEE TAES*, 39(3):760–776, Jul. 2003. [4, 12]
- Ulander, L. M. H., Pierson, W., Lundberg, M., Follo, P., Frörlind, P.-O. & Gustavsson, A., “Performance of VHF-band SAR change detection for wide-area surveillance of concealed ground targets”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 5427, Apr. 2004. [83]
- Ulander, L. M. H. & Martin, T., “Bistatic ultra-wideband SAR for imaging of ground targets under foliage”, in *Proc. IEEE Radar Conf.*, May 2005. [84]
- Ulander, L. M. H., Lundberg, M., Pierson, W. & Gustavsson, A., “Change detection for low-frequency SAR ground surveillance”, *IEE Proc. - Radar, Sonar & Navig.*, 152(6):413–420, Dec. 2005. [83, 87]
- Ulander, L. M. H. & Lundberg, M., “Modeling of change detection in VHF- and UHF-band SAR”, in *Proc. EuSAR Conf.*, Jun. 2008. [87]
- Ulander, L. M. H., Barmettler, A., Flood, B., Frörlind, P.-O., Gustavsson, A., Jonsson, T., Meier, E., Rasmusson, J. & Stenström, G., “Signal-to-clutter ratio enhancement in bistatic VHF-band SAR images of truck vehicles in forested and urban terrain”, *IET Radar, Sonar & Navig.*, 4(3):438–448, Jun. 2010. [84]
- Ulander, L. M. H., Flood, B., Frörlind, P.-O., Gustavsson, A., Jonsson, T., Larsson, B., Lundberg, M., Murdin, D. & Stenström, G., “Change detection of vehicle-sized targets in forest concealment using VHF- and UHF-band SAR”, *IEEE A & E Systems Magazine*, 26(7):30–36, Jul. 2011a. [83]

- Ulander, L. M. H., Baqué, R., Cantalloube, H., Dreuillet, P., Flood, B., Fröling, P.-O., Gustavsson, A., Jonsson, T., Larsson, B., Murdin, D., Ragnarsson, R., du Plessis, O. R. & Stenström, G., “Observations of clutter suppression in bistatic VHF/UHF-band SAR”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 8051, Apr. 2011b. [84]
- Ulander, L. M. H., Flood, B., Fröling, P.-O., Gustavsson, A., Jonsson, T., Larsson, B., Stenström, G., Ragnarsson, R., Baqué, R., du Plessis, O. R., Bonin, G., Oriot, H. & Dreuillet, P., “Bistatic VHF/UHF-band airborne SAR experiment”, in *Proc. IET Radar Systems Conf.*, Oct. 2012. [84]
- Van Trees, H. L., *Optimum array processing: Part IV of detection, estimation & modulation theory*, John Wiley & Sons, 2002. [2, 30, 85, 99, 101, 103, 145]
- Van Veen, B. D., “Adaptive convergence of linearly constrained beamformers based on the sample covariance matrix”, *IEEE TSP*, 39(6):1470–1473, Jun. 1991. [101]
- van Zyl, J. J., Zebker, H. A. & Elachi, C., “Imaging radar polarization signatures: Theory and observation”, *Radio Science*, 22(4):529–543, Jul. 1987. [177]
- van Zyl, J. J., “Unsupervised classification of scattering behavior using radar polarimetry data”, *IEEE TGRS*, 27(1):36–45, Jan. 1989. [3, 81]
- van Zyl, J. J., “Calibration of polarimetric radar images using only image parameters and trihedral corner reflector responses”, *IEEE TGRS*, 28(3):337–348, May 1990. [3, 7, 118, 176, 180, 188]
- VandenBerg, N., Sheen, D., Shackman, S. & Wiseman, D., “P-3 ultra-wideband SAR: System applications to foliage penetration”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 2757, Jun. 1996. [81]
- Vickers, R. S., Lowry, R. T. & Schmidt, A. D., “A VHF radar to make terrain elevation models through tropical jungle”, in *Proc. IEEE Radar Conf.*, Apr. 1988. [79]
- Vickers, R. S., “Design and applications of airborne radars in the VHF/UHF band”, *IEEE A & E Systems Magazine*, 17(6):26–29, Jun. 2002. [89]
- Wackerly, D. D., Mendenhall, W. & Scheaffer, R. L., *Mathematical statistics with applications*, 5th edn, Duxbury Press, 1996. [59]
- Wahl, D. E., Yocky, D. A. & Jakowatz, C. V., “An implementation of a fast back-projection image formation algorithm for spotlight-mode SAR”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 6970, Apr. 2008. [12]
- Wahl, D. E., Yocky, D. A., Jakowatz, C. V. & Simonson, K. M., “A new maximum-likelihood change estimator for two-pass SAR coherent change detection”, *IEEE TGRS*, 54(4):2460–2469, Apr. 2016. [87]
- Walker, J. L., “Range-doppler imaging of rotating objects”, *IEEE TAES*, 16(1):23–52, Jan. 1980. [11]
- Weber-Hoen, E. & Zebker, H. A., “Penetration depths inferred from interferometric volume decorrelation observed over the Greenland ice sheet”, *IEEE TGRS*, 38(6):2571–2583, Nov. 2000. [14, 81, 120, 128]
- Wheeler, K. & Hensley, S., “The GeoSAR airborne mapping system”, in *Proc. IEEE Radar Conf.*, May 2000. [78, 81]
- Williams, M. L., *PolSARproSim: A coherent, polarimetric SAR simulation of forests for PolSARpro* [PolSARproSim_Design.pdf], ESA, Dec. 2006. [6, 78, 80, 82, 157, 158, 159, 160, 203]
- Willis, N. J., “Bistatic radar”, in Skolnik, M. I., ed., *Radar Handbook*, 3rd edn, McGraw-Hill, 2008. [76]
- Winter, E. M., Schlangen, M. J. & Hendrickson, C. R., “Comparisons of target detection in clutter using data from the 1993 FOPEN experiments”, in *Proc. SPIE Alg. SAR Im. Conf.*, vol. 2230, Apr. 1994. [1, 82]

- Xiao, S., Munson, D. C., Basu, S. & Bresler, Y., “An $N^2 \log N$ back-projection algorithm for SAR image formation”, in *Proc. Sig., Sys. & Comp. Asilomar Conf.*, pp. 3–7, Oct. 2000. [12]
- Yitayew, T. G., Ferro-Famil, L., Eltoft, T. & Tebaldini, S., “Tomographic imaging of fjord ice using a very high resolution ground-based SAR system”, *IEEE TGRS*, 55(2):0–1, Feb. 2017. [23]
- Yueh, S. H., Kong, J. A., Barnes, R. M. & Shin, R. T., “Calibration of polarimetric radars using in-scene reflectors”, *J. Electromagnetic Waves & Applications*, 4(1):27–48, Jan. 1990. [3, 176]
- Zebker, H. A., van Zyl, J. J. & Held, D. N., “Imaging radar polarimetry from wave synthesis”, *J. Geophysical Research*, 92(B1):683–701, Jan. 1987. [3, 175]
- Zebker, H. A. & Villasenor, J., “Decorrelation in interferometric radar echoes”, *IEEE TGRS*, 30(5):950–959, Sep. 1992. [13, 14, 81, 90, 91, 119, 120]
- Zhang, L., Li, H.-L., Qiao, Z.-J., Xing, M.-D. & Bao, Z., “Integrating autofocus techniques with fast factorized back-projection for high-resolution spotlight SAR imaging”, *IEEE GRS Letters*, 10(6):1394–1398, Nov. 2013. [12]
- Zhang, Q., Mercer, J. B. & Cloude, S. R., “Forest height estimation from INDREX-II L-band polarimetric InSAR data”, in *Proc. ISPRS Congress*, Jul. 2008. [78]
- Zhang, Q., Huang, Y., Schwäbisch, M., Mercer, B. & Wei, M., “Forest height estimation using single-pass dual-baseline L-band PolInSAR data”, in *Proc. IGARSS*, Jul. 2012. [77, 87, 106]