Performance factors for airborne short-dwell squinted radar sensors

By:

Gavin Beard

A thesis presented to the University of London for the degree of Doctor of Engineering.

Communications Engineering Doctorate Centre

Department of Electronic and Electrical Engineering

University College London

September 2010

Abstract

Millimetre-wave radar in a missile seeker for the engagement of ground targets allows all-weather, day and night, surface imaging and has the ability to detect, classify and geolocate objects at long ranges. The use of a seeker allows intelligent target selection and removes inaccuracies in the target position. The selection of the correct target against a cluttered background in radar imagery is a challenging problem, which is further constrained by the seeker's hardware and flight-path. This thesis examines how to make better use of the components of radar imagery that support target selection.

Image formation for a squinted radar seeker is described, followed by an approach to automatic target recognition. Size and shape information is considered using a model-matching approach that is not reliant on extensive databases of templates, but a limited set of shape-only templates to reject clutter objects. The effects of radar sensitivity on size measurements are then explored to understand seeker operation in poor weather.

Size measures cannot easily be used for moving targets, where the target signature is distorted and displaced. The ability to detect, segment and measure vehicle dimensions and velocity from the shadows of moving targets is tested using real and simulated data.

The choice of polarisation can affect the quality of measurements and the ability to reject clutter. Data from three different radars is examined to help to understand the performance using linear and circular polarisations.

For sensors operating at shorter ranges, the application of elevation monopulse to include target height as a discriminant is tested, showing good potential on simulated data.

The combination of these studies offers an insight into the performance factors that influence the design and processing of a radar seeker. The use of shadow imagery on short-dwell radar seeker imagery is an area offering particular promise.

Declaration

I, Gavin Spencer Beard, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed:

Date:

Acknowledgements

I would like to thank many of my colleagues for their assistance when I was undertaking the projects that formed this thesis, and during the process of completion under pressing timescales.

Thanks go to Dr Adrian Britton for providing guidance, useful suggestions, and asking many awkward, but valuable, questions. I would also like to thank Dr Jonathan Collier for his support in the final stages of this work and Patrick Beasley, who has provided information and answers to my questions on radar hardware. I would like to thank John Whitehead and Robert Gibson for their advice, feedback and contrasting review styles. Many thanks to Andrew May, now of Dstl, for answering questions over the years, and for many interesting discussions.

I would like to thank Professor Chris Baker, now of the Australian National University, for accepting me as an EngD student, and I am indebted to Professor Hugh Griffiths of UCL, who accepted me as a student at a late stage in my EngD, and has provided much encouragement and support.

I would like to thank the MoD for funding this work and allowing publication, as well as the Engineering and Physical Sciences Research Council (EPSRC) for sponsorship.

Finally, thanks must go to Naomi for putting up with my work, and to my parents for their support and for never forgetting to ask when this would be complete.

List of contents

List of abbreviations

ADC	Analogue to Digital Converter
ADU	Air Defence Unit
APC	Armoured Personnel Carrier
ATR	Automatic Target Recognition
BBC	British Broadcasting Corporation
CAD	Computer-Aided Design
CDAA	Context Data Aided Acquisition
CEP	Circular Error Probable
CFAR	Constant False Alarm Rate
CIF	Composite Image Formation
CLTG	Closed-Loop Terminal Guidance
CNR	Clutter to Noise Ratio
COTS	Commercial Off The Shelf
CW	Continuous Wave
DBS	Doppler Beam Sharpening
DDS	Direct Digital Synthesis
DMS	Dual Mode Seeker
DRO	Dielectric Resonant Oscillator
DTED	Digital Terrain Elevation Data
EM	Electromagnetic
EO	Electro-Optic
EPSRC	Engineering and Physical Sciences Research Council
ESR	Enhanced Surveillance Radar
FFT	Fast Fourier Transform
FIBUA	Fighting In Built-Up Areas
FM	Frequency Modulation
FMCW	Frequency Modulated Continuous Wave
FOV	Field of View
FPGA	Field Programmable Gate Array
GPS	Global Positioning System
GUI	Graphical User Interface
HRR	High Range Resolution
IC	Integrated Circuit
INS	Inertial Navigation System
IR	Infra-Red
ISAR	Inverse Synthetic Aperture Radar

JDAM	Joint Direct Attack Munition
LDA	Linear Discriminant Analysis
MBT	Main Battle Tank
MEMPHIS	Multi-Frequency Experimental Monopulse High-resolution Interferometric SAR
MMW	Millimetre Wave
MRMS	Multi-Role Modular Seeker
MTI	Moving Target Indication
NATO	North Atlantic Treaty Organization
NSDP	Normalised Standard Deviation of Power
PCA	Principal Component Analysis
P_d	Probability of detection
P_{fa}	Probability of false alarm
PRF	Pulse Repetition Frequency
PSF	Point Spread Function
RCS	Radar Cross Section
RDM	Range-Doppler Map
RFOD	Residual False Object Density
ROC	Receiver Operating Characteristic
ROE	Rules Of Engagement
SADARM	Search And Destroy ARMor
SAR	Synthetic Aperture Radar
SFM	Sensor Fused Munition
SFW	Sensor Fused Weapon
SNR	Signal to Noise Ratio
SPG	Self-Propelled Gun
SPTA	Salisbury Plain Training Area
STALO	Stable Local Oscillator
TCR	Target to Clutter Ratio
TEL	Transporter Erector Launcher
TRE	Telecommunications Research Establishment
UCL	University College London
U.S.	United States
UAV	Unmanned Aerial Vehicle
UHF	Ultra High Frequency
USA	United States of America
VHF	Very High Frequency
WRF	Waveform Repetition Frequency
WW	World War

List of symbols

Symbol	Meaning	Units
A	Rain attenuation	dB/km
A_C	Resolution cell area	m^2
A_E	Antenna effective area	m^2
B	Radar bandwidth	Hz
b(θ)	Smear factor	Unitless
c	Speed of light	ms^{-1}
Corr	Correlation score	Unitless
d	Antenna diameter	m
D	Difference signal	V
d	Radar footprint width	m
d_G	Radar footprint ground extent	m
F	Receiver noise figure	Unitless
f_d	Doppler frequency	Hz
f_r	Received frequency	Hz
G	Antenna gain	Unitless
h	Platform height	m
H	Object height	m
K	Boltzmann constant	J/K
k	Monopulse slope	V/V/degree
L	Radar loss factor	Unitless
L_T	Object length	m
n	Number of pulses	Unitless
$\emptyset$	Depression angle	degrees
P_A	Reflected power per unit area at antenna	W
P_C	Clutter power	W
P_I	Intercepted power	W
P_N	Noise power	W
P_R	Received power	W
P_S	Signal power (from target)	W
P_T	Transmitted power	W
R	Range	m
R_G	Ground range	m
R_{MAX}	Maximum range	m
S	Sum signal	V

τ	Pulse duration	s
T	Dwell time	s
T_0	System temperature	K
τ_L	Extended pulse length	s
v	Platform velocity	ms^{-1}
V_Z	Vertical component of velocity	ms^{-1}
W	SFM rotation rate	Hz
W_{CR}	Object cross-range extent	m
W_T	Object width	m
x	Position	m or pixels
y	Position	m or pixels
Δ_{CR}	Cross-range resolution	m
Δf	Change in frequency	Hz
Δf_{cell}	Change in frequency over one range cell	Hz
Δf_{trip}	Change in frequency over two-way path	Hz
Δ_R	Range resolution	m
Δt	Change in time	s
η	Backscatter coefficient	m^2/m^3
θ	Squint angle	degrees
θ_B	Antenna beamwidth	degrees
λ	Wavelength	m
ρ	Rainfall rate	mm/hr
σ_0	RCS per unit area	m^2/m^2
σ_C	Clutter RCS within cell	m^2/m^2
σ_T	Target radar cross section	m^2/m^2
$\dot{\theta}$	Antenna scan rate	degrees/s
$\dot{\theta}_{radar}$	Effective scan rate	degrees/s
$\delta\theta$	Monopulse measurement accuracy	degrees

List of Figures

List of Tables

1 Introduction

1.1 Missile Sensors

The focus of the work contained in this thesis is on the performance of radar sensors in missiles. While missiles in their simplest form have been used to engage targets since pre-historic times, it is only in the last century that active sensors were added to improve their capabilities.

Fuses for the detonation of artillery shells, based on time elapsed from firing, have been in existence since around 1400, and percussion fuses were first used in 1650, but did not become reliable until the copper percussion cap in 1818.

The first active sensor fitted to an artillery shell was developed by the British in the 1930s, a photo-electric fuse for use in anti-aircraft rockets. In World War II radar fuses were used, with detonation based on strength of returned signal.

While rockets had been used in China since the 12th century, and by the British in the 18th and early 19th century, conventional guns were favoured until the 1920s as they were more accurate. The creation of the first efficient liquid-fuelled rocket motor by Robert Goddard in 1926 started a rapid period of development in the US, Germany and Russia. World War II saw the use of the V-2 rocket and the V-1 flying bomb, which used a pulse-jet engine, and was an early example of a modern cruise missile. Both used inertial guidance systems to reach a target area and detonated on impact, with no form of terminal guidance to engage a specific target or proximity fuse to explode above the ground.

Following WW II, rocketry developed for large applications such as space programs and ballistic missiles, with their current battlefield use being mainly limited to helicopters and light aircraft for ground attack.

Advances in sensor technology from the 1950s focused primarily on the development of air-to-air and ground-to-air missiles using Infra-Red (IR) and radar seekers. Millimetre-wave (MMW) radar seekers offer superior adverse weather performance compared with IR or electro-optic (EO) technologies for precision guided weapons.

The MMW seeker forms part of the missile guidance and control system and must be compatible with the power available on the missile, the aerodynamic capabilities of the airframe, and not cause any excessive impediment to the operation of the warhead.

The requirement for a seeker is to address any uncertainty in the missile launch point, its flight path and the target location, where it would be impossible to achieve the desired accuracy without a seeker [1]. Artillery-launched weapons may also benefit from use of a radar seeker to refine the target location. These weapons are known as sensor-fused munitions (SFMs) as the seeker is used to trigger the warhead.

The use of active radar for air-to-surface engagements has lagged behind its use in anti-air operation or anti-ship applications. Semi-active radar seekers tend to be used for anti-air applications, such as the AIM-7 Sparrow [2] and related BAE Skyflash [3], with some fully fire and forget active radar seekers as used on the AIM-54 Phoenix [4] and AIM-120 AMRAAM [5]. Numerous anti-ship missiles use fully active radar such as the MBDA Exocet missile [6], Boeing Harpoon [7] and Russian KH-35 [8].

Active air-to-surface radar seekers have to overcome the technical challenge of achieving a high probability of target acquisition whist maintaining a very low probability of mistaking discrete clutter objects, such as buildings, as valid targets.

There are a limited number of air-to-surface missiles which use an active radar seeker. Two that are known to be in-service are the MBDA Brimstone missile, which has a 94 GHz MMW seeker [9] [10] and which entered service in 2005, and the Lockheed Martin AGM-114L Longbow Hellfire missile [11]. Both of these offer fire and forget capability, allowing the launch platform to relocate after launch. Numerous research programmes have worked on the development of anti-surface radar seekers for ground attack, but the challenges of target selection in a cluttered environment and the risk of collateral damage, when compared to anti-air and anti-ship missiles, place challenging constraints on the required performance.

Radar has also been used in SFMs such as SADARM and SMArt [12] to provide target selectivity over a limited search area.

Global Positioning System (GPS) guidance may be suitable for some targets and mission objectives. The use of GPS information allows very precise positioning of a weapon regardless of weather conditions and visibility of the target. The Boeing Joint Direct Attack Munition (JDAM) [13] is a kit which converts unguided bombs into all-weather smart munitions, guided by an inertial navigation system (INS) coupled with a GPS receiver. This system has achieved an accuracy of less than 10 m Circular Error Probable (CEP) [14], far better than could be expected from

unguided munitions. However, while appropriate for some fixed targets, relying solely on GPS guidance is not suitable for cases where targets may have relocated in the time between the last update to the weapon (if that capability is implemented) and the time of arrival. Any case where the target may have moved would require a search of a localised area during the terminal engagement and an intelligent selection of the correct target given knowledge of its original position. An approach has been tested to guide a JDAM munition onto a moving target [15]. However, it was resource intensive to provide a high update rate on the target's position.

Consideration must also be given to the possibility of a GPS system being jammed [16]. Localised jamming near a target's location would require the weapon to be guided solely on the INS, generating an increasing level of error with the time travelled under jammed conditions. While there is research being undertaken on anti-jamming technologies, reliance on GPS alone does not appear to be suitable for engagement of a broad target set.

For longer-range missiles, designed to engage moving or relocatable targets, a wide area search is required, and to provide the capability to engage a broad target set with a low probability of selecting a collateral object the sensor must be capable of producing high-resolution imagery. To avoid ambiguity, high / fine resolution indicates resolution on a point target of typically of 0.5 m or less in the down-range and / or the cross-range directions. A new generation of missiles with high-resolution imaging radars, potentially combined with other sensors and with data links for communication during flight, are proposed as a solution.

1.2   Millimetre Wave (MMW) radar for missile seekers

Millimetre-Wave (MMW) radar in a missile seeker for the engagement of ground targets allows all-weather, day and night, surface imaging and has the ability to detect, classify and geolocate objects at long stand-off ranges.

Autonomous air-to-ground missiles or Unmanned Aerial Vehicle (UAV) sensors are technically challenged by the difficulties in adequately searching a large area for targets, correct target selection and potential engagement with a low probability of misclassification. First generation MMW seeker sensors which fill this gap [9] use real-beam radars, operating with a high range-resolution waveform to image targets albeit with low cross-range resolution. The radar used primarily in this work operates with an imaging waveform, and the application of synthetic aperture radar (SAR) allows the production of data with high resolution in both the down and

cross-range directions. This form of high-resolution imagery makes possible the exploitation of techniques developed for radar surveillance and targeting purposes, which can be of benefit in developing radar seeker Automatic Target Recognition (ATR) processes.

Key to the acceptance of a MMW seeker in operational use is the verification of a capability to engage specific targets in a scene following the provision of information from a targeting sensor. To fulfil this target selection requirement, the MoD has funded a package of seeker research, jointly developing the necessary hardware and algorithms.

This thesis presents work which addresses aspects of the algorithmic side of the target acquisition problem, examining and developing techniques to extract information on objects in imagery to improve target selection.

1.3   MMW Doppler Beam Sharpening (DBS) and SAR

Early imaging radars used Doppler beam sharpening, a form of unfocused SAR, to produce imagery with greater detail than could be achieved with a real-beam radar. Rather than illuminating the ground directly in front of the seeker, by squinting the beam away from the velocity vector, the difference in Doppler shift across the illuminated area on the ground provides information that can be used to create imagery with cross-range resolution finer than the beamwidth of the radar.

The radars considered in this report are a further development of the early DBS radars, with the inclusion of image focusing to operate a close to the velocity vector, highly-squinted when compared to normal sideways-looking synthetic aperture radars with fine cross-range resolution. The techniques used are detailed in chapter 3, and the experimental systems used to gather data are described in chapter 4.

1.4   Analysis of imagery and Automatic Target Recognition

In order to understand the imagery produced using synthetic aperture techniques, and to quantify what information may be obtained from that imagery, a number of tools are used. An outline of these tools, and of the basic ATR algorithms, is contained in chapter 5.

1.5 Motivation

The aim of this work is to improve the ability of airborne radar seekers, through the identification of the important hardware and processing factors, to allow intelligent decisions to be made for the recognition of vehicles. The seekers may be cued by a long-range surveillance sensor or other prior information.

1.6 Objectives

This work aims to develop a greater understanding of the properties of radar imagery that affect the ability of air-to-ground, highly-squinted radar sensors to detect, and provide information on, vehicle targets. This is achieved through the development of theoretical understanding as well as proof of concept tests on real data to demonstrate techniques that can assist with the target detection and classification problem.

1.7 Methodology

In this thesis we explore the factors that affect the capability of ATR and their dependencies on the capability of the radar system.

Chapter 2 provides background and context for the research, building from the early days of radar to the technologies which make ATR possible.

Chapter 3 develops analytical expressions for missile seeker radar resolution in the real-beam, and imaging domains. This develops an understanding of what can be expected from any later experimental work.

Chapter 4 describes the process of gathering real radar data to support the algorithmic work, and the radar modelling used to explore the performance space for additional radars, targets and backgrounds.

Chapter 5 outlines an approach to ATR, then investigates a robust shape-based template matching technique, followed by an investigation of how size information is affected by reduced radar sensitivity.

Chapter 6 looks at ways to use shadow information from a short-dwell squinted SAR sensor, and measures the additional information which can be gathered compared to the bright signature, primarily for the case of moving targets.

Chapter 7 considers the effect of different radar polarisations for the purposes of ATR, examining data of military vehicles in linear and circular polarisations from three different radar systems.

Chapter 8 examines the case of a real-beam sensor in the context of sensor fused munitions, and the extra information on potential targets that can be gained through the use of enhanced hardware to provide fine range-resolution and elevation monopulse.

Chapter 9 contains conclusions on the research and suggestions for future work.

1.8 Novel aspects of the work

The major original contributions from this thesis are believed to be, for highly squinted radar sensors:

- The use of shadow information from moving-vehicles to inform target size and velocity estimation where the target's bright signature is displaced and blurred. Techniques developed showed good potential for improving vehicle size and position estimation, Chapter 6, Section 6.6.9;

- For radars with a steep imaging geometry, the application and modelling of monopulse information to measure target height, and therefore improve selectivity, for radar sensor fused munitions, Chapter 8, Section 8.5.10;

- Investigations into the implications of radar hardware and processing options to identify factors which influence the performance of automatic target recognition for a range of vehicles. These include polarisation choice (Section 7.9), reduced radar sensitivity (Section 5.6) and robust template matching (Section 5.3.11).

1.9 Publications

G. S. Beard, "Multi-Spectral Target Detection Fusion", *Proceedings of the London Communications Symposium*, September 2005 (Annexe A.1).

Aspects of the work have been presented to students from the Defence College Shrivenham as "Millimetre Wave Air-Ground Seekers", given as part of the 'Defence College Guided Weapons Systems Course' Number 58, 12th November 2008, Malvern. This presentation is not publically available.

G. S. Beard, "Seeker De-Risking", Complex Weapons Centre for Defence Technology Research Conference", QinetiQ Malvern, July 2010

University College London Postgraduate Poster Sessions:

- 2005 – Radar and infra-red detection fusion

- 2006 – Radar polarisation effects on armoured vehicle detection (Chapter 7)

- 2008 – Vehicle detection by radar from a steep elevation angle (Chapter 8)

All work was originally produced as technical reports for the MoD customer community.

2 Research context

2.1 Early History

The detection and ranging of objects using electromagnetic waves has been an area of research for more than 100 years, although the term radar (radio detection and ranging) was introduced in the 1940s by the U.S. Navy [17].

There were a number of early contributors to the concept of the detection of objects using radio waves:

- Heinrich Hertz, in 1887, was the first person to demonstrate the existence of electromagnetic (EM) waves through development of a receiver that consisted of a coil and a spark gap, where a spark was seen when EM waves were received. Hertz demonstrated important principles such as reflection and polarisation, but felt there was no practical use for EM waves;

- Nikola Tesla, in 1900, published thoughts on the detection and measurement of position and speed of an object using electrical waves [18], although a practical system was not demonstrated;

- In 1904, Christian Hueslmeyer demonstrated equipment which used radio echoes to detect ships [19]. The transmitter used a spark gap fed by an induction coil, with radiated pulses beamed by a funnel-shaped reflector which could be rotated in the azimuth direction. A separate antenna fed a coherer receiver [20] which then sounded a bell if a signal was detected. Despite good results, no military or commercial interest was shown in the system and no further experimentation was carried out.

The further advancement of radar does not appear to have started until after World War I, where military research produced a number of advances:

- In June 1922, Marconi addressed the American Institute of Electrical Engineers and proposed the use of beams of electrical waves which would be reflected back to a receiver screened from the transmitter and reveal the presence and bearing of a ship in fog or bad weather [21];

- In September 1922, Taylor and Young at the U.S. Naval Aircraft Laboratory carried out Very High Frequency (VHF) radio propagation experiments, and noted the ability to detect the passage of a nearby wooden ship [19]. They

then proposed that, with suitable antennas and a more focused beam, detection of steel vessels could take place at a much longer range irrespective of darkness, fog or smoke. The use of pulsed techniques was suggested by Young, to allow determination of range, but the timing equipment to make this possible needed to be developed;

- In 1934, building on Young's work, Robert Page demonstrated a pulsed transmitter and constructed a system capable of detecting aircraft flying up and down the course of the Potomac river;

- In 1935 the Daventry Experiment in Great Britain, using a BBC short-wave radio station as the transmitter, detected the interference signal from an aircraft at ranges up to 8 miles. The success of this test led to the construction of a pulsed system, which was demonstrated at Orfordness by detecting a flying boat at ranges out to 17 miles.

This work led to the development of the Chain Home network of radar stations in Britain for the detection of aircraft, which operated at frequencies between 20 and 30 MHz, requiring very large antennas.

The development of the cavity magnetron allowed significant power levels to be generated at microwave frequencies, which then allowed the use of far smaller antennas, with improved directivity, that could be fitted to aircraft for navigation in bad weather and as a bombing aid. An example of this is the H2S system, developed at the Telecommunications Research Establishment (TRE) Malvern, which dramatically improved the effectiveness of British bomber forces in World War II following its introduction to service in 1943 [22].

2.2 2D imaging radar

Following World War II, radar had been developed for use both on the ground and in the air which allowed long-range detection of targets, and imaging of the ground for navigation and bombing purposes. Whilst these radars were capable of detection of smaller targets (ships, armoured vehicles), the lack of cross-range resolution at longer ranges limited the information on the target signature that could be used to aid recognition.

Synthetic-Aperture Radar (SAR) was invented by Carl A. Wiley in 1951. SAR is a technique to improve resolution in the cross-range direction through synthesis of a large antenna, using the perpendicular motion of the platform relative to the ground.

By recording the amplitude and phase of a signal returning from a given area of ground, the results can be combined as if they were formed at a single time from a much larger antenna.

Initial SAR systems used holographic techniques to produce SAR images. However, advances in digital signal processing and computing power allowed the radar information to be recorded and processed either directly after recording or at a later time. A thorough description of SAR is given in [23] and [24].

The use of a SAR approach to produce high-resolution imagery for missile seekers is detailed in chapter 3.4.

2.3   Radar ATR

While many radar systems have resolutions of tens of metres, which may be suitable for certain applications, with the fine resolution of a SAR it is possible to detect and locate a target, and also measure information on that target which may provide information about its class, or even specific type.

This type of recognition has traditionally been carried out by human image analysts, familiar with the nature of the imagery and the appearance of targets and clutter, which may be quite different to their visual appearance.

However, there are cases where this process may be better carried out by a computer:

- where large quantities of data are produced, a human operator is prone to fatigue, errors, and can only examine a limited quantity of data;

- if the data is produced on an unmanned platform without a loitering capability, such as a missile, and the timescales of the engagement are short, then there is no scope for human examination of the data.

These are cases where computerised ATR is required. The role of ATR is to detect objects (targets) in a background of clutter given a sensor image containing noise and other artefacts. Depending on the application, the problem may be the detection of targets against noise rather than clutter. However, for ground-imaging radar it is nearly always the case that, for stationary targets, the target is against a clutter background. The definitions of target and clutter depend on the case, with the target category as loose as an armoured vehicle (recognition), or as specific as a certain type of tank (identification) [25]. Clutter refers to imaged objects that are not classified as targets. This can be broad-area clutter such as grass or tarmac, or

specific objects such as trees, buildings and vehicles. When large areas are imaged, the ability to discriminate between targets and clutter is essential to avoid the processing being swamped with false alarms, which are clutter objects declared to be targets.

Factors which improve the capability of an ATR process in a missile seeker radar to discriminate between targets and clutter, and to improve classification capability, are the focus of this work and a review of these is presented in the following chapter.

2.4   Factors influencing ATR performance

There is a limited set of radar data with which ATR decisions can be made. The data available when examining the returned signal in the case of a monostatic radar comprises:

- Amplitude

- Time

- Frequency

- Phase

- Polarisation

- Bearing

Factors which influence the collection of this information are:

- The motion of the radar platform;

- Atmospheric propagation effects between the radar and imaged area;

- Duration of imaging or repeated imaging.

When these inputs are used to produce data suitable for the application of ATR, whether real-beam or high-resolution SAR [26], a set of these factors will be available. Radar capabilities can vary significantly, from having much detailed information available to the ATR process, to little more than range and reflected power at a given bearing.

All processes are ultimately dependent on the quality with which these parameters can be measured, and the capabilities of the platform from which the measurements are taken.

The current understanding of aspects of this parameter space, relevant to the topics covered in this thesis, is outlined in the following sections.

2.4.1 Amplitude

Given a radar input for an ATR process, whether it is real-beam or SAR data, the amplitude of the returns is the crucial factor, and all subsequent information relies on sufficient amplitude above receiver noise to function. For the purposes of ATR, there are three distinct stages in the problem: detection, discrimination and recognition [27].

For detection with an ATR process, the signal from the target should be sufficient to separate it from the local noise and clutter without an excessive probability of false alarm [28]. For this, a simple locally adaptive constant false alarm rate (CFAR) process [29] can be used. This process sets the detection threshold for a specific cell based on the nearby background. The greater the level of background clutter, the higher the required threshold. The assumption of CFAR processing is that the samples used to set the threshold represent homogeneous clutter of a certain type, for example exponential, Gaussian or K-distributed. If, however, the samples contain returns from other targets or bright clutter objects, then the threshold may be set too low or too high respectively. An example of where this may happen is when the clutter samples include returns from an area of grass adjoining an area of tarmac, where the change in background reflectivity is significant.

Given the initial detection of an object, a range of processes can be used to discriminate target signatures from clutter. A decision first needs to be made on exactly which data points form the object, and which are considered to be clutter. A CFAR process produces a binary image, and the detections from this image can be clustered using prior knowledge of expected target size to reject clusters which are too small or too large. Following clustering, a bounding-box may be applied to the cluster to select a rectangular shaped area, or the clustered pixels alone may be used. An alternative is to use bright CFAR detections (see Chapter 5.1) to seed a segmentation process on an area of the image, with the aim of identifying the outline of the object. Active contour or 'snake' models [30] have been used for this process as a more robust method of selecting the pixels belonging to an object, being less susceptible to effects that may affect a cluster based solely on CFAR returns.

Having separated the returns that are considered to come from potential target objects rather than the background, feature-based techniques can be used to perform recognition or identification by mapping data of an object, such as its statistical features, into a declaration of identity [31]. Depending on the quality of the measurements available, this stage can be used as a target / clutter discriminator, or to perform recognition / identification on the target.

A simple approach to this process is through simple constraints on a number of feature measures to decide whether a target falls into a specific class, or through a more optimised process that reduces the number of dimensions in the classification features.

Principal components analysis (PCA) [32] is a commonly used technique for feature extraction and dimensional reduction. The aim of PCA is to use linear combinations of the original variables to derive new and uncorrelated variables called principal components.

Linear discriminant analysis (LDA), finds a linear combination of features which best separate two or more classes of objects. The approach taken by Fisher [33] to separate two classes (targets and clutter in our case) was to find a projection onto a 1D line that optimally separated the two classes.

The approach used in this work, described in section 5.1, focuses on a simple threshold approach for each feature, since the ability to measure target attributes is of interest rather than optimal separation of two classes. Whilst alternative thresholding schemes have been considered, such as the use of a Fisher Classifier or Quadratic Classifier [34], the consideration of different classification approaches is not of primary importance in this work.

Both PCA and LDA require knowledge of the likely target characteristics as well as those of the clutter. A simple thresholding approach requires only target characteristics to permit detection, but is not optimised for the rejection of clutter.

Model or template-matching approaches rely on direct comparison of the target's signature in the imagery with a database of sample signatures. This database needs to encompass sufficient examples of the target's signature through a full rotation in azimuth and at the elevation angles at which the sensor may operate. The generation of these templates requires extensive data gathering or simulation, and can be very sensitive to exact target configuration, with different configurations of vehicles of the same type potentially producing very different signatures. Novak

[35] presents an evaluation of template-based classification against a broad set of vehicles, and showed that target configuration variability can decrease inter-class separability and degrade performance. Additional templates are suggested to counteract this. However, generation of additional templates would be both computationally expensive and require increased data gathering / simulations.

Whilst a range of approaches are suitable for analysis of a target's bright returns, the information available from the analysis of shadows in radar imagery can provide important target information for ATR. Shadow information has also been used in the area of synthetic aperture sonar to aid the detection of mines [36]. Jahangir, Blacknell, Moate and Hill [37] demonstrate the segmentation of shadow information for stationary and moving targets as well as for buildings. From this information, the height and shape of the objects casting the shadows can be inferred and, for moving targets, multi-frame processing allows estimation of velocity as well as detection where the bright signature has blurred to such an extent that it is no longer detectable. Jahangir and Rollason [38] demonstrate a 'track-before-detect' process for shadows of moving vehicles given multi-frame imagery, which improves detectability for reduced shadow contrast.

Sparr, Hansen, Callow and Green [39] propose an approach for re-focusing SAR imagery to enhance the clarity of shadows and reduce smearing effects during the dwell. Such an approach would aid extraction of information from shadows at the expense of increased processing time to generate the refocused image. This approach shows greatest benefit for focusing the shadows from tall objects and when the image has been formed over a wide synthetic aperture, hence a high resolution. For shorter targets, and for shadows of moving targets, the approach would not be of benefit.

2.4.2 Polarisation

The choice of polarisation can be a critical factor in radar ATR, due both to the different scattering characteristics from targets and clutter as polarisation changes, and also to the fundamental effects it can have on propagation and target detectability. In radars that can measure different polarisations, the different scattering characteristics at each polarisation can be used as additional information in target discrimination and recognition. Extensive multi-polarisation studies have been performed on targets and clutter [40], from which it can be seen that the addition of certain polarisation channels can aid in target / clutter discrimination.

The choice of a single polarisation is also important for system performance [41], to optimise the strength of scattering from targets, to minimise that from clutter and to mitigate the effects of weather. At the frequencies used in radar seekers, atmospheric attenuation and backscatter due to rain can seriously degrade performance. Circular polarisation has been shown to be of benefit for rain-clutter rejection, since a circularly polarised wave incident on a sphere is reflected as a circularly polarised wave of the opposite rotation and can be rejected by the antenna [42]. However, raindrops, especially at high rain rates, are not perfect spheres, and the distorted shape favours the use of an elliptical polarisation, the ellipticity ratio of which is dependent on the rain rate.

2.4.3 Use of target measurements in a seeker system

The quality of data gathered, and the ability to separate target and clutter classes, gives a radar seeker in a missile the capability to engage a broader set of problems than has previously been possible.

Rather than rely solely on the ATR in the seeker to select a specific target, a challenging task given strict rules of engagement (ROE) [43], a missile with a radar seeker is likely to be fired with prior knowledge of target location and a requirement to engage a specific target. For early radar seekers, this could have been through the pilot knowing the target's approximate location, or having visually identified the target before firing [44]. For engagements where the missile is fired from a greater distance, perhaps due to hostile air-defence units, identification of a target from the launch platform may not be possible. The target may also relocate during the fly-out. This relocation would then require a large search area, increasing the probability of a clutter object being declared a target.

Prior knowledge can be used to improve the probability of the radar seeker acquiring a specified target through fusion of the radar measurements with those from another source. Salmond [45] proposes an approach which applies a Bayesian fusion process between prior information and the attributes of targets of interest in the scenario and measurements from an observer (in this case the radar seeker) to predict the posterior probability of selecting the correct target.

Such approaches are robust to bulk position changes in the scene (due to geolocation errors in the measurements) and to the relocation of specific vehicles in the scene. This approach to target selection does not necessitate the challenges of detailed templates for a large range of vehicles, but rather aims to combine the

information available in an optimal manner. Thus, a radar seeker could be used to engage a target whose selection on the basis of ATR alone would not be guaranteed with sufficient certainty to permit a weapon launch.

2.5 Summary

Some of the important factors involved in radar seeker ATR have been described, along with an approach to target selection, using prior information, which offers the ability to engage a wide target set under tight ROE. The work contained in this thesis on radar seeker performance aims to reduce the number of false targets passed to a fusion process and enhance the quality of attribute measurements to improve the probability of correct target selection. This is achieved through studies into target shape, seeker sensitivity, target shadows, polarisation and the use of monopulse to measure height information.

3 MMW SAR image formation and analysis

3.1 Introduction

This chapter formulates the theory of MMW SAR image formation and properties, as used for data from the airborne radar systems such as that subsequently described in chapter 4. The detectability of a target is defined, followed by the processes for achieving fine resolution in the range and cross-range directions, together with effects that may degrade the resolution.

3.2 Radar sensitivity

The initial requirement for a radar system is that sufficient signal level be returned such that subsequent processing can be used. The signal could be composed of the return from the ground for SAR image formation, or the returns from a target object. The derivation of the radar equation can be found in many books on the subject, for example [24].

Given a radar with a transmit power of P_T watts, with an antenna gain of G, the power per unit area, P_I incident on an object at a range R is given by

$$P_I = \frac{P_T G}{4\pi R^2}$$

(3.1)

Energy is reflected from a target proportional to its Radar Cross-Section (RCS) σ_T. If the system is monostatic, i.e. the same antenna is used for transmission and reception, the energy received is multiplied by the antenna gain, G. The received power at the radar per unit area, P_A, is

$$P_A = \frac{P_T G}{(4\pi)^2 R^4} \sigma_T$$

(3.2)

For an antenna of effective area A_E, this can be simply calculated. However the antenna gain is related to area and wavelength, λ, through

$$G = \frac{4\pi A_E}{\lambda^2}$$

(3.3)

The received power through the antenna can then be formulated as

$$P_R = \frac{P_T G^2 \lambda^2 \sigma_T}{(4\pi)^3 R^4}$$

(3.4)

In order to make use of the received energy, P_R, the signal level must exceed the thermal noise in the system by a certain margin. The receiver noise power is given by

$$P_N = KT_o BF \qquad (3.5)$$

where K is Boltzmann's constant, T_0 is the temperature of the system in kelvin, B is the receiver bandwidth in hertz and F is the receiver noise factor representing the difference in noise available from a practical and ideal receiver. T_0 is defined to be 290 K, giving a value for the factor KT of 4 x 10^{-21} W / Hz of bandwidth.

Attenuation, representing losses during two-way transmission and in the radar hardware, is represented as a loss factor, L.

The resultant Signal to Noise Ratio (SNR) is given by

$$SNR = \frac{P_R}{P_N} = \frac{P_T G^2 \lambda^2 \sigma_T}{(4\pi)^3 R^4 KT_o BFL} \qquad (3.6)$$

This equation can then be reformulated to show the maximum detection range at a given minimum SNR

$$R_{MAX} = \left(\frac{P_T G^2 \lambda^2 \sigma_T}{(4\pi)^3 KT_o BFL(SNR_{MIN})} \right)^{1/4} \qquad (3.7)$$

This is the simplest form of the equation for target detection using a monostatic radar. It assumes detection against a noise background, with attenuation during the round-trip between radar and target included in the loss term.

For a SAR-type sensor, the SNR for point targets benefits from integration gain during the formation of the image. When considering this gain, point targets and distributed targets need to be treated in a different manner.

The received power for a single pulse was given by equation (3.4). If n pulses are used to form the synthetic aperture, the noise will add incoherently, giving a gain increase by a factor n, and the echoes will add coherently, giving an gain increase of n^2. This results in an increase in the overall SNR of a factor n compared to a single pulse.

$$P_R = \frac{nP_T G^2 \lambda^2 \sigma_T}{(4\pi)^3 R^4} \qquad (3.8)$$

For distributed targets, there is equivalent gain during the processing, however the gain is compensated for by a reduction in resolution cell area, leaving no net compression gain for a distributed target.

Should the target be imaged against a clutter background, as is normally the case for an air-to-surface radar seeker, then equation (3.6) and equation (3.7) need to be modified to represent detection against clutter. Clutter can be considered as any unwanted radar returns which may come from land, sea, rain, birds or countermeasures such as chaff.

If a radar is designed for the detection of aircraft, then the presence of surface clutter in the radar returns is an unwanted signal, and clutter in the side-lobes can make detection challenging. For a ground-imaging radar, high levels of clutter are expected, but certain clutter types can make the detection of targets more challenging. However, for a SAR sensor, as will be described in chapter 3.4, ground clutter is not essential to the formation of an image, but provides contextual information in the image to accompany the returns from any targets. Clutter can be distributed spatially such that its extent is much greater than the resolution cells of the radar, or can take the form of discrete objects which may be of comparable size to a desired target class, and hence act as a possible confuser in target selection.

For surface clutter, rather than an RCS as was used for a conventional target in equation (3.2), a cross-section per unit area, σ_0, is defined as the RCS of the clutter within a resolution cell, σ_C, divided by the area of that cell, A_C.

$$\sigma_0 = \frac{\sigma_C}{A_C}$$

(3.9)

This allows use of a clutter density per unit area, hence the total backscatter in a resolution cell can be calculated. For the case where the noise level is small compared with the clutter, and can be ignored, the signal to clutter ratio (SCR) for a point target is defined as the power from the target, i.e. signal, P_S, to the power from the clutter, P_C.

$$SCR = \frac{P_S}{P_C} = \frac{\sigma_T}{A_C \sigma_0}$$

(3.10)

In this case the SCR increases as resolution is improved until the point where the resolution cell is the same size as the target, after which the target's returns will be

spread over multiple cells. The range dependency cancels out as long as the clutter level is much higher than the noise level.

The value of σ_0 is dependent on the terrain type, elevation angle, frequency and polarisation. Man-made clutter can give more intense echoes than those from rural areas due to the presence of flat-sided buildings and corner-reflectors. Rivers and lakes or hard surfaces such as roads give very little backscatter and can appear dark or at the noise level in the radar imagery. Clutter with a range of heights can also cast shadows in the scene, which may appear similar to the less reflective surfaces.

The reflectivity of clutter increases as the depression angle, Ø, changes from a low angle towards normal incidence. This can be approximated as a *1/sinØ* relation. However, clutter reflections can increase at very low angles due to the returns from tall objects in the scene that are then normal to the radar [46].

Land clutter can be represented by a range of statistical distributions including the Rayleigh, log-normal and Weibull, with different parameter choices dependent on the clutter type [47]. For high resolution imagery, a compound model such as the K-distribution [48] can be used, which contains two separate distributions, one for the RCS of a cell, i.e. the texture of the scene, and another for the speckle which is a characteristic of coherent imaging. While originally developed to represent sea clutter, this model has proven to be a good representation for homogeneous patches of clutter in synthetic aperture radar [49].

However, detection against homogeneous clutter is not a primary concern for the high-resolution radar discussed here. The targets observed in imagery are not point returns, with many target pixels observed against the background. At the CFAR stage of the ATR process, the pixel detections are sensitive to the choice of background model, but are then used to seed a segmentation routine which produces a very small segmentation if seeded on single bright clutter or speckle pixels, allowing easy rejection of that type of clutter at the feature stage. An accurate clutter model is valuable to reduce the number of CFAR false detections, and to reduce load on the subsequent processing. The type of clutter which causes a problem for the ATR process, as described in Chapter 5.1, are discrete clutter objects such as man-made structures, vehicles and occasional pieces of rural clutter. At fine resolutions, these are not easily modelled by a statistical distribution.

Also of interest are the two distinctive effects that rain can have on radar imagery; backscatter from rain drops and attenuation of the signal. These two effects are dependent on rain intensity, drop size, frequency and polarisation. The effects of rain clutter are discussed further in section 5.4.

3.3 Range resolution

As described in chapter 2, radar has typically been good at measuring range to a fine resolution, but fine cross-range resolution requires more complex radar hardware, information formation requirements, and processing.

Early radars used short pulses to reduce the distance at which two objects become separable in range, or the range resolution. For a pulse of duration τ, the range resolution, Δ_R, is equivalent to

$$\Delta_R = \frac{c\tau}{2}$$ (3.11)

In order to achieve the resolutions necessary for target recognition (ideally less than 0.5 m), a pulse length of 3 ns would required. While possible, this limits the amount of energy that can be transmitted.

Rather than using a very short pulse, a much longer pulse, τ_L can be transmitted, which is modulated in such a way as to spread the energy over a bandwidth B

$$\Delta_R = \frac{c}{2\beta}$$ (3.12)

Hence, to achieve a resolution of 0.5 m, a bandwidth of 300 MHz is required, which is practical both in terms of the necessary hardware and the carrier frequency.

The total energy of the transmitted pulse is now $P_T\tau_L$. For the short-pulse case, the transmitted energy $P_T\tau$ is equivalent to P_T/B. Hence, the SNR for a given target at a given range has increased by a factor of $B\tau_L$. This is known as the time-bandwidth product of the waveform, and can be of great benefit for the detection of targets above noise.

The bandwidth can be achieved using a number of techniques such as a linear Frequency Modulated (FM) chirp, or a stepped-frequency waveform. For the purposes of much of the work detailed in this document, the exact nature of range-compression does not affect the results.

3.4 Cross-range resolution

While the down-range resolution can be measured accurately through use of a waveform with a large bandwidth, the cross-range resolution of a non-coherent, conventional radar is limited by the beamwidth of the antenna. The finest possible beamwidth in radians, θ_B, can be calculated from Rayleigh criterion for diffraction, for a uniformly-illuminated circular antenna of diameter d and wavelength λ.

$$\theta_B = 1.22 \frac{\lambda}{d}$$

(3.13)

For a radar seeker, a beamwidth of 1° would be considered a fine beam. However, this results in a cross-range resolution of 17 m at a range of 1 km, resulting in target signatures which contain no detail in the cross-range direction, as well as combining the returns from the target with those from background clutter. Figure 3-1 shows the radar returns from a number of vehicles, imaged from an aircraft using a radar with a narrow beam and mapped to a square-pixel grid. The vehicle signatures have limited extent in the down-range direction due to the fine-range resolution of the radar, but are significantly extended in the cross-range resolution.
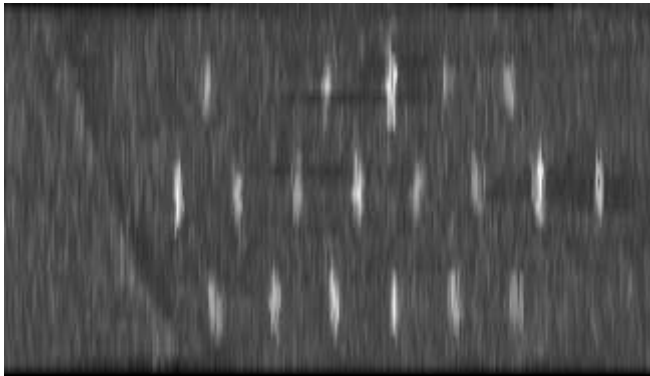
Figure 3-1 Imagery of vehicles mapped to a square-pixel grid. Down-range from left to right, cross-range in the vertical direction

The coarse cross-range resolution limits the information that can be gathered about the target from its radar returns. Given the requirement to form an image at long range, a real-beam radar may not provide adequate detail for an ATR process.

The detailed derivation of cross-range resolution using Synthetic Aperture Techniques is covered in sources such as [24]. A derivation which accounts for the forward squint of the antenna in a missile seeker is given here.

A synthetic aperture uses the sensor's motion in the cross-range direction to simulate an aperture over an extended dwell time. This synthetic aperture can be much larger than would be physically possible for a real antenna in this application. The resolution is then no longer dependent on the size of the real antenna, but is determined by how accurately the radar can measure the Doppler frequencies returned from the ground and how widely the Doppler frequencies that the imaging geometry provides are spread.

The projection of the movement of the radar to the ground plane is depicted in Figure 3-2, where v is the velocity of the platform, θ is the azimuth angle from the velocity vector, Ø is the depression angle of the radar and λ is the transmitted wavelength.
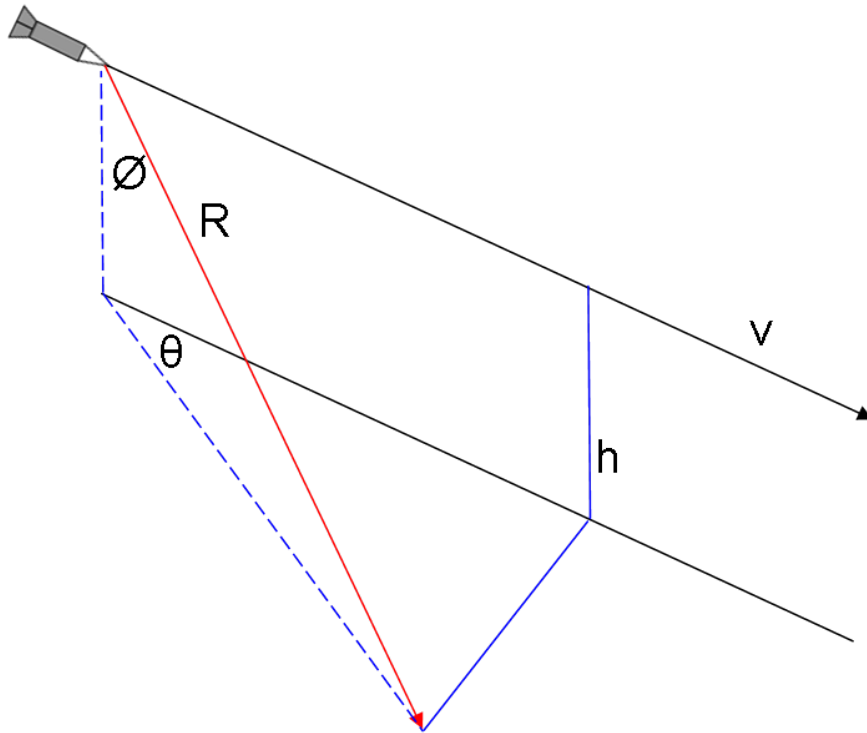
Figure 3-2 DBS Imaging Geometry

The minimum frequency difference that the radar can measure (Doppler Resolution) is given by *1/T* (Hz), where *T* is the dwell time in seconds. For this simple case, the dwell time is considered to be the duration of the aperture formation, and the point on the ground remains in the radar beam for all of that time. The Doppler frequency, f_d, returned from a point on the ground is given by

$$f_d = \frac{2 \, v \cos(\theta) \cos(\phi)}{\lambda}$$

(3.14)

The rate of change of the Doppler frequency with angle is given by

$$\frac{df_d}{d\theta} = \frac{-2\,\mathrm{v}\sin(\theta)\cos(\phi)}{\lambda}$$

(3.15)

Rearranging equation (3.15) to find d(θ), substituting d(fd) with *1/T* and multiplying by *R* to convert an angle to distance yields equation (3.16) which is the cross range resolution equation for a DBS sensor in straight and level flight.

$$\Delta_{CR} = \frac{\lambda\,R}{2\,\mathrm{v}\,\mathrm{T}\sin(\theta)\cos(\phi)}$$

(3.16)

Equation (3.16) applies only for a spotlight mode of illumination where the amplitude received by a scatterer is constant throughout the dwell. If the amplitude received from a scatterer varies throughout a dwell this amplitude taper will degrade the resolution from its theoretical value.

The next section will discuss how a finite beamwidth, coupled with a scanning beam, affects the maximum achievable resolution. This will then be developed to discuss the effects of the antenna illumination pattern. A comparison of the theory developed will be made to real radar data. Finally, the variation of resolution within the illuminated patch will be discussed.

3.5   Scanning effects

Should there be a requirement on the missile seeker to image a wide ground area during its flight, one approach to improve the area coverage, as used for the system described in Chapter 4, is to scan the antenna in the azimuth direction around the squint angle and form multiple Range-Doppler Maps (RDMs). This can allow the seeker to operate in a stripmap-like mode whilst accommodating the small size of the beam's footprint on the ground, an example of which is shown in Figure 3-3.
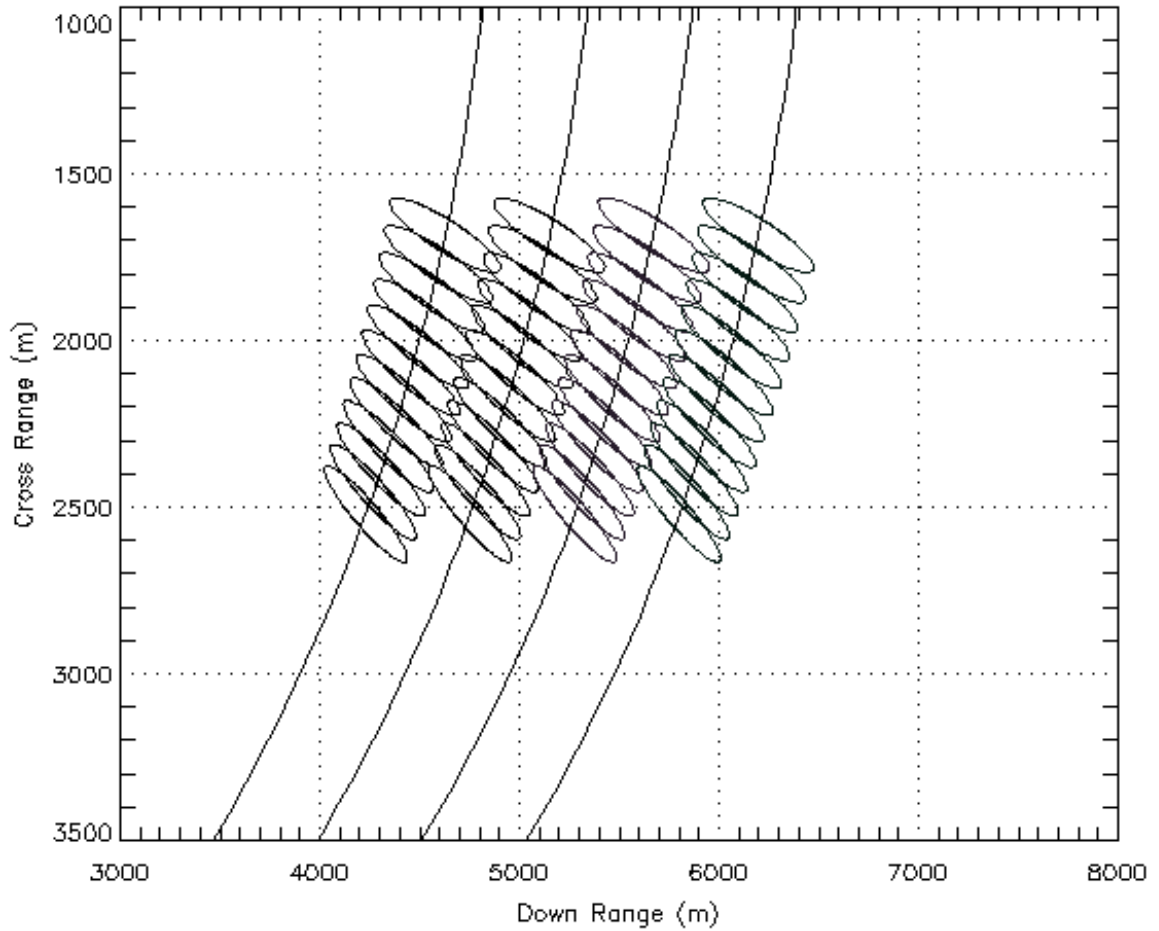
Figure 3-3 Example scan pattern showing formation of individual RDMs

The beam transmitted by a radar has a finite width, and this is usually quoted as the 3 dB width (half-power width of the 1-way beam) in terms of angle. For an initial investigation into the effects on imagery of the antenna scan, the beam can be approximated to a 'top-hat' function with width equivalent to the 3 dB width. It should be clarified that the pixel spacing in an image, and the actual resolution in the imagery, are independent values, although ideally the pixel spacing should be finer than the resolution.

It follows that, if the radar antenna is scanning across the ground, a scatterer may not be illuminated for the entirety of the dwell time. As the resolution of a scatterer is inversely proportional to the dwell time, the resolution of the scatterer will decrease. The effective time, T_{eff}, for which a scatterer is illuminated is given by equation (3.17), where θ_B is the 3 dB width of the beam, $\dot{\theta}$ is the scan rate about the vertical axis and Ø is the depression angle, which relates to the component of the rotation rate perpendicular to the beam.

$$T_{eff} = \frac{\theta_B}{\dot{\theta}\cos(\phi)} \qquad (3.17)$$

The physical scanning of the antenna is one of two factors that can reduce the time a scatterer is illuminated. For an antenna at a fixed angle, the forward motion of the platform causes an effective scan of the beam across the ground. The effective scan rate from the platform motion, $\dot{\theta}_{radar}$, is given by

$$\dot{\theta}_{radar} = \frac{v\sin(\theta)}{R} \qquad (3.18)$$

From equation (3.17) and equation (3.18), it can be seen that there is a coupling between the antenna scan rate and the platform motion which can alter T_{eff} depending on the direction of scan. If the antenna is scanning away from the platform velocity vector (positive direction), the scan counters the platform motion, prolonging the time for which a scatterer is illuminated. A special case of this is when the scan rate is equal to the effective scan rate from the platform's motion, and can result in spotlight illumination. If the antenna is scanning towards the velocity vector (negative direction) the time a scatterer is illuminated in the beam can be greatly reduced, degrading imaging resolution.

The resolution of a point scatterer in a scanning beam can be calculated by equation (3.16), where T is replaced by T_{eff} which is defined in equation (3.19) and gives the maximum time for which a scatterer is illuminated.

$$T_{eff} = \frac{\theta_B}{\left| \dot{\theta}\cos(\phi) - \frac{v\sin(\theta)}{R} \right|} \qquad (3.19)$$

By varying the scan rate and dwell time, a resolution surface can be produced as shown in Figure 3-4. Here, it is possible to see where the *1/T* theoretical resolution curve is intercepted by a linear slope determined by the scanning antenna. From Figure 3-4 it is also possible to see how the direction of the scan affects the resolution from the point where the slope intercepts the curve. To accentuate this, Figure 3-5 is a plot of how resolution changes with scan rate for a constant dwell time of 0.5 s.

Resolution surface for varying scan away from the velocity vector and dwell time
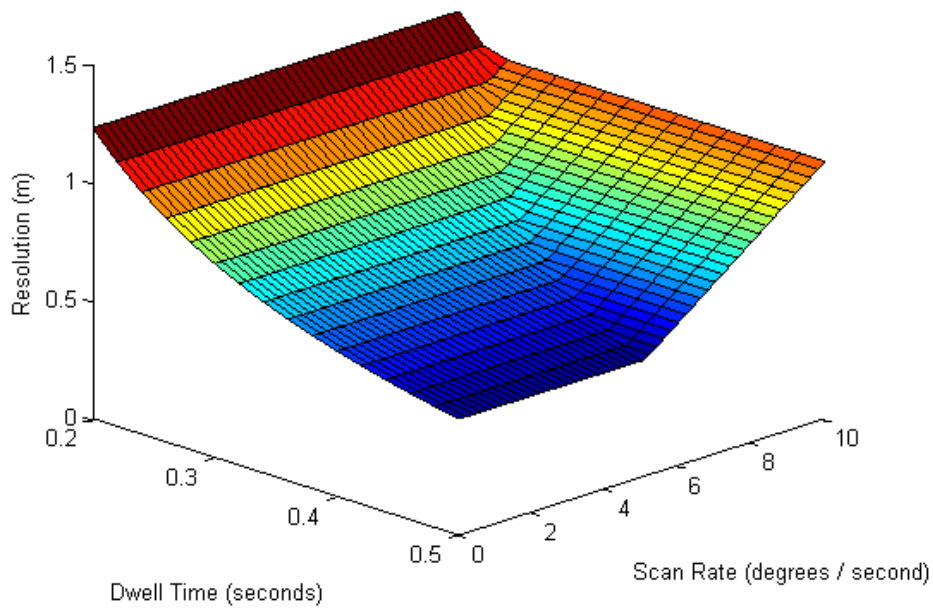
Resolution surface for varying scan towards the velocity vector and dwell time
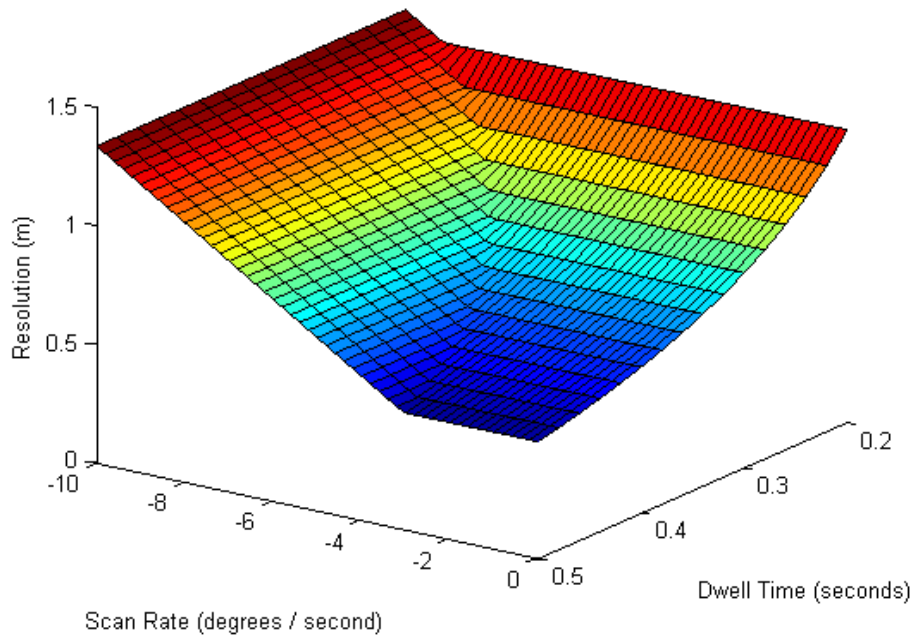
Figure 3-4 Resolution surface for varying scan rate and dwell time. This is for a 35 GHz radar with a 2^o beamwidth, a velocity of 80 ms^{-1}, looking at an azimuthal angle of 30^o from the velocity vector and at a range of 2300 m
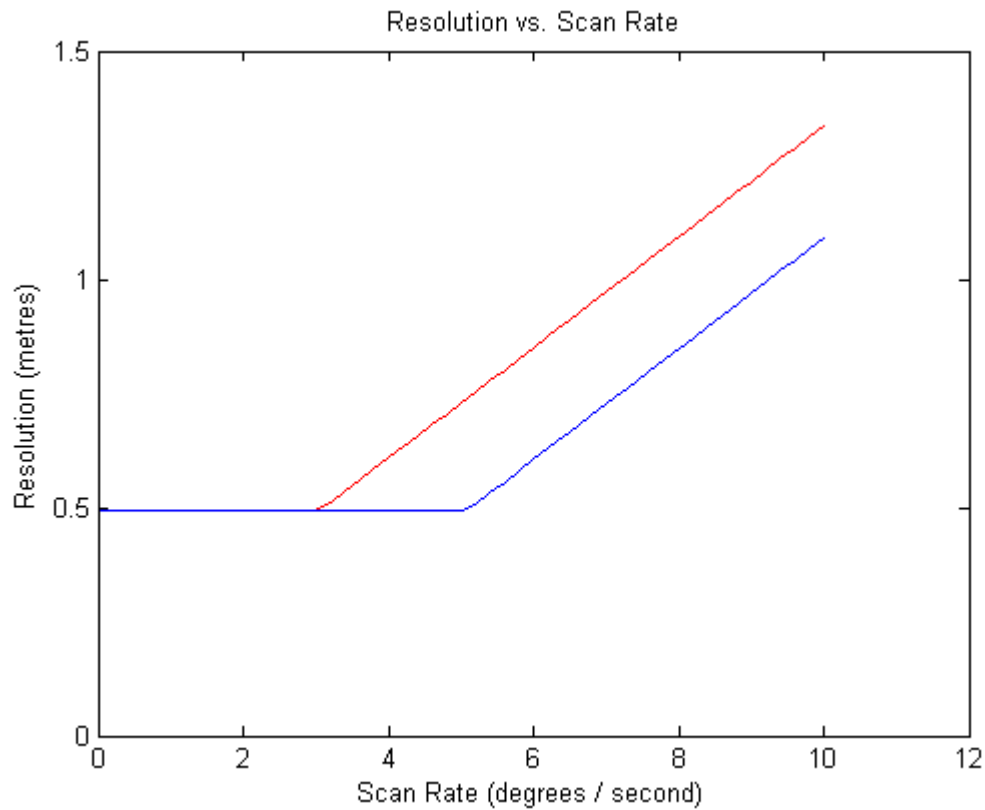
Figure 3-5 Resolution vs. Scan Rate for a constant dwell time of 0.5 s. Outboard scan in blue, inboard scan in red. 2° beam, 80 ms^{-1}, 30° squint, 10°/s scan

At low scan rates, the resolution on the outward and inward scans is only limited by the synthetic aperture length, and the resolution is the same. As the scan rate increases, the inward scan is the first to experience a reduction in the time spent imaging the target and the fine resolution degrades as the scan rate increases. As the outward scan has a reduced rate of beam movement on the ground due to the forward velocity of the platform, a higher scan rate can be reached before the reduction in imaging time on the target affects the resolution. This is similar to the difference between spotlight and stripmap SAR imagery, where the resolution limit is due to the time taken for the target to move through the beam.

The resolution of a point scatterer and the resolution of the system must be distinguished. The frequency spacing (grid spacing) of the system is always *1/T*; this is the frequency resolution of the radar and is independent of scan rate. The resolution of a scatterer is the point spread function that an ideal point scatterer has on the frequency (*1/T*) grid. By scanning the beam, the point spread function of a scatterer will change and it is this change in resolution that must be accounted for.

3.6 Antenna beam shape effects

In the previous section the radar beam was approximated to a 'top hat' function within the confines of the 3 dB width. This is a reasonable assumption. However, the beam is approximately Gaussian in shape and includes sidelobes. The following section explores how the resolution is affected by the beam shape and the scanned angle.

It is assumed that, at the centre of the dwell, the scatterer is at the centre of the beam shape profile i.e. the scatterer is at the centre of a cross-range slice. For spotlight mode (or when a very small angle is scanned through) the amplitude modulation is minimal and equation (3.16) still applies. As the angle through which the beam is scanned during the dwell increases, the amplitude variation becomes more pronounced and the shape of the beam also becomes more pronounced. When this happens, the resolution achieved will gradually decrease. When a large enough angle is scanned during a dwell, the dwell time is artificially limited as in the previous section. Figure 3-6 shows how the amplitude weighting changes when the angle through which the bean is scanned increases.
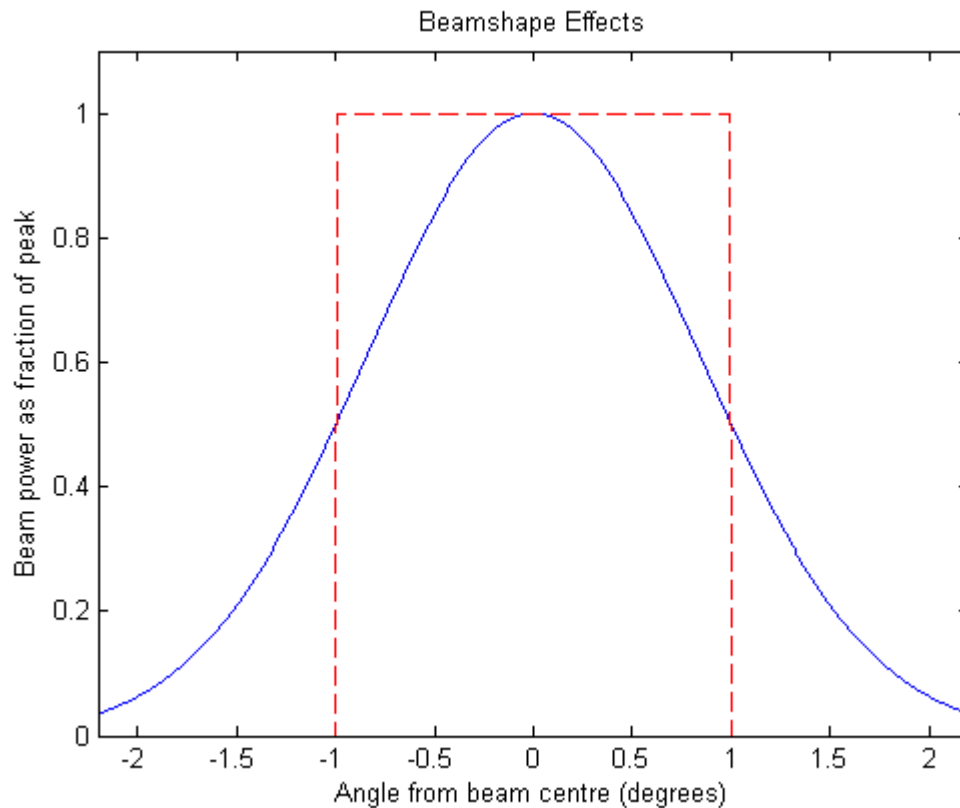
Figure 3-6 Effect of increasing scan angle on amplitude weighting

Figure 3-6 shows the beam shape assumed previously as a red dotted line. It can be seen that taking the 3 dB width as a 'top-hat' function is an approximation as there is still energy illuminating the target beyond the 3 dB width.

As the beam shape has a smooth profile, its effects on T_{eff} are not as simple to understand as with the 'top-hat'. If a small angle is scanned through (less than the 3 dB width) there will be some variation in the amplitude weighting, reducing a point scatterer's resolution. Using equation (3.19) for the same angle scanned, the resolution would not have changed from the theoretical case. By having a smooth amplitude profile, the change in resolution with the angle scanned through will also be smooth. The effect of beam shape as opposed to the 3 dB top hat function is a smoothing of the change in gradient in the plots displayed in Figure 3-5.

This effect was tested by May [50], who calculated a 'smearing factor', b(θ), representing the resolution achievable with a scanning antenna compared to that from a spotlight antenna for the same range, velocity and squint angle. This smearing factor can then be used in conjunction with equation (3.16) to estimate the degraded resolution of a scanning system, giving equation (3.20).

$$\Delta_{CR} = \frac{\lambda R}{2 \, v \, T \sin(\theta)} b(\theta)$$
(3.20)

3.7 Validation with real data

Through examination of reference (trihedral) reflectors in air-carry seeker imagery, the point spread function resolution of the reflector was measured. To do this, the data was over-sampled by a factor of eight, and the 3 dB width measured manually. The measured resolution was found to be 0.67 m, and the expected resolution was calculated to be 0.48 m. This gave a smearing factor of 1.4. For the trials scan-rate, using equation (3.19), a smear factor of 1.5 was calculated. This difference of 0.1 in the smearing factors may be accounted for by small errors in the measured platform velocity and radar look angle.

The dwell time in the seeker data processing software was varied to produce imagery with increasing levels of scan during a dwell. A reflector's point spread function (PSF) was measured and compared with the theoretically predicted resolution. In Figure 3-7, the solid red line represents the predicted smearing factor for an inward scan, the dotted blue line represents the predicted smearing factor for an outward scan and the crosses represent measured data from the radar on an outward scan.

Good agreement is seen between the predicted smear values and those measured from imagery, although a test with a larger data set would be beneficial to confirm the relation on an inward scan.
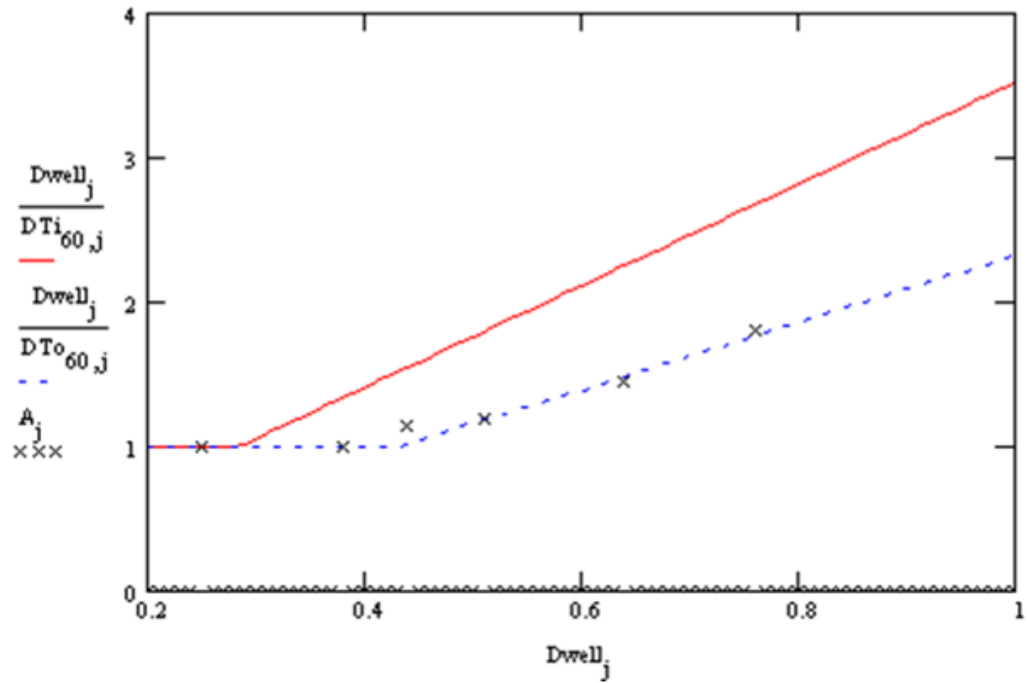
Figure 3-7 Comparison of measured and theoretical smearing factors, image courtesy of A. May

3.8   Image resolution conclusions

It has been shown that the resolution of a radar system may not necessarily be the theoretical resolution predicted by equation (3.16). The platform motion and antenna scan can reduce the dwell time which will, in turn, degrade the resolution.

A prediction of how the dwell time will be reduced for a scanning antenna on a moving platform is shown in equation (3.19), which can limit the cross-range resolution of the imagery.

To support stable ATR, it is important to understand effects that may vary the resolution and to compensate for these variations when measuring attributes of objects in the scene. Scaling errors in the imagery will directly relate to measurement errors in object dimensions and position. The resolution must remain sufficient to support the ATR techniques described in the following chapter, so careful consideration should be given to the design of a scan for searching an area of interest.

A scan of uniform angular rate in both the inboard and outboard direction has been modelled for this example, and is used by the radar system described in Chapter 4. Given the ability to select an ideal scan, a 'stop-stare' approach, operating in a similar manner to spotlight SAR with dwell-times varying with squint angle, would allow the generation of imagery with reduced resolution variability.

4 Experimental data gathering and radar simulation

4.1 Introduction

Work reported in this thesis is supported by a range of experimental data which is necessary to develop and test algorithms. This data has come from air-carry trials of static and moving vehicles (Section 4.2) and quarry trials which use vehicles on a turntable (Section 4.3). Data has also been simulated to cover a broader range of scenarios than can be achieved using real trials data. Under this work a tool was created to compare the image properties of the simulated data with real radar data, and the model modified to improve the simulation. This is addressed in Section 4.4.

4.2 Data gathering trials – Air carry

Data gathered by the experimental Multi-Role Modular Seeker (MRMS) radar was used to design and test many of the algorithms and work reported in this thesis. Of the novel aspects listed in Section 1.8, shadow information, polarisation and model matching all use data from this radar. The radar has been developed jointly between QinetiQ and MBDA [51] over several years. The radar data used was gathered during the 2006 flight trials of the MRMS radar.

The MRMS sensor was flown in a pod beneath the QinetiQ Andover trials aircraft, Figure 4-1.

Figure 4-1 MRMS Phase 2 radar mounted in pod below Andover trials aircraft

The MRMS radar has undergone two different builds during research programmes thus far. Phase 1, in 2002, allowed the cross and co-polar channels to be recorded

with a range of bandwidths to change down-range resolution, and different scan-rates to vary cross-range resolution. For Phase 2, in 2005 and 2006, a monopulse mode was added to the antenna, generating sum and difference beams in the azimuth direction. This allowed two different output modes, a polarimetric mode as before which used only the monopulse sum beam, and a mode which recorded both sum and azimuth difference beams, but only in the 'odd' circularly polarised channel (cross-polar return). These modes are shown in Figure 4-2.
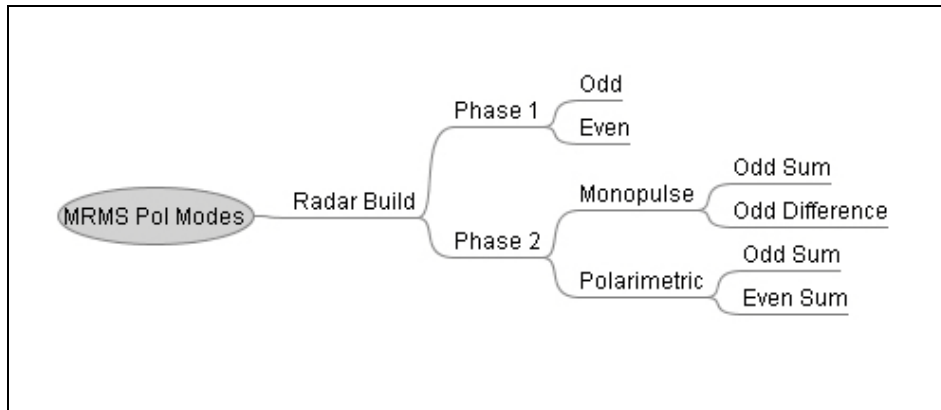
Figure 4-2 MRMS Polarisation Modes

The radar had previously been flown in 2002 (Phase 1), however, the new build standard allowed an increased dynamic range to avoid saturation effects which had been seen previously, as well as the addition of a monopulse-capable antenna. The addition of amplitude monopulse angle estimation (see Chapter 8) was a large step in the system's capability, as it allowed confirmation that use of monopulse information in radar seeker imagery was possible, providing accurate geo-location of objects and the development of monopulse based moving target indication (MTI) algorithms.

The radar operated at a nominal range of 2.3 km, a 30^{o} squint angle from the velocity vector and a 10^{o} depression angle. The imagery gathered has a cross-range resolution of 0.5 m and a down-range resolution of 0.5 m or 0.25 m depending on the choice of waveform bandwidth.

The 2006 data gathering trials flew mainly over the Copehill Plantation area of Salisbury Plain Training Area (SPTA). Here, data was gathered of stationary targets, decoy vehicles, moving targets and the Fighting In a Built Up Area (FIBUA) village to gather urban area data. Imaging runs were also performed over the Defence Academy at Shrivenham to obtain imagery of static targets in an urban area.

The static targets and decoys used were:

- Two Main Battle Tanks (MBT);
- Two Self-Propelled Guns (SPG);
- One mobile Air Defence Unit (ADU);
- Large and small Tractor;
- Land rover pickup;
- Cars;
- Decoy constructed from 4 trihedral reflectors;
- Decoy constructed from several omni-directional reflectors.

The moving targets imaged were:

- Range Rover augmented with a reflector array;
- MBT.

The target motions were at various speeds (5 – 25 kph) in directions parallel and perpendicular to the aircraft's track.

Example imagery from individual range-Doppler maps is shown in Figure 4-2 to Figure 4-7.
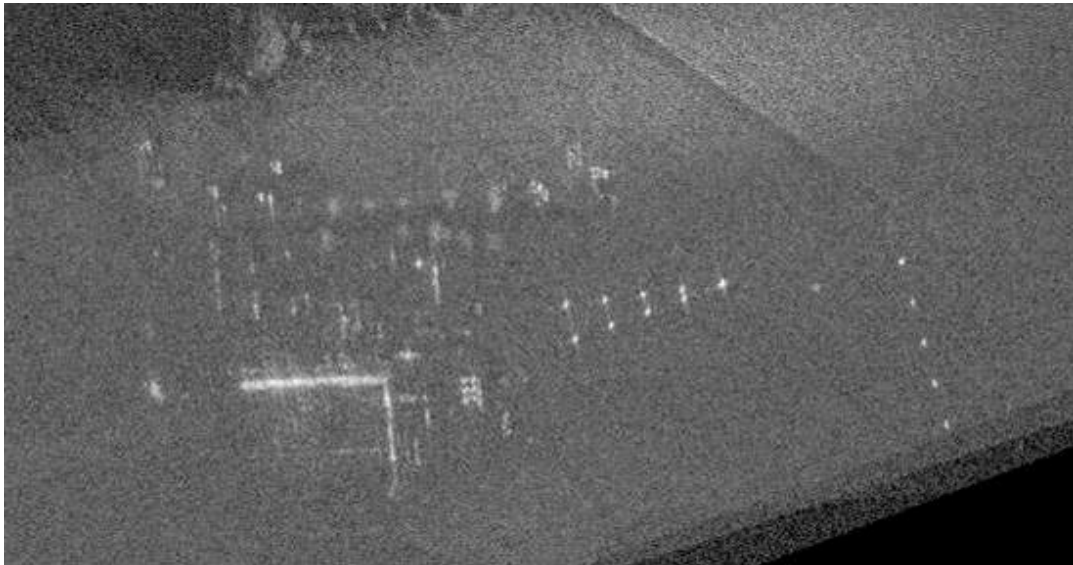
Figure 4-3 Section of RDM containing reflector arrays for measurement of monopulse performance and image resolution

Figure 4-4 Section of RDM containing buildings
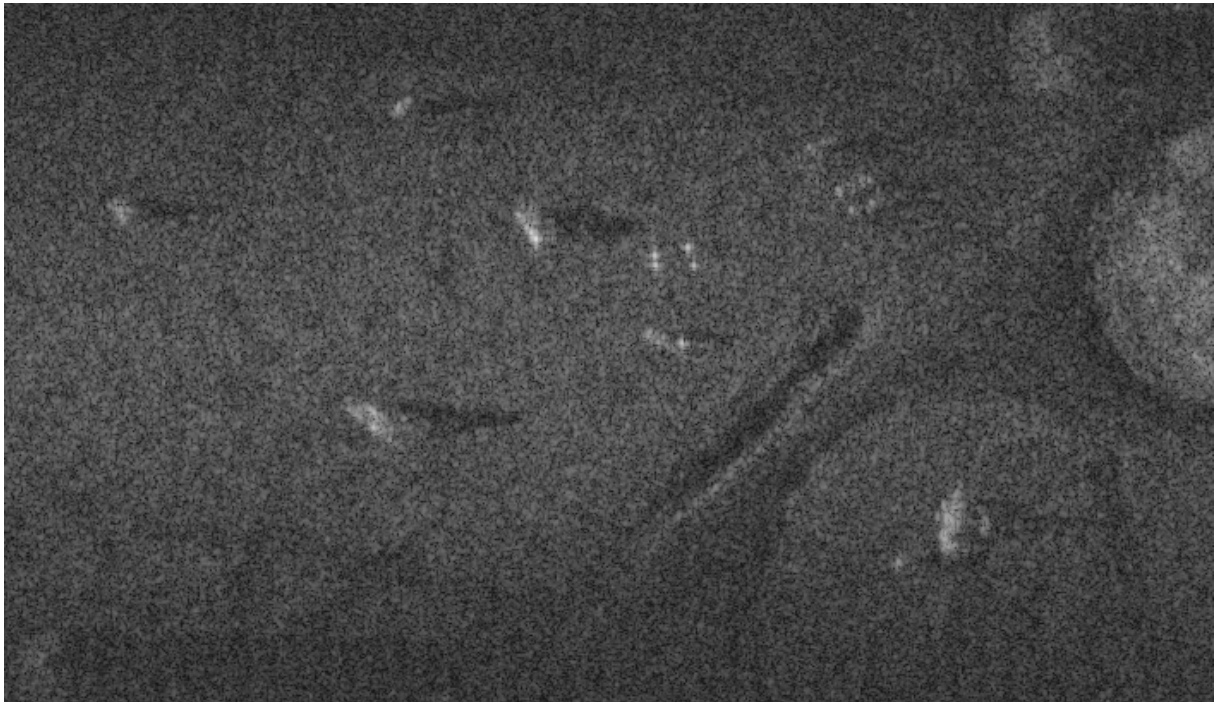
Figure 4-5 Section of RDM with rural clutter, vehicles and radar reflectors
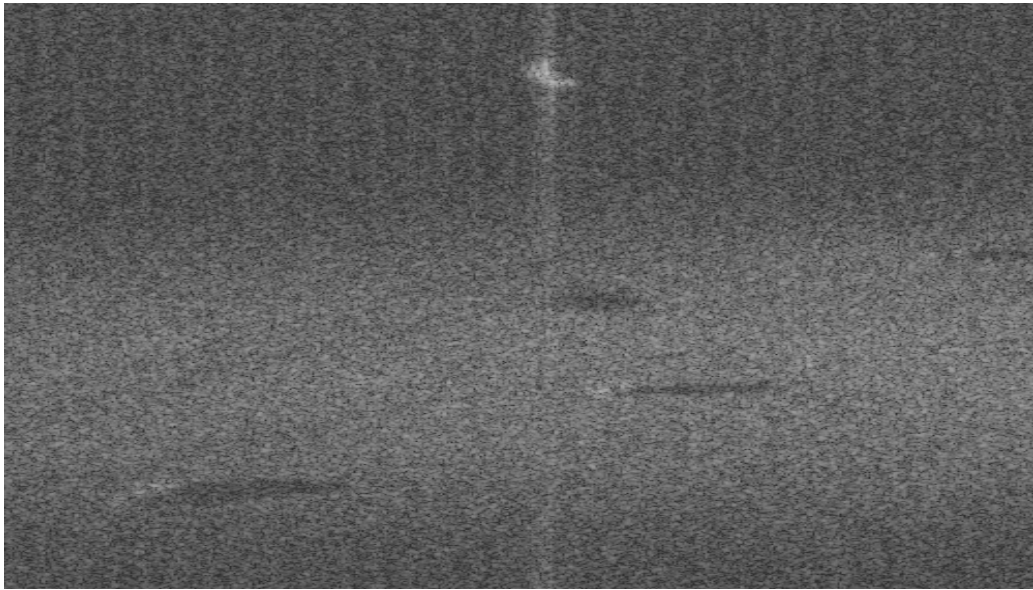
Figure 4-6 Section of RDM showing displaced moving target and its shadow, in the presence of rural clutter
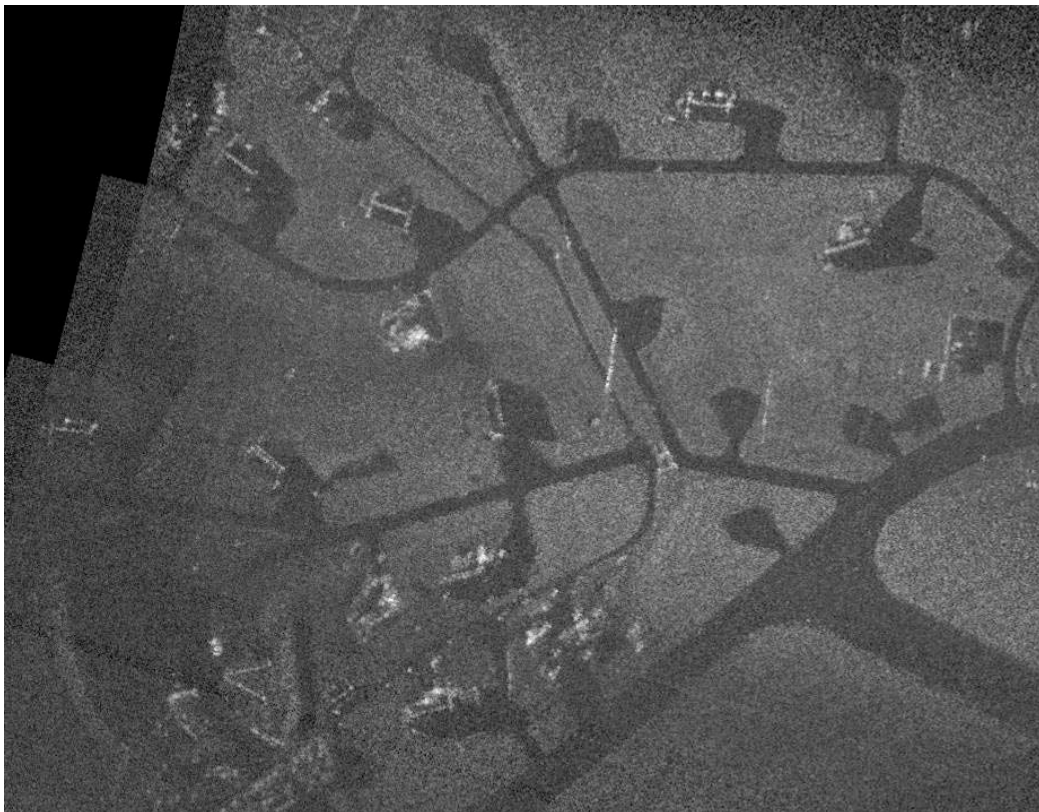
Figure 4-7 Composite image formed from multiple RDMs of the taxiways and aircraft shelters at Boscombe Down

4.3 Data gathering trials – Turntable

In addition to air-carry radar data, use has been made of Inverse Synthetic Aperture Radar (ISAR) data to provide detailed information on target signatures. ISAR relies on the rotation of a target to provide the Doppler information required to produce a fine cross-range resolution. Data from a target can be gathered for a complete rotation, and the sub-sections of the data can be processed to generate target images through a complete rotation, examples of which are shown in Figure 4-8 at 30° steps from 0° to 90°. The cross-range resolution is proportional to the angle through which the target rotates during the integration time, so the choice of integration angle can be varied to produce different image resolutions.
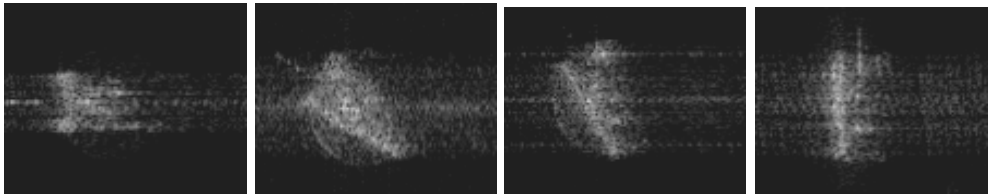
Figure 4-8 Example ISAR images of armoured vehicle at 30° angular steps

It is also possible to back-rotate ISAR imagery of a target and sum the responses to produce a detailed map of the target's scattering over a full rotation, as shown for two vehicles in Figure 4-9.
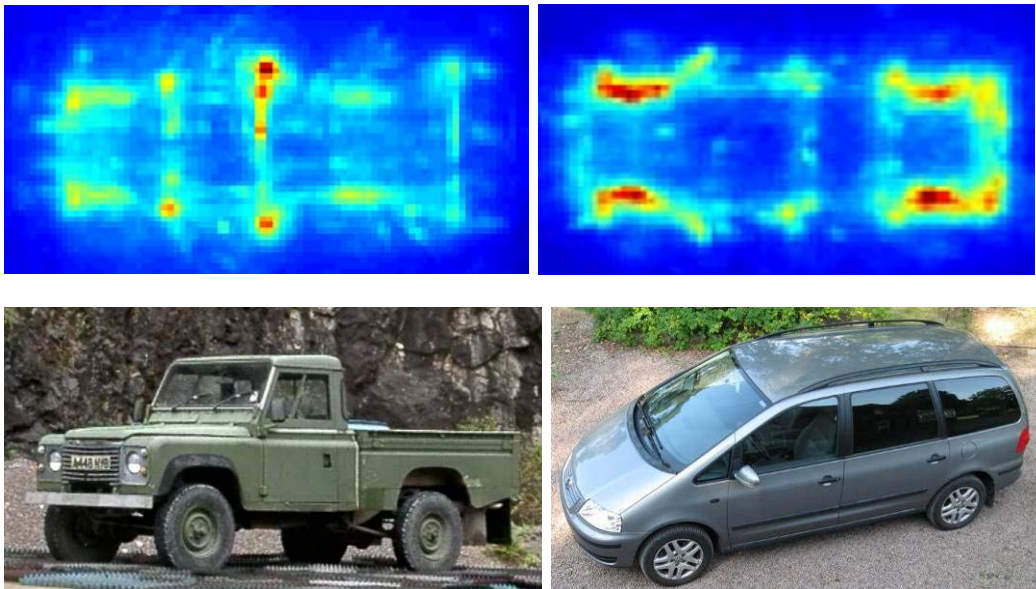
Figure 4-9 360° back-rotated ISAR images of Land Rover pickup and VW Sharan

A primary use of ISAR data has been to validate and improve the vehicle models used in radar simulations, which are further described in section 4.4.

4.4 Radar Simulation

4.4.1 Introduction

The Closed Loop Terminal Guidance (CLTG) software is a Matlab radar simulation, developed by Thales and QinetiQ, which generates synthetic millimetre wave (MMW) seeker imagery. The CLTG is capable of producing imagery of large areas, including detailed target signatures.

The CLTG can produce realistic RDM imagery from a seeker trajectory, which has been used to generate data for tests of template matching and performance at reduced sensitivity (Chapter 5), for shadow analysis (Chapter 6) and for unconventional imaging geometries (Chapter 8).

| CLTG | Air-carry trials data |
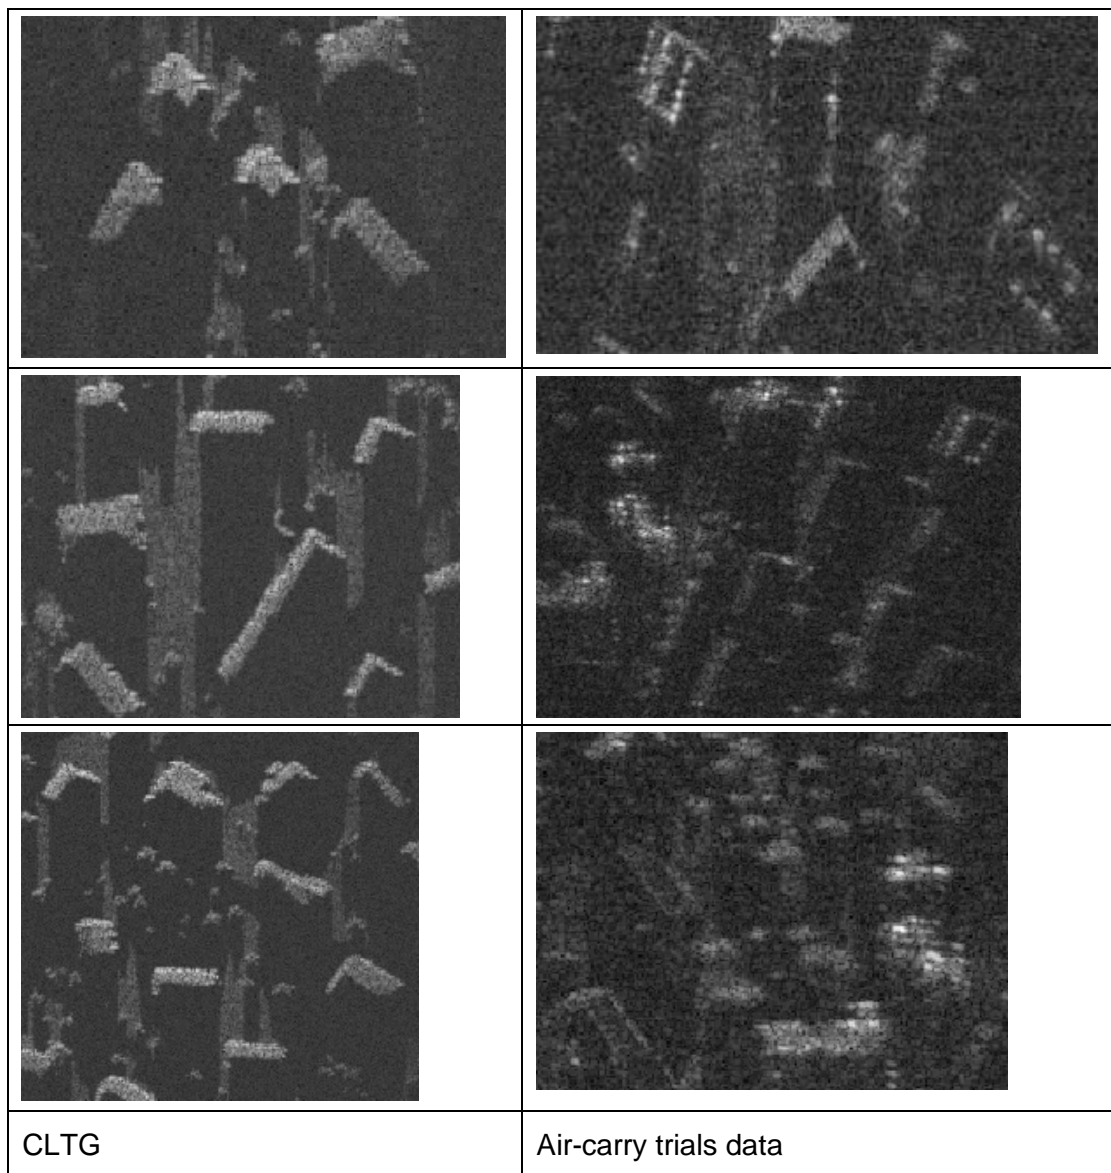
Figure 4-10 Exemplar CLTG (left) and trial radar imagery (right) of urban areas
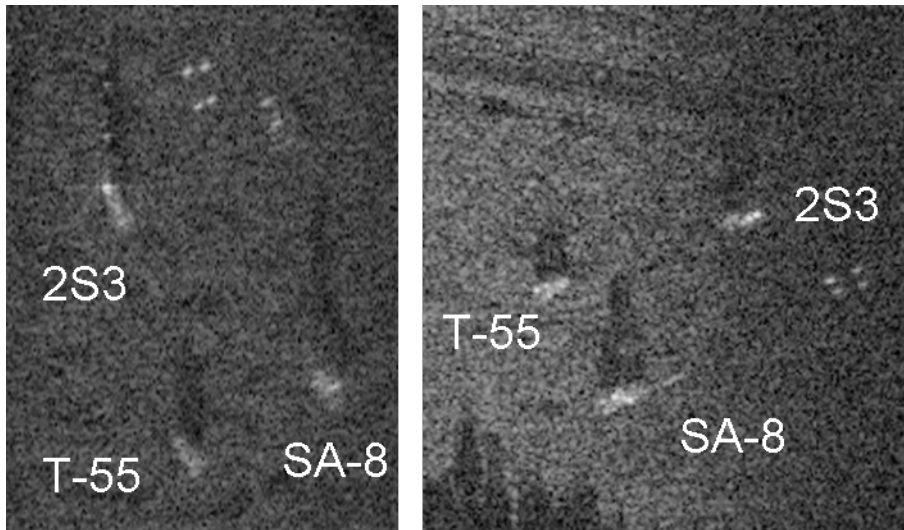
Figure 4-11 Vehicles from two air-carry trials sorties
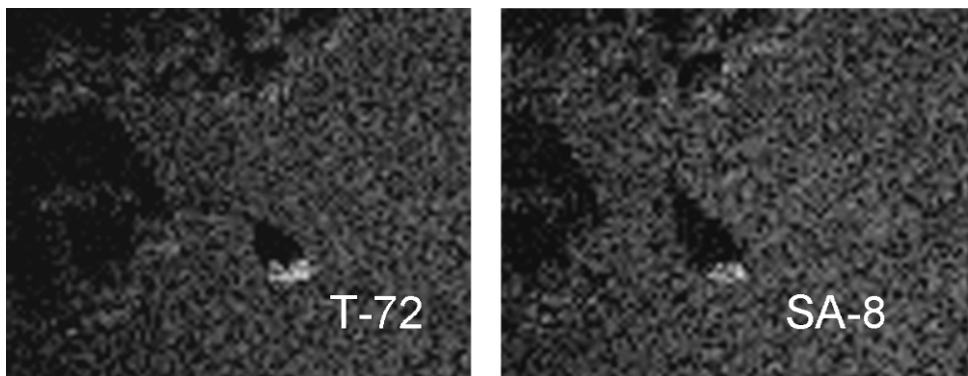
Figure 4-12 Vehicles from CLTG data, MBT (left) and Mobile ADU (right)
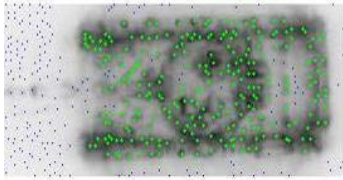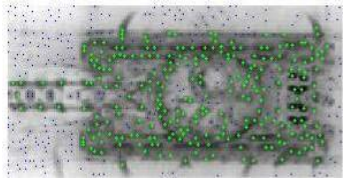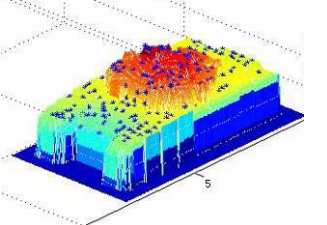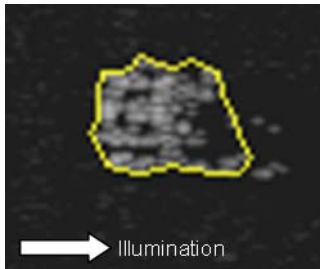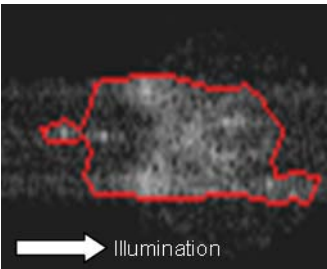
4.4.2 Modelling validation

The CLTG has been the subject of extensive data validation, resulting in adjustments to the modelling process to allow the generation of more realistic target and background signatures. The detailed validation is not reproduced here, but this section gives a brief overview of some of the work that has taken place to improve the fidelity of the modelling, and to understand its limitations.

The statistics of target objects were made more representative of trials data through the use of 'scatterpoint' models. The modelled targets consist of a large number of individual points spread over a height map, and each point has a reflectivity for each degree of a 360° rotation. The scatterpoints can be generated either from ISAR data as described in Chapter 4.3, or from a high-fidelity simulation of the target. A high-fidelity simulation of the entire scene would require an extremely detailed model of a large area and be computationally prohibitive.

A number of scatterpoint models for various targets were generated, using both real and simulated data to provide the scatterpoint information. An example of scatterpoint generated targets is shown in Figure 4-13.

Turntable MBT	Simulated MBT	Main MBT Scatterers
CLTG simulated MBT using ISAR derived scatterpoints	CLTG simulated MBT using high-fidelity modelled scatterpoints	ISAR turntable MBT imagery

Figure 4-13 A comparison of target scatterpoint distributions and example output imagery with segmentation outlines. From left: CLTG MBT with scatterpoints derived from ISAR data, CLTG MBT with scatterpoints derived from high-fidelity model, ISAR turntable image of MBT

A comparison was made between real targets and simulated targets which were generated either using scatterpoints or relying on simple backscatter calculations. A range of size-based and statistical features were considered, which showed that the use of the scatterpoint approach better represented the type of internal structure seen in high-resolution target imagery, without increasing the processing overhead of the modelling.

The simulated clutter background was also compared in detail against real air-carry trials imagery, looking at both the relative mean values of clutter areas, and the distribution of power in those areas. While the clutter in real and simulated data appeared visually similar over a broad area, the aspects of the clutter to which an ATR algorithm is sensitive needed to be compared.

A software tool was created for this work which allowed the examination of areas of RDMs and comparison of different clutter types. Figure 4-14 is an example of the

Graphical User Interface (GUI) of the tool in use, where an area of woodland and a grass area have been selected.
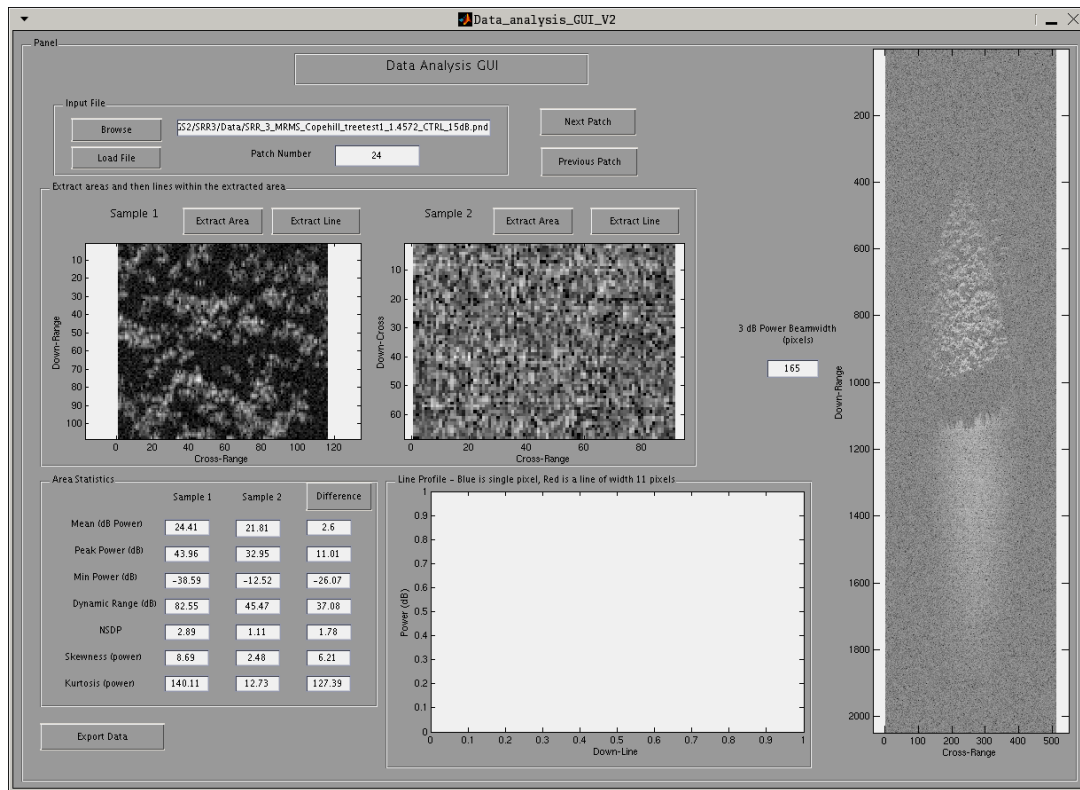
Figure 4-14 RDM analysis GUI, showing extracted woodland (left), grass (middle) and the full RDM (right). Down-range is down the page

This analysis indicated that the reflectivity levels were consistent between simulated and real imagery, but that the internal structure of simulated woodland clutter was not representative of the tree-types in the bulk of real imagery. Figure 4-15 shows woodland and grass returns from two real and one simulated RDM.

The majority of wooded areas imaged are similar to the top picture in Figure 4-15, with a spatial variation in intensity, but without many shadowed areas. The middle image shows a very structured plantation of trees, where the area around each tree is shadowed, containing only noise returns. This makes the bright tree responses more challenging to discriminate from targets as each tree can show good contrast relative to its local background. The third image shows simulated data, with similar structure to the top example, but with greater contrast due to the distinctly shadowed areas. Whereas real trees, dependent on foliage density, will have a diffuse structure, the modelled trees do not allow for that, and show solid edges. This was accounted for using a process which smoothed much of the height

variation of the tree canopy and introduced varying levels of reflectivity for the clutter which better represented the spatial structure seen in the majority of trees.
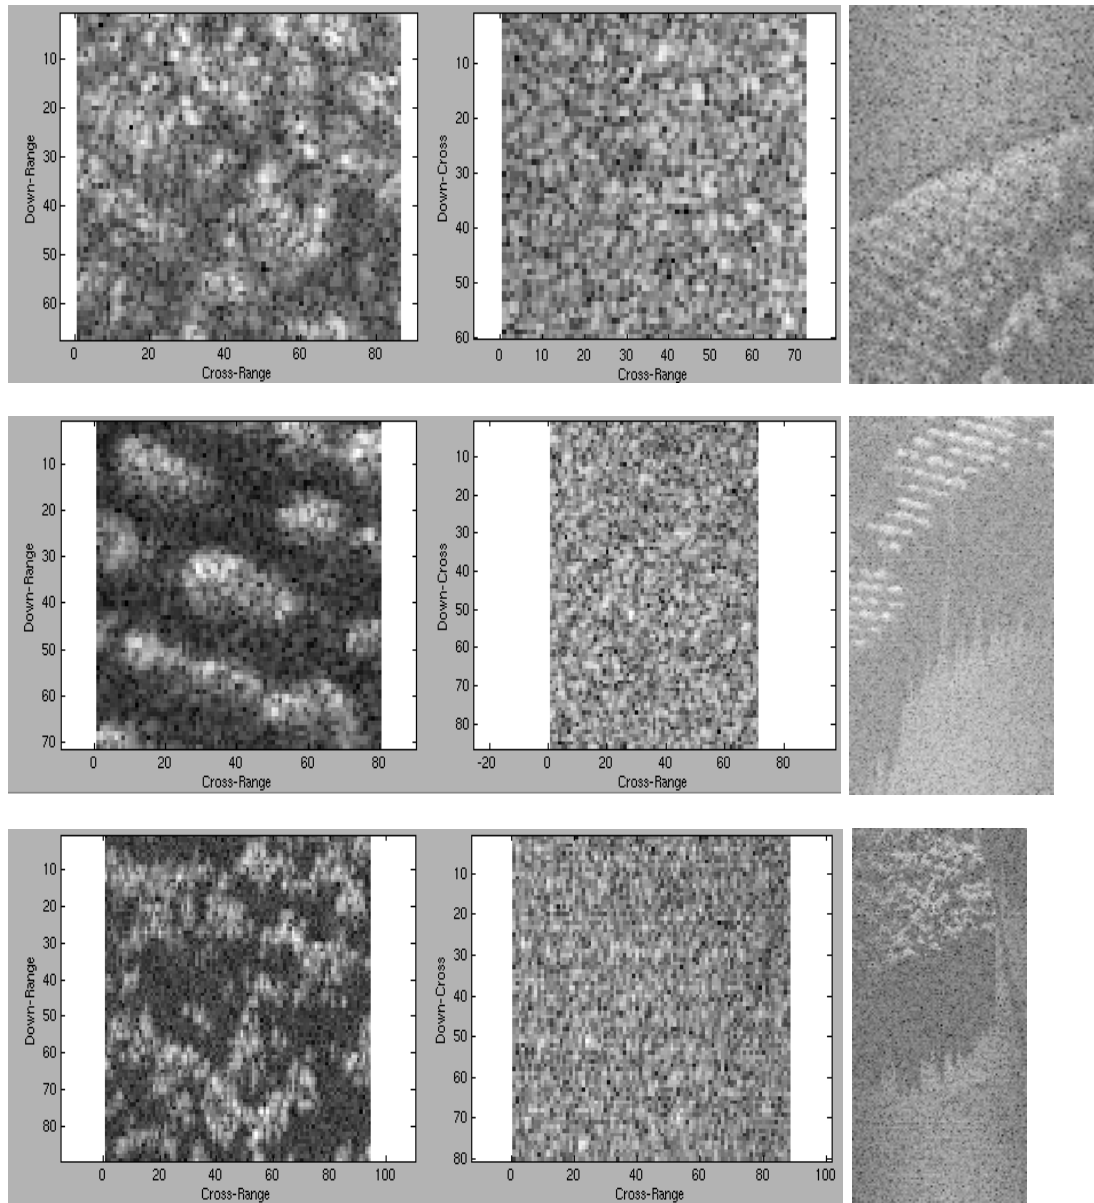
Figure 4-15 Tree clutter (left), grass clutter (middle) and extract from scene (right) for MRMS Trees (top), MRMS 'awkward' trees (middle) and CLTG trees (bottom)

Also of interest was the interface between clutter types. A strong discontinuity in the imagery can produce many detections from the CFAR stage of the ATR, and these were seen on the edges of woodland in synthetic data, shown in Figure 4-16.

The images on the left show a wooded area with a white line superimposed which highlights the region chosen to analyse the pixel returns in the down-range direction as the clutter type changes from grass to tree. The graphs on the right show the pixel relative power (in dB) for each down-range bin, with the blue plot representing

a single line of pixels, and the red plot the average of 11 pixels in the cross-range direction.

Input	Image	Profile
Original		
Smoothed internal areas		
Edge-smoothed		
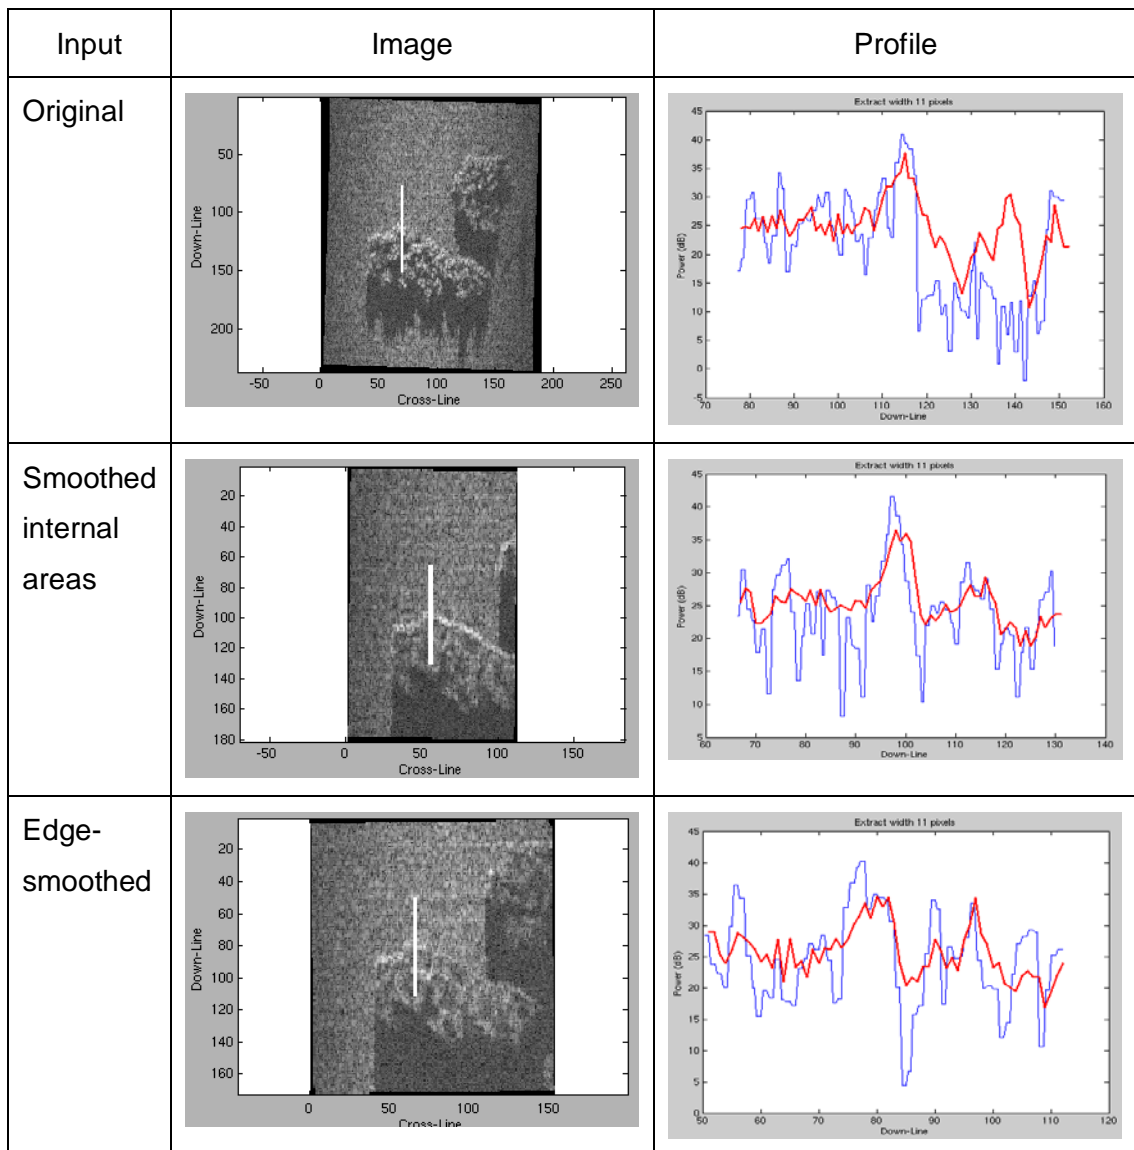
Figure 4-16 Line profiles of a wooded area for three versions of the input scene. White lines on the images show the pixels under test. The blue plot shows the returns from a single row of pixels, and the red plot the average across 11 pixels. Down-line units are measured at a sample spacing of 0.5 m.

The reflectivity of these border areas was then adjusted such that the step-change in returns was reduced to a level consistent with that seen in trials imagery, Table 4-1.

Scene Type	Grass-Tree Power Step
CLTG Original	12.2 dB
CLTG Smoothed internal areas	12.1 dB
CLTG Edge Smoothed	9.0 dB
Real air-carry	9.4 dB

Table 4-1 Average grass-tree interface change in brightness

4.5 Conclusions

Good quality air-carry SAR and ISAR quarry trials data have been gathered, and used as the basis for all the performance studies in this thesis. However, the quantity of data, and range of targets, target aspects and scene backgrounds are limited, as is the geometry from which it is gathered. The real data are therefore accompanied, where appropriate by synthetic data to expand on the cases available. CLTG data has been used alongside real trials imagery to:

- Generate target signatures over a complete rotation for model-matching studies (Chapter 5.3);

- Generate numerous target signatures for attenuation studies (Chapter 5.4);

- Generate synthetic data of moving targets to expand the data available to assess the performance of shadow information extraction (Chapter 6);

- Produce data from unconventional imaging modes to investigate potential upgrades to sensor-fused munitions (Chapter 8);

5 Object signal and shape information

Traditional radar seeker ATR has used high-resolution range profiling with a poor cross-range resolution due to the use of a small real aperture. Statistical features and measures of target length and structure from the range profiles, or a series of range profiles, have been used to provide information on objects within the radar beam. Imaging seekers, capable of producing imagery with high resolution in both down-range and cross-range directions, have much more information available for the measurement of size and statistical parameters.

This chapter examines the ability to measure the size of a range of target classes using a variety of techniques and for different levels of image quality. The sections cover:

- An investigation into the use of the shape of the returned target signature to match against a database of simple signatures generated from basic CAD models;

- Size measurement with reduced sensitivity;

- Size measurement through the use of the target's shadow.

5.1 MMW ATR Theory and Outline

Given the ability to generate high-resolution data, examples of which are given in chapter 4 where data-gathering trials are described, an ATR process has been developed to detect likely targets and extract information about their properties that could be used for target selection. This autonomous capability is necessary given the short time window between target detection and engagement, and there being insufficient time to consult a human operator unless the weapon has a loitering capability. Depending on algorithm configuration, this detection and attribute measurement software can run in real time on desktop computers. While not a focus of this work, computational efficiency is an important factor in selection of an ATR algorithm, since delays while data are processed will reduce the time available for the weapon to manoeuvre following target acquisition.

Section 2.4.1 described some background to an algorithmic approach to the detection, segmentation and attribute measurement of objects in imagery. This section details the approach developed for the majority of the work reported in this thesis.

The level of ATR which can be applied is dependent on the type of sensor and to what degree the structure of the target is resolved, and can be split into three classes [25]:

- If the target is a single point, it may be possible to rely on the target being brighter than the background, and a simple detector can be used. For most targets under consideration, their radar returns are stronger than those from the background, so that detection is possible. If target brightness could not be relied on, then a more complicated approach would need to be used which searched for motion, expected signal fluctuations, or a pattern of points representing an expected layout of vehicles;

- Should the target not be a point, and have a resolvable pattern of returns, then the spatial extent of the signature, along with the statistics of the pixels within the signature can be used as effective discriminants. This approach may be capable of recognition, separating targets into broad classes, but not identifying specific target types;

- For high-resolution imagery, with many pixels on target, identification of a target as a specific class of vehicle may be possible [52]. However, this can require an extensive training set and may be very configuration dependent, with small changes in the state of the target affecting the classifier output [35].

The ATR approach implemented here focuses on the second class, where there are sufficient pixels to gather information on the attributes of objects but, in most cases, insufficient detail to perform identification.

The output of the ATR process is considered to comprise a list of objects, each accompanied by their location and measurements of size and statistical features, alongside a decision as to whether it is believed that object may be part of a desired target class or a clutter object. These can be summarised for one configuration of scene, sensor and processing as:

- A probability of detection (P_d) for the desired target class;

- A residual false object density (RFOD) of objects which are declared as potential targets;

- The quality of attribute information which can be measured from the target.

Figure 5-1 is a functional diagram of the algorithmic chain, showing the major components through which the data flows.
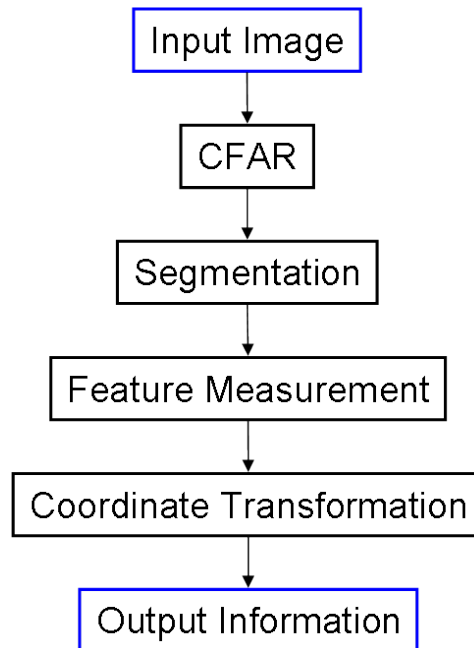
Figure 5-1 Functional diagram of algorithmic chain. Processing modules in black, input and output in blue

The function of each block is defined below.

The initial detection stage is a Constant False Alarm Rate (CFAR) filter, which applies an adaptive threshold to the image to detect areas with brighter returns than the local background. A window is passed over the data, illustrated in Figure 5-2, that compares the value within a single test pixel to the statistical distribution of a sample set of pixels. These sample pixels are separated from the test pixel by a number of pixels, the guard cells, in order to avoid contamination from the returns of an extended target. If the value on the test cell exceeds the chosen threshold, it is set to 1 in the CFAR output, if not, it is set to 0.
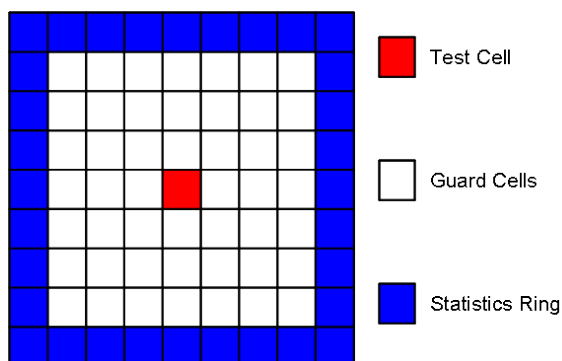
Figure 5-2 Example CFAR window

The choice of window, size, clutter distribution, and threshold can be selected, as appropriate, for the data under test.

The CFAR detections are used to cue a segmentation algorithm, based on that described by Chesnaud et al. [53]. This is a region-based approach which considers a target region and its background, separated by a connected set of nodes. The nodes are moved and a minimisation criterion applied to judge whether the move improves the separability between the two regions. This is often referred to as 'snake' segmentation. A more advanced segmentation routine was developed by Moate and Denton [54], offering improved performance over the original version through a number of modifications.

The improved segmentation automatically adapts the number of nodes to suit the target size. As nodes are moved around, if the distance between any two nodes exceeds a certain number length, an extra node is created at the midpoint. The version in [53] used a fixed number of nodes which did not vary with the size of the segmented object. To improve the speed of this process, a number of iterations are performed with a limit of four nodes, which finds the basic target outline, and then a further set are calculated with an adaptive number of nodes. An example of this is shown in Figure 5-3.
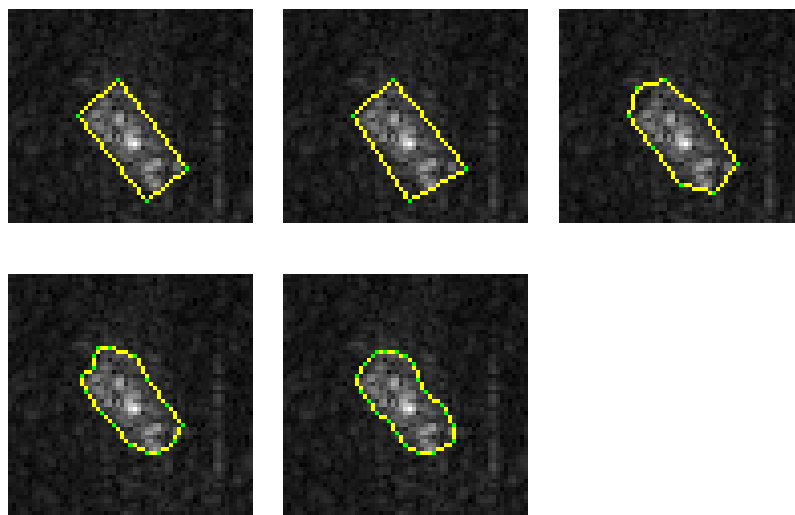
Figure 5-3 Example of delineating an image first with the rectangle model (top left) and then using the multi-resolution polygon model with increasing nodes. Node positions are shown in green

It is also possible to run multiple snake hypotheses in parallel and choose the best option when segmenting a target. While this increases processing load, it helps to avoid the situation where the minimisation criteria for a single segmentation could fall into a localised minimum and does not find the best segmentation.

The segmentation outlines are then used in a feature measurement process, whereby a set of size-based and statistical features are calculated from the pixels within the outline, and from the surrounding local clutter.

Warping, or axis transformation, of the segmentations to a flat ground plane is applied to account for the variable resolution within an RDM. This terminology is used in the ATR process for an image with axes of range and Doppler that has not been corrected for resolution variation with range. The seeker imagery gives a constant Doppler and range resolution that, when mapped to the ground, introduces a range dependent cross-range resolution. Using corner coordinates from an RDM, illustrated in Figure 5-4, the resolution at near range is approximately 0.56 m / pixel, and the resolution at far range is 0.75 m / pixel. The resolution also varies significantly depending on the antenna scan, as shown in Chapter 3.5. A second example taken on an outboard scan produced a resolution of 0.34 m at near range and 0.49 m at far range. The mapping process, therefore, prevents what can be up to a factor of 2 in scale from affecting measured size.

For the first example from the paragraph above, the mapping of an RDM to the ground plane is shown in Figure 5-4. The left part of the figure shows an RDM with the corners and the centre of the near and far range extent as blue circles, the target segmentation node positions as green points, and the shadow segmentation nodes as red points. The right image shows the same RDM warped into a square pixel imaging geometry on an assumed flat ground plan, with a transform based on the seeker navigation information and antenna pointing angle. The change in the width of the RDM on the ground (the radar is imaging from the top-left), is visible in the warped image.
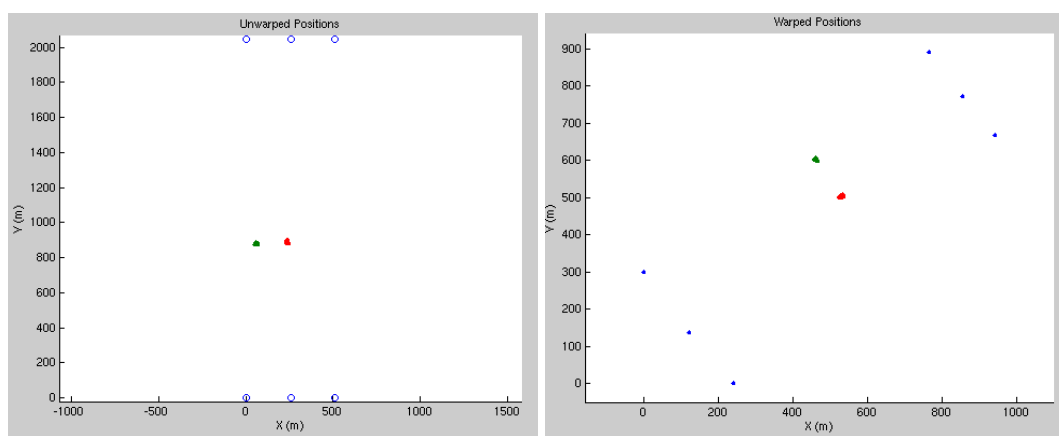
Figure 5-4 RDM (left) and warped coordinates (right). RDM corners are shown as blue points, target nodes as green, and shadow nodes as red

Through mapping to the ground plane, which is an affine transform of rescaling and rotating, multiple RDMs can be registered to produce a large-area composite image with increased contextual information.

Studies in this thesis investigate performance at different parts of the chain that was shown in Figure 5-1. The processed RDM data is used (Chapter 7, Polarisation, Chapter 5.4, Attenuation), and also at the segmentation level (Chapter 5.3, model matching and Chapter 6, Shadows), or at the range-profile level without use of a synthetic aperture (Chapter 8, real-beam ATR).

An example of part of an RDM being passed through this chain is shown in Figure 5-5 to Figure 5-8. Range is shown from left to right and Doppler in the vertical direction. In Figure 5-5, three areas have been highlighted: the blue area contains two armoured vehicles, the amber area contains four trihedral reflectors and a soft-skinned vehicle, and the red area contains two construction vehicles. The output from the CFAR process can be seen in Figure 5-6. Many pixels on the vehicles and reflectors have been passed. However, there are also returns due to rural clutter. Segmentation and some basic size constraints reduce the number of areas of interest in the scene to those shown in Figure 5-7, where the segmentation outlines are visible and superimposed on the original RDM. In Figure 5-8, thresholds have been placed on feature measurements from the objects, with potential targets highlighted in red, and those considered to be clutter objects in yellow.
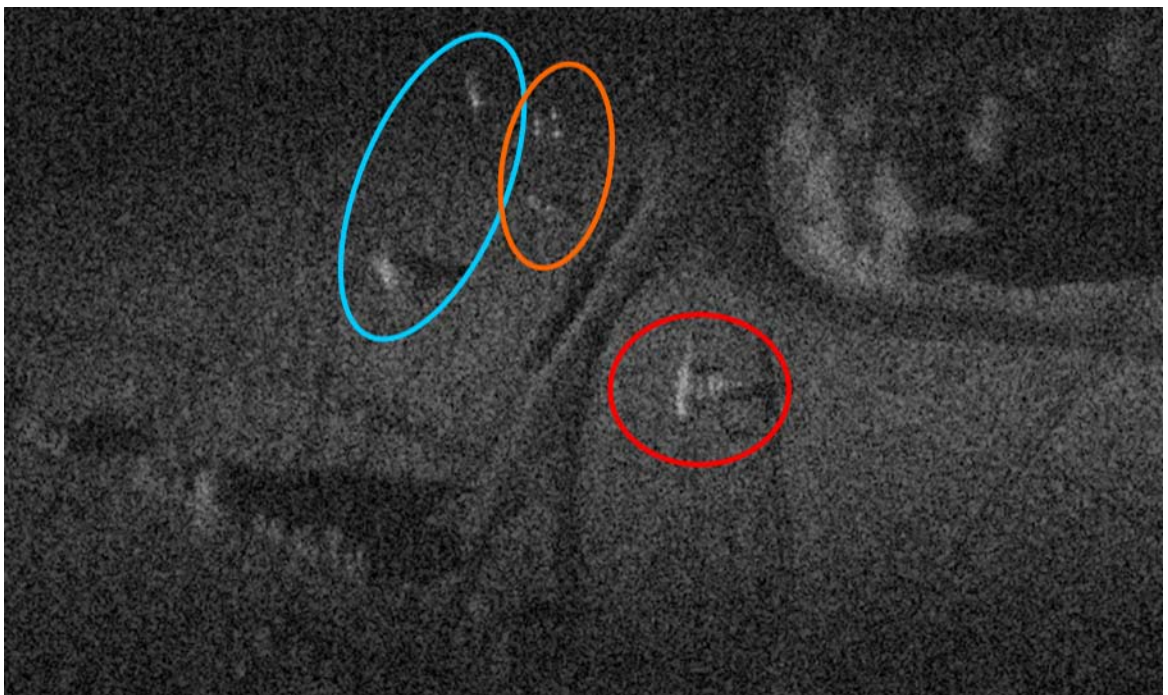
Figure 5-5 Input RDM with sets of objects highlighted
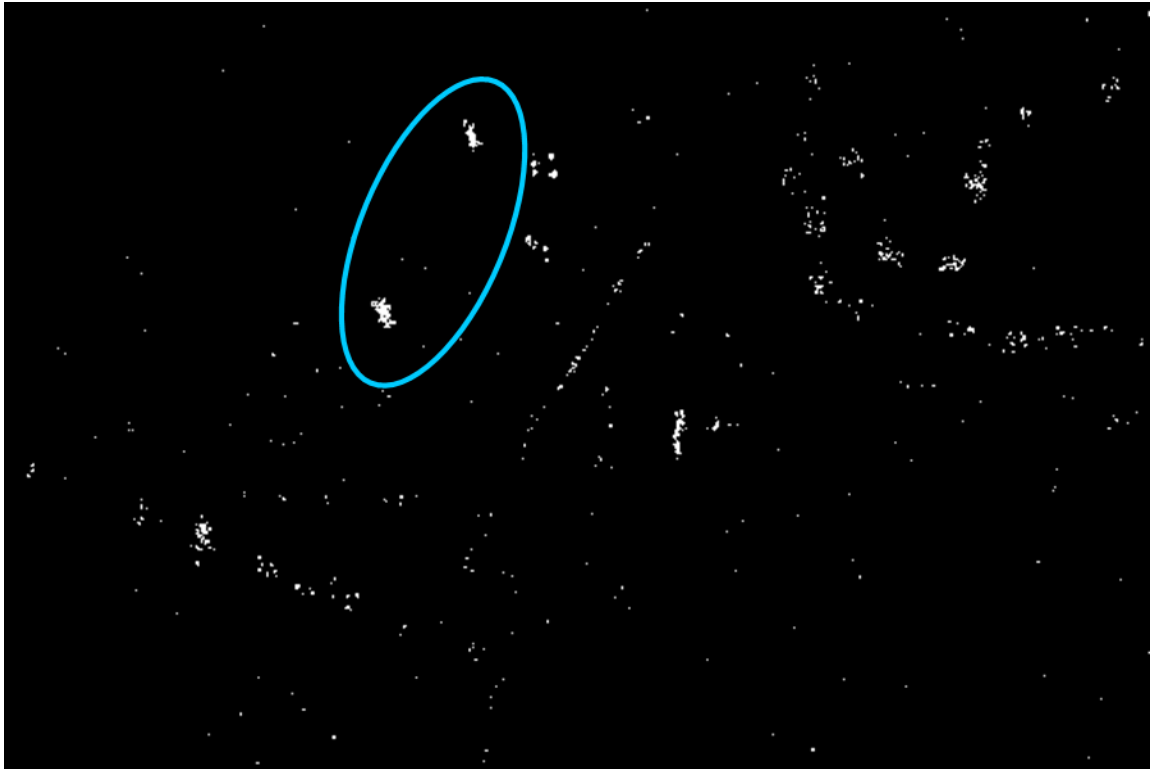
Figure 5-6 RDM following CFAR

Figure 5-7 RDM following segmentation
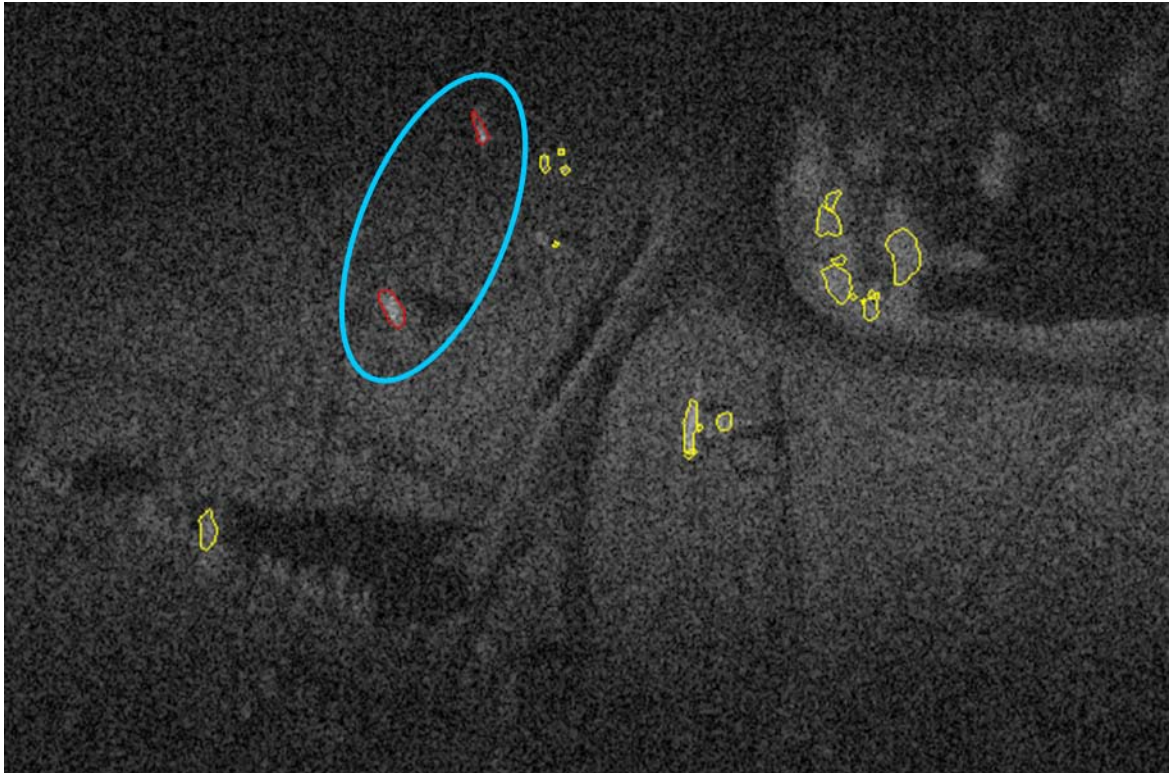
Figure 5-8 RDM following feature thresholding

Rather than rely solely on the ATR capabilities of the seeker to select a target, it is envisaged that the output information could be used in a fusion process [45], with prior information provided before launch or during flight to the target area, to allow Context Data Aided Acquisition (CDAA). This changes the requirement on the seeker processing from being one of finding what is perceived as probably a suitable target to providing the information that will allow selection of a specific target. The exact target selection can then be performed in a final stage of on-board processing (not detailed here).

5.1.1 ATR Summary

The generation of fine-resolution data offers greater information on target attributes than can be achieved with a real-beam system. An approach to the detection, segmentation and attribute measurement of objects in radar imagery has been described. This approach is used in the following chapters, including a specific study of performance against changes in the signal to noise ratio of the imagery (Chapter 5.4). Many of the performance measures are calculated as a single value that incorporates returns from a number of different viewing angles against the target, averaging performance over multiple aspects.

5.2　Challenges in measuring target size and shape

The 'radar size' of an object in an RDM image can vary significantly from its true size due to a number of contributing factors which are outlined here.

Self-shadowing, in which the structure of an object obscures parts of the object from the radar's line of sight, can significantly change perceived size. Figure 5-9 shows this effect for an armoured vehicle at different viewing angles. The perceived shape of the object varies significantly depending on the imaging angle.
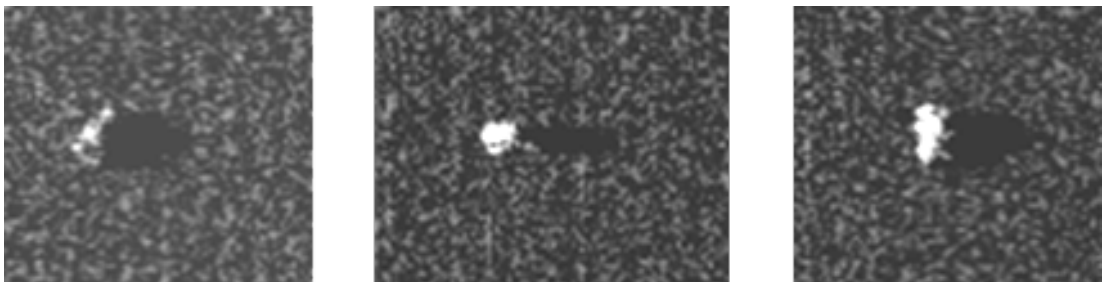
Figure 5-9 Self-shadowing in radar imagery

The effects of self-shadowing on perceived size using an object segmentation and measurement approach are further illustrated in Chapter 5.5.

This effect can be a problem for all sensor domains, affecting optical techniques as well as all types of radar. Although, in the optical domain, no shadowed regions are visible since returns fill every azimuth and elevation pixel.

Solutions to this problem include repeated looks at a range of azimuth angles to build-up a better understanding of the object, with the greatest level of detail available through a series of high-resolution back-rotated images from all angles as shown in Figure 5-10.
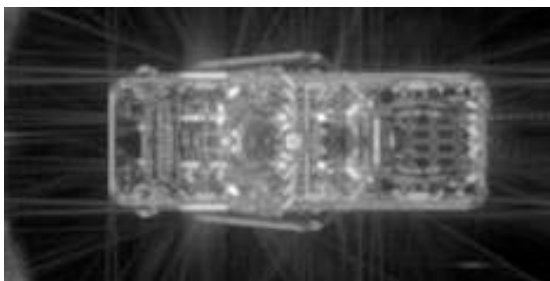
Figure 5-10 Back-rotated series of high-resolution images

However, depending on the radar platform's capabilities, this may not be possible. An approach which uses knowledge of likely self-shadowing is detailed in Section 5.3.

The variability of scattering within an object can result in signatures which change significantly over small changes in viewing and elevation angles. These variations add noise to any size measurement technique since parts of the target may be included / excluded depending on the exact signature. The overall level of the signature against the local background also affects how well size can be measured, and this is explored in Section 5.4.

Should the object of interest be moving, the signature is both displaced and blurred [55], making size measurement from the high-resolution image impossible. Techniques for focusing the target signature to such a level that size can be measured have been explored [56], but with limited success on seeker data because of the difficulties in compensating for acceleration components. Information from the object's shadow, if available, can contribute to the estimation of the size of the object, and this work is presented in Section 6.

Further effects in the radar imagery can make size measurement difficult including spill-over of energy into adjacent cells and incorrectly focused imagery, which may smear the signature.

5.3 Simple CAD model signature prediction

5.3.1 Introduction

The decision as to whether an object detected in the static target detection processing is declared as a potential target is based on the analysis of a number of size-based and statistical features. These features contain useful information about the nature of the object, including length and width measurements. However, information about the target's shape in the imagery should also be considered.

Any constraints based on target length and width need to be broad, owing to the wide variation in target signature, primarily with imaging angle. This limits the ability of simple size features to reject clutter objects.

The shape of the object and its gross internal structure have the potential to act as a further filter to reduce the RFOD, the number of false declarations per km^2.

The approach to this work was to use simple CAD models of targets, generated using the CLTG signature simulation software (Section 4.4), for correlation with real air-carry imagery to produce a matching metric.

The model matching routine is designed to be added to the static detection chain (described in Section 5.1) as an extra clutter filter, processing objects which are currently assigned as targets. As well as clutter rejection, a likelihood of match value is produced, which could be passed to a target selection process.

To assess the potential of this technique, data from two sources were cross-compared. Section 5.3.2 describes the real data gathered from air-carry trials, used as the test data. Section 5.3.3 details the synthetic radar data generated using CAD models in the CLTG, which was used to create a database of templates.

5.3.2 Trials data

Data using the airborne radar, described in chapter 4.2, was used to provide real images of targets with which the algorithm could be tested.

Two targets selected to use for this work are shown in Figure 5-11. These were the Mobile ADU and the SPG. The ADU was chosen since a detailed CLTG model already existed, and the SPG was the most frequently imaged target in the trial. The simple shape of the SPG made it ideal for the generation of a basic model.

Figure 5-11 ADU (left) and SPG (right)

Imagery of the targets was available with a resolution of approximately 0.5 metres in the cross-range direction and 0.25 or 0.5 metres in the down-range direction, dependent on the radar waveform.

The resolution within an RDM is not uniform. As shown in Chapter 3, the cross-range resolution is dependent on the scan rate, range from the radar and angle from the velocity vector. The down-range resolution varies depends on the transmitted bandwidth.

To allow correlation between real and synthetic imagery, it was necessary to ensure a uniform pixel spacing was applied. A spacing of 0.6 metres was chosen since this was the finest pixel spacing that accommodated the worst cross-range resolution.

To achieve this, the composite image formation (CIF) mode of QinetiQ's radar seeker processing software was used to transform the RDM image to the ground plane. A database of SPG and ADU images was then extracted from these warped images, as well as a large number of clutter objects to use for RFOD analysis. The database was formed of 'image chips', each chip being a small extract from the main image, of sufficient size to ensure that the complete signature of the object was included.

Much of the data gathered in the 2006 trial was of a higher quality than that gathered in previous trials. Improvements made to the radar resulted in a higher signal to noise ratio and the target shape became more apparent in the imagery. This improvement in the imagery should be a key facilitator in the performance of this model matching technique. Figure 5-12 shows the SPG imaged next to the array of four reflectors at the Copehill plantation. The shape of the body of the target is clearly visible, with some self-shadowing of the lower right hand side because of the turret. The extra height of the turret is apparent in the length of the shadow, where a step-change is visible from the hull shadow to the turret shadow.
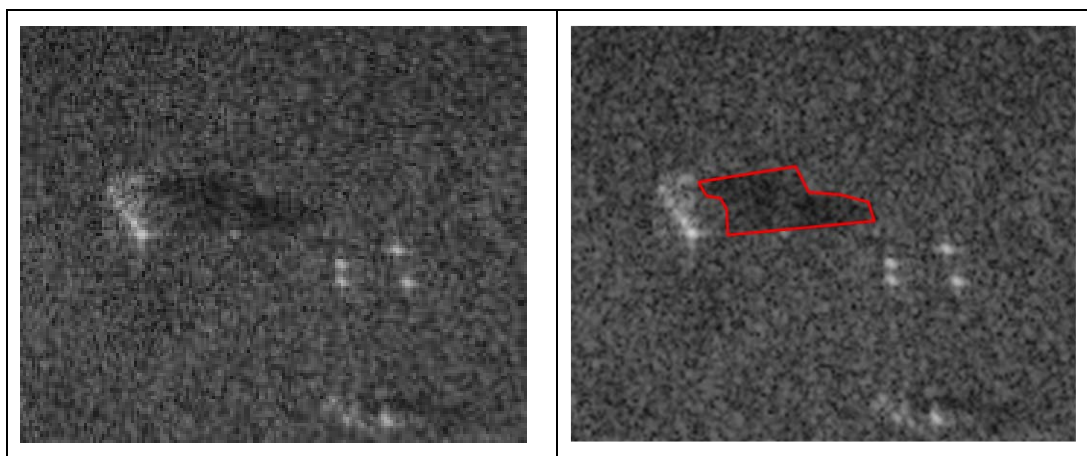
Figure 5-12 SPG, image (left) and with shadow manually highlighted (right). Down-range from left to right, cross-range vertical

These types of effects were not seen in data gathered from previous trials and, whilst self-shadowing effects are accounted for in this study, the full shadow information is not currently processed, but forms a component of the work reported in Chapter 6.

5.3.3 Synthetic CLTG data

As described in Chapter 4.4, the CLTG is capable of producing RDM imagery of large areas including target signatures. The signature representation is simple but is considered sufficient for the representations of targets required for this shape-

based work. Whilst signature prediction tools more powerful than the CLTG are available, they tend to either work on very small areas or a single object and require highly detailed models and significant computing power.

The CLTG simulation is sufficient to:

- Capture the gross structure of the target in the RDM;
- Replicate the self-shadowing effects seen in trials data.

However, due to the nature of the CLTG and the simplicity of the models it is not possible to create images with a reflected power distribution identical to that measured in real data, as has been used for high-resolution template matching studies [35]. While a mode using 'scatterpoint models' has been developed, this requires very detailed knowledge of the desired target, and does not produce an exact signature match.

The idea behind this work is not to require the detailed information necessary to perform such matching, but to use less specific information to allow robust acquisition of a broad target set. The use of detailed templates can increase the sensitivity to changes in target configuration [35], and require what may be a prohibitively large set of templates. This information may not be available for a desired target, or may not be practical to use within the time frame in which a seeker operates, and hence an approach which is not sensitive to minor changes in target configuration would be of benefit.

In order to produce a set of templates, the CLTG was configured to replicate the geometry of the air-carry trials, with targets placed against a grass background, and imaged using the same range, depression angle and resolution. The output of each simulation was an image containing the target, from which a small section or 'target chip', typically a square extract from the image with a side length of 20 m would be selected.

Each target model was imaged every 5 degrees over a full rotation. A step of 5 degrees was chosen as a compromise between processing load and accuracy. A larger angular step produces a set of fewer templates, saving on space and processing, while a smaller step is better for capturing changes in the target's response. It was felt that 5 degrees was sufficiently small to ensure that changes between sequential images, for self-shadowing purposes, were not significant.

5.3.4 CLTG models

The ADU model used in this study has a reasonable level of detail and should be adequate to reproduce the gross structure in the synthetic imagery. The CLTG target model is represented by two 2D arrays, one containing height information and the other the reflectivity for each point. An example of the CLTG height model, viewed from the side, is shown in Figure 5-13. For comparison, a line drawing of the target is shown in Figure 5-14.

The CLTG simulation only considers height above the surface for each point, therefore all parts of the object extend down to the ground. Such modelling can result in different radar scattering characteristics from those expected from a real target.

Figure 5-13 ADU model side view
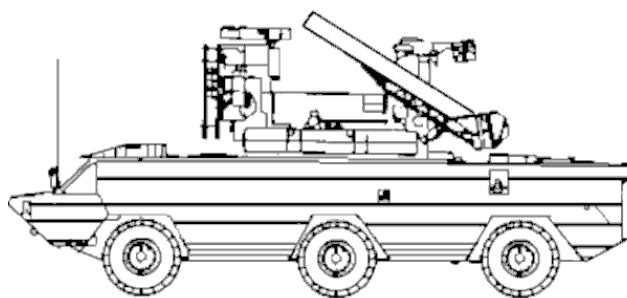
Figure 5-14 ADU Line drawing

An example of the CLTG imagery for the ADU is shown in Figure 5-15. To the left, the CLTG target chip can be seen with a bright return from the target in the centre, and the shadow extending down the image which indicates an imaging direction from above. To the right are the results of the segmentation of the target outline and the shadow respectively.

| CLTG ADU | Segmented Target Outline | Segmented Shadow Outline |
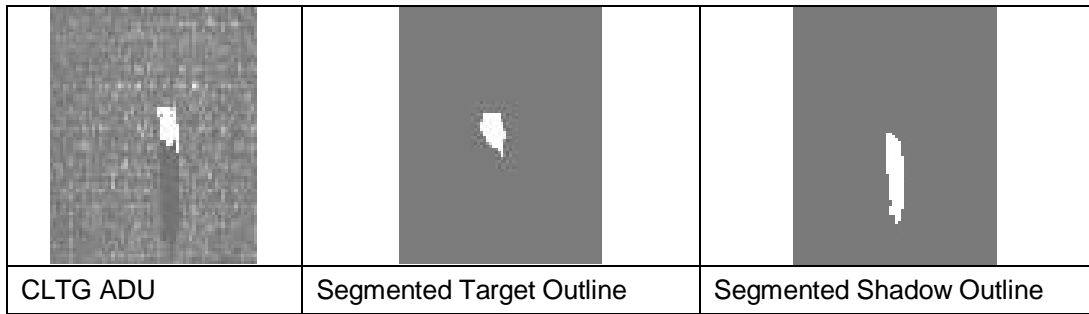
Figure 5-15 CLTG ADU Image

The synthetic image can be compared with that from a real ADU as displayed in Figure 5-16. This target is at a different orientation, and self-shadowing is evident on the left of the image. Since the vehicle is approximately 45° from head-on to the radar, the shadow displays both the height of the vehicle hull, and a longer shadow cast by the mounted radar and missiles.

| Real ADU | Segmented Target Outline | Segmented Shadow Outline |
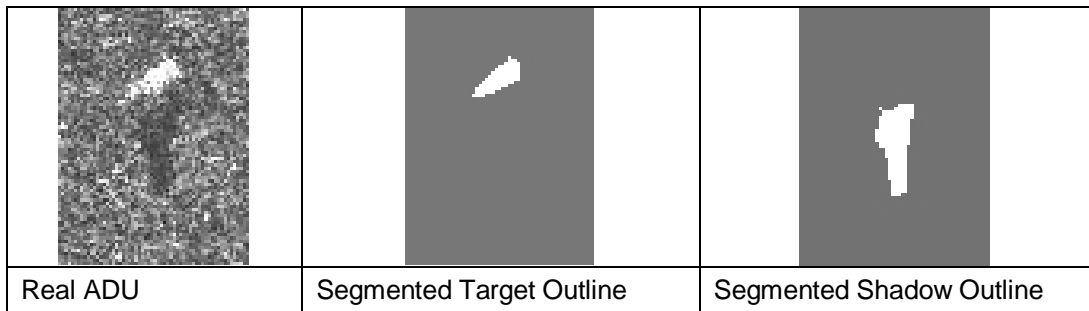
Figure 5-16 Real ADU Image

Owing to the nature of the CLTG modelling process, there may be little to be gained by producing a more detailed model. The reflectivity calculations of the CLTG differ significantly from those of a real target when using this type of model and extra detailing will not rectify this problem. The aim of this work is to correlate gross target shape, size and self-shadowing, therefore a simple target representation should be adequate. Figure 5-17 shows an extremely simple SPG CLTG target model compared with a line-drawing of the SPG. The basic shape and dimensions of the model are comparable with those from the drawn vehicle.
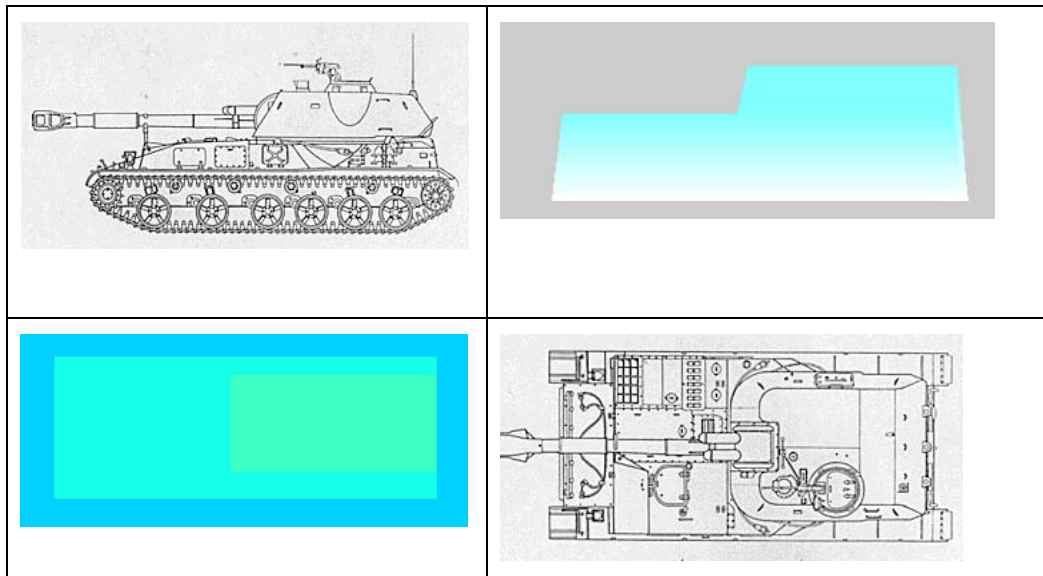
Figure 5-17 Simple SPG model of basic target shape

Owing to the nature of the modelling, details such as the barrel of the gun will detract from the accuracy of the output. The height-map nature of the imagery means that the inclusion of a gun barrel, as shown in Figure 5-18, will lead to the formation of a very large vertical surface that would produce a significant reflection at certain orientations. The returns from the gun are also rarely visible in the trials imagery at representative resolutions. Consequently, the gun barrel was not modelled.
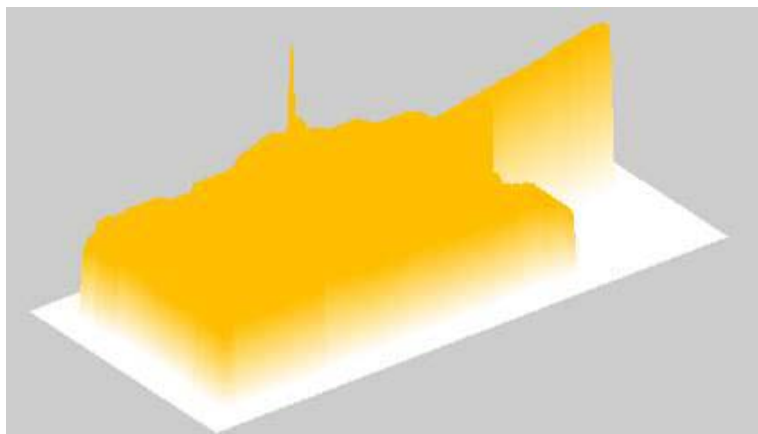
Figure 5-18 CLTG MBT with barrel modelled

5.3.5 CLTG comparison with real imagery

In order to verify that the CLTG imagery was broadly comparable with real data, target images from the same illumination angles were compared. Figure 5-19 shows a CLTG SPG image and an image of the SPG from the air-carry trials. The

output target shape is very similar, with a bright leading edge on the top of the target since the radar illumination is from above. There is a small notch of self-shadowing visible on the down-range side of the images, although this is less clear in the real image. The shadows also have a common shape, with the longer shadow from the turret visible on the right of the image, which is distinct from the shorter shadow cast by the hull of the vehicle. The clutter to noise ratio (CNR) in the real image is lower than in the synthetic image, hence the shadow is less distinct from the clutter.

| CLTG SPG | Real SPG |
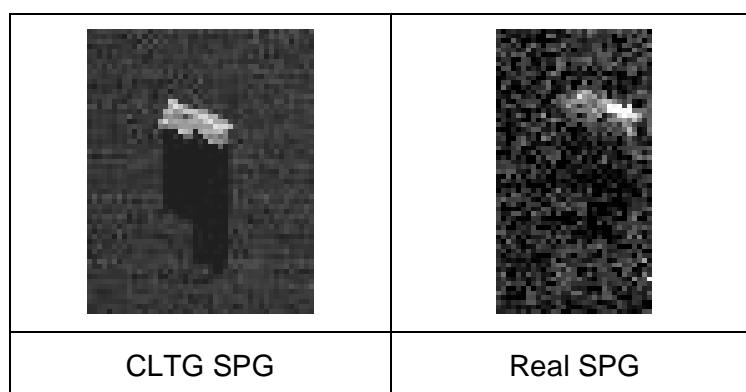
Figure 5-19 Comparison of CLTG and real SPG images

CLTG and real images of the ADU target are shown in Figure 5-20. The match between the model and real data does not appear to be as close as for the SPG. The variation of brightness within the target signature of the real image appears to make the outline shape slightly different to that of the CLTG image, in which the brightness is fairly uniform over the target. It can be seen that useful information is again available within the target shadow, although the clutter to noise ratio for the real target is lower than for the CLTG image.

| CLTG ADU | Real ADU |
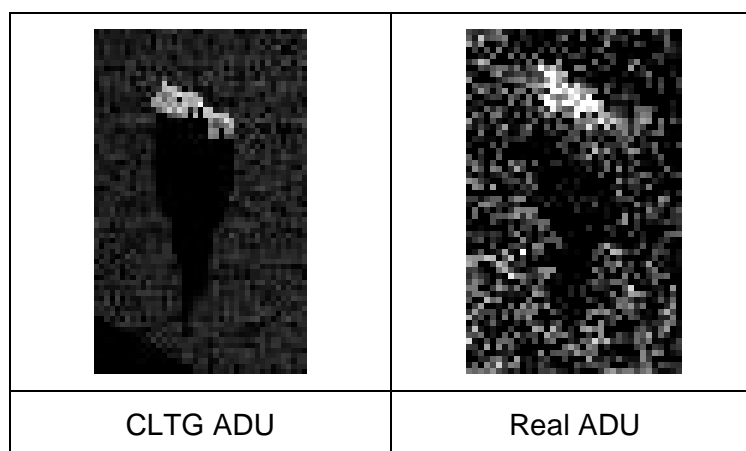
Figure 5-20 Comparison of CLTG and real ADU images

82

5.3.6 Snake segmentation

Before any correlation processing was performed, all target chips from both the trials and the CLTG required a segmentation routine to be applied to separate the target from the background, and hence provide information on the target shape, as shown in Figure 5-21.
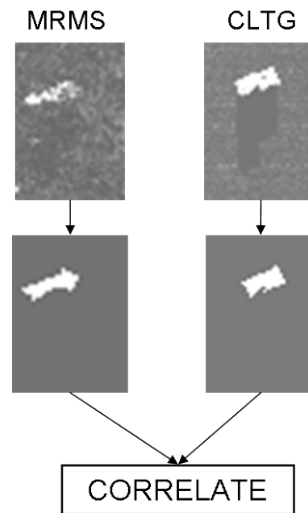
Figure 5-21 Segmentation of real and synthetic data before correlation

The segmentation used is clearly a key stage in the process, since any correlation is dependent on the segmentation result and the 'snake' segmentation, described in Chapter 5.1, was used.

The outputs from the multi-hypothesis segmentation were found to be consistent, with the segmentation converging on very similar solutions when run repeatedly on a single template. The use of multiple hypothesis segmentation helped prevent the selection of a localised solution. The selected outlines also closely matched the outline that would be chosen by a human operator with few noticeable errors when the segmentations were checked visually.

5.3.7 Model matching technique

The model matching process operates on single image chips which have been passed as belonging to the target category by the existing static algorithm.

A phase correlation technique was used to match trials data with sets of synthetic images. This technique is described in [57].

The technique was applied using 2D fast Fourier transforms (FFTs) in Matlab. The quality of the correlation between the two images was taken as the magnitude of

the maximum absolute value of the output image, and the position of the maximum indicated the offset required to align the images to achieve that correlation.

This is a correlation in position only; there is no need to include a rotational element in the correlation process as this is accounted for by the multiple target templates over 360°. While it may seem intuitive to include rotation to find the best match, for a given viewing angle of the target there will be a synthetic image within the matching set formed at a very similar angle. For the set of 72 signatures, the difference will be less than 2.5° as there is a template image every 5°. Introducing a rotation term would increase the processing power demands, as well as potentially confusing the classifier.

The matching process is shown in Figure 5-22. The top images show the targets from the real data and CLTG after segmentation, where only the target shape should remain, with background clutter having been rejected. The segmented target is placed into a zero-padded 256 x 256 array, which is much larger than the target size. The large array size increases the precision in calculating the matching coefficients, but could be reduced should it be necessary to increase the speed of the process. All pixels that are within the segmented area are considered to be target pixels and assigned a value of 1; all other pixels are given a value of 0. The middle images show the Fourier transformed 2D targets, and the lower image shows the correlation that results. The position of the correlation peak, towards the top right, indicates that the second image chip needs to be moved up and to the right to overlap the first chip for optimum correlation. The amplitude of the peak denotes the quality of the correlation.
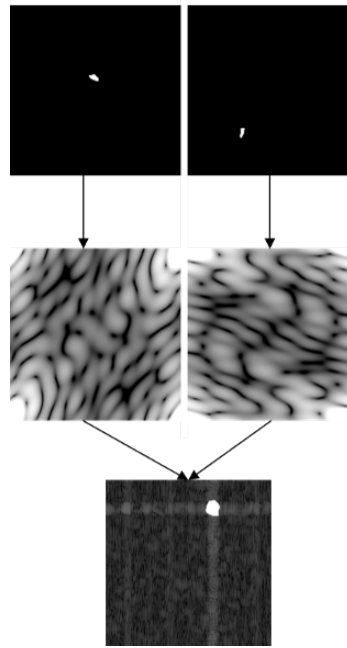
Figure 5-22 Model matching Fourier transform process

The output value of the match for this specific example, 0.53 is fairly low compared to that seen in other cases, indicating that, for this orientation, the synthetic target is not a good match for the real target. The highest correlation value for this real target, 0.85, is produced when the real and CLTG targets are at a very similar orientation, shown in Figure 5-23.
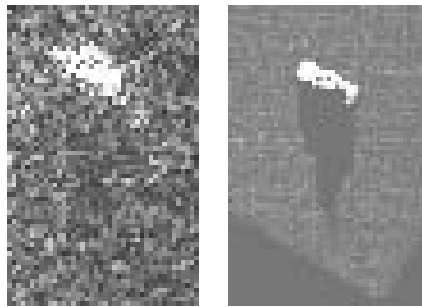
Figure 5-23 Real ADU and CLTG ADU with best match

5.3.8 Amplitude model matching

Initial testing of the correlation process used input images which contained their full dynamic range. The performance in this mode of operation was compared with that achieved when the power in the chips was normalised, and also when power was excluded completely and the correlation was performed on segmented shape alone, with the image having binary values, 1 for target and 0 for background.

When correlation was applied to segmented images with their full dynamic range, since the variability within a target signature can be extremely large, the bright

points significantly bias the match. Figure 5-24 illustrates this process; the highest correlation can be where the two bright points align, and since the magnitude of the radar returns from those points is much higher than in the rest of the image, a good correlation value is returned, despite the fact that the true correlation is poor.
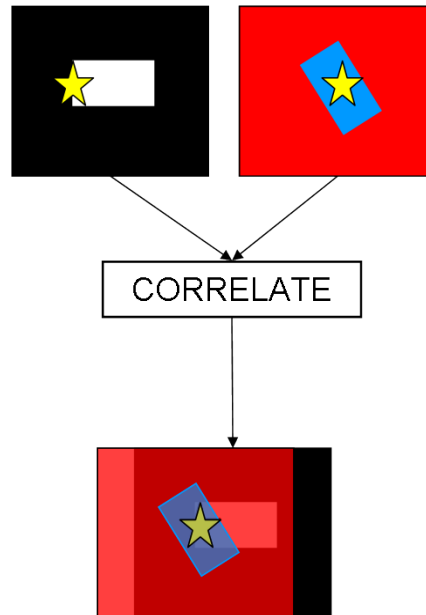
Figure 5-24 Illustration of correlation of full dynamic range images, with bright scatterers represented by a star

Normalised cross-correlation may provide the benefit of reducing the impact of these effects [58]. However, since the CLTG does not capture detailed scatterer information, it is unlikely that a full amplitude approach would be successful. Work on the CLTG, described in Chapter 4.4, which involved the extraction of significant scatterers from ISAR data to create better target representations in synthetic imagery may partially solve this problem. However, that approach moves away from the benefits of using simple models.

Comparison of the results of the binary and full amplitude techniques over a small dataset showed that, for a fixed probability of detection, the binary matching resulted in a better match and the ability to reject clutter objects for both rural and urban clutter. The binary approach was subsequently used to produce the results in this chapter.

5.3.9 Application to MRMS data

The model matching scheme was applied to the output from the ATR algorithm as a technique to reduce the RFOD. The model matching was implemented as a stand-alone routine, operating at the end of an existing ATR chain. The inputs to the

model matching were objects that the current algorithm considered to be potential targets following feature discrimination. The outputs were a correlation value for each object which could then be thresholded for a further target / clutter decision.

The P_d from the model match routine is the P_d from the clutter discriminator multiplied by the probability of a target passing through the model match routine. The RFOD on the output of the model matching scheme is the RFOD from the existing algorithm multiplied by the probability of a clutter object being passed by the matching routine.

Figure 5-25 shows the points at which these target / false target calculations are performed.
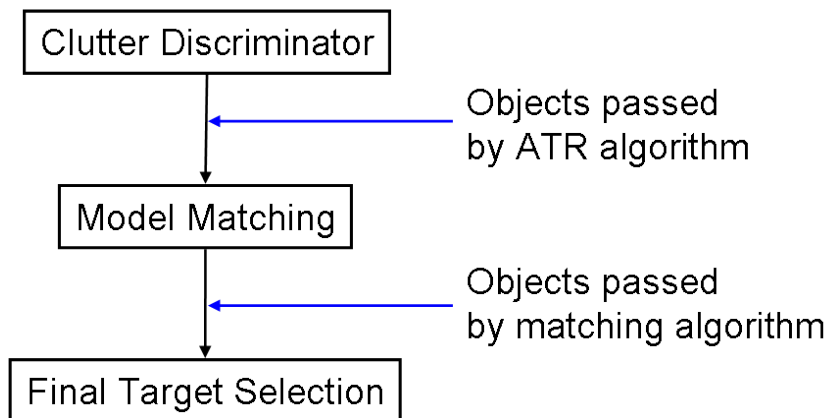
Figure 5-25 Implementation of model matching routine

To assess the performance on real targets, a set of ADU and SPG trials imagery was passed through the model matching routine. A threshold was applied to the correlation value such that the probability of a target being passed was close to 1 to avoid lowering the P_d.

To gather clutter information, two complete imaging runs over the Salisbury Plain trials area were analysed. These included; arrays of radar reflectors, the FIBUA village, moving targets, static targets and decoys. The radar imagery was passed through the ATR algorithm and all false detections were used to measure the RFOD.
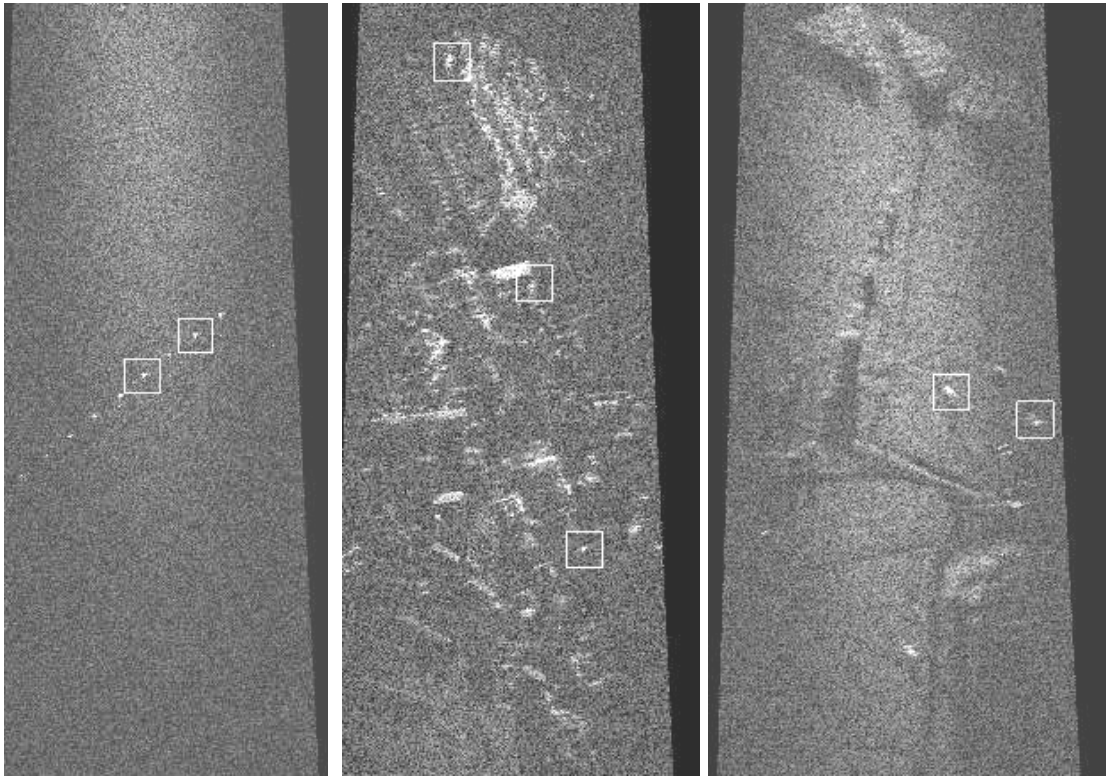
Figure 5-26 False detections on reflector array (left), FIBUA (middle) and target detections at the Copehill plantation (right)

Following tests with the SPG and the ADU models, it was found that the synthetic SPG imagery produced a lower RFOD than the ADU model, and gave a good P_d on both real targets. Since the performance was higher on both target sets, it was decided to only use the simple SPG model for the calculation of results. That the simple model produced the better results indicates that increasing model complexity does not necessarily improve performance for this technique, and that a small database of basic model types may be sufficient for a wide range of targets. Key parameters for the success of this technique are a comparable target size and similar levels of self-shadowing. This result, albeit on a limited dataset, indicates that robustness to changes in configuration would be high.

5.3.10   Clutter rejection

The results of applying the model matching routine to a dataset of target and clutter objects are shown in Table 5-1. For a range of object types the columns show the:

- number of objects labelled as a target by the ATR algorithm;
- number of objects which have subsequently been declared as targets by the model match process;
- probability of detection (targets only).

The radar point reflector detections, whose bright response coupled with their returns spreading into adjacent resolution cells caused them to be occasionally labelled as targets are often rejected by the algorithm with only 13% being passed though. Half of all urban false detections are rejected, and slightly more than half of rural detections.

For the two selected targets, which had a P_d of 1 before matching, the ADU is passed through each time it was imaged, and the SPG passed 90% of the time. It should be noted that this does not necessarily imply that performance against one target type is worse than for the other. The quantity of data available is limited, and further data would need to be examined. The failure on a SPG occurred when matching to a faint image at the edge of the beam, and should not be an issue for a case where the seeker has a comprehensive scan pattern.

Decoy targets, designed to confuse observers in one or more of the visual, radar or IR domains are still passed by this approach. Since the approach uses shape information, which a decoy can replicate relatively easily, it does discriminate against them, and reliance would have to be placed on other aspects of the ATR algorithm.

Type	Objects passed	Targets declared	Probability passed	P_d
Point reflectors	31	4	0.13	
Urban	121	58	0.48	
Rural	8	3	0.38	
Civilian Vehicles	25	12	0.48	
Moving Target	5	2	0.40	
Decoy	6	5	0.83	
ADU	3	3	1.00	1
SPG	10	9	0.90	0.9

Table 5-1 Probability of clutter and target objects passing the model matching

The results from the model-matching technique appear to be promising, with a significant improvement in clutter rejection performance over that provided by the existing algorithm. Fewer false detections should result in a smaller processing load for subsequent fusion algorithms, and it will make available an additional feature of matching likelihood to aid selection of the correct target.

The RFOD for these clutter types can be calculated through knowledge of the imaged area and the false alarm count. Whilst the reduction in rural and urban false alarms is significant, it should be noted that:

- the rural RFOD was already low, hence only a small number of samples were passed forward for testing;

- the urban RFOD remains relatively high, of the order of tens per km^2. Dependent on search area this may be too high for a simple target selection process. Fusion with contextual data can allow acquisition with a far higher RFOD.

5.3.11   Model matching summary

The use of simple CAD based target prediction models has been shown to reduce the RFOD for a wide range of clutter types while retaining a high probability of detection.

Very limited target knowledge is required to achieve this level of performance, with the simple SPG model outperforming the detailed ADU model. The techniques have been tested on a small sample of data, and more comprehensive tests would be required to better understand the performance.

Further work should be performed on this algorithm to reach a level of development where it can be included in an ATR algorithm. Sufficient target contrast is a pre-requisite for reliable segmentation, as is a suitable number of pixels on the target. The number of pixels is dependent on resolution and target size. A resolution of 0.6 m was used, which for medium-sized military vehicles gave a good level of performance. Should the image resolution fall below 0.6 m, or smaller targets with less structure be considered, performance is likely to fall.

5.4   Performance effects of radar sensitivity / attenuation

A short parametric study into the performance impact of changing the seeker sensitivity, as could be varied through power levels or atmospheric attenuation, has been undertaken to understand the effect on performance characteristics such as detection and size measurements.

5.4.1 Seeker sensitivity

The sensitivity of a radar seeker is considered to be the ability to detect a target of a specific RCS against a noise background at a defined range. For an imaging radar seeker, ground targets are seen against the surrounding clutter, which may have a high reflectivity (RCS/m^2) and be relatively bright (e.g. grassland) or have a low RCS per unit area and be noise dominated (e.g. smooth roads).

Noise can be considered to consist of thermal noise, plus any noise / interference from other sources, such as the radar waveform.

The path between the radar and target will also affect the sensitivity of the radar. For the frequencies used in many common radars, atmospheric attenuation may be relatively low, even at long ranges. However, for radar seekers which operate at higher frequencies, such as 35 or 94 GHz, attenuation from the atmosphere, especially due to precipitation can reduce performance. In addition to the attenuation effects, backscatter from rain will contribute to clutter returns and could affect target detectability.

The backscatter at higher frequencies is first considered, followed by consideration of the effects of attenuation.

5.4.2 Rain clutter backscatter

In order to investigate whether backscatter is likely to adversely affect the performance of a radar seeker, it is necessary to calculate the expected level of returns in the volume relating to a range-Doppler cell on the ground.

A set of measurements of rain backscatter were performed by the US Ballistic Research Laboratory in 1973 with radars at 9.375, 35, 70 and 95 GHz [59]. Rain backscatter and attenuation were measured over a wide range of rain rates while recording raindrop size. A simple relationship was fitted to the results, equation (5.1).

$$\eta = a\rho^b \qquad\qquad (5.1)$$

where η is the backscatter coefficient in square metres per cubic metre, ρ is the rainfall rate in mm / hr while *a* and *b* are two frequency-dependent parameters.

For a real-beam radar, which does not simulate a synthetic aperture, the expected backscatter from rain can be derived by calculating the volume of a resolution cell based on the antenna beamwidth and range-resolution. However, for an airborne range-Doppler radar, the volume which contributes to a single range-Doppler cell needs to be calculated. Figure 5-27 shows an annulus of constant Doppler at an angle θ around the velocity vector of an airborne radar.
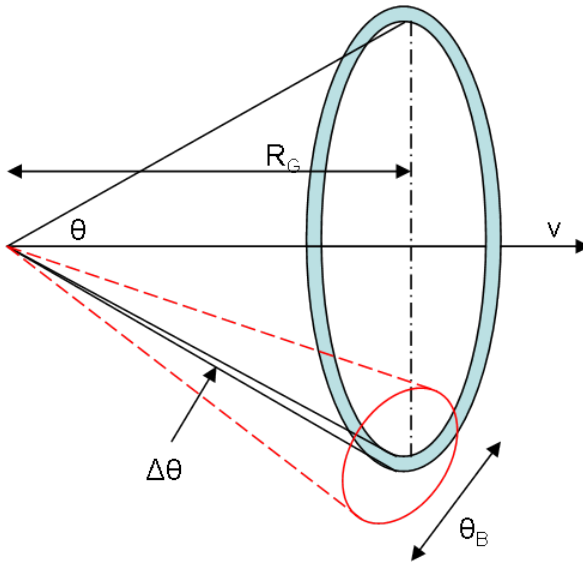
Figure 5-27 Illustration of annulus of constant Doppler

Here R_G is the range along the ground in the direction of travel of the radar, *θ* the squint angle off the velocity vector, *Δθ* the Doppler resolution and θ_B the beamwidth.

The area, *A*, of the annulus of constant Doppler can be found by combining the circumference of the annulus with the cross-range dimension of the radar beam

$$A = 2\pi.R_G \tan\theta.\frac{R_G \Delta\theta}{\cos\theta} = \frac{2\pi.R_G^2 \Delta\theta \sin\theta}{\cos^2\theta} \qquad (5.2)$$

The volume within a range Doppler cell can then be calculated using the resolution in the down-range direction, Δ_{CR}. The volume is then further reduced due to two factors:

- The 3 dB width of the beam, θ_B, limiting the sampling of this annulus

- One half of the beam will be below the ground plane, so will not produce clutter returns at that range

The fraction of the annulus sampled due to the 3dB width is

$$\text{Fraction Sampled} = \frac{\dfrac{R_G \theta_B}{\cos \theta}}{2\pi.R_G \tan \theta} \tag{5.3}$$

When combined with the components of the signal blocked by the ground, the volume of rain clutter for a single range-Doppler cell is

$$V = \frac{2\pi.R_G{}^2 \Delta\theta \sin\theta}{\cos^2 \theta} \cdot \frac{\theta_B}{2\pi.\sin\theta} \cdot \frac{1}{2}.\Delta_{CR} \tag{5.4}$$

$$V = \frac{R_G{}^2 \Delta\theta \theta_B \Delta_{CR}}{2\cos^2 \theta}$$

The rain-clutter backscatter can then be calculated for a typical imaging seeker geometry, with a squint angle of 30°, a range of 5 km, and resolutions in the down and cross-range directions of 0.5 m. This gives a clutter cell volume of 38 m^3.

For linear polarisations, the RCS of the rain in a resolution cell is given in Table 5-2.

Freq (GHz)	Volume RCS (m^2 / m^3)	Cell RCS (dB m^2)	Cell RCS (dB ^{m2})
		1 mm/hr rain	4 mm/hr rain
9.375	1.04E-08	-64.0	-54.9
35	2.21E-05	-30.8	-24.5
70	1.44E-04	-22.6	-19.1
95	8.98E-05	-24.7	-21.2

Table 5-2 Rain clutter returns from the volume of one resolution cell

These clutter reflectivities can be compared with that which may be expected from typical rural clutter at a representative geometry with an elevation angle of 10°, where a value of -15 dB m^2/m^2 can be used for rural clutter [60]. Given a clutter cell area of 0.25 m^2, for the frequencies above 35 GHz rain clutter backscatter is as large or larger than for the ground clutter.

Since raindrops are near spherical in shape, the backscattered energy will be affected differently compared with that from most clutter or man-made objects. If a circularly polarised wave is incident on a sphere, it will be reflected as a circularly polarised wave with the opposite sense of rotation. For a complex target, there is a

more balanced distribution of polarisations in the returned wave. If the receive channel is not responsive to the opposite sense of rotation (cross-polar) to that which was transmitted, and has good isolation between the channels, then it is possible to reject much of the clutter backscatter from rain.

Raindrops are not perfect spheres, so there will be some backscatter in the co-polar channel. An elliptical polarisation may also be preferable for optimum rain-clutter rejection, though the exact choice of ellipticity will be dependent on the drop size and may not be an easy parameter to vary. The cancellation by an optimised antenna may be as high as 40 dB, but is more likely to be of the order of 20 dB for most moderate rain rates [61].

Given 20 dB of clutter rejection for circular polarisation, the ratio of clutter to rain is such that, for all frequencies, the rain should not have a significant effect on the ability to produce good quality imagery. Table 5-3 shows the ratio of returned power from clutter relative to rain within a resolution cell at a range of frequencies. A high clutter to rain ratio is desirable.

Freq (GHz)	Clutter / Rain (dB)	
	1 mm/hr rain	4 mm/hr rain
9.375	63.0	54.0
35	29.8	23.4
70	21.6	18.1
95	23.7	20.2

Table 5-3 Background clutter relative to rain backscatter for two rain rates

5.4.3 Atmospheric attenuation

The Clutter to Noise Ratio (CNR) in radar imagery is dependent on the large number of factors that form the radar equation but, for a given radar and waveform, the two factors that would typically cause CNR variation are the operating range and atmospheric attenuation. For a radar with a given CNR in clear conditions at a range of 5 km, the CNR as a function of atmospheric attenuation is shown in Figure 5-28. It can be seen that, for high levels of attenuation, the CNR can fall well below 0 dB. The CNR values plotted vary from a level where noise is not considered to be an issue (15 dB) to a level where noise will dominate the imagery (-15 dB CNR).
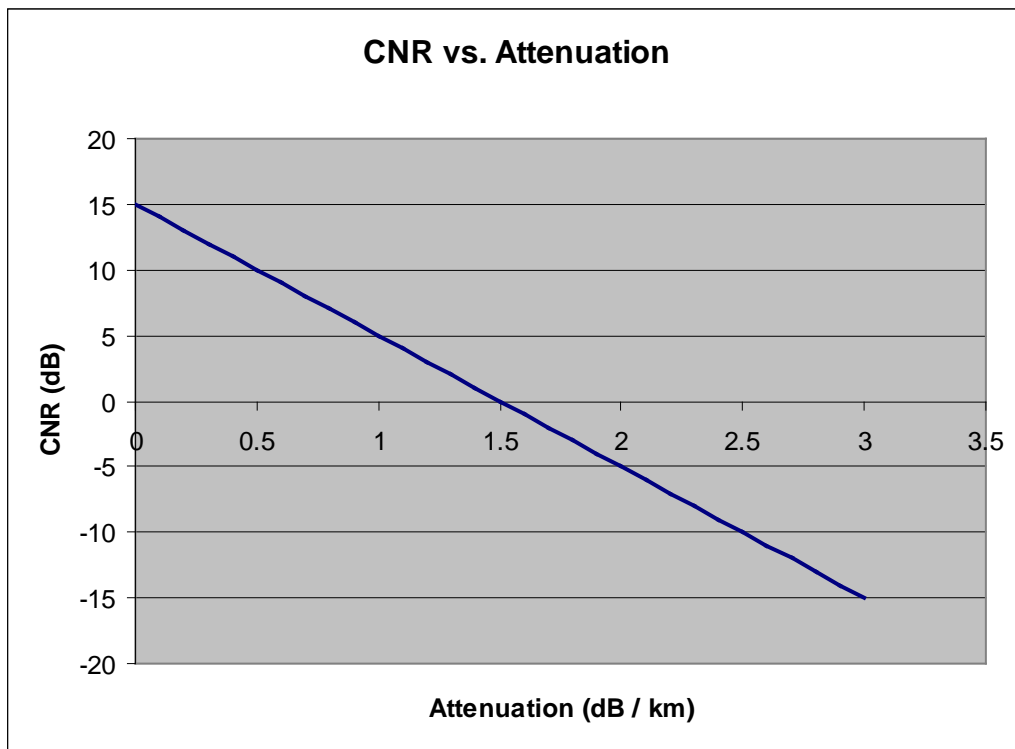
Figure 5-28 CNR vs. attenuation for a generic imaging radar at 5 km range

Rather than consider the impact of a number of factors on sensitivity, this study examines where reduced sensitivity has an impact on key performance areas for ATR and a scene matching algorithm. The sensitivity levels can then be related back to defined rain and range values as appropriate.

Attenuation as a function of frequency and rain rate is described in [27]. Here the attenuation per km, A, can be calculated as a function of rain rate in mm / hr, ρ, and two frequency dependent factors, c and d.

$$A = c\rho^d \tag{5.5}$$

The attenuation, in dB / km for the four frequencies that were considered was calculated and is shown in Table 5-4.

Rain Rate (mm/hr)	Attenuation (dB / km)			
	10 GHz	35 GHz	70 GHz	94 GHz
1	0.009	0.27	0.63	1.6
2	0.021	0.54	1.2	2.5
4	0.046	1.1	2.1	3.9

Table 5-4 Rain attenuation (dB / km) for varied frequency and rain rate

At low frequencies, such as 10 GHz, attenuation is minimal, even over a long two-way path. However, at higher frequencies, attenuation can be a significant factor, and an understanding of how performance as sensitivity is reduced is valuable due to the increased likelihood of significant attenuation.

5.4.4 Sensitivity effects in imagery

Imagery with a high CNR is more easily interpretable to the human eye than imagery with reduced sensitivity. Figure 5-29 shows synthetic data of multiple RDMs generated using the CLTG model. The model was been configured to represent the MRMS radar operating with different levels of rain in the scene, which represent increasing attenuation and therefore decreasing sensitivity.

Multiple targets are present, seen as objects noticeably brighter than the rural clutter. A network of roads is clearly visible in the image on the left where CNR is approximately 15 dB, but much of the contrast is lost with the reduction in sensitivity in the image on the right, where the CNR is approximately 6 dB.
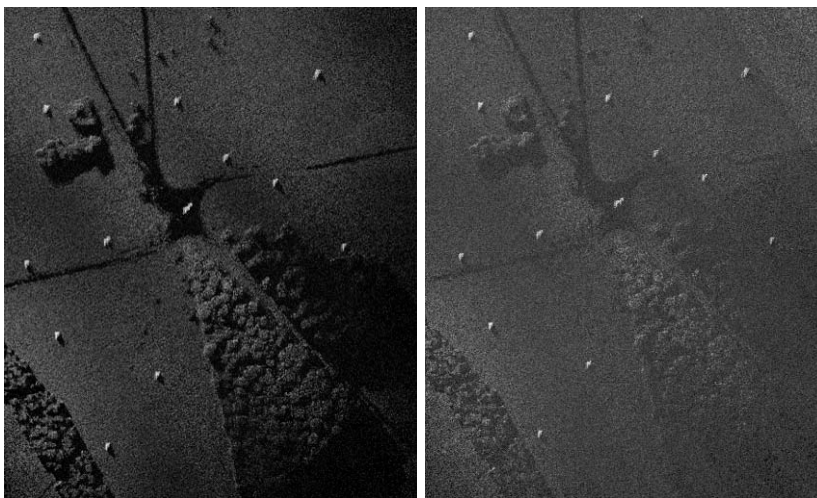
Figure 5-29 Synthetic imagery of multiple MBT targets with 0.5 mm/hr rain (left) and 4.0 mm/hr rain (right)

As sensitivity decreases further, the clutter background can be lost entirely, covered by noise in the imagery. Figure 5-30 shows a single RDM with 15 dB CNR on the left and 0 dB CNR on the right. The brightest part of the clutter is faintly visible in the scene, but the tree structure has been lost against noise.
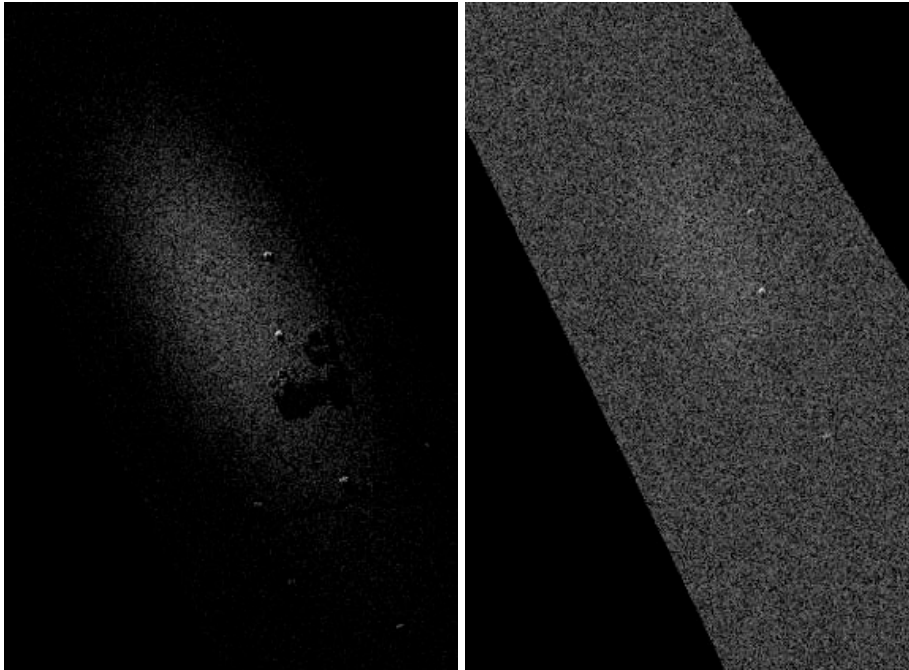
Figure 5-30 Synthetic imagery with 15 dB CNR (left) and 0 dB CNR (right)

5.4.5 Modelling reduced sensitivity

In order to study performance at varying levels of sensitivity, there was a requirement to use both real and synthetic data that was representative of a range of radar sensitivities.

To produce suitable data, a process was adopted in which noise of random amplitude and phase was added to complex radar imagery before conversion to power values:

- An array of complex noise was generated using two, normally distributed, random distributions;

- The real data was 'converted' to complex data through multiplication with a uniformly distributed random phase term (appropriate if phase is considered to be random from pixel to pixel in the image);

- The real data was combined with the appropriately scaled noise array to give a desired level of CNR, and then converted back to absolute values;

- Checks were performed on the output imagery to ensure that the CNR was correct.

This approach was developed under previous seekers research programmes, and used when generating the example imagery shown in Figure 5-29 and Figure 5-30.

The effects of sensitivity on acquisition performance are discussed in Section 5.4.6. The ability to measure attribute information as sensitivity degrades is presented in Section 5.5.

5.4.6 Sensitivity effects on acquisition performance

The MRMS sensor was simulated with a reduced CNR relating to a number of different rain rates to a maximum range of 5 km. The attenuation levels chosen represented a variety of rain rates, and a P_d and RFOD were calculated for each.

The target under test was an MBT, and the scene considered was the rural scene, shown earlier in Figure 5-29. This is a simple scene for ATR processing, and typically a very low RFOD was reported using simple constraints on target size, brightness and internal statistics.

Table 5-5 shows the CNR values for which data was generated and the P_d and RFOD values that resulted from the data having passed through the ATR algorithm. Some quantisation of the results is visible in the RFOD measures due to the sample size.

CNR (dB)	P_d	RFOD
15.0	1.00	0.23
13.9	1.00	0.00
12.7	1.00	0.23
10.4	0.95	0.00
5.9	0.96	0.00
-1.0	0.97	0.23
-11.0	0.82	14.2

Table 5-5 P_d and RFOD for differing CNR levels, rural data

It can be seen that performance remains high until there is a large drop in CNR such that the noise floor is 11 dB above the clutter level. This may not be intuitively expected, since one of the key factors in discriminating targets from clutter is the target's brightness compared with the local background. Target detectability should remain unchanged until the noise floor becomes higher than that of the background. Given the typically observed Target to Clutter Ratio (TCR) from trials data of approximately 18 dB for an armoured vehicle, even a noise floor of 11 dB above the clutter should leave a detectable target presence.

Figure 5-31 shows imagery with a 15 dB CNR (left), and a -11 dB CNR (right). The target in the centre of the beam, where the gain is highest, is still retained as a bright point in the noise-dominated imagery, although targets away from the beam centre have been lost, along with all clutter details.

The segmentation process is seeded on the brightest points in the image that have passed the CFAR detector. As long as the targets are brighter than their local background, they should be detected and subsequently segmented.
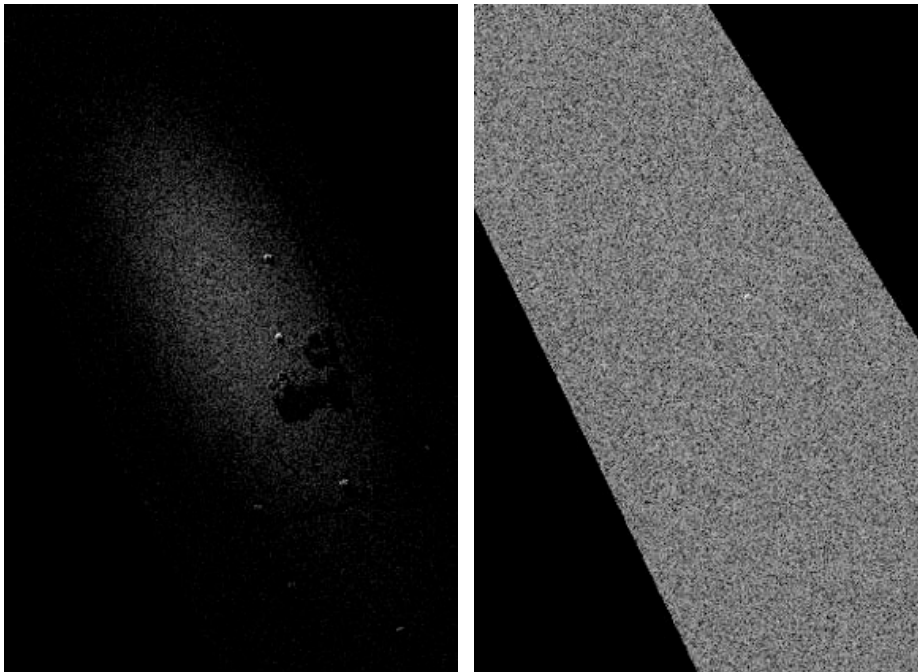
Figure 5-31 0.5 m resolution data with 15 dB CNR (left) and -11 dB CNR (right)

While acquisition performance is not significantly reduced until sensitivity is severely reduced, there are other factors in the ATR process beyond simple P_d and RFOD measures that can be affected, and are discussed in Section 5.5.

5.5 Attribute measurement capability

When viewing the target data generated under previous programmes, degradation of the size measurement capability was noticeable at high levels of attenuation. However, the limited number of targets, and the generation of data at only seven different target headings relative to the radar (head-on to tail-on in 30° steps), did not provide a sufficient number of images for a statistically valid analysis.

To increase the range of imagery available, CLTG data was generated of a single target viewed over a full rotation in 1 degree steps. The imagery was generated with a down-range and cross-range resolution of 0.5 m.

Figure 5-32 shows example imagery of the MBT as the CNR is lowered from 15 dB to -6.8 dB. It is not until the lowest contrast level that the change in target outline becomes noticeable.

| 15 dB CNR | 6.5 dB CNR | 0 dB CNR | -6.8 dB CNR |
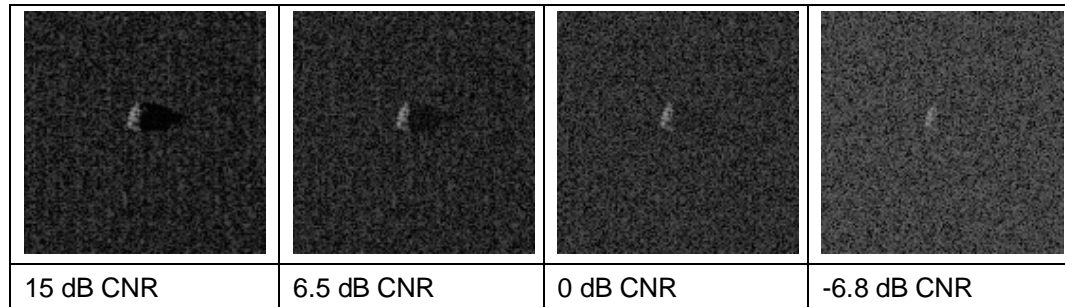
Figure 5-32 Example MBT imagery at varied levels of CNR

Having generated sets of 360 images at a range of CNR levels, the ATR algorithms were applied to the data to provide a set of feature measurements. The features of interest are the target's TCR and size attribute measurements.

The retention of the ability to detect a target as noise dominates clutter can be seen by examining the TCR of a target in the scene for different noise levels, Figure 5-33. Large peaks in the TCR are visible when the target is broadside to the radar (at 90° and 270°), with a narrower peak at tail-on (180°) and a small peak at head-on (0° / 360°). Even with a noise floor 10 dB higher than the clutter level, the target retains detectability, despite a drop in mean TCR of 7 dB.

However, many of the target signatures are then showing a TCR of less than 10 dB, at which point the segmentation applied can suffer degraded performance.

Note that the 0 dB CNR case has only reduced the TCR by about 3 dB. This is because the noise returns are the same amplitude as the clutter returns.
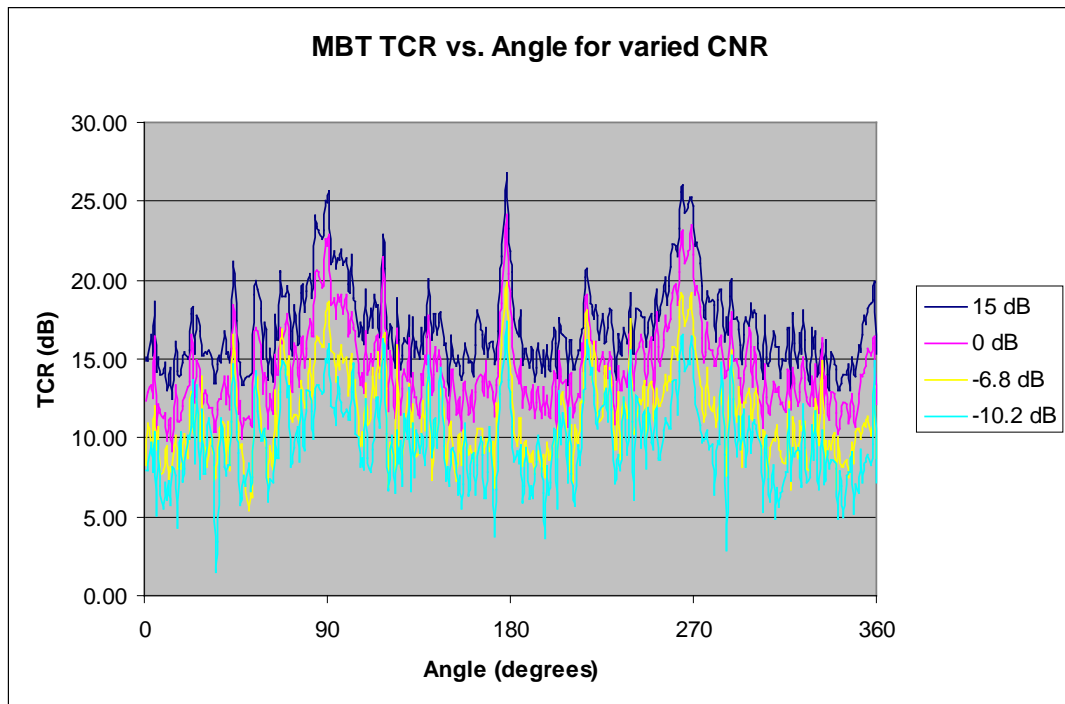
Figure 5-33 TCR vs. angle for an MBT target at varied CNR

However, a degradation in size measurement capability for both the length and width features is observed. Figure 5-34 to Figure 5-37 show the length and width measures of the MBT throughout a 360 degree rotation for varying levels of CNR.

For the cases where the target retains good contrast relative to the local background, 15 dB and 0 dB CNR, the size measurement with angle is in agreement with what would be expected from the target's appearance in the scene and self-shadowing. Length is greatest towards a broadside viewing angle, and at head-on and tail-on, self-shadowing reduces the measured length such that it is comparable with width. The target's true length and width are 7 m by 3.4 m.
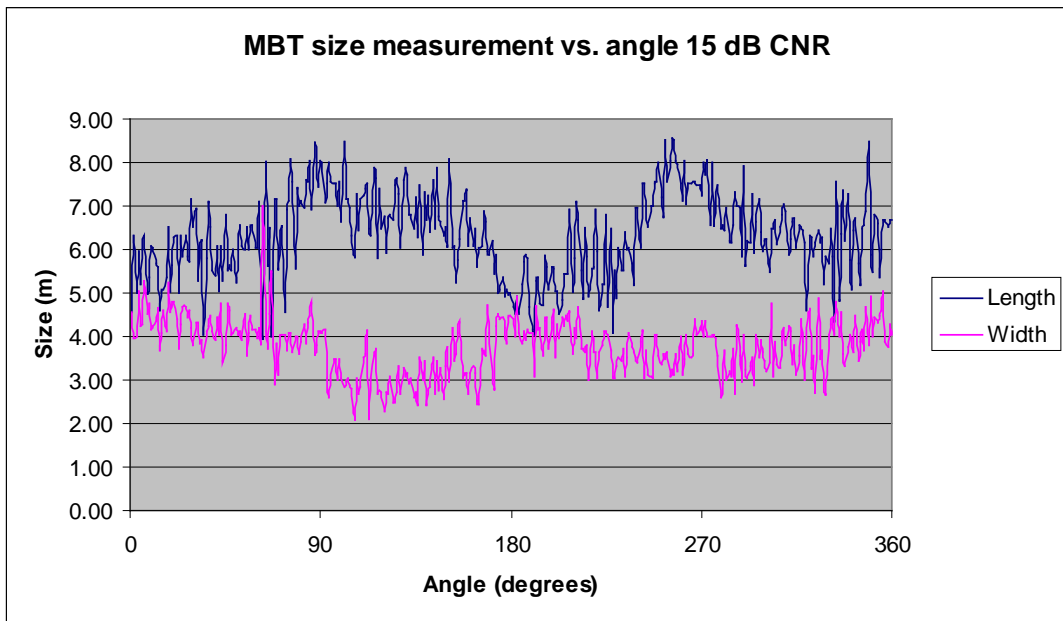
Figure 5-34 Size vs. angle for an MBT target with 15 dB CNR
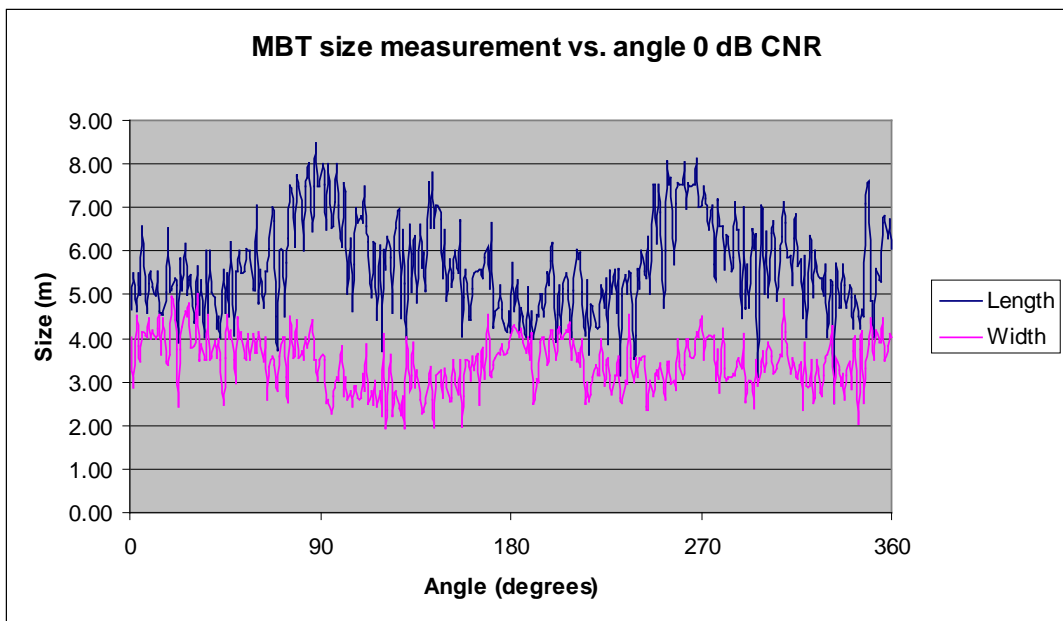
Figure 5-35 Size vs. angle for an MBT target with 0 dB CNR

For the very high noise cases, - 6.8 dB and -10.2 dB CNR, the target begins to lose its defined outline, and the length and width measures become similar. There are still exceptions at the cardinal points, where the increased target brightness results in a clear separation between length and width.
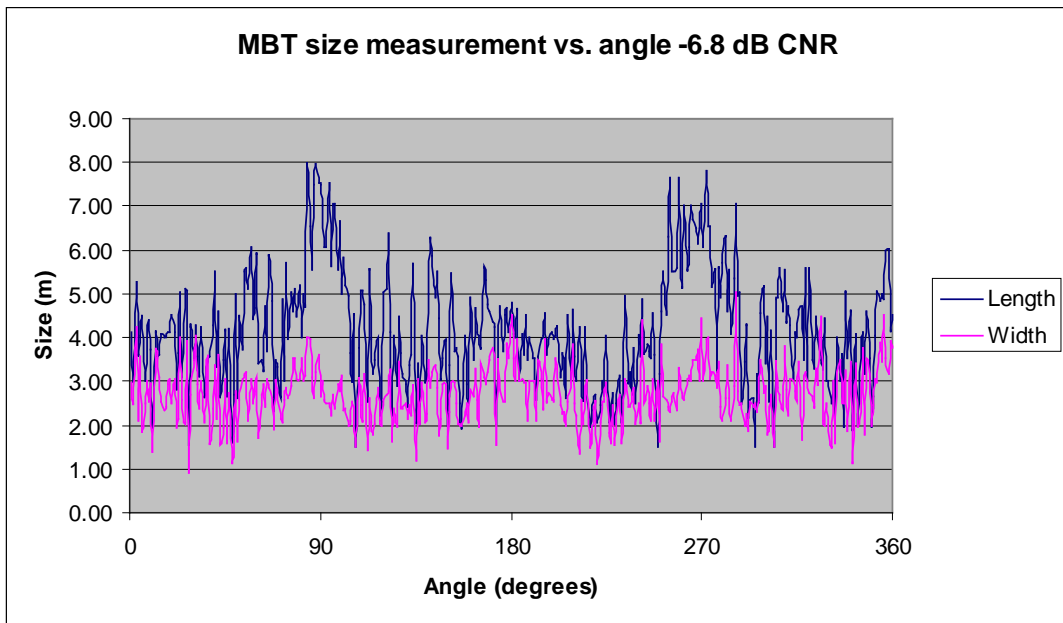
Figure 5-36 Size vs. angle for an MBT target with -6.8 dB CNR
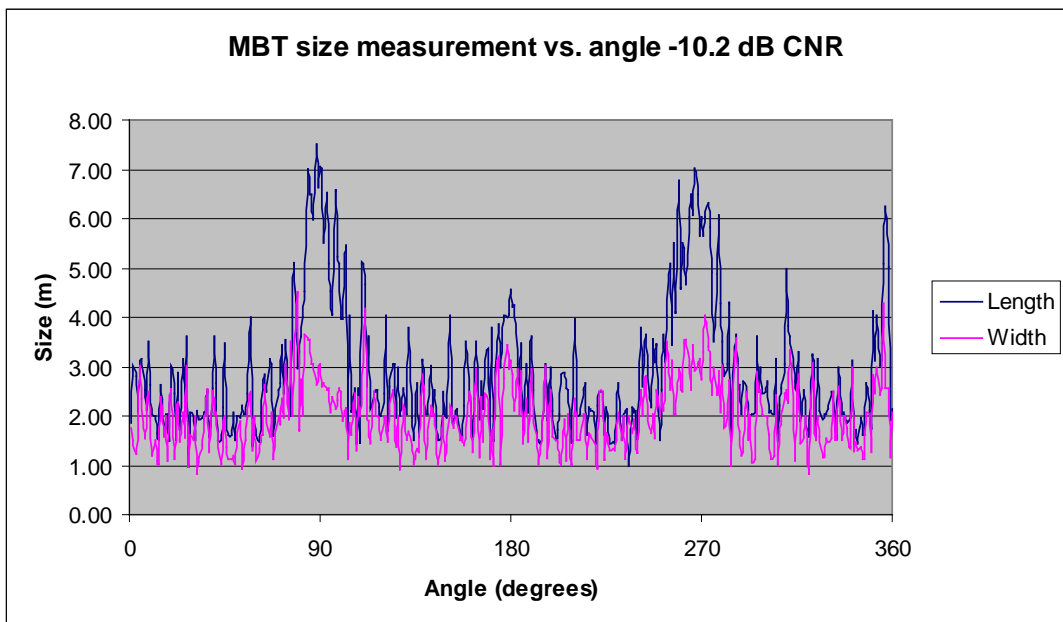
Figure 5-37 Size vs. angle for an MBT target with -10.2 dB CNR

Less detailed size information gathered from target attributes will reduce the quality of output passed to the target acquisition process. Depending on the level of targeting information available, a detection and its position may be sufficient information to select a target. Comparing size information directly across the CNR levels, at poor CNR levels the target is effectively a small bright object with little discernible shape, as in Figure 5-38 and Figure 5-39.
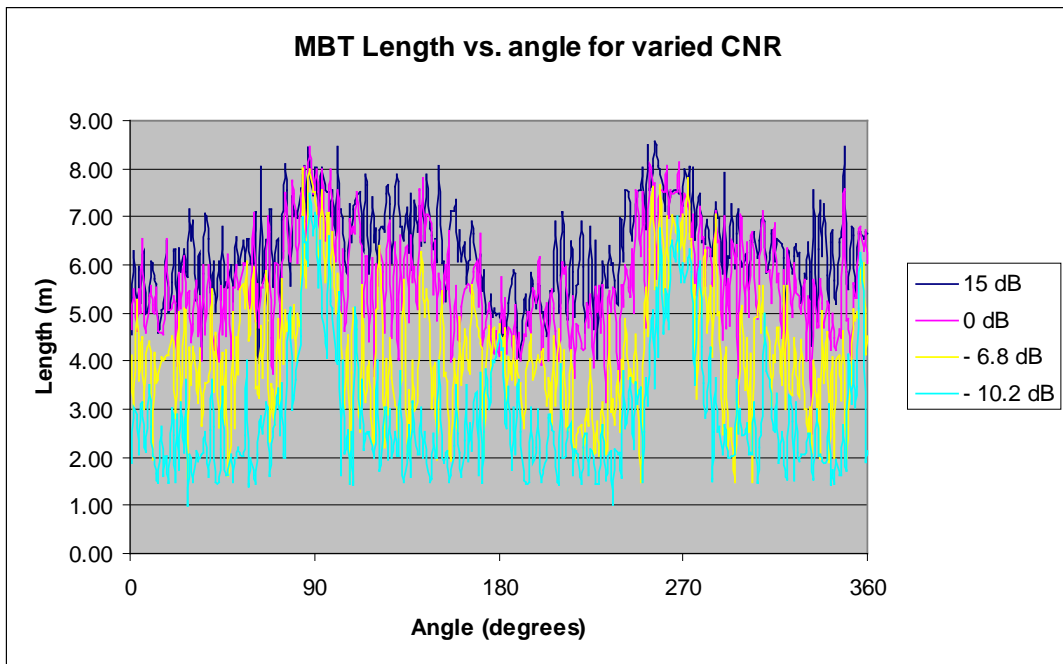
Figure 5-38 Length measurement vs. angle for MBT, varied CNR
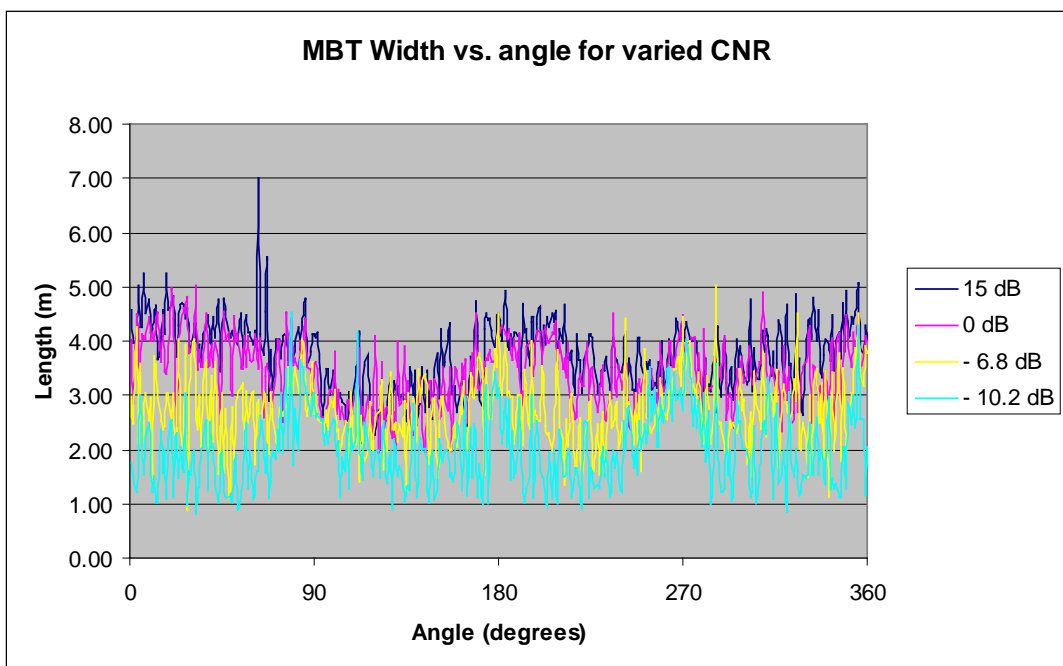
Figure 5-39 Width measurement vs. angle for MBT, varied CNR

This process was repeated with a Land Rover target. This target has a lower average TCR than the armoured target (approximately 14 dB), and a smaller size at 4.5 m by 1.8 m. The Land Rover target shows an asymmetric power profile, with a peak towards tail-on.
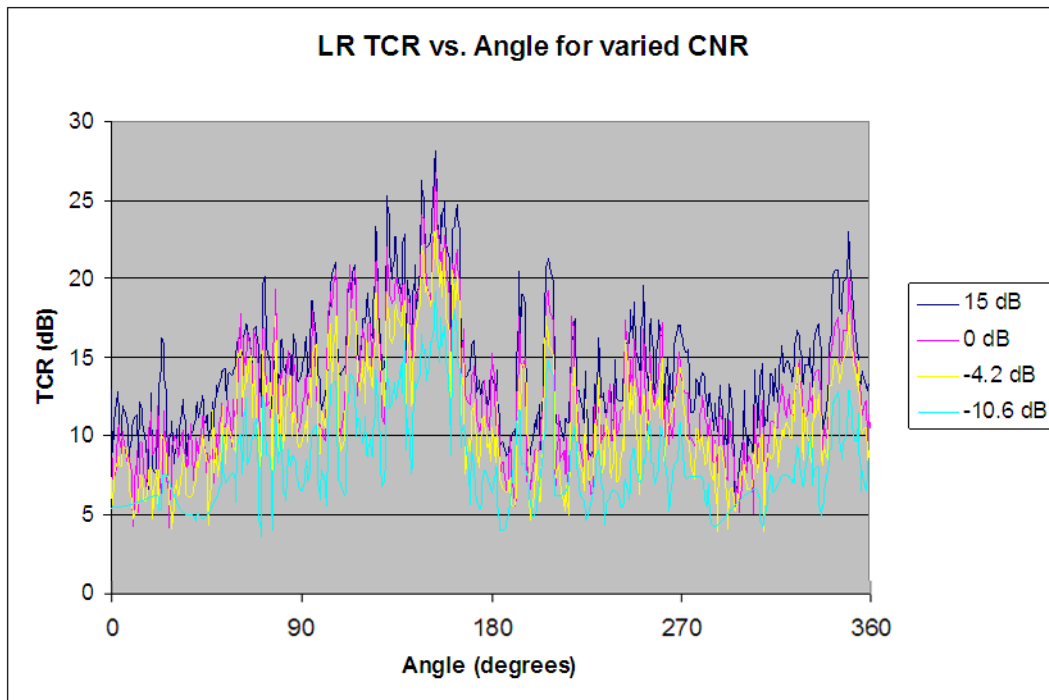
Figure 5-40 TCR vs. angle for a Land Rover target at varied CNR

With this smaller and lower-contrast target, appreciable numbers of signatures were missed by the detection process as the noise floor increased, as shown in Table 5-6. As noise levels rise above the clutter floor, many examples of the target are not detected despite the simple scene, and at very high noise levels, 35% of the signatures are not detected, as well as having much reduced segmented size information.

TCR (dB)	Not detected	Detections
15	1	359
5.4	1	359
0	5	355
-4.2	20	340
-6.9	59	301
-10.6	127	233

Table 5-6 Target detection capability at variable CNR

Size measurements for this target as the CNR degrades are shown in Figure 5-41 to Figure 5-44. Where measurements are missed due to a failure of the algorithm to detect the target at high noise levels, the line connects the two closest data points.

For the Land Rover, the size measurement capability is less accurate than for the MBT target, and with its lower TCR the degradation with increased noise is more severe.
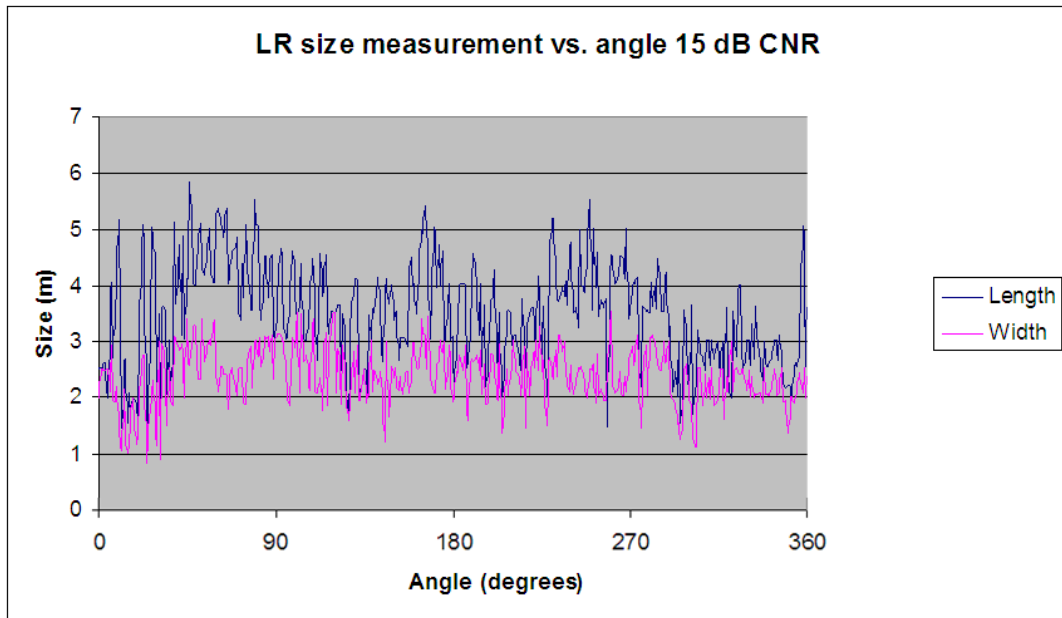
Figure 5-41 Size vs. angle for a Land Rover target with 15 dB CNR
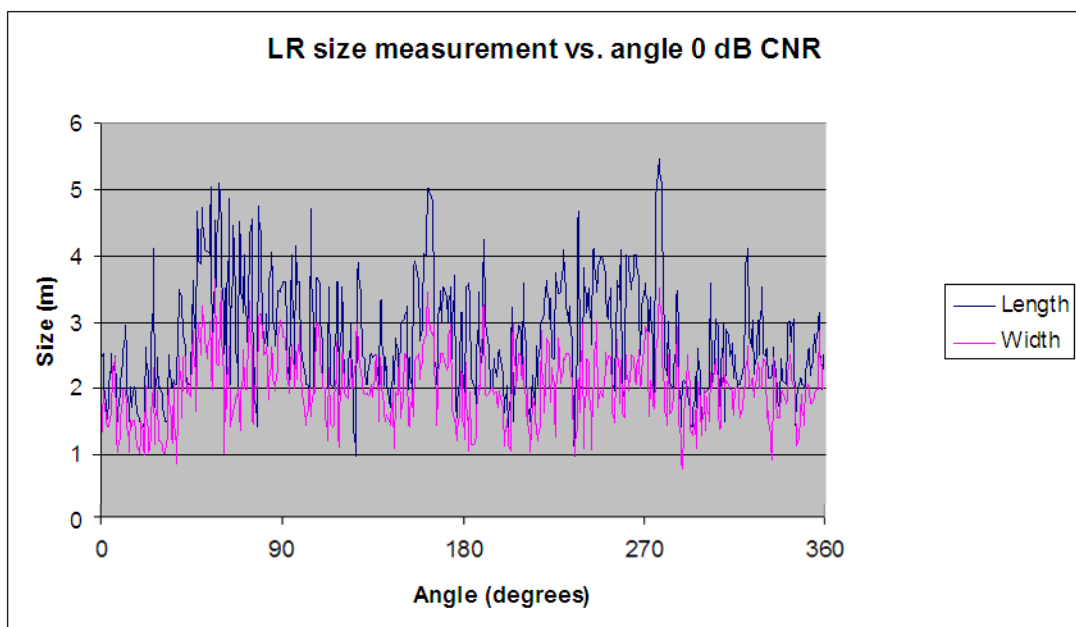
Figure 5-42 Size vs. angle for a Land Rover target with 0 dB CNR
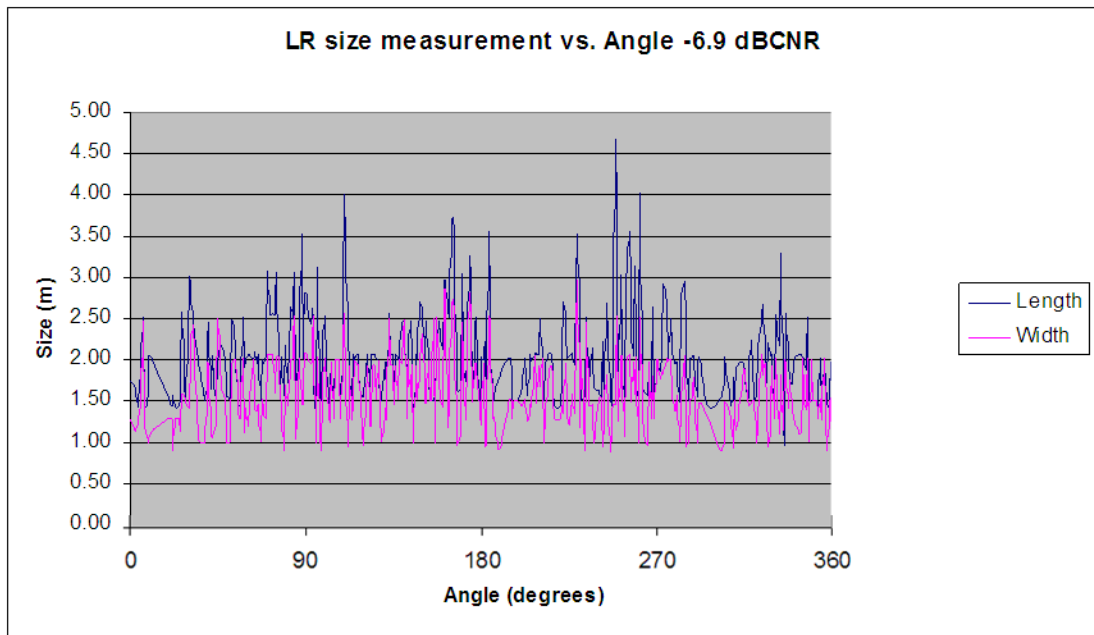
Figure 5-43 Size vs. angle for a Land Rover target with -6.9 dB CNR
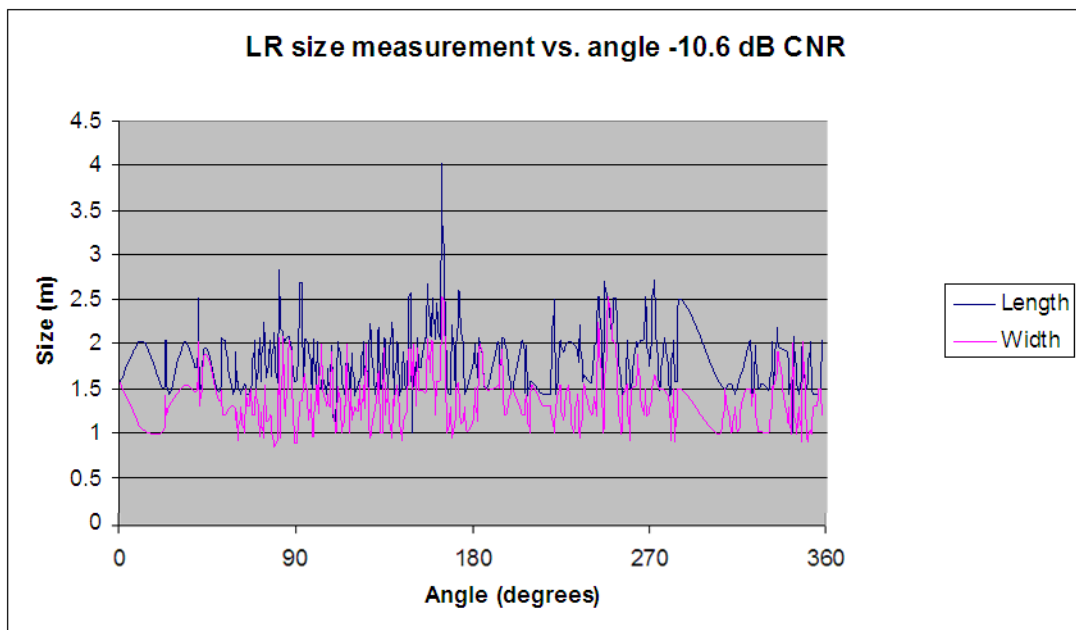
Figure 5-44 Size vs. angle for a Land Rover target with -10.6 dB CNR

5.6 Seeker Sensitivity Conclusions

The effects of decreased sensitivity affect two aspects of the ATR algorithms:

- The detection processing, i.e. whether the target is still bright relative to the local background;

- The segmentation and attribute measurement, i.e. whether the characteristics that separate targets from clutter are retained.

If the seeker is required to acquire a target with reasonable contrast against the background, detection performance remains strong down to levels where the noise floor exceeds the clutter level. However, the quality of attribute information rapidly degrades before this point. Should acquisition of a lower contrast target such as a soft vehicle be required, reductions in sensitivity are less acceptable.

If the requirement is simply for detection and position measurement, with little emphasis on attribute information, then low seeker sensitivity levels may be tolerable.

6 Object Shadow Information

6.1 Introduction

While the use of the measurements of static targets size when imaged with a radar seeker to provide information to a CDAA process is a well-developed technique, the equivalent process for moving target signatures poses a greater challenge.

Synthetic Aperture Radar (SAR) techniques assume the ground is stationary, and the Doppler frequency shift due to the platform velocity varies across the width of the beam. Additional Doppler shift from the radial motion of a target results in the signature being displaced from its true position in the RDM image. Blurring of the target's signature also occurs. This is due to motion during the dwell as it moves between range and frequency cells, and from radial acceleration which smears the signature over multiple frequency cells.

While the target is displaced, its shadow will remain at the target's true location in the scene. Unlike the target returns, this shadow will not be blurred due to motion effects. The shadow, therefore, has no need to be focused since it is an absence of information, and assuming the nearby clutter is well focused, the shadow will exhibit a clear outline. However, the motion of the object during the imaging period can degrade the contrast towards the edges of the shadow.

A shadow detection process has been designed to operate following cueing from a moving target indication algorithm, capable of detecting a moving target that may produce a shadow. Work performed on shadow detection for moving targets in SAR imagery needs to begin with a search for any suitable shadows, since the duration of the imaging dwell can be such (around 30 s) that the target signature effectively disappears through blurring. However, due to the short dwell of a radar seeker (typically less than 0.5 s) the target, while blurred, remains detectable and, hence, shadow analysis can be an additional post-detection technique.

Information available from the position and shape of the shadow can potentially aid the target acquisition process in a number of ways:

- Shadow width and length relate to the target's cross-range dimension and height;
- Shadow location can give the true position of the moving target in the scene with greater accuracy than a monopulse estimate [47];
- Shadow location from multiple images can be used for velocity estimation.

This extra information would be of benefit to seeker performance. However, shadow information can only be used in scenarios where the background is sufficiently bright relative to the noise or interference level to produce a shadow with good contrast, for example grassland clutter rather than tarmac. Because of this, the use of shadow information can only be an additional approach in the seeker processing, rather than a required aspect of the detection process.

6.2   Shadow segmentation and measurement

Shadowed areas in a seeker RDM image, created by moving targets, need to be identified in order to analyse the shadow and infer information on the attributes of the shadow's source.

When approaching this detection problem, two options exist upon which a shadow search can be based. These are:

- To search for all shadows in the scene and identify whether each lacks an obvious source (the target is displaced);
- To cue a search from MTI detections.

At an early stage in the research a decision was made to cue the search from MTI detections. A set of image-based MTI algorithms for radar seeker imagery have been developed by QinetiQ, and are able to achieve a high probability of target detection for a reasonable false alarm rate. Image-based MTI schemes have also been developed by FGAN [62]. Using these detections, a shadow search can then be cued to investigate the scene further. Starting the detection process with a search for any shadows in the imagery gathered by the seeker would be more computationally expensive. Should any shadows be detected where an MTI detection was not present, such as from dead ground or faint rural clutter, the uncertainty in determining that a valid, but undetected, target is the source of the shadow would be too high for the information to be of use. Equally, a target could be displaced back to its true position in Doppler, directly associating the shadow with an object. An approach has been developed by Jahangir and Rollason [38] which uses multiple images over a period of time to detect movement of the shadow. However, an extended image formation process may not be available to a seeker sensor, hence the detection operates on a single image.

The moving target detection processes are able to identify the true position of the target to the accuracy of the monopulse measurements from its signature. The nature of the exact processes used to detect the moving targets initially are not further described herein. All that needs to be provided to the shadow detection

routine is a target range, possibly supplemented by an approximate cross-range position from the monopulse capability. The RDM imagery can then be investigated at this range.

Shadow analysis is considered to be equally applicable to providing information on static targets, for example [37], [63] and [64]. However, since static targets already provide good attribute information, the bulk of the research described here has focused on moving target cases.

6.2.1 Shadow analysis process

Figure 6-1 shows the shadow analysis processing chain that has been developed. A shadow search can be cued by an MTI detection, selecting potential shadows before the application of a segmentation algorithm. A coordinate transformation process accounts for the variable resolution of the seeker and converts from slant-range to ground-range before features of the shadow are measured.

```
┌──────────────────┐
│ Potential Targets │
└──────────────────┘
         │
         ▼
┌──────────────────┐
│ Shadow Search    │
│ and Selection    │
└──────────────────┘
         │
         ▼
┌──────────────────┐
│ Shadow           │
│ Segmentation     │
└──────────────────┘
         │
         ▼
┌──────────────────┐
│ Coordinate       │
│ Transformation   │
└──────────────────┘
         │
         ▼
┌──────────────────┐
│ Shadow Feature   │
│ Analysis         │
└──────────────────┘
```
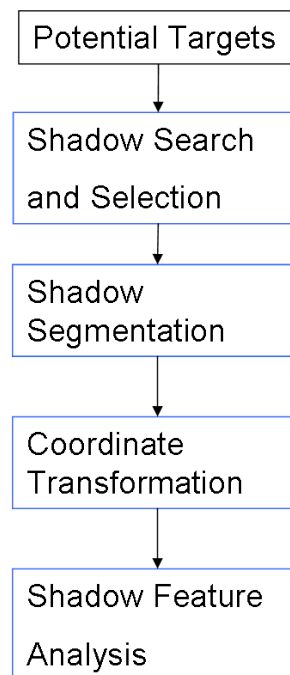
Figure 6-1 Processing chain for shadow analysis

The search, segmentation, coordinate transformation and feature analysis stages are described in sections 6.2.2 to 6.2.6.

6.2.2 Shadow search

The shadow detection process starts with the assumption that shadow analysis will be cued from a bright target detection, and uses the location of that detection to initiate a localised search within the image.

When an MTI detection is declared at a coordinate within an RDM, a section of the image with a set extent up-range and down-range of the target is extracted, such as that shown in Figure 6-2 (Real data from 2006 MRMS air-carry trials data). This section includes the complete cross-range or Doppler bandwidth of the RDM along the *x*-axis. The initial target detection is visible to the right of the image at approximately position 450 on the *x*-axis and 40 on the *y*-axis (down-range), with shadows visible both at the target's true position (140, 50), a tree (250, 55-110), and two trees for which the bright returns are above the sampled area, (340, 0) and (375, 0). The target has been displaced out of the antenna beam, so is observed against a noise background.
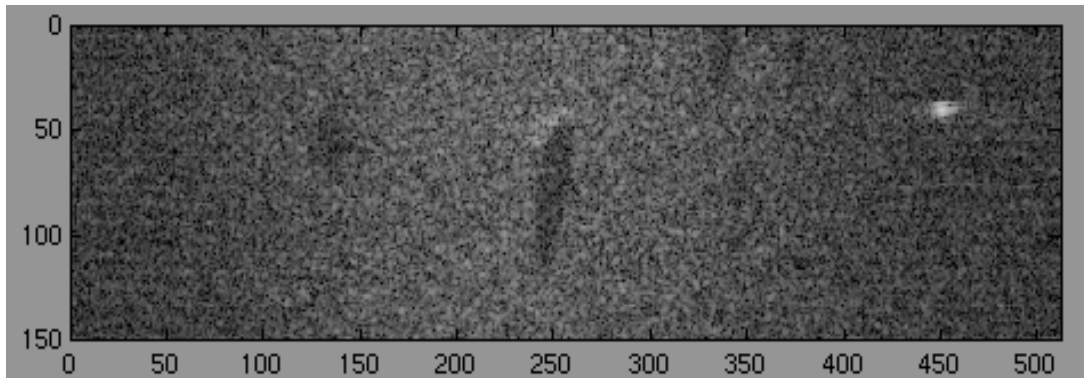
Figure 6-2 Sub-section of RDM for shadow detection. Range down the image, cross-range from left to right

The detection processing begins with the segmentation and removal of the bright target returns to prevent them from affecting the adaptive thresholding that follows. This thresholding is designed to pass almost all pixels that stem from clutter, leaving gaps in the detection map where the radar returns are low relative to their local background.

A process of dilation and erosion then fills small gaps in the detection map, before it is filtered to leave a number of 'dark' clusters where the returns are likely to come from noise (i.e. shadow). Clusters that are very small or large are rejected, leaving only those shadows that could potentially be created by vehicles of a set range of sizes. The result of this process being applied to Figure 6-2 is shown in Figure 6-3.
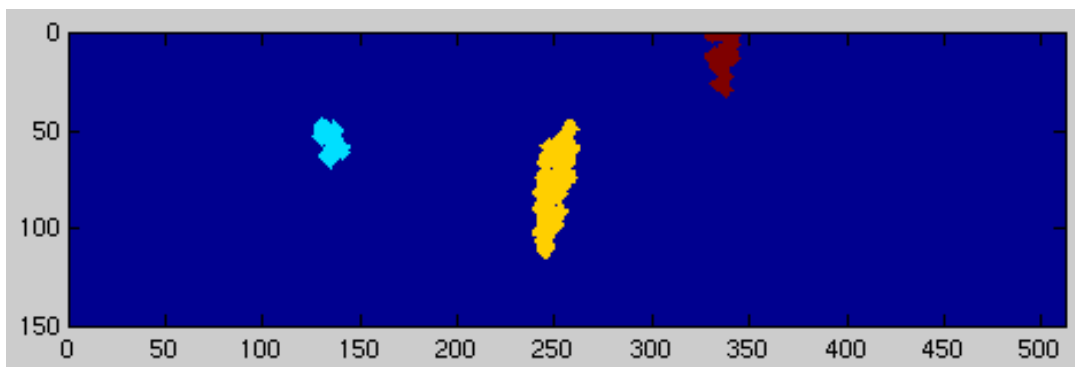
Figure 6-3 Potential shadow objects

Given knowledge of the target's location, the red shadow, detected closer in range than the target, can be rejected. Of the two remaining objects, the size of the blue return is a closer match to the range of expected target heights than the yellow shadow, and is correctly selected as the appropriate shadow to analyse, shown in Figure 6-4. The shape and position of this shadow can then be used to provide information on the target.
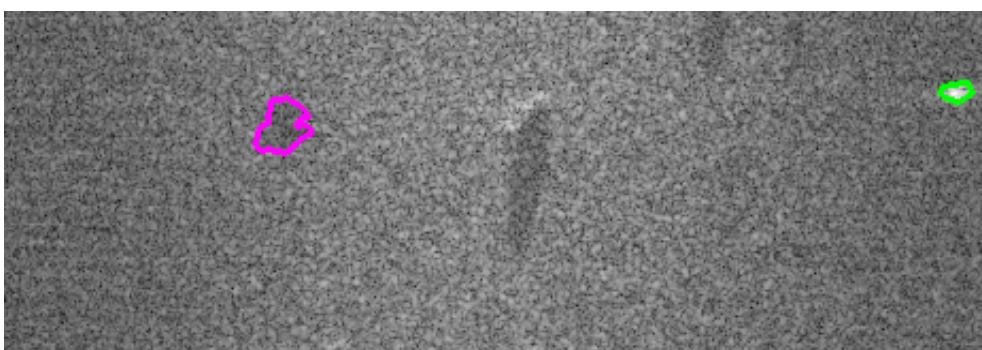
Figure 6-4 Correct shadow selected for analysis

For the range of moving targets against a rural clutter background imaged in air-carry trials, where the potential confusers are trees, bushes, roads and tracks, this shadow detection approach works well in all cases where a reasonable CNR exists (i.e. the shadow retains good contrast and is not too close to the edge of the beam).

In this example, the entire Doppler (cross-range) width of the RDM has been searched for shadows since, for this study, the monopulse information contained in the imagery that can relocate the target to its true position has not been used. When searching the full extent, the shadow detection performance is high.

The angular accuracy with which the moving target position can be measured, $\delta\theta$ is related to the antenna beamwidth, θ_B, and the signal to noise ratio of the measurements.

$$\delta\theta = \frac{\theta_B}{\sqrt{2.SNR}}$$

<div align="right">(6.1)</div>

The use of monopulse sum and difference channels allow this to be improved such that $\delta\theta/k$ may be obtained where k is the gradient of the monopulse difference / sum slope near the beam centre [65] [66]. A term can then be applied to represent a loss in accuracy as the target moves away from beam centre, D/S which represents the normalised ratio of the target's returns.

$$\delta\theta = \frac{\theta_B}{k\sqrt{2.SNR}}[1 + D/S]^{0.5}$$

<div align="right">(6.2)</div>

A well designed system, with reasonable values for SNR should be able to calculate the true position of the displaced MTI signature with a standard deviation of 5 to 10 metres in the cross-range direction, thereby providing a second source of information to aid the shadow search that will be beneficial in more cluttered environments.

6.2.3 Shadow segmentation - Adaptive contour

A key focus of this work is an assessment of the ability to segment shadows from the local background using techniques that have been developed during previous research programmes. The segmentation technique currently implemented in the ATR process is an adaptive contour technique, referred to as 'Snake segmentation' [53], [54]. This has proven to be a robust technique for segmentation of bright objects in an RDM from the local clutter background, and has the capacity to segment areas that are dark compared with their local background.

This approach has been tested with shadows of static targets, indicating an ability to segment both the target and shadow. The segmentation of a static self-propelled gun (SPG) and its shadow is shown in Figure 6-5. Here, the stages of segmentation are visible from the initial seeding on a bright CFAR detection, followed by a bounding rectangle and then a full segmentation. Down-range is from left to right in the image.
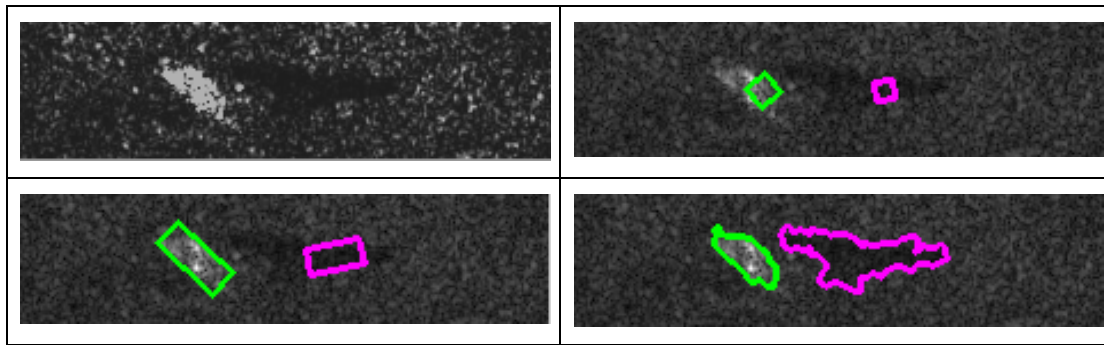
Figure 6-5 Raw image of SPG (top left), three stages of segmentation – initial seed (top right), bounding box (lower left), full segmentation (lower right)

A configuration suitable for segmentation of shadows in MRMS imagery was developed, and applied to trials imagery of moving targets and their shadows.

6.2.4 Shadow segmentation – Model based

A second approach to segmentation was briefly investigated. It uses some of the functionality of the snake segmentation from section 6.2.3, combined with a simple wire-frame target model with variable dimensions to produce simulated shadows. This approach was originally designed by QinetiQ for providing height, length and width information for buildings in high-resolution SAR imagery [67], but the same principles are applicable to the seeker imagery.

In this approach, a simulated shadow from a model is compared with a selected shadow in the image. The model that created the shadow represents the object in the scene, and can provide information on its dimensions, particularly height which is not normally directly measurable from the bright signature.

The parameters of the model are moved in a similar manner to the nodes of the regular segmentation, and with each move the quality of the match between the simulated shadow and the real imagery is checked until the minimisation criterion indicates that the optimum delineation for the image has been obtained.

A solution is found when further node movements do not produce any improvement in the correlation between simulated shadow and the scene shadow, although that can be a local solution and is not necessarily optimum.

The left image in Figure 6-6 shows an example SAR image of a building and its shadow. The image on the right is an illustration of the simple model used when analysing building shadows and the shadow it may cast.
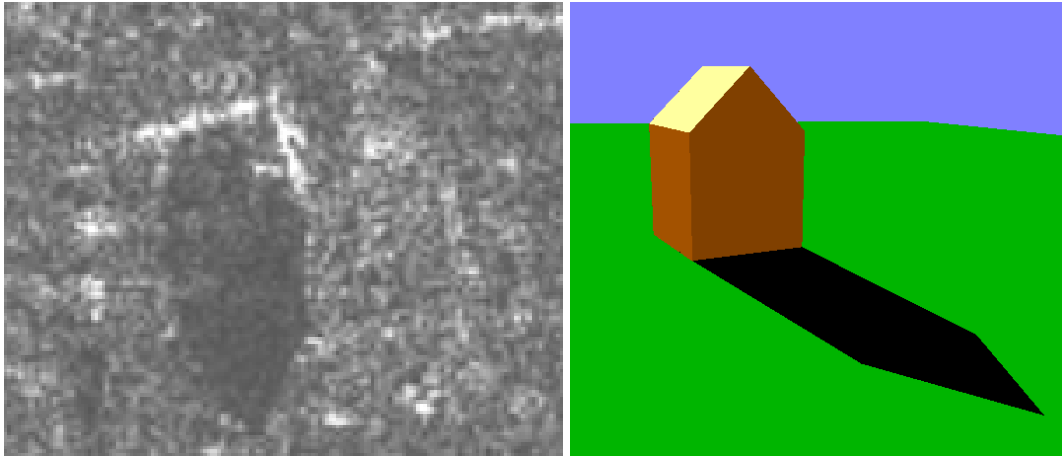
Figure 6-6 Building and shadow from SAR image (left), model for urban scene analysis (right). Down-range from top to bottom, cross-range from left to right

This approach was tested on data from the MRMS air-carry trials. An example target segmentation is shown in Figure 6-7.
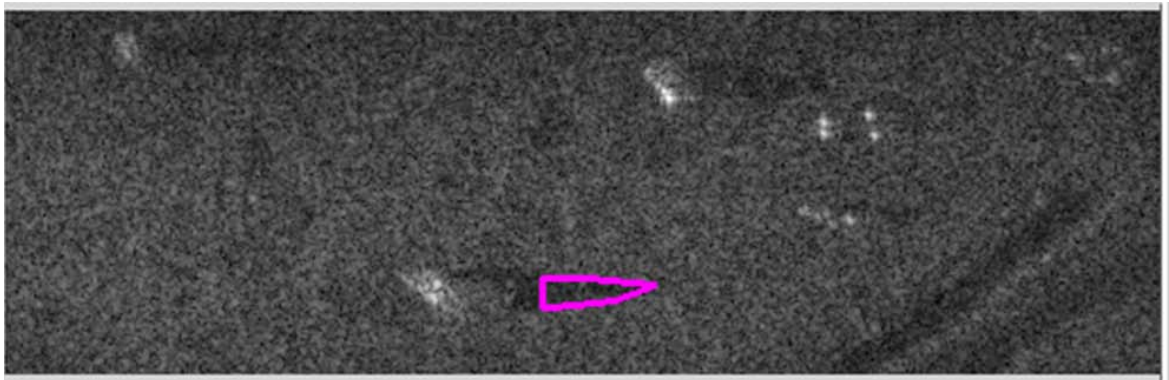
Figure 6-7 Example shadow segmentation using model-based technique. Down-range from left to right, cross-range in vertical direction

The segmentation has excluded the near part of the shadow that may include overspill from the bright target returns, but has segmented the far extent of the shadow reasonably well.

The model that best suits this segmentation was generated, and the parameters that define that model are shown in Figure 6-8. The model used had the degrees of freedom associated with the building model for which the code was developed, which is why the segmentation is more pointed than would be expected. For optimal performance, modification to better suit a target class would be beneficial.

Length (m)	9.4
Width (m)	4.0
Height (m)	2.0
Roof Height (m)	2.5
Orientation (deg)	70
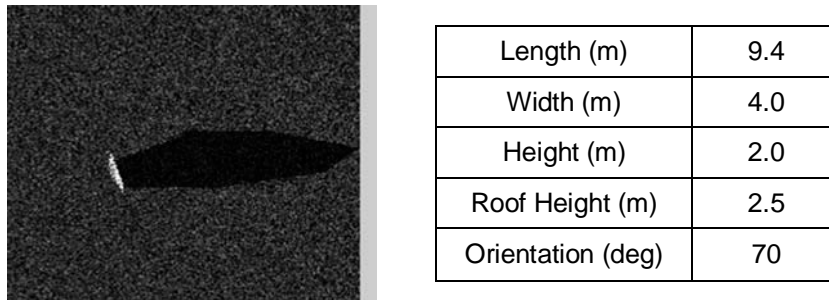
Figure 6-8 Model and simulated shadow (left), model parameters (right). Down-range from left to right

While the parameters produced for the target are not far from the true values, the orientation of the object is incorrect. For the static target case this could be constrained through the use of information on the bright target signature but, for the moving target case, orientation measurement from a single image is not possible, and multi-frame tracking would be required, during which the target's shadow may vary.

The model-based approach was designed for high-resolution SAR imagery in which multiple images of a static target can be produced over a range of aspect angles. For the seeker case, imagery is only considered at a single angle or at best a narrow range of angles based on the trajectory, and the ability to be able to produce multiple images is not guaranteed.

The lack of high-resolution data from multiple angles results in the model being over-fitted to the data, having more degrees of freedom than the data can support. For the moving target case and simple static target cases, the information measures that will be described in Section 6.2.6, in which the shadow is used to determine the target's cross-range extent and height, contain all the information that can sensibly be extracted from the signature.

Should the seeker engage a static target at a known location, and have the capability to fly a trajectory that would give a broad range of angles, then this technique may provide additional information. However, the additional processing and demands on the trajectory are unlikely to warrant the benefits of the approach, and so it was not investigated further during this research.

6.2.5 Coordinate transformation

To correct the shadow for the variable resolution, the node positions of the segmentation were passed through the transformation process described in section

5.1. Figure 6-9 shows the segmentation node positions in white applied to an extract from the full RDM. It is clear that the segmentation agrees with what the user would perceive as the blurred target and shadow outlines.
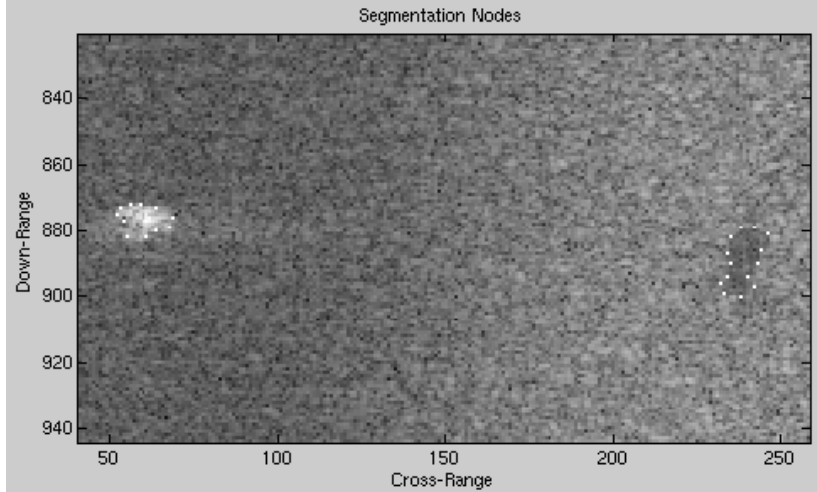
Figure 6-9 Segmentation nodes highlighted on target and shadow areas

Within the constraints of the flat earth assumption, which is reasonable for the trials data, shadow segmentations are converted to ground coordinates, thereby correcting for the slant range and variable resolution.

6.2.6 Shadow analysis

Once shadow segmentation has been applied, estimates of the cross-range width and height of the source object can be derived.

The cross-range width represents the dimensions of the target in the cross-range direction, rather than the actual width of the target. Converting this measurement to real dimensions would require either the object to be at a known head-on or broadside angle to the radar, or for both the orientation and one of the other dimensions to be known. Figure 6-10 shows the contribution of the two sides of a rectangular target to the perceived cross-range shadow width for a target that is facing in the same direction the seeker is flying. The shadow's cross-range width, W_{CR} is the sum of the target length, L_T, and width, W_T, components.

$$W_{CR} = L_T \sin(\theta) + W_T \cos(\theta) \tag{6.3}$$
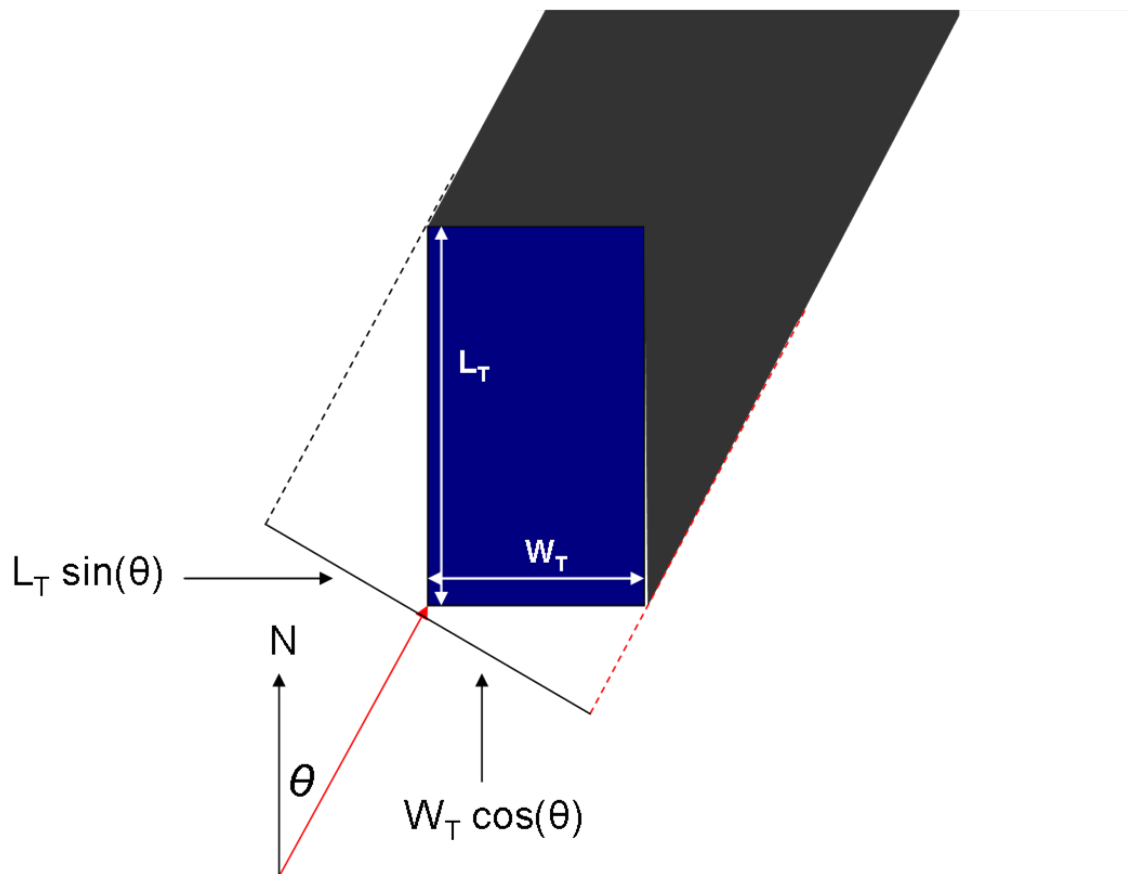
Figure 6-10 Combination of target length and width to measured shadow width

The target height is the greatest distance measured when sampling the down-range dimension of the shadow across its width, illustrated in Figure 6-11 for a target side-on to the radar. Assuming the terrain is flat and the depression angle is known, the height of the target can then be calculated. The effect of terrain slope is discussed in section 6.6.7.
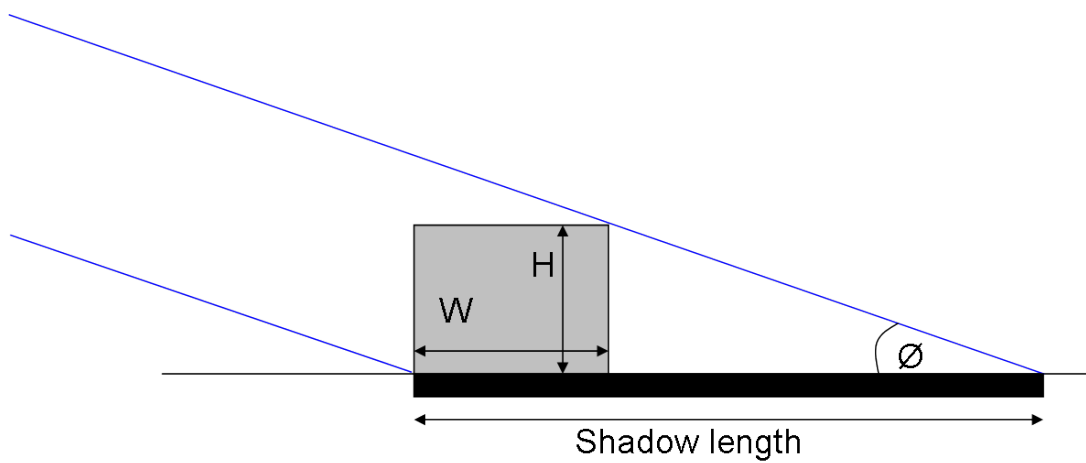
Figure 6-11 Combination of target width and height to measured shadow length

The contribution to the shadow length (and, therefore, target height) of the target's obscuration of the ground was removed. For calculation of the shadow length, the near-range extent of the shadow (pink outline) was limited in line with the far-range extent of the returns from the bright signature (green outline), illustrated in Figure 6-12, where range is down the page.
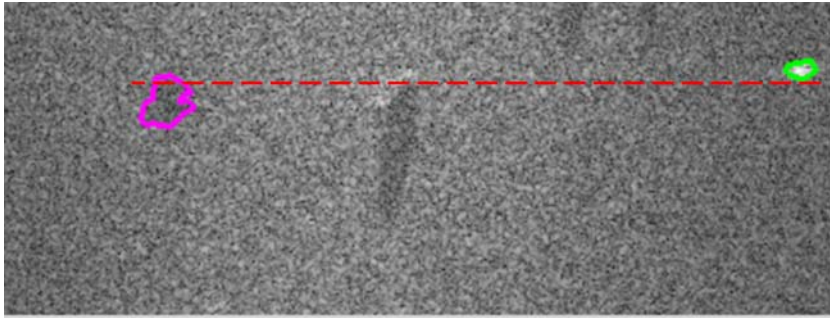
Figure 6-12 Illustration of bright signature extent used to limit shadow length measurement. Range down the image, cross-range across

6.3 Software implementation

In order to be able to analyse all available data, it was necessary to build a tool in which shadow detection, segmentation and measurement techniques could be tested.

Matlab code was constructed to perform the following functions:

- Input data from a range of sources;
- Examine regions of imagery in detail;
- Segment dark shadows and bright targets;
- Compensate for variable image resolution by transforming coordinates to a ground plane;
- Measure attributes of shadow regions;
- Apply the CFAR-based shadow detection and clustering scheme (Section 6.2.2);
- Search for potential shadows and select the one most likely to relate to an MTI detection;
- Measure clutter to noise ratio and compare shadow contrast in different polarisations.

The interface is shown in Figure 6-13, containing an example RDM on the left, an extract from that map (top right), and target and shadow segmentation applied to the extract (lower right).
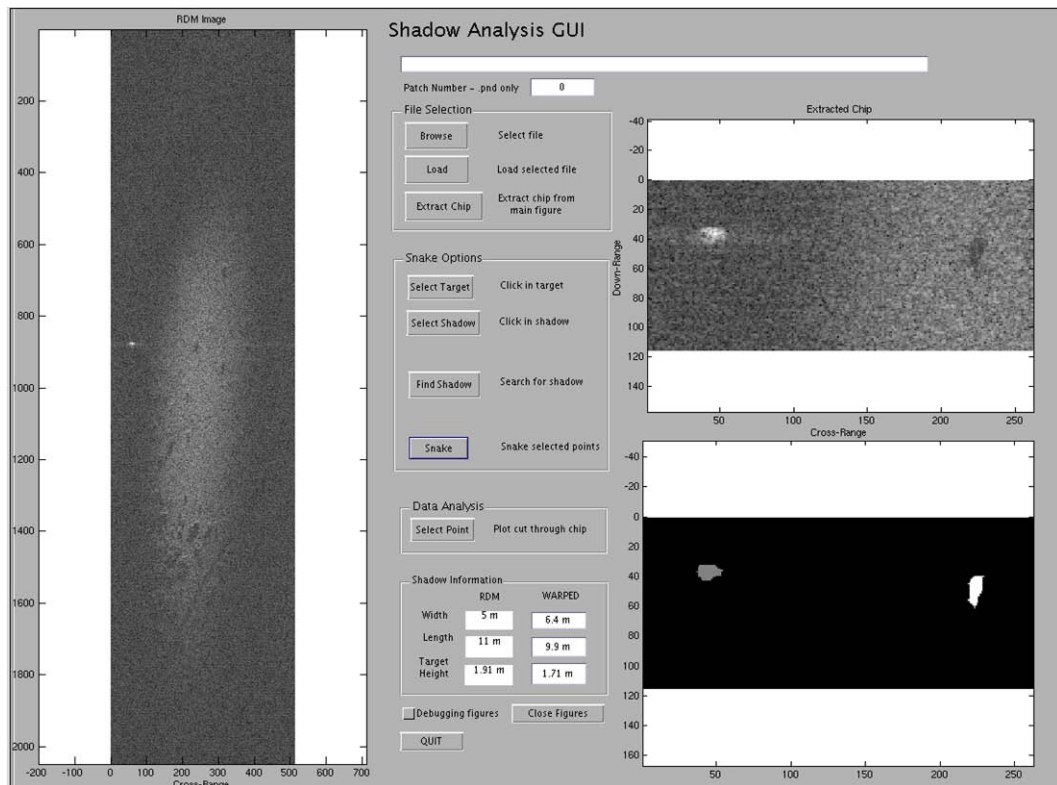
Figure 6-13 Shadow analysis GUI

Through the combination of these functions, the software was used to generate the shadow attribute information detailed in sections 6.4 and 6.5.

Signal power across the shadow in the cross-range and down-range directions can be examined as shown in Figure 6-14. The blue lines represent an extract of single pixel width, whereas the red lines are averaged over 9 pixels. This example shows a shadow with a CNR of approximately 11 dB. Empirical evidence suggests that a CNR of 8 dB or greater is necessary for accurate segmentation of a shadow outline.

121
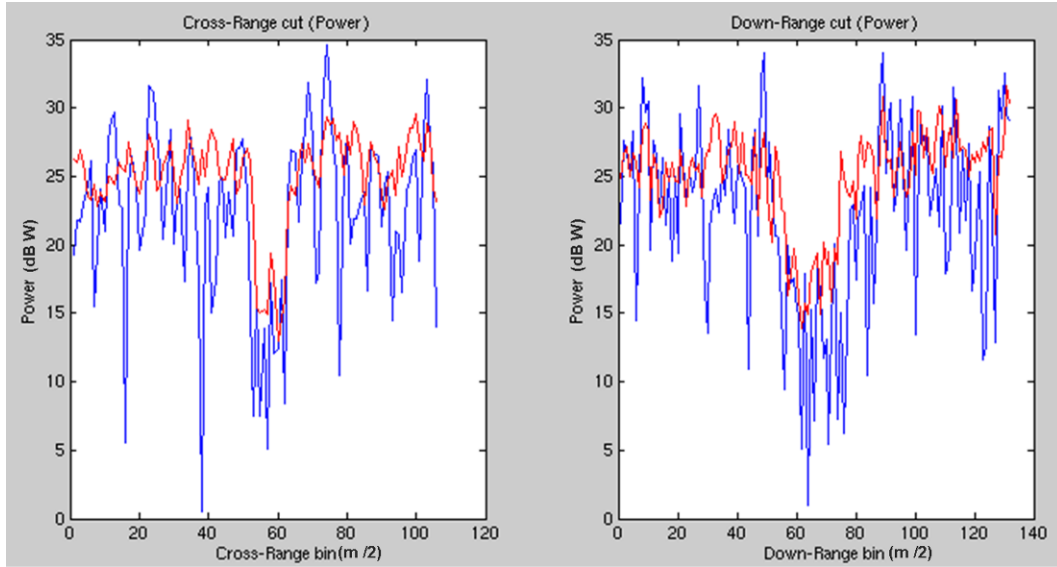
Figure 6-14 Cross-range and down-range signal power through a shadow left by a moving target. Blue plot is for single pixel width, red plot is an average over 9 pixels

6.4 Data analysis and results

6.4.1 Air carry trials data

Data gathered with the MRMS radar during the air-carry trials of 2006 provided a representative dataset to develop and test shadow analysis algorithms. The ability to segment shadows and gain useful information from this data was investigated to demonstrate the utility of the approach for future seekers with a similar or better CNR.

Two moving vehicles were imaged a number of times during the trial, a Main Battle Tank (MBT), Figure 6-15, and a civilian Range Rover configured with an array of trihedral reflectors on a substantial roof rack, Figure 6-16. These reflectors are not relevant to this study on shadows, but were placed there as a reference for use in the development of moving target focusing algorithms.

Figure 6-15 Moving MBT with Range Rover visible in the background to the left

Figure 6-16 Range Rover with trihedral reflector array and MBT visible in the background

As shown, the targets were located in an open area of grassland. Uninterrupted line-of-sight was always available to the vehicles, and some rural clutter was present in the form of scattered trees and bushes.

Three sorties were flown on separate days over the trials location, with a total of 29 imaging flights over the moving target locations. The data gathered in these flights were analysed to extract all available examples of moving targets for the test of shadow analysis algorithms.

Example imagery of the moving MBT and Range Rover is shown in Figure 6-17. The displacement of the bright return from the target's true position where the shadow remains can be observed, as well as blurring due to the target's movement and acceleration.
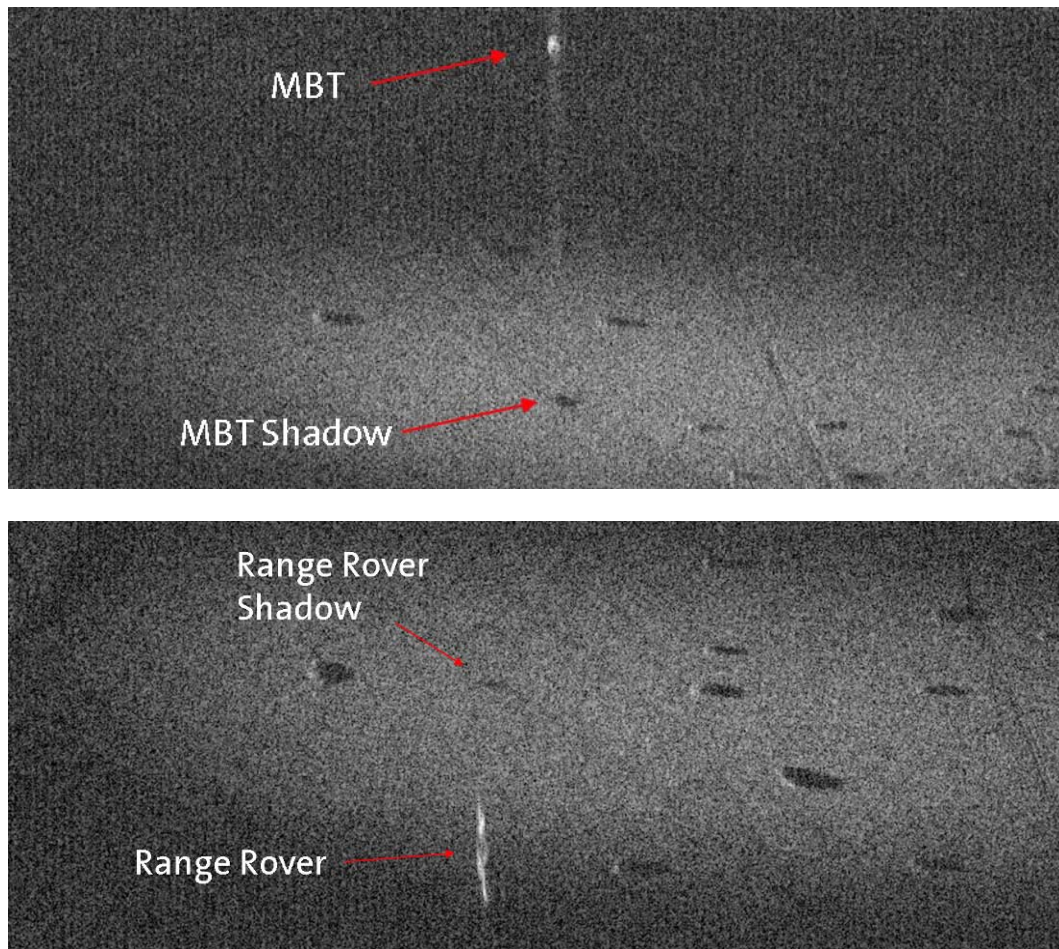
Figure 6-17 Moving MBT (top) and Range Rover (bottom). Down-range from left to right

In the case of the Range Rover, the array of trihedral reflectors on the vehicle roof strongly reflect the 'odd bounce' or cross-polar channel of the circularly polarised radar, producing extremely bright returns in this cross-polar imagery. When combined with the vehicle's movement, this leads to greater blurring of the signature than is seen for the MBT target. These bright signature effects are not considered here, having no effect on the shadow in the imagery.

6.4.2 Data available

While the seeker over-flew the target array on 29 imaging runs, inaccuracies in the aircraft flight path, and incomplete ground coverage in the scan, reduced the data yield. In many cases the target was imaged too far from the beam centre, resulted in a reduced yield of data suitable for analysis. Only 15 imaging runs provided useful data of either the MBT or the Range Rover, where the target was within the main beam of the radar. While this may be seen as a concern for the use of this approach, a system with adaptive scan control would have a scan pattern sufficient

to produce RDMs of the search area that provide 100% coverage within the 3 dB beamwidth.

At the centre of an RDM containing grassland clutter, the CNR is approximately 15 dB, and good shadow contrast is visible.

Moving away from the centre of an RDM in the cross-range direction, the clutter returns fall with the shape of the beam while the noise remains constant. For static targets that are likely to have a TCR of the order of 12 to 20 dB, detectability remains high even at the point where the CNR has fallen to zero (Chapter 5.4). For the segmentation of shadows, however, a requirement for a CNR ratio of greater than 8 dB is required. Given an appropriate seeker scan, with image overlap at the point where antenna gain has dropped by 3 dB, this should ensure that all parts of the ground will be adequately imaged. The seeker scan pattern used for the air-carry trials did not achieve such complete ground coverage, and changes in the height of the aircraft often placed the moving target towards the near or far range extents of the beam, where gain was again reduced.

Figure 6-18 shows the relationship between target and clutter power for objects moving away from the centre of the radar beam. The antenna pattern chosen for this example has a one-way Gaussian shape, with a beamwidth typical of that which may be expected from a millimetre-wave seeker. The expected signal level from clutter is shown in black, and for the trials imagery the 3 dB point is 35 m each side of the centre of the beam. The noise level is set at 15 dB below the peak clutter response, and a simulated target response is included at 15 dB above the clutter level. Given the requirement of approximately 8 dB of contrast for segmentation, the width across the beam at which accurate segmentation is possible for targets and for clutter is shown.
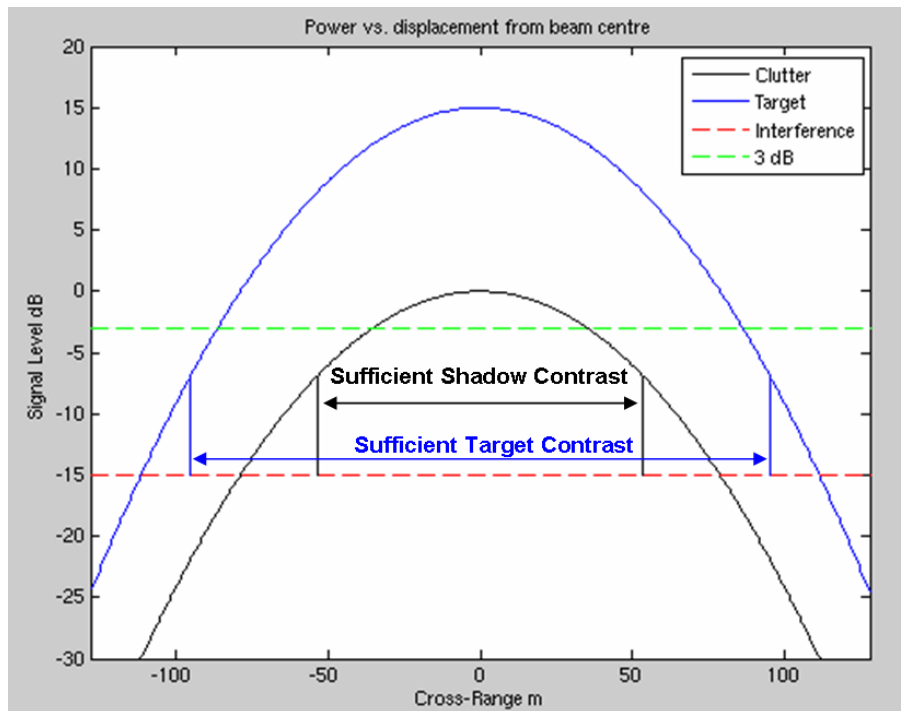
Figure 6-18 Power levels of target and clutter returns with cross-range position within the radar beam

For this system configuration, typical armoured targets should have adequate contrast to be segmented from the background in a strip 190 m wide, 23 dB from the beam centre, or approximately ± 2.1° from boresight. However, there will be sufficient contrast to segment shadows accurately in a strip only 107 m wide, 7 dB down from the beam centre, approximately ± 1.3° from boresight.

An example of a shadow imaged away from the main beam is shown in Figure 6-19. For this image, where the local contrast is reduced, the shadow is still detectable, and can be segmented from the background with reasonable accuracy. However, there are a large number of images where the location of the shadow is farther from the beam centre, and detection and segmentation become unreliable.
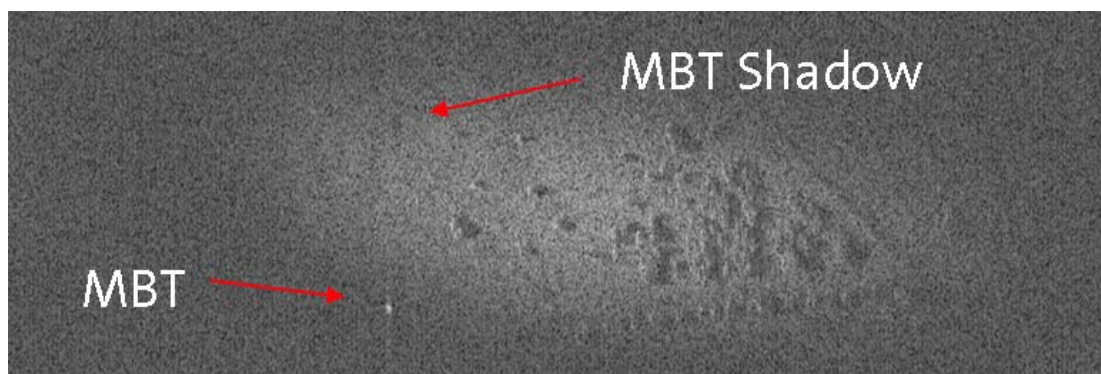
Figure 6-19 RDM with target shadow away from beam centre

To account for the variable quality of shadows, the moving target data was split into three classes:

- Not used   – no shadow, or shadow barely visible
- Acceptable   – reasonable shadow, reduced contrast
- Good   – well placed in the image, high contrast

The numbers of images that fell into each category are listed in Table 6-1. It is to be expected that a large number of potential images should be rejected for shadow analysis based on the reduced extent over which shadows remain detectable.

Target	Total MTI	Not Used	Acceptable + Good	Good
MBT	35	24	11	5
Range Rover	42	22	20	9

Table 6-1 Data yield for MBT and Range Rover from 2006 air-carry trials

6.4.3 Analysis of air-carry imagery

Target data was passed through the shadow analysis software to produce target width and height measures when transformed to the ground plane. The results for the combination of acceptable and good data are shown in Figure 6-20.
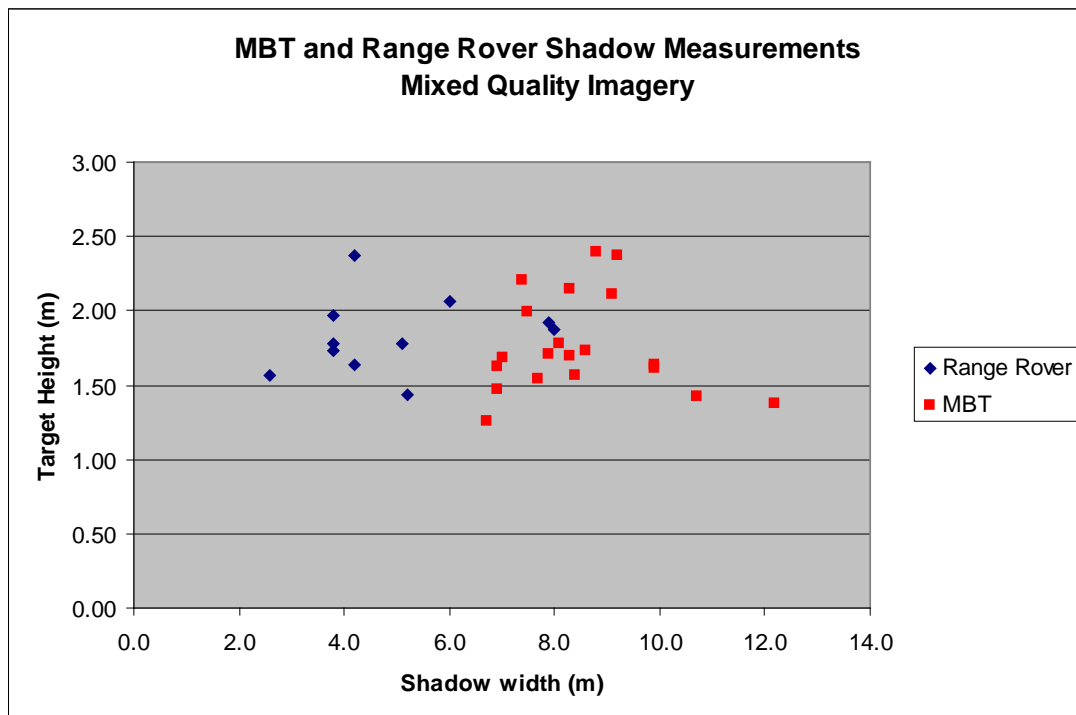
Figure 6-20 Measured dimensions of MBT and Range Rover targets, mixed quality imagery

If only the high-quality imagery is considered, the separation between target classes becomes more distinct, Figure 6-21.
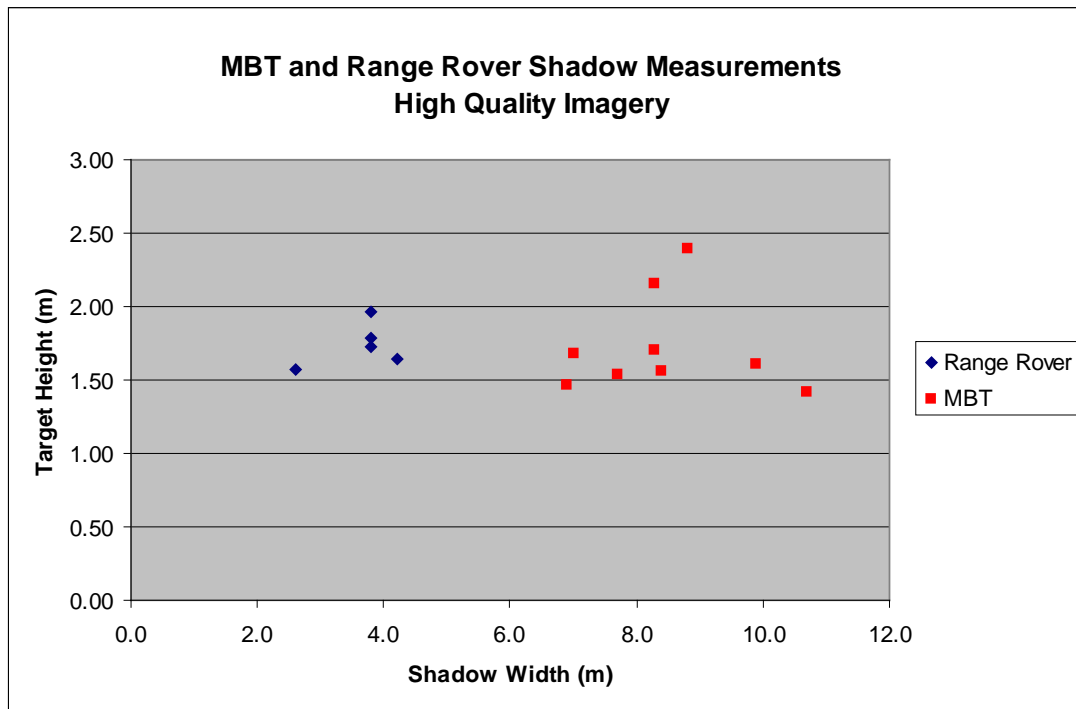
Figure 6-21 Measured dimensions of MBT and Range Rover targets, high quality imagery

The quality of size estimates taken directly from unfocused target signatures, measured in separate work undertaken by QinetiQ, is shown in Figure 6-22. It can be seen that there is a significant spread in the size measurements, especially for the Range Rover target where the reflectors on the roof can cause significant blurring. The Range Rover is more frequently measured as larger than the slower-moving MBT, which exhibits less blurring, but still includes some very large size measurements. These size measures are considered to be of little use for providing information on moving targets.
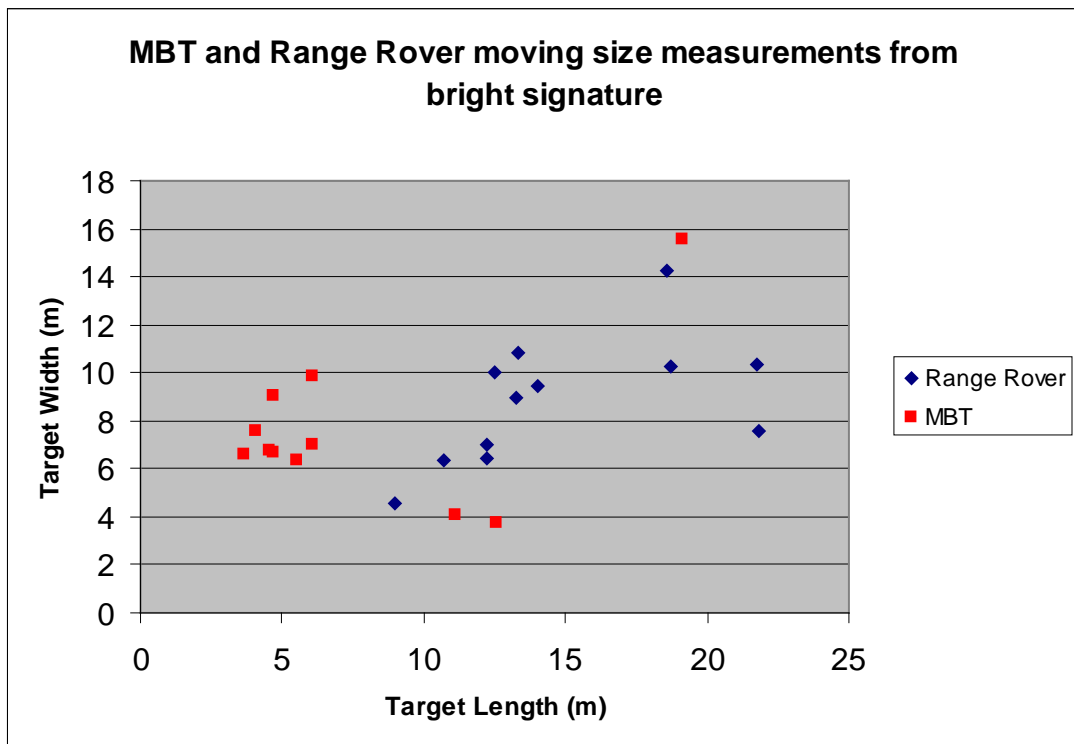
Figure 6-22 Size measurements taken from unfocused bright signatures of moving target.

For the measurements taken from target shadows, Table 6-2 and Table 6-3 display statistics for the measured dimensions of these two targets for both the mixed quality and high quality cases.

Dimension	Mixed Quality		High Quality	
	Width	Height	Width	Height
Mean	4.9	1.9	3.6	1.7
Std Dev	1.82	0.23	0.61	0.15

Table 6-2 Range Rover dimensions, mean and standard deviation

Dimension	Mixed Quality		High Quality	
	Width	Height	Width	Height
Mean	8.5	1.8	8.4	1.7
Std Dev	1.40	0.33	1.25	0.33

Table 6-3 MBT dimensions, mean and standard deviation

The true values for the two vehicles are shown in Table 6-4.

Dimension	Range Rover	MBT
Height	2.1 (roof rack)	2.4 (turret), 1.65 (hull)
Length	4.5	6.5
Width	1.8	3.4

Table 6-4 Range Rover and MBT true dimensions

The MBT turret is sloping and does not occupy the full length of the hull, hence the shadow cast from its peak is more likely to have its contrast reduced by target motion effects (Section 6.6.2), and does not appear clearly as a change in height of the target profile. Hence, lower height measures are recorded than would otherwise be expected.

While the measured dimensions do not agree directly with the known target size, a number of factors contribute to the mismatch:

- The combination of object length and width in the shadow width measure (Figure 6-10);
- Terrain slope effects on shadow length and therefore target height (Section 6.6.7);
- Shadow 'spread' due to vehicle motion during the dwell (Section 6.6.2);
- Variability of the segmentation algorithm, especially for reduced contrast imagery (Section 6.6.2).

The separation between the two classes, and the spread of measurements seen in the results compare favourably with that seen when the length and width are measured from static bright targets in air-carry trials imagery. Table 6-5 lists measurements on the SPG and ADU that were present as static targets in the 2006 air-carry trials. The measurements cover a limited range of angles, excluding head-on and broadside-on to the radar. Although these direct measurements are made of the static target's signature, there is still a notable difference between the true values and the measured values, as well as variability in the measurements that is comparable to that seen for the shadow dimensions.

Dimension	SPG		ADU	
	Length (m)	Width (m)	Length (m)	Width (m)
True	7.77	3.25	9.14	2.8
Mean	10.0	3.7	8.7	4.2
Std Dev	1.6	0.7	1.7	0.6

Table 6-5 Bright target dimensions from MRMS imagery

6.4.4 Aspect angle considerations

Whilst there is only a limited quantity of high-quality data of moving targets from air-carry trials, this has been separated into the two directions travelled; one in-line with the aircraft velocity vector, and the other crossing the velocity vector. Due to the squint angle of the antenna, the moving targets were therefore either imaged at 30° or 60° from head or tail-on respectively. Figure 6-10 showed the imaging alignment for a target moving in-line with the aircraft velocity, and the contributions of the target's dimensions to the shadow width measurement.

Dividing the high-quality data into sets based on the direction of target movement allows comparison between expected shadow width and measured shadow width. A very limited quantity of data is available for this, hence, work was undertaken with the CLTG model to simulate a wider range of target motions, which are described in Chapter 6.4.5.

Figure 6-23 shows the expected shadow width for imaging angles from head-on to tail-on for the MBT and Range Rover targets. The target movement in the trials relate to the dashed blue line for in-line and the dashed green line for crossing targets. For the range of angles measured in the air-carry trials, a significant difference between the in-line and crossing width measurements for each target would not be expected. Between classes, good separability should be visible, as was demonstrated in Figure 4-32.
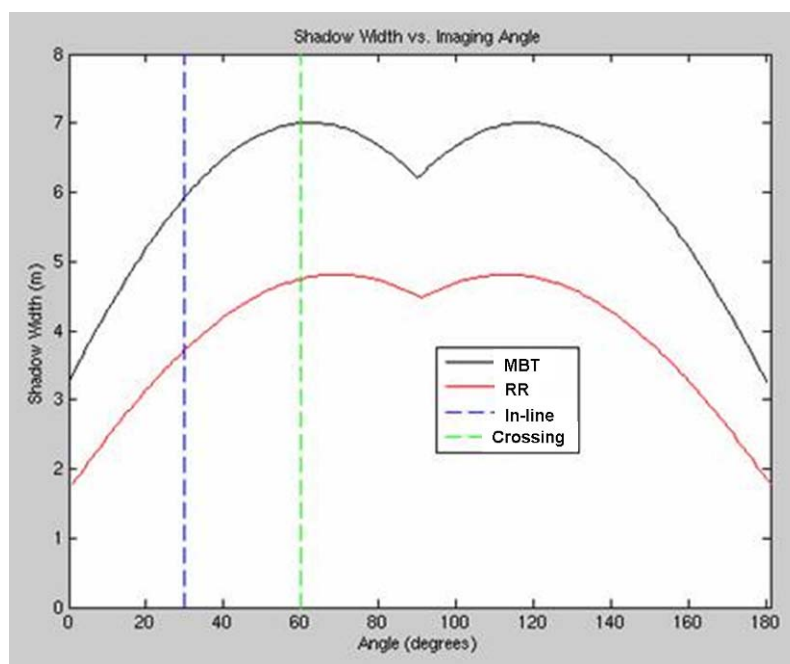
Figure 6-23 Predicted shadow width for MBT and Range Rover targets

Table 6-6 shows the expected and measured shadow widths for the two targets based on the direction of target movement. For both targets the crossing results are larger than the in-line results as expected, although results are subject to the limited data.

Target	Width Expected (m)	Width Measured (m)
MBT in-line	5.93	8.3
MBT crossing	7.04	9.9
Range Rover in-line	3.77	3.4
Range Rover crossing	4.76	4.2

Table 6-6 Shadow width by target direction of movement

Components of this extra width can be caused by the target's movement. Vehicles travelled at a set speed for each run, which was recorded by GPS. However, time synchronisation between the radar and vehicle records is not exact. Hence, it is possible to catch the vehicle while it is accelerating or decelerating, but there were no instances when it was stationary.

Considering the in-line cases, of which there are 8 for the MBT, four at low-speed (0.6 ms^{-1}) and four at a higher speed (2.8 to 4.7 ms^{-1}), the average measured width for the higher speed cases is 0.9 m greater than for the low speed cases, roughly the increase expected due to target movement broadening the shadow. For the Range Rover, there were insufficient data of a high quality to draw appropriate conclusions.

6.4.5 CLTG Data

Good results were achieved through examination of air-carry data gathered with the MRMS radar; however, the range of available data is limited, and investigation of the techniques on a broader set of targets was considered beneficial.

The CLTG model (Chapter 4.4) was used to generate imagery of moving targets, including displacement and blurring effects for the bright signature, and contrast reduction for the shadow.

Targets in three different size categories were simulated against a rural background with the CLTG (Figure 6-24). The targets are:

- Land Rover Discovery
- MBT
- Transporter-Erector-Launcher (TEL)
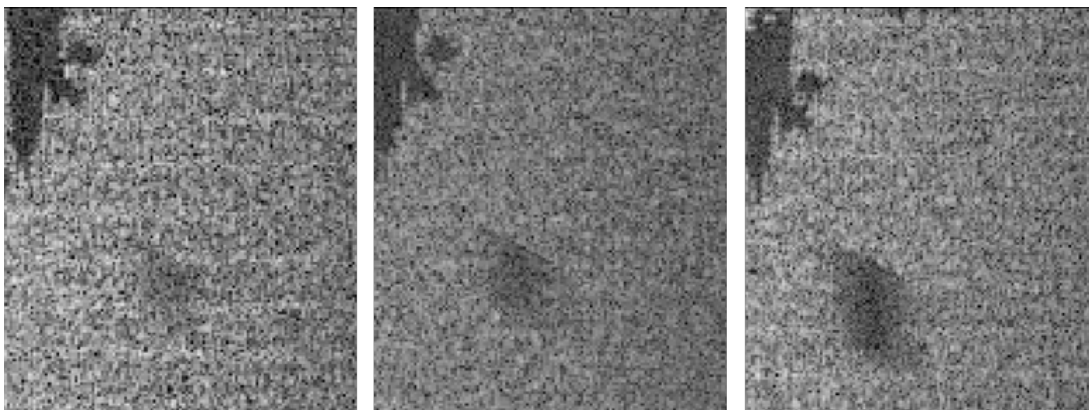
Figure 6-24 CLTG simulated shadow images of (from left to right) a Land Rover, MBT and TEL moving at 22ms^{-1}, heading 270°. Down-range from top to bottom of the image

Examples of a CLTG MBT shadow and air-carry trials MBT shadow for similar ground speeds (about 7 ms^{-1}) are shown in Figure 6-25. The shadows are of comparable size and contrast, although the clutter from the real imagery shows slightly more spatial structure.
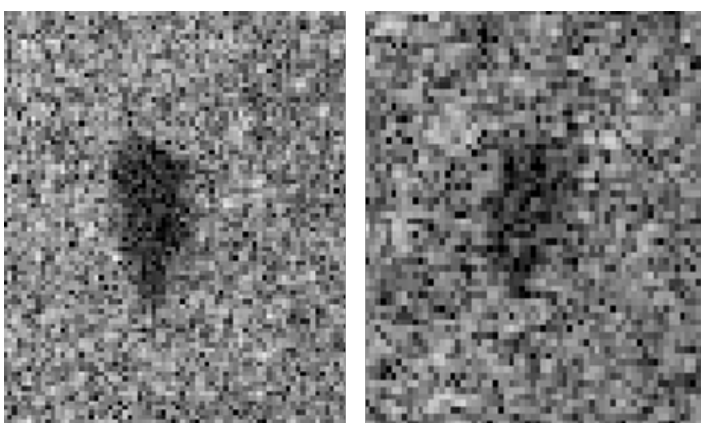
Figure 6-25 CLTG MBT signature (left), MRMS MBT signature (right)

Size measurements of the CLTG shadows from moving and static targets indicate the high level of separability between the classes that would be expected following the results from the air-carry data. When imaged with reasonable CNR (i.e. within the beam), the segmentation techniques work well on the CLTG data.

Figure 6-26 displays the size measurements for the three target types at a heading of 0°, for static targets and for speeds of 7, 13 and 22 ms^{-1}, which are 16, 29 and 50 mph respectively. The separation of the three target types is visible, and the greater height of the TEL provides separability in that feature, which is not apparent between the Land Rover and MBT, nor in the real data of Figure 6-20.
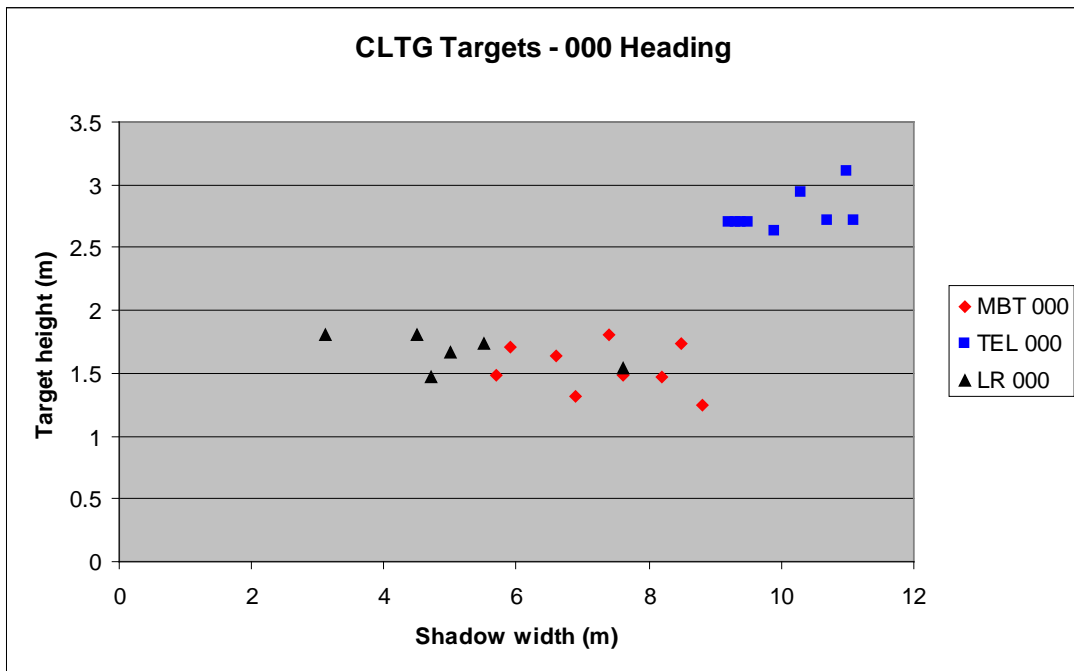
Figure 6-26 CLTG shadow size measurements – 0° heading

For a heading of 0°, the target will be viewed at 210° from head-on as shown in Figure 6-10. The expected and measured values for the three targets are in reasonable agreement. Table 6-7 contains the expected and measured widths and heights for the three targets modelled using the CLTG.

Target	Expected Width (m)	Measured Width (m)	Expected Height (m)	Measured Height (m)
TEL	9.4	9.3	2.85	2.64
MBT	6.5	6.4	1.57 (hull) 2.23 (turret)	1.63
Land Rover	5.0	3.8	2.0	1.73

Table 6-7 Shadow widths and target heights from CLTG 0° data

When combining all target speeds and headings, the distribution shown in Figure 6-27 is produced. There is notable overlap between the MBT and Land Rover classes. However, given the combination of different imaging angles, this is to be expected. One erroneous measurement has been reported for the Land Rover vehicle where the segmentation failed.
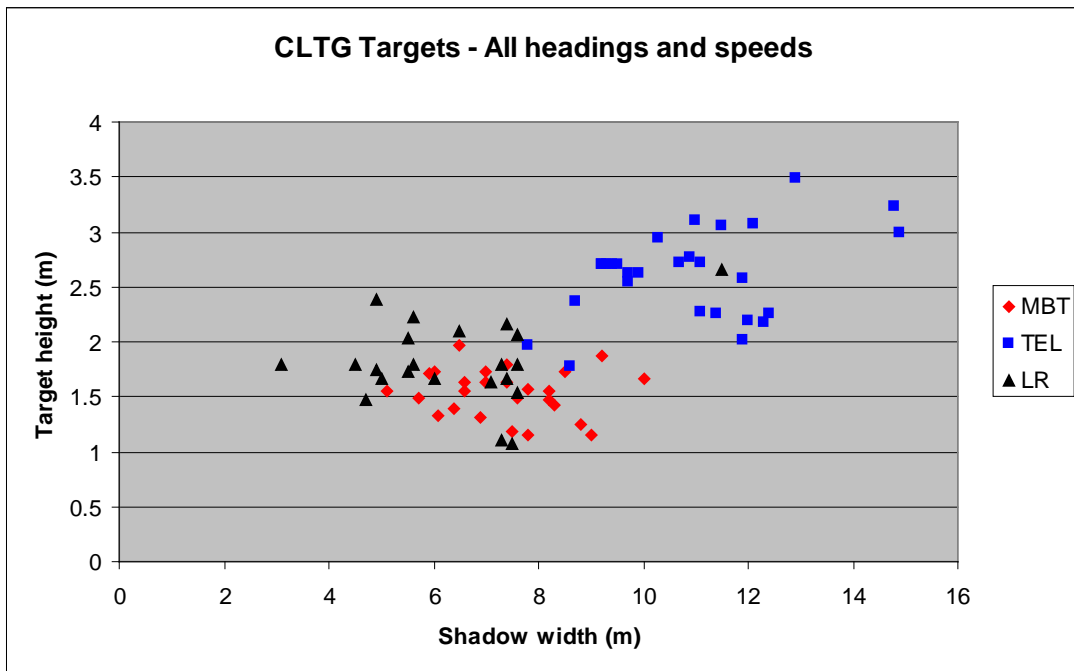
Figure 6-27 CLTG shadow size measurements – all headings

6.5   CLTG shadow velocity estimation

The behaviour of shadows in CLTG imagery closely relates to that of shadows in real data in the key aspects of contrast and shape. This makes data from the CLTG suitable for a short study into the ability of a radar seeker to measure moving target velocity through the change in position of the target shadow when re-imaged after an appropriate delay. The ability to measure velocity from the moving target's bright signature alone is initially discussed, followed by the accuracy with which velocity can be estimated from the shadow's position.

6.5.1 Velocity measurement using bright target signature

A study into measuring the velocity of a target from its bright signature was undertaken by QinetiQ under a separate research package. This considered two techniques, an MTI burst, and a multi-look approach.

An MTI burst is a high Pulse Repetition Frequency (PRF) imaging mode with resolution cells in the down- and cross-range directions of the order of the target size. The rapid PRF allows the Doppler frequency of a target to be measured unambiguously, and the large cell size helps to concentrate the returns from the target. When augmented with monopulse information, an accurate estimate of the radial velocity can be made.

The second technique of interest here uses knowledge of position when geolocating a target to the ground plane. This can be used to give a velocity estimate from two images containing a target separated by a time interval. Image registration can be used to remove bulk errors between the two scenes. However, inaccuracies still occur from the need to use a monopulse technique to place the target signature in its correct position. The standard deviation of the angular accuracy of a monopulse measurement was defined as a function of the monopulse slope of the antenna, and the signal to noise (or background) ratio of the desired target in equation (6.2).

For an example monopulse slope of 0.40 V/V/degree, which is within the range of values which could be expected for a seeker antenna, the angular measurement accuracy at one standard deviation is as shown in Figure 6-28.
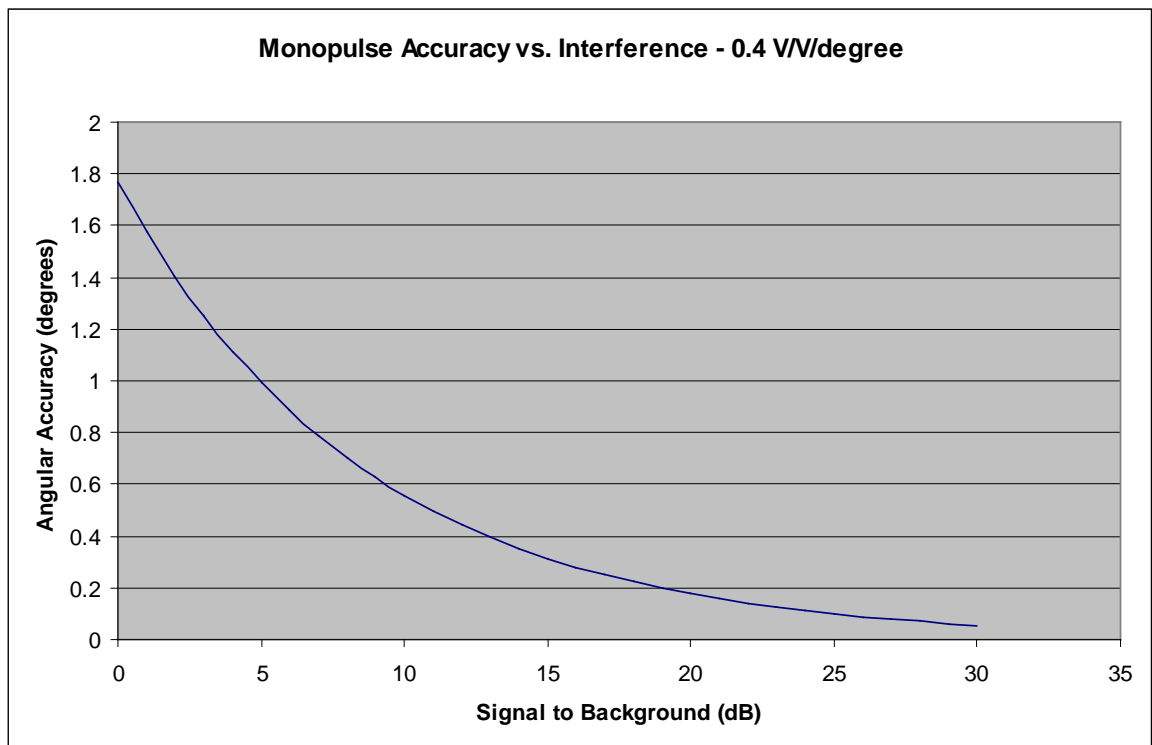
Figure 6-28 Angular accuracy (one standard deviation) of monopulse measurement vs. signal to background

For the geometry and ranges that were used in the trials and synthetic data generation, this monopulse error is predicted to have a standard deviation of 5 m for an MBT and 12 m for a 4x4 vehicle due to the reduced signal strength.

When predicting the performance of velocity estimation on the bright target signature, ideally the data points would be passed to a tracker. However, as this

work only examined short time intervals and few example images, a straight line was fitted between two positions with errors generated using a Monte Carlo method. Figure 6-29 illustrates this process. Two measurements of a moving target are made. The target's true position for each measurement is marked with an 'X' and the measurement errors are represented by the ellipses, the solid blue lines show possible trajectories between the measurement ellipses.

Due to the fine down-range resolution, the ellipses are narrow in the down-range direction, but errors in the monopulse cross-range position result in broad-ellipses in the cross-range direction, and, hence, less accurate measurements.

This technique produces speed estimates with an average value that can be significantly higher than the true speed. For the simple radial case, this is shown in Figure 6-29, where monopulse position errors will make the distance the target has travelled appear to be larger than the truth. The speed errors for tangentially moving targets are larger than radial, as the monopulse errors are in the same direction of travel.
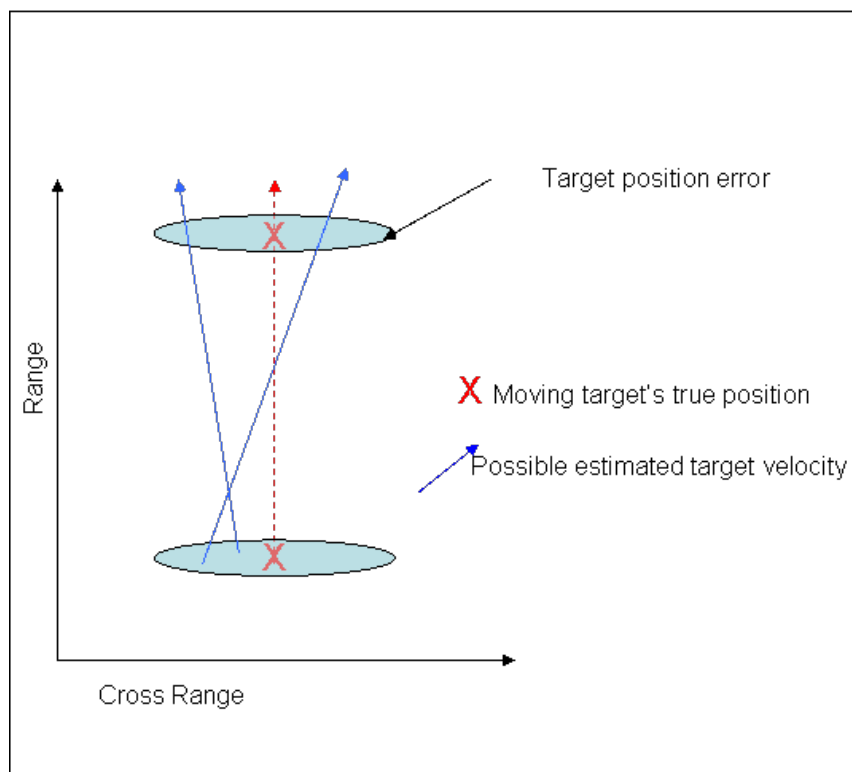
Figure 6-29 An example of multiple look detection based velocity estimation

6.5.2 Velocity measurement using shadow information

While the errors in identifying a target's true position using its bright signature can be large, the position of the centroid of a target's shadow in the cross-range direction is subject to a much smaller level of error than the monopulse estimate of the true position from the bright return. This error is comparable to finding the centroid of a bright static target signature, with the advantage of a relatively uniform power level within the signature. However, at low elevation angles, the signature can be significantly extended (Section 6.6.7).

The reduced positional error should result in a more accurate target velocity measure. This hypothesis was tested using simulated data.

6.5.3 Modelling

The CLTG model has been used to create imagery of a target moving at a range of speeds and in different directions. The directions used were 0°, 90°, 180° and 270° for all targets, plus some additional angles for the MBT.

For each configuration, two RDMs containing the target were generated, with a one second time interval between the images, during which the target's position changed relative to the background. The RDMs are then mapped into the ground plane, shown in Figure 6-30. The use of a common ground plane allows the difference in shadow centroid position to be measured.
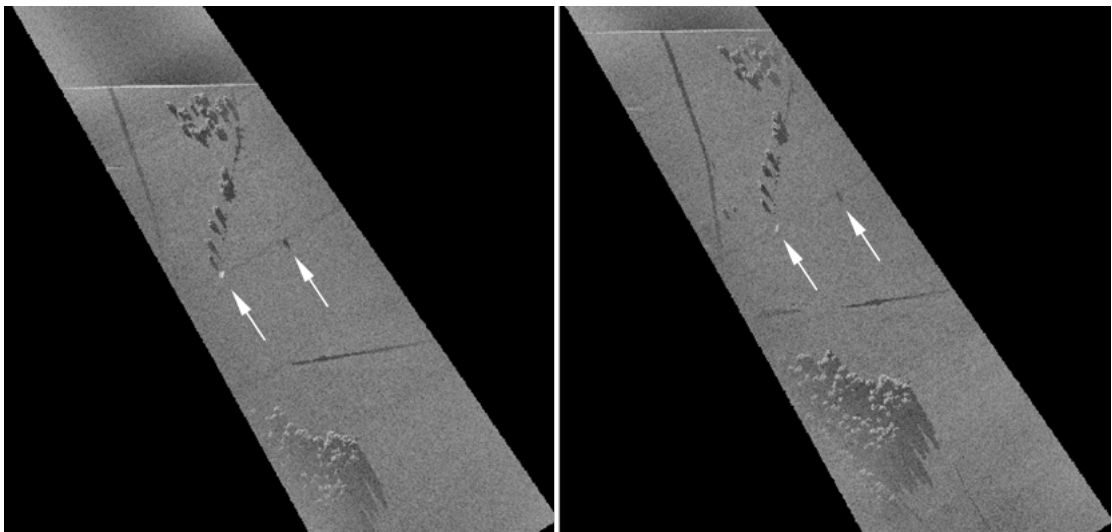
Figure 6-30: CLTG scene with target moving at 13 ms^{-1} on a heading of 000 degrees, one second interval between left and right images. Target and shadow indicated

Through accurate alignment of the two scenes, as shown in Figure 6-31, the change in shadow position and the centroid, which is marked as a bright point on the shadows, can be seen. The centroid was calculated as the central moment shadow shape.
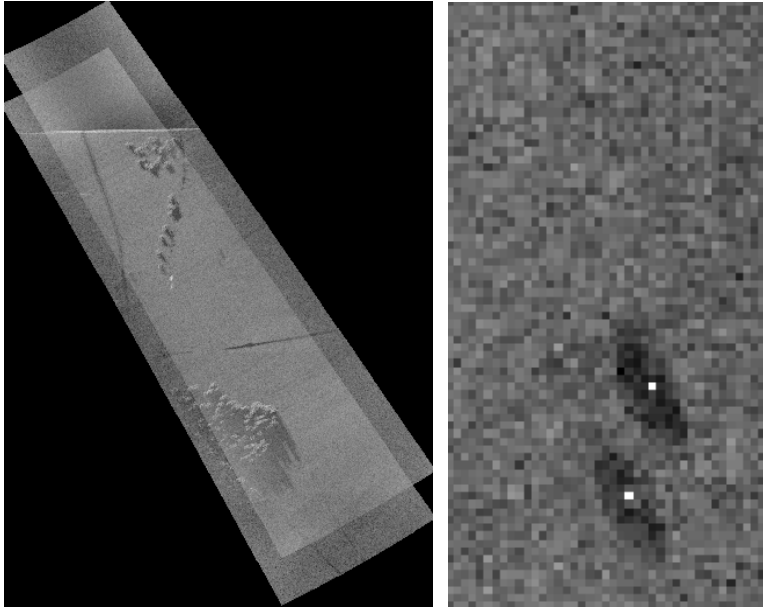
Figure 6-31 CLTG scene with target moving at 13 ms^{-1} on a heading of 0°; combined images (left) and extracted shadow positions with centroid marked (right)

For this example, the true speed was 13 ms^{-1} and the heading 90°. The change in measured centroid position indicates a speed of 15.03 ms^{-1} and a heading of 86.19°, which is close to the true values.

6.5.4 Shadow velocity measurement performance

A table of measurements from simulated data of an MBT is shown Table 6-8. For all cases, the time delay between measurements is one second. For each case there is a single simulation followed by the shadow detection and measurement process. This was repeated for the three target types.

Speed (ms^{-1})	True Heading (degrees)	Speed (ms^{-1})	Speed Error (ms^{-1})	Heading (degrees)	Heading Error (degrees)
0	0	0.10	0.10	0.00	0.00
7	0	4.90	-2.10	0.00	0.00
13	0	10.80	-2.20	0.00	0.00
40	0	38.01	-1.99	-9.39	-9.39
7	90	9.06	2.06	83.66	-6.34
13	90	15.03	2.03	86.19	-3.81
7	120	13.63	6.63	107.50	-12.50
13	120	15.70	2.70	116.89	-3.11
7	210	8.44	1.44	216.33	6.33
13	210	13.89	0.89	210.26	0.26
22	210	21.18	-0.82	211.29	1.29
7	270	8.00	1.00	269.10	-0.79
13	270	13.00	0.00	269.52	-0.48
22	270	25.02	3.02	272.06	2.06
7	315	6.28	-0.72	307.23	-7.77
13	315	13.45	0.45	311.99	-3.01

Table 6-8 Simulated MBT. Target speed and heading, true vs. measured, one second delay

In addition to the simulated MBT target, similar tables of performance were generated for a moving TEL, and a moving Land Rover. These vehicles represent two extremes of vehicle size that could be of interest for detection.

While the size and brightness of the returns from the targets varies significantly between these classes, the shadows remain consistently detectable. The accuracy of vehicle position measurements does not appear to be affected when locating the centroid of large shadows such as for the TEL.

The accuracy of velocity and heading measures, using data combined over the three target types and a range of headings, is shown in Table 6-9.

Speed (ms^{-1})	Speed Error (ms^{-1})		Heading Error (degrees)	
	Mean	Std Dev	Mean	Std Dev
0	1.09	1.03	/	/
7	0.34	0.92	-1.17	5.08
13	0.39	1.06	0.16	3.42
22	1.63	1.43	1.16	3.79

Table 6-9 CLTG shadow simulations, velocity and heading errors

The measurements from simulated data show an ability to estimate a target's speed and heading from the change in shadow position with greater accuracy than can be obtained from the use of monopulse to relocate the bright signature.

Should the capability to re-image the scene with a delay of greater than one second be available, the accuracy of the velocity estimates would improve. Should the time delay be such that re-acquisition of the target may be problematic, then an additional interim image may be required to allow shadow tracking. A combination of bright return and shadow tracking over multiple images would lead to an accurate speed and heading measure.

6.5.5 Bright signature velocity performance

The shadow-based velocity performance reported in Table 6-9 can be compared with that predicted using knowledge of the accuracy of the bright signature relocation and a monopulse velocity error model developed by QinetiQ. The results from using bright target measurements are shown in Table 6-10.

Speed (ms^{-1})	Speed Estimate (ms^{-1})		Heading Estimate (degrees)	
	Mean	Std Dev	Mean	Std Dev
0	10.7	6.2	/	/
7	12.7	6.85	31	74
13	16.2	6.75	20	60
22	24.1	9.05	6.6	49

Table 6-10 Bright target signature velocity prediction results

These errors are significantly larger than for the shadow measurement approach. For this range of speeds and time delays, the standard deviation of the results is at a level where it will limit the utility of any measurements.

6.5.6 Shadow centroid measurement accuracy

To check that this level of performance was realistic, the velocity error model was adapted to match the setup used in the shadow simulation, and the mean and standard deviation of velocity and heading errors were generated.

Given the speed and heading errors measured from the imagery (Table 6-9), the model was configured with position accuracy values such that it produced similar speed and heading errors. The best-fit between the simulated data and the model occurs for the following errors in shadow centroid accuracy, shown in Table 6-11.

Speed (ms^{-1})	Centroid Measurement Error Standard Deviation (m)
7	0.5
13	0.7
22	1.2

Table 6-11 Error in target location for CLTG simulated moving target data

The ability to measure the centroid position accurately appears to degrade as speed increases. This is expected as the shadow becomes extended and contrast is reduced due to the target's movement during the dwell. For the low-speed cases, the consistency of position measurement appears to be very good, i.e. of the order of a single pixel. Given a clear shadow this consistency is expected as the shadow changes little in the time between images. The accuracy of velocity measurement, as shown when comparing performance using the bright signature against that from the shadow, is proportional to the position measurement error.

6.6 Shadow Properties

6.6.1 Introduction

While the shadow signature has shown its usefulness in providing both size and velocity information for moving targets, there are a number of factors which need to be accounted for. These factors determine whether the target motion is such that the shadow contrast would be lost (Section 6.6.2), whether the background will provide necessary contrast for shadows (Section 6.6.5), and the effects of terrain slope on the shadows (Section 6.6.7).

6.6.2 Target motion effects on shadows

The effect on shadows of the target velocities observed from air-carry trials (typically less than 10 ms^{-1}) and modelled with the CLTG (low to 22 ms^{-1}) do not appear to result in noticeable changes between the size estimates for the static and moving cases.

For the 7 ms^{-1} speeds, the target will move 3.5 m during the 0.5 s dwell that has been simulated. This will have the effect of extending the shadow, leaving areas of lower contrast where the target has not shadowed the ground for the entire dwell.

In the data produced by the CLTG, this effect can best be illustrated on a simulation of the TEL moving with the missile erect. Whilst this is not a physically possible scenario, it produced a distinct missile shadow in the image where the effects of changing contrast were visible, Figure 6-32.
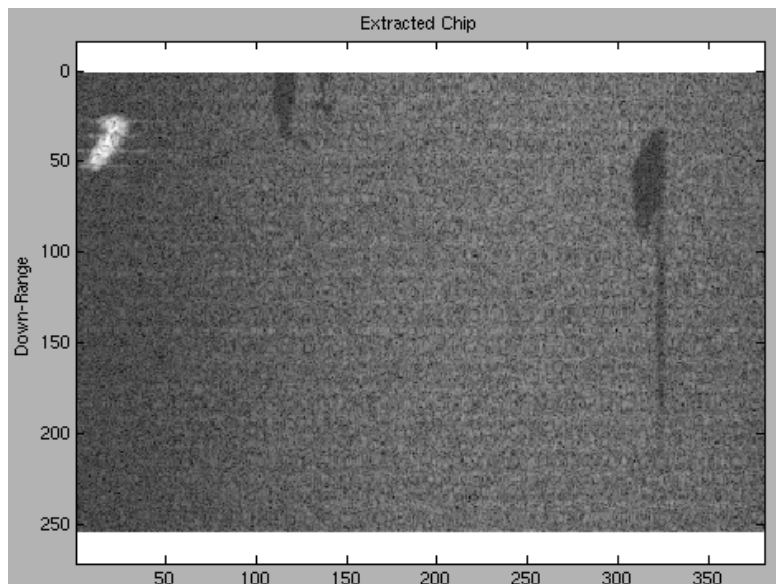
Figure 6-32 Displaced TEL with missile erect, and shadow. Down-range from top to bottom of image

The target's velocity of 7 ms^{-1} on a 0° heading results in a velocity in the cross-range direction of 3.4 ms^{-1}, and for a half-second dwell, a displacement of 1.7 m. This displacement during the dwell manifests itself as visible effects in the imagery in Figure 6-33.
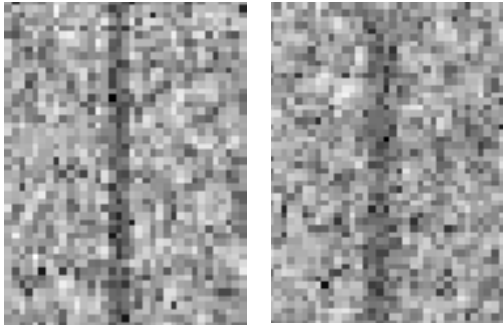
Figure 6-33 Shadow of missile from stationary TEL (left) and moving (right)

Owing to the additional shadowing from the TEL missile support structure, the shadow is expected to be roughly 2 m wide. The shadow width is measured as 1.75 m in the static case, and 3 m in the moving case.

6.6.3 Contrast calculations

The contrast variation over a shadow was calculated for a range of target velocities and radar dwell times. While the calculations in this section focus on tangential velocity, radial velocity would be treated in the same manner. Under typical viewing geometries, due to the shadow length exceeding its width, the effects of radial velocity on contrast are less pronounced, and hence performance will be retained over a broader range of scenarios.

Figure 6-34 shows the loss of contrast expected for a 2 m wide target moving tangentially at 3.5 ms^{-1} for a variety of dwell times. Any shift in the position of a line below the 0 dB level (no loss) indicates that the shadow contrast is reduced. It can be seen that, for the same velocity, the contrast decreases as dwell duration increases owing to the reduced fraction of the dwell in which ground returns are obscured.
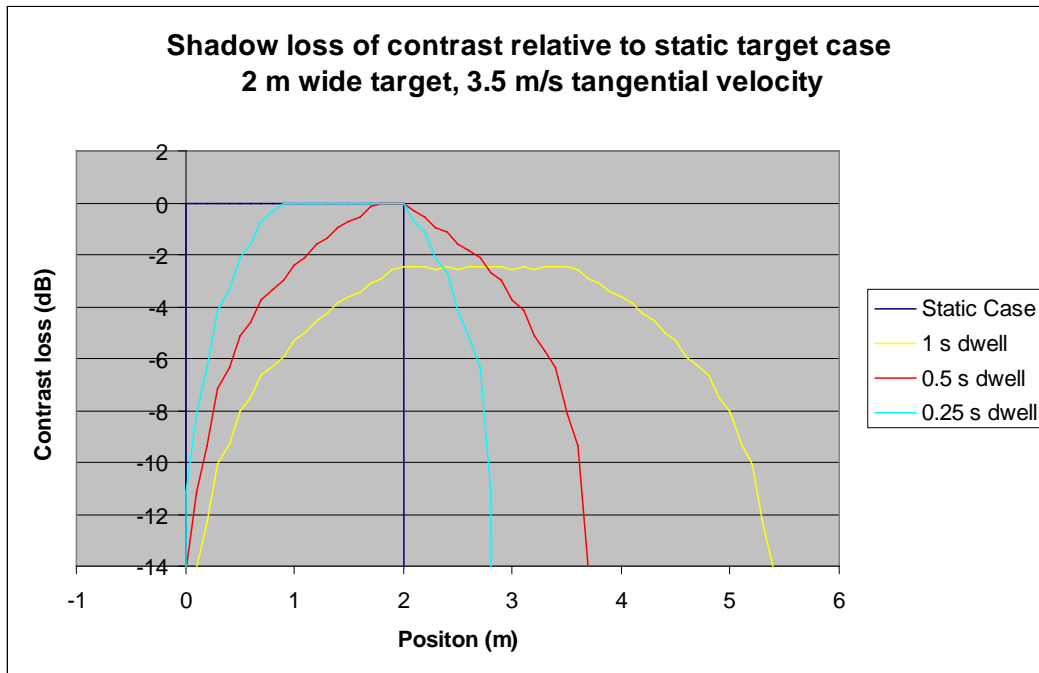
Figure 6-34 Shadow contrast loss with tangential target movement

For a larger object, even at a higher speed, the reduction in target contrast and fractional increase in width can be less pronounced (Figure 6-35).
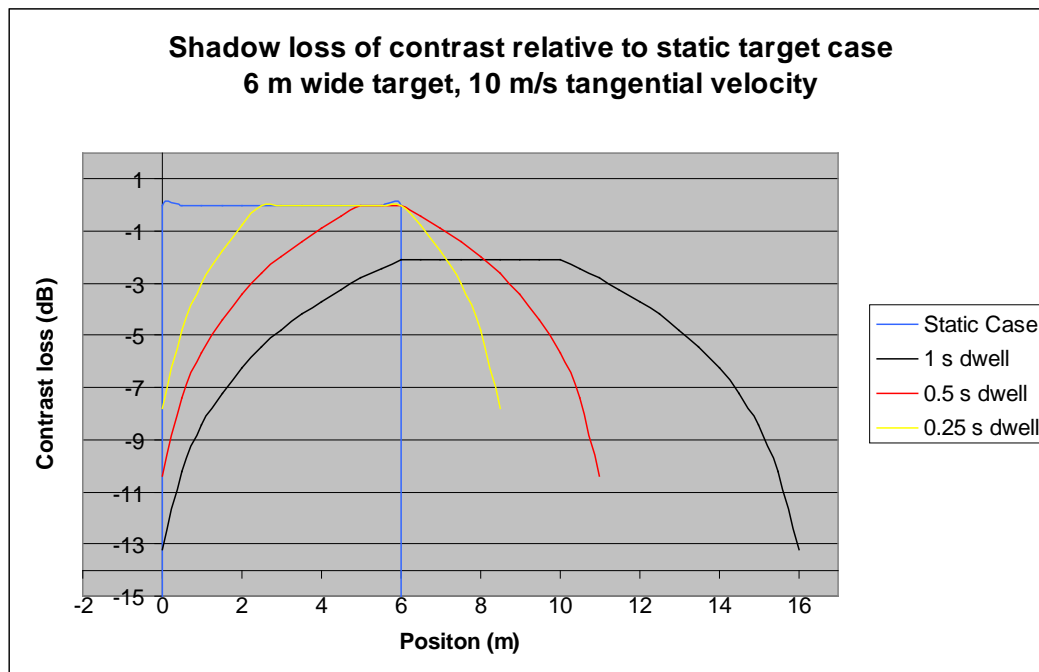
Figure 6-35 Shadow loss of contrast with tangential target movement

6.6.4 Contrast prediction tool

To visualise the effects of target motion quickly, and judge where it may cause problems for shadow analysis, a tool was created using Matlab to plot shadow contrast. The following parameters were allowed to vary:

- Target length
- Target width
- Target speed
- Target heading
- Noise floor
- Contrast required to segment

This allowed a range of target sizes and speeds to be tested, noting the levels of contrast required for the segmentation to operate, and the noise floor observed in trials imagery.

The contrast was plotted in a similar manner to that used for examination of the shadow contrast formed when taking cuts through the imagery in the cross-range and down-range direction, shown previously in Figure 6-14.

An example contrast prediction for an MBT sized target is shown in Figure 6-36. For this target moving at 10 ms^{-1} (22 mph), it can be seen that there is some signature spread, but the bulk of the shadow is limited by the radar noise floor, so performance will remain high.
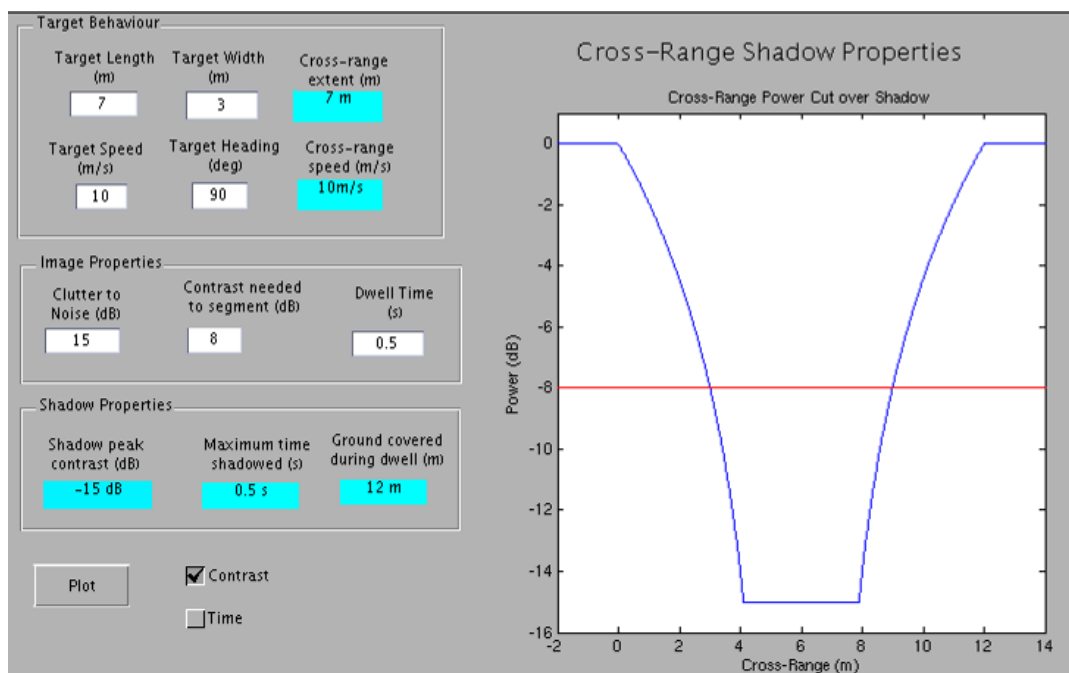
Figure 6-36 Contrast prediction for MBT sized target moving tangentially to the sensor at 10 ms^{-1}

The predictions from this code were tested against CLTG shadow data. Figure 6-37 and Figure 6-38 show cross-sections through the shadow in the simulated data against cross-sections through the predicted shadow for speeds of 7 ms^{-1} and 22 ms^{-1} respectively. The *y*-axis power values do not align, but the scaling is consistent.

For the plots using real data on the left of the figures, the red line represents a smoothed cut through the shadow, averaging the returns from a band of pixels, and the blue line is the return from a single line of pixels. For the plots on the right of the figures, the blue line shows the ideal contrast against a uniform background, and the red line the level of contrast required for the segmentation to operate reliably.
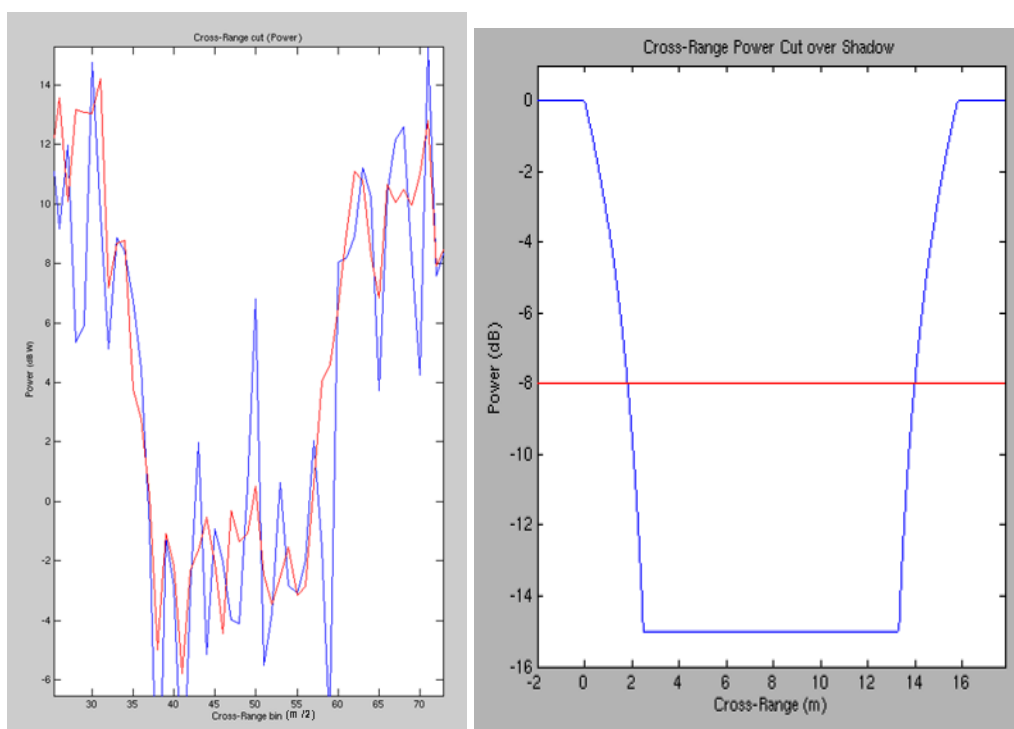
Figure 6-37 TEL, 7ms^{-1} tangential, CLTG (left) and predicted (right). Blue plot is for single pixel width, red plot is an average over 9 pixels
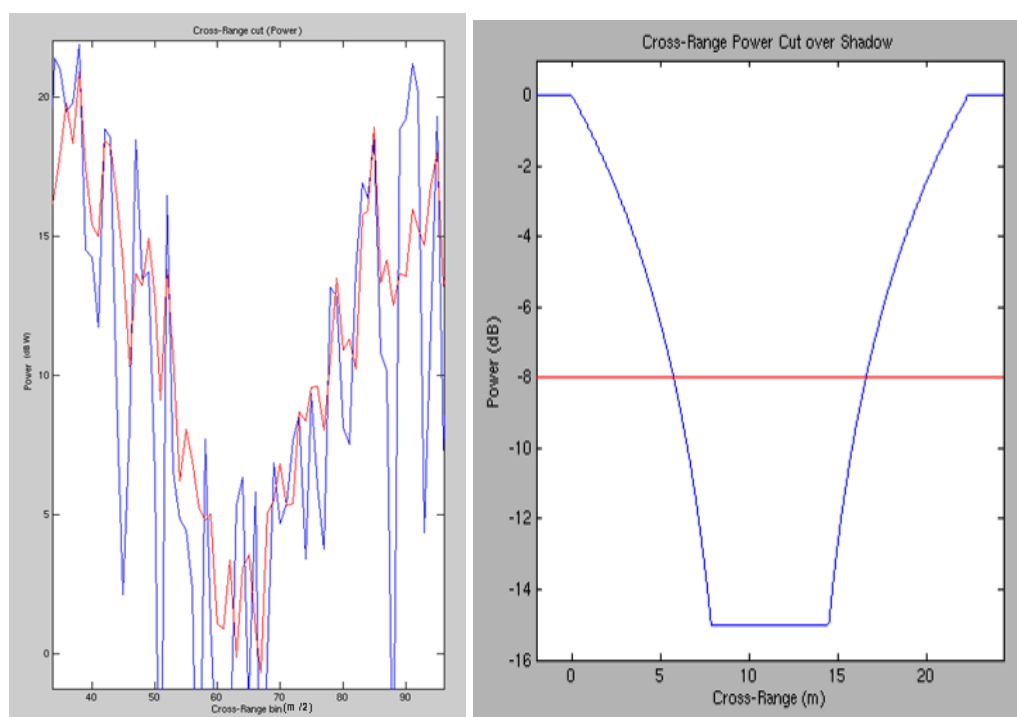
Figure 6-38 TEL, 22ms^{-1} tangential, CLTG (left) and predicted (right). Blue plot is for single pixel width, red plot is an average over 9 pixels

The quantity of real data containing the MBT target within the beam at speeds above walking pace was limited. However, Figure 6-39 is an example of the MBT shadow cross-section from an RDM, and the prediction of how the contrast should appear.
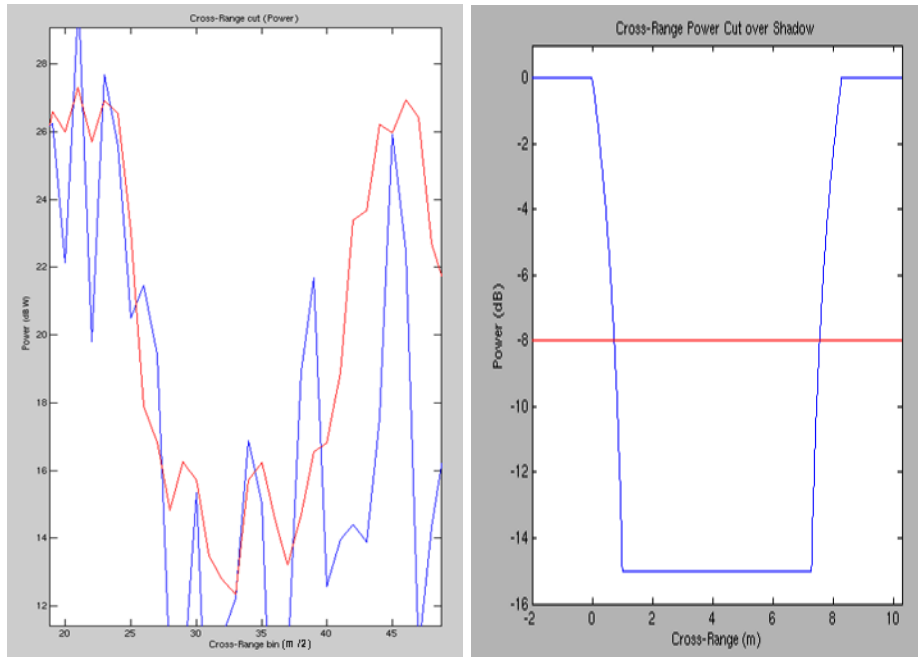
Figure 6-39 MBT, 2.8 ms⁻¹, 60 degree heading, Trials data (left) and predicted (right). Blue plot is for single pixel width, red plot is an average over 9 pixels

Accounting for the noise present in the imagery, the contrast within the shadows from real and simulated imagery appears to match the predictions at this level of observation.

The tool can be used to examine at what point shadow information will be lost due to the reduction in contrast resulting from target motion. For the case of targets moving tangentially to the radar, Table 6-12 shows the upper velocity limit for a range of dwell times for the segmentation of shadow information when the noise floor is 15 dB below the clutter level.

	Target		
	Land Rover	MBT	TEL
Dwell Time (s)	Approximate maximum velocity for segmentation (ms⁻¹)		
1.0	7.2	10	19.5
0.5	14.5	20	39
0.25	29	40	78

Table 6-12 Maximum tangential speed to allow shadow analysis for a range of target types and dwell times

Based on each vehicle's capabilities, the colour scheme indicates whether the target's shadow should be sufficient for analysis (green), whether the target may travel too fast (amber), and where the target is likely to travel too fast (red).

6.6.5 Clutter background

All real and simulated shadow data have considered targets on a grassland background, as illustrated in Figure 6-15 and Figure 6-16. This terrain has sufficient reflectivity to appear in radar imagery as significantly brighter than the noise floor.

Grassland in the air-carry trials data, dependent on polarisation, has been measured at levels comparable with those in [68], with an RCS typically in the range of -15 to -20 dB m^2/m^2. This surface is expected to be the most reflective surface upon which a vehicle will be imaged.

Hard surfaces such as tarmac or concrete for roads or urban areas have a significantly reduced diffuse reflectivity, and a reduction in contrast of the order of 10 dB can be expected [68]. An extremely low level of noise would be required for shadow analysis to be effective.

If a target is travelling on a surfaced road, the road is narrow, the surrounding area has favourable clutter and the imaging geometry supports shadow analysis, there is potential that the technique could be used. This technique has shown potential in SAR imagery [38], but a sequence of multiple images was available for processing, and the low elevation angle left long shadows.

CLTG simulated imagery of an MBT moving on a road is shown in Figure 6-40. It can be seen that, for this grazing angle and background, the shadow appears as a small extension to the dark road. Use of shadow information, in this case, would require either a multiple image change-detection approach, or segmentation using knowledge of the likely target size and tightly constrained to prevent areas of the road from inclusion in what is considered to be the shadow area.
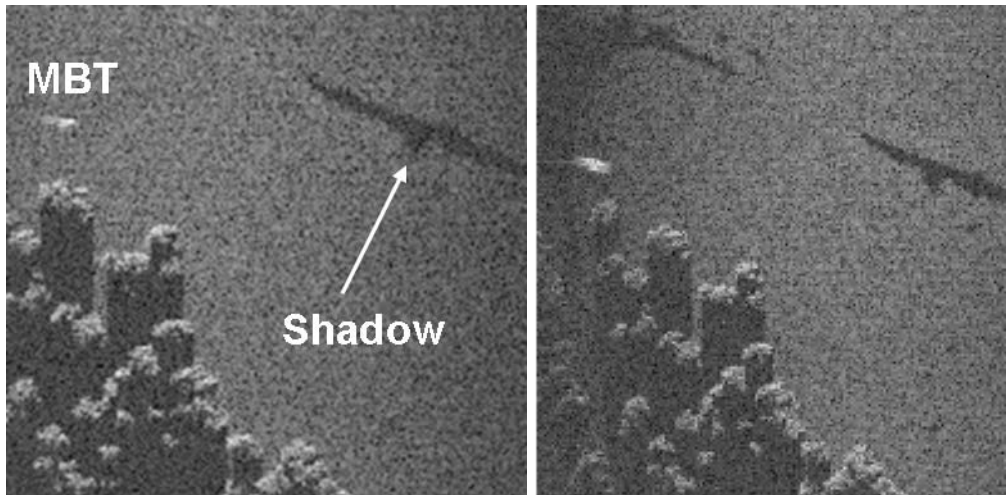
Figure 6-40 Target moving on road at 7 ms^{-1}, 1 second interval between images

Rough surfaces, such as unmade tracks and rocky desert areas are likely to have sufficient contrast for shadows to be observed. For the 2006 MRMS trials, an SPG was imaged in the car park outside the FIBUA village, see Figure 6-41.

Figure 6-41 SPG in FIBUA car park

Example imagery of the SPG in different positions against the rough car park background is shown in Figure 6-42. For two of the cases the turret has been rotated from the forward position. In all cases, the shadow is visible against the background. However, while the shadow is visible, the CNR in the imagery is of the order of 8 dB or lower and the shadow segmentation does not work reliably.

Against this type of clutter background, a radar with a lower noise floor may provide a more robust capability.
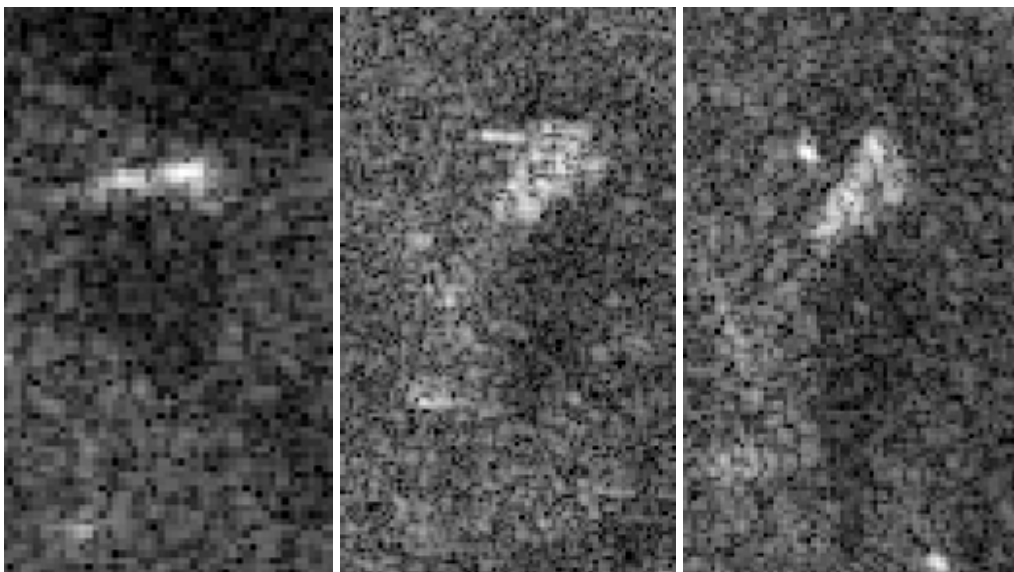
Figure 6-42 SPG imaged in FIBUA car park

6.6.6 Polarisation

The choice of polarisation can affect the clutter reflectivity from a scene and, therefore, also affect the CNR, which is a key factor in the ability to segment shadow outlines. This is further discussed in Chapter 7.

For the circular polarisations used in the air-carry trial, the use of the choice co- or cross-polar channel makes a relatively small difference in the reflectivity of background clutter. Should a linear cross-polarisation be used, this difference can be of the order of 10 dB or higher [60] and the subsequent reduction in CNR would degrade shadow analysis performance.

The MRMS air-carry trials data was gathered almost exclusively in the cross-polar circular polarisation. A few runs were performed which recorded both the cross- and co-polar channels, but these did not include any moving targets.

Due to the nature of the seeker waveform, the noise floor is limited by the mean clutter level in each image. Hence, comparisons of CNR between the odd and even channels were effectively normalised for clutter reflectivity and a noticeable difference was not observed.

Calibration against sky-noise was used to normalise the radar channels, and clutter levels could then be directly compared between the two channels. This indicated that the grassland clutter in the cross-polar radar channel was approximately 3 dB brighter than that in the even channel. For a given noise floor, this would indicate that the cross-polar channel is preferable for shadow analysis.

Information on clutter reflectivity for linear polarisations from [68] indicates that expected levels of backscatter in the HH and VV polarisations are greater than has been measured from the 2006 trials data for circular polarisation and, therefore, would be suitable for shadow analysis.

While the choice of polarisation may make a small difference to the contrast, it should be considered as one component in the overall radar seeker design to ensure that the noise floor is sufficiently far below the clutter levels of likely scenes to permit information from shadows to be used.

6.6.7 Effects of terrain slope

As described in Section 6.4.5, in order to broaden the amount of data available with which shadow analysis techniques could be assessed, the CLTG seeker model was used to generate imagery of moving targets.

A limited quantity of imagery was generated with static targets distributed over a broad area with small terrain slope variations dependent on position. In this imagery, the relationship between shadow size on imaging geometry and terrain slope was clearly visible. The signature of three identical vehicles located at different positions in the scene is shown in Figure 6-43.
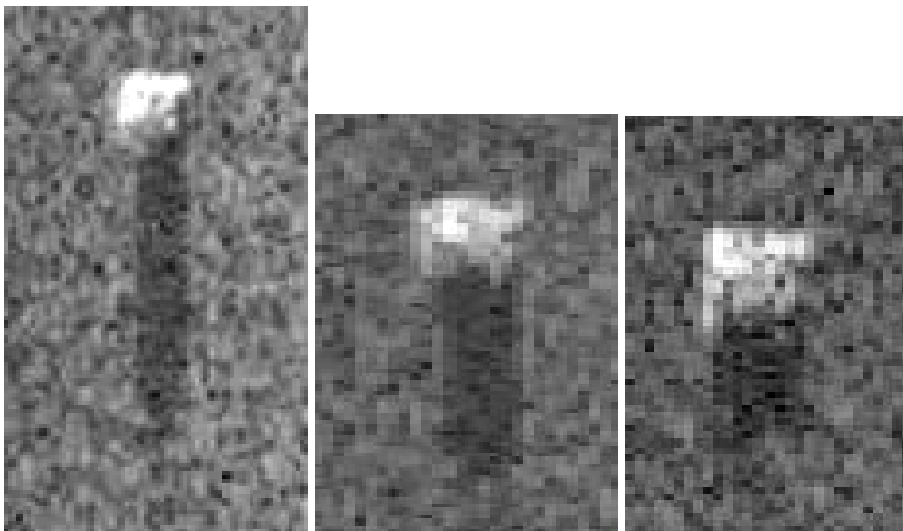
Figure 6-43 CLTG imagery of the same vehicle class positioned in different locations at the Copehill plantation

Owing to the sensor elevation angle used in the CLTG data of approximately 10°, slight variations in the slope of the imaged ground can result in a significant variation in the shadow length. This effect is illustrated in Figure 6-44, although to a lesser extent due to the choice for illustrative purposes of a radar depression angle

(Ø) of approximately 20°. The slope used for the example is approximately 4 degrees, and highlights the lengthening of the shadow for small variations in terrain.
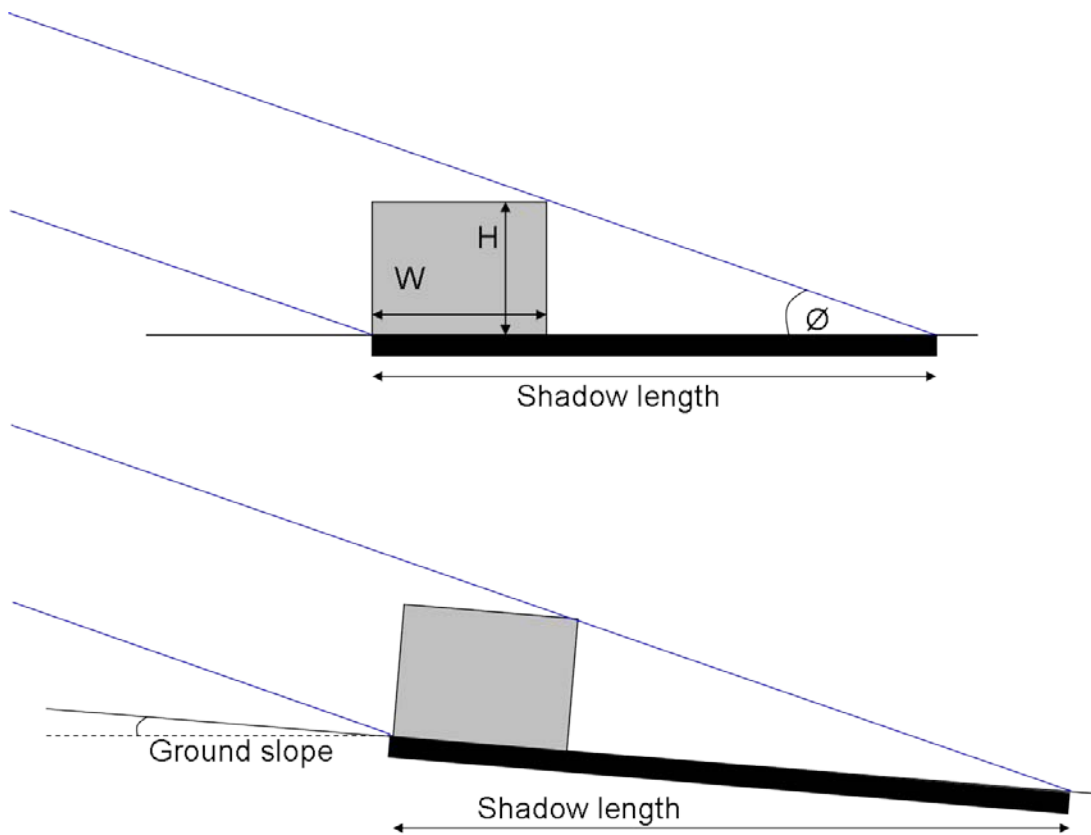
Figure 6-44 Effect of terrain slope on shadow length

The shadow length effects from terrain slope and from different seeker depression angles are interchangeable, hence it is the relative depression angle which is of interest. It should be noted that at steeper angles it can be expected that the clutter reflectivity will increase, improving shadow contrast. For grassland, and paved / unpaved roads, should the effective depression angle change from 10° to 20° the increase in σ_0 could be of the order of 3 to 5 dB m^2/m^2 [60]. A trade therefore exists between long shadows at low angles, which may give more accurate height estimates, and shorter shadows at steeper grazing angles, where the clutter RCS should be higher and hence shadow contrast improved.

For a target of height 2 metres and width in the down-range direction of 3 metres, the shadow length for a range of elevation and slope angles is shown in Figure 6-45. The choice of elevation angles represent the lowest expected angle for a likely seeker trajectory (10°) increasing in 5° steps to a steeper angle (25°) that could either be achieved through an increase in seeker height or through forming an image later in the trajectory.

The shadow lengths shown in Figure 6-45 cover a range of slope angles from 20° away and towards the sensor. Low elevation angles present a challenge when the slope is away from the seeker. As the slope angle away becomes comparable in magnitude to the elevation angle, the shadow length increases until the point at which the entirety of the ground is shadowed and no measurement is possible. It is clear that, even for small changes in angle, terrain slope needs to be known to convert these length measures into a meaningful target height.

For increasing slopes towards the target (positive slope), the shadow length decreases, and the relative difference between the range of grazing angles becomes less pronounced. Steep slopes towards the sensor are not necessarily beneficial as a reduction in shadow length increases the effect of segmentation errors on measuring size.

The ± 20° range represents approximately a 1 in 3 slope towards and away from the seeker, which is likely to be at the most extreme range of observed terrain containing moving targets. Figure 6-46 shows in greater detail the effects on shadow length of ± 6°, or a 1 in 10 slope.
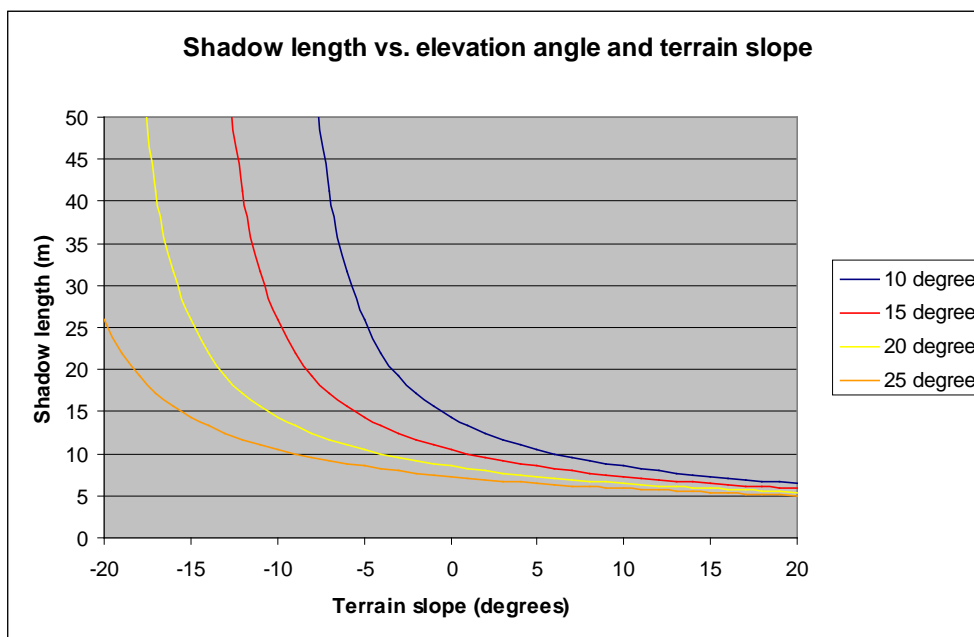
Figure 6-45 Shadow length vs. terrain slope for a range of elevation angles
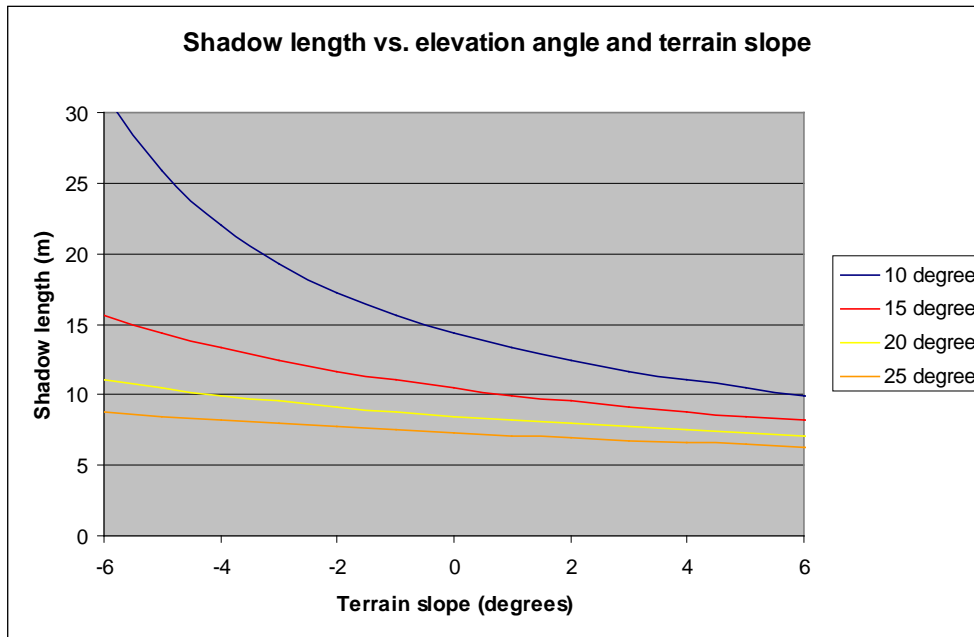
Figure 6-46 Shadow length vs. terrain slope for a limited range of elevation angles

The effective shadow length, a combination of the ground beneath the target, as well as that shaded by the target's height varies as shown in Table 6-13.

Elevation Angle (degrees)	Shadow Length (m) for varied terrain slope						
	-20°	-10°	-5°	0°	5°	10°	20°
10°	∞	∞	25.9	14.3	10.5	8.5	6.5
15°	∞	25.9	14.3	10.5	8.5	7.3	5.9
20°	∞	14.3	10.5	8.5	7.3	6.5	5.4
25°	25.9	10.5	8.5	7.3	6.5	5.9	5

Table 6-13 Shadow length vs. elevation angle

It is clear that, even for small terrain slope changes of ± 5°, at lower elevation angles there is a significant change in shadow length and, therefore, the ability to estimate target height.

6.6.8 DTED Information

Knowledge of the relative height of terrain imaged by the radar seeker will allow correction of shadow length and, therefore, height measurements to be made with a known level of accuracy. To do this, seeker data could be accurately geolocated with maps of the ground or Digital Terrain Elevation Data (DTED).

Uncertainty in the seeker position (potentially due to GPS denial) during image formation will prevent any correction taking place. There are techniques such as

scene-matching algorithms that could compensate for this, given prior knowledge of the terrain, and information from those maps could be used to provide information on ground slope.

In addition to the coarse height information which may be available for the seeker, DTED data may be available with sufficient resolution to assess terrain slope. Level 1 DTED data has a sample spacing of 90 m and is used for military activities and systems that require slope, elevation and gross terrain roughness. This sample spacing may be adequate to correct terrain slope effects.

Level 2 DTED data has a sample spacing of 30 m, and therefore for shadows that may be of the order of 15 m in length, should be adequate to correct terrain slope effects.

6.6.9 Conclusions

The ability to segment the shadows cast by vehicles in radar seeker imagery from their local background has been demonstrated using representative air-carry data gathered with the MRMS seeker. The ability to detect shadows in an automated manner, cued from a bright target detection, has been shown, and information on the dimensions of the object casting the shadow can be calculated.

Through analysis of the available data, the size information from the shadow cast by a moving vehicle appears to offer significantly improved discrimination between armoured and civilian classes than is possible using the blurred bright signature alone. Synthetic data was generated using the CLTG model and used to supplement the size measurement capability testing of shadow segmentation, and to test the quality of velocity prediction. The reduced location errors offer improved velocity and heading measurement accuracy compared with techniques that use the target's bright signature.

While shadows can provide useful additional information, there are a number of challenges to overcome to ensure that the results are not distorted by terrain slope and to understand where shadow information will not be available.

The use of shadow information for short-dwell SAR sensors offers good performance benefits. Should an application allow for an extended imaging time, such as for a surveillance UAV, the ability to generate an extended series of images, potentially with a change of aspect angle, would further enhance the capability.

7 Polarisation

7.1 Introduction

The use of a radar seeker enables a designated target to be selected under day/night, all weather conditions. A key factor in the radar seeker design is the polarisation used as this affects both the performance and cost of the system.

The choice of polarisation for the data gathering radar seekers used in the air-carry trials was circular polarisation rather than the linear polarisations traditionally used in surveillance radars. The received signals from the target are decomposed into cross- and co-polar components relative to the circular transmission polarisation. These components are often referred to as the odd and even channels which, as a simplification, can relate to the number of reflections the signal has experienced before returning to the radar for flat plates (or spheres), dihedral, and trihedral reflectors. This does not cover all scattering examples, with a dipole returning a linear polarisation that will have equal components in the co-and cross-polar channels [60].

Use of polarimetric features provides additional information on the scattering from the target, forming an important part of its signature. This information has potential to be used as a clutter discriminator if multiple channels are available, or to allow selection of a single channel that is best suited to the required task.

The use of two polarisation channels, coupled with a monopulse angle measurement capability gives rise to an expensive, multiple channel radar. By using a single polarisation, the radar would be less complex and the required number of receiver channels would be reduced.

This chapter assesses the performance implications of using a single circular or linear polarisation for an imaging radar seeker through the use of data from three radar sensors at 35 and 10 GHz.

The data from these sources is described in Section 7.3, and the process to convert linear polarisations to circular is detailed in Section 7.4. Sections 7.5, 7.6, and 7.7 assess the performance implications of polarisation choice for each of the three radars. Section 7.8 summarises the performance measurements and Sections 7.9 and 7.10 contain the conclusions and recommendations respectively.

7.2 Polarisation properties

The polarisation of a wave is a description of the shape traced out by the E-vector in a plane perpendicular to the direction of propagation. Under specific conditions this may trace a line or a circle, resulting in linear or circular polarisation, but the more general case is an ellipse.

The sense of rotation of the tip of the E-vector describes the 'handedness' of the wave. A wave is defined as right-handed if the rotation is clockwise for an observer looking down the direction of propagation [69].

The polarisation state is specified by the orientation angle Ø and ellipticity angle χ.
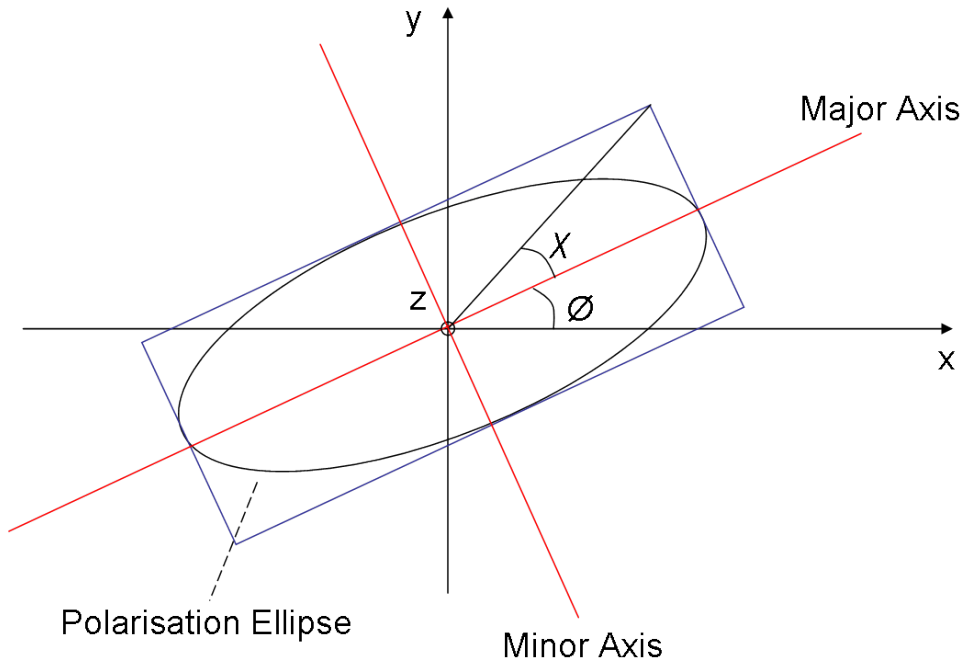
Figure 7-1 Polarisation ellipse in the x-y plane, ellipticity angle X and orientation angle Ø

The polarisation state can also be defined in terms of the Stokes' Parameters, I_0, Q, U and V which can then be written in vector form [69].

$$F = \begin{bmatrix} I_0 \\ Q \\ U \\ V \end{bmatrix} = \begin{bmatrix} |E_v|^2 + |E_h|^2 \\ |E_v|^2 - |E_h|^2 \\ 2\,\mathrm{Re}\,E_v E_h^* \\ 2\,\mathrm{Im}\,E_v E_h^* \end{bmatrix} = \begin{bmatrix} I_0 \\ I_0 \cos 2\phi \cos 2\chi \\ I_0 \sin 2\phi \sin 2\chi \\ I_0 \sin 2\chi \end{bmatrix} \tag{7.1}$$

Only three of the four parameters are independent since they are related by:

$$I_0^2 = Q^2 + U^2 + V^2 \tag{7.2}$$

The state of polarisation of a wave can then be mapped onto the surface of a sphere of radius I_0, called the Poincaré sphere, using the parameters Q, U and V as the Cartesian coordinates. The angles 2χ and 2ψ define the latitude and longitude of the point. The sign of χ defines the handedness, with the upper hemisphere relating to left-handed polarisations and the lower to right-handed polarisations. The two poles represent the circular polarisations. Linear polarisations are represented by points in the equatorial plane.
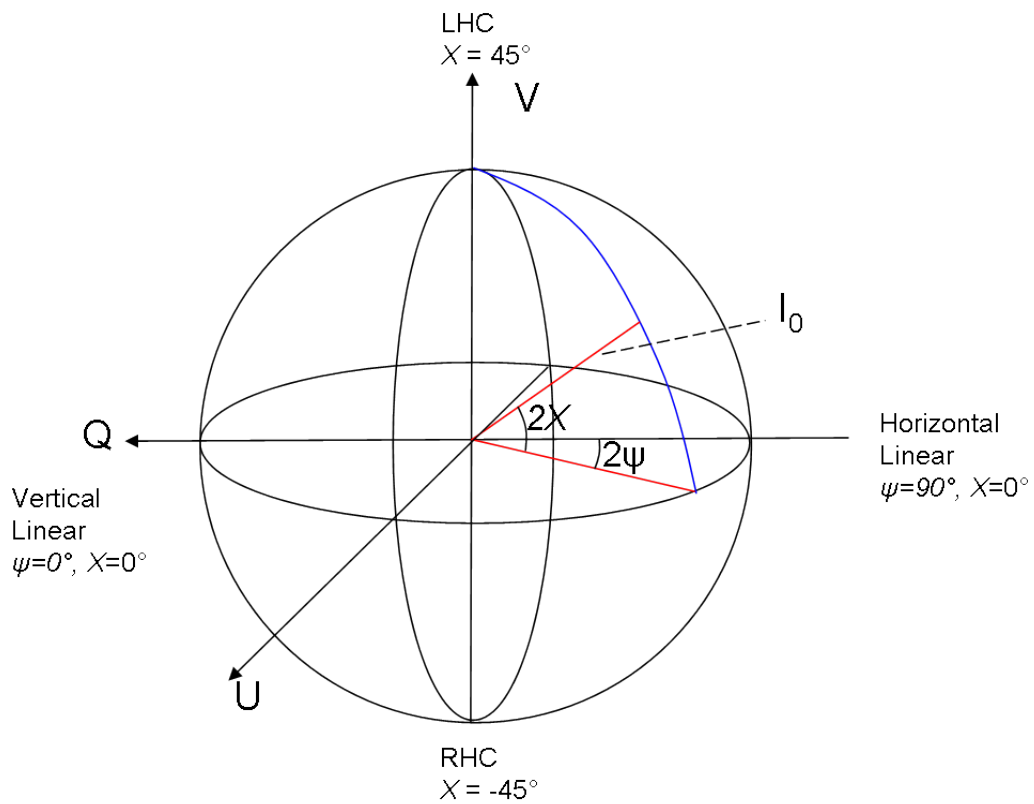
Figure 7-2 Poincaré sphere with Stokes parameters as axes

All information about the scattering properties of a target, imaged at a specific geometry and frequency, can be contained in a 2 x 2 complex matrix. Each element of the matrix represents the amplitude and phase of the returned signal for an orthogonal polarisation.

$$S = \begin{pmatrix} \sqrt{\sigma_{11}}e^{j\phi 11} & \sqrt{\sigma_{12}}e^{j\phi 12} \\ \sqrt{\sigma_{21}}e^{j\phi 21} & \sqrt{\sigma_{22}}e^{j\phi 22} \end{pmatrix} \tag{7.3}$$

To gather this information on a target would require two pulses transmitted with the polarisation switched between them, and the ability to receive both polarisations simultaneously. This may be appropriate for an experimental data-gathering system, but the extra complexity may exclude its use in more cost-sensitive applications.

For the scattering matrix, σ_{mn} is the amplitude of the reflectivity for the transmitted polarisation m and the received polarisation n, and $\emptyset_{mn}$ is the phase.

The polarisations typically considered are horizontal, vertical, left circular and right circular, but can be any orthogonal set. The polarisation matrix can be theoretically transformed into any other set.

Since the radars under discussion in this work are monostatic, the cross-polarisation terms can be considered equivalent, $\sigma_{12} = \sigma_{21}$ and $\emptyset_{12} = \emptyset_{21}$.

$$S = \begin{pmatrix} \sqrt{\sigma_{11}}e^{j\phi 11} & \sqrt{\sigma_{12}}e^{j\phi 12} \\ \sqrt{\sigma_{12}}e^{j\phi 12} & \sqrt{\sigma_{22}}e^{j\phi 22} \end{pmatrix} \tag{7.4}$$

Information on the specific scattering matrices for a range of object types such as cylinders, spheres and other di- and tri-hedral reflectors can be found in [27] and [60].

It is the difference in scattering properties between certain types of clutter and man-made objects which make the use of polarisation an important factor in achieving good seeker performance.

7.3 Available Data

Three sources of data were available for use within this study:

- 35 GHz data from the MBDA MRMS radar;

- 35 GHz data from the fully polarimetric MEMPHIS radar owned by FGAN of Germany, supplied under NATO Set 069;

- 10 GHz SAR data gathered by the Enhanced Surveillance Radar (ESR) owned by QinetiQ.

Data from MRMS is best suited to this study since the radar is most representative of the likely capability of a future seeker. The MEMPHIS and ESR radars are sideways-looking SAR sensors.

The FGAN and ESR radars are linearly polarised, and the HH, HV, VH and VV channels are recorded. Since this data is fully calibrated [70], it is possible to create circularly polarised images, as will be discussed in Section 7.4.

7.3.1 The MRMS seeker

The analysis of MRMS data in this report uses data gathered in the 2006 trial (Section 4.2). Whilst there is a limited quantity of data, the imagery is the highest quality currently available for a highly-squinted SAR sensor, representing the most developed build of the radar.

An example of the MRMS imagery used is shown in Figure 7-3, which is a rural scene with a number of vehicles present as well as decoy reflectors. A crossroads is visible slightly to the right of centre, and a plantation of trees is to the right of the image.

Figure 7-3 MRMS RDM Image of vehicles in a rural scene – odd sum polarisation

The MRMS data was recorded at a depression angle of 10° and a height of 400 m, giving a range to the imaged ground of 2300 m. Two military vehicles are present in Figure 7-3; an Air-Defence Unit (ADU) and a Self-Propelled Gun (SPG). Figure 7-4 shows these vehicles. Also present, but for only one of the two sorties considered, were an MBT, large and small tractors, and a Land Rover pick-up.

| ADU | SPG |
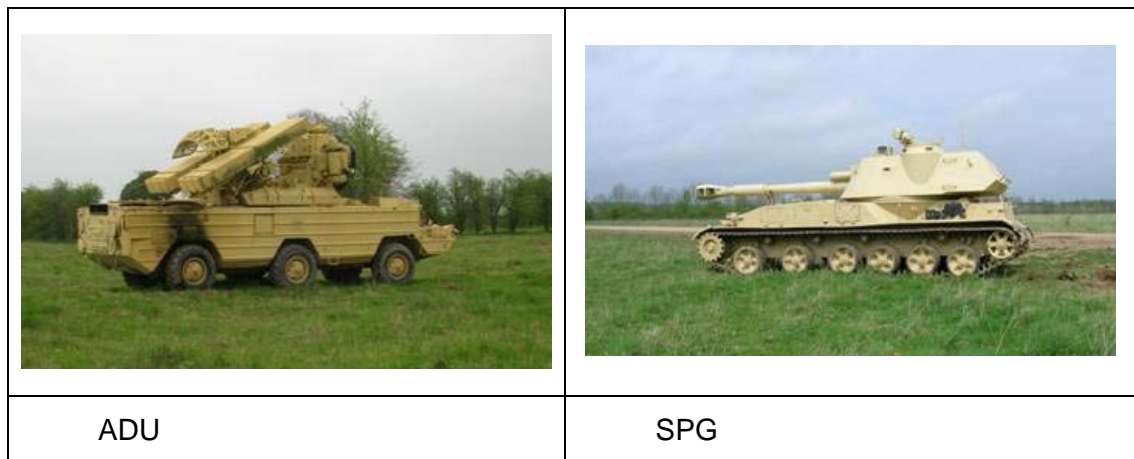
Figure 7-4 Military vehicles imaged at Copehill Plantation

7.3.2 The MEMPHIS radar

Data including a range of targets imaged with the FGAN Multi-frequency Experimental Monopulse High-resolution Interferometric SAR) MEMPHIS [70] (Multi-Frequency Experimental Monopulse High-resolution Interferometric SAR) radar was provided for use in this work by FGAN.

MEMPHIS is an airborne, sideways-looking radar with the ability to operate at 10, 35 and 94 GHz. The data examined here is 35 GHz, the same frequency as that used for the MRMS radar. The data were recorded with HH, VV and the cross-polarisations.

A sample image of targets from the MEMPHIS radar is shown in Figure 7-5. Only a small amount of FGAN data were available. The data were recorded from a height of 400 m at a depression angle of 20°, giving a slant range to the ground of 1200 m.

Figure 7-5 MEMPHIS stripmap SAR patch

Three flights over an array of targets were made available. The target array contained seventeen objects. The known vehicles were an MBT, an armoured personnel carrier (APC) and a civilian coach. Figure 7-6 displays photographs of the known vehicle types. Four different types of decoy were also present in the

array, constructed from trihedral, dihedral, and gridded trihedral reflectors as well as a short tower structure. There were a further 9 vehicles in the array for which the designation is not known.

In a single run there are six or seven looks at each target. Over the set of three runs there are 20 images of the MBT, 20 of the APC and 21 of the coach. A similar number of images were available of the decoy reflectors and the tower.

MBT	APC
Coach	Reference Point (tower)

Figure 7-6 Objects imaged with the MEMPHIS radar (images from [70])

7.3.3 The Enhanced Surveillance Radar

ESR was flown by DERA / QinetiQ providing high-quality fully polarimetric SAR data. This radar operates at 10 GHz rather than the 35 GHz at which MEMPHIS and MRMS operate.

Since the antenna can operate in a spotlight SAR mode, a target can be illuminated for an extended period as the platform flies past. By forming repeated SAR images at different angles, it is possible to achieve many individual looks at a target.

Data are available from ESR airborne trials of military vehicles. Targets A, B, D and G are military vehicles. Decoy rubber vehicles, a tractor and trailer and a water tower were also imaged.

The SAR imagery was collected in spotlight mode and processed at 0.4 by 0.4 m resolution with a pixel spacing of 0.3 by 0.3 m. Sets of target 'chips' were extracted from the imagery for further analysis. Figure 7-7 is a multi-angle composite of SAR images from the trial with the object types labelled.
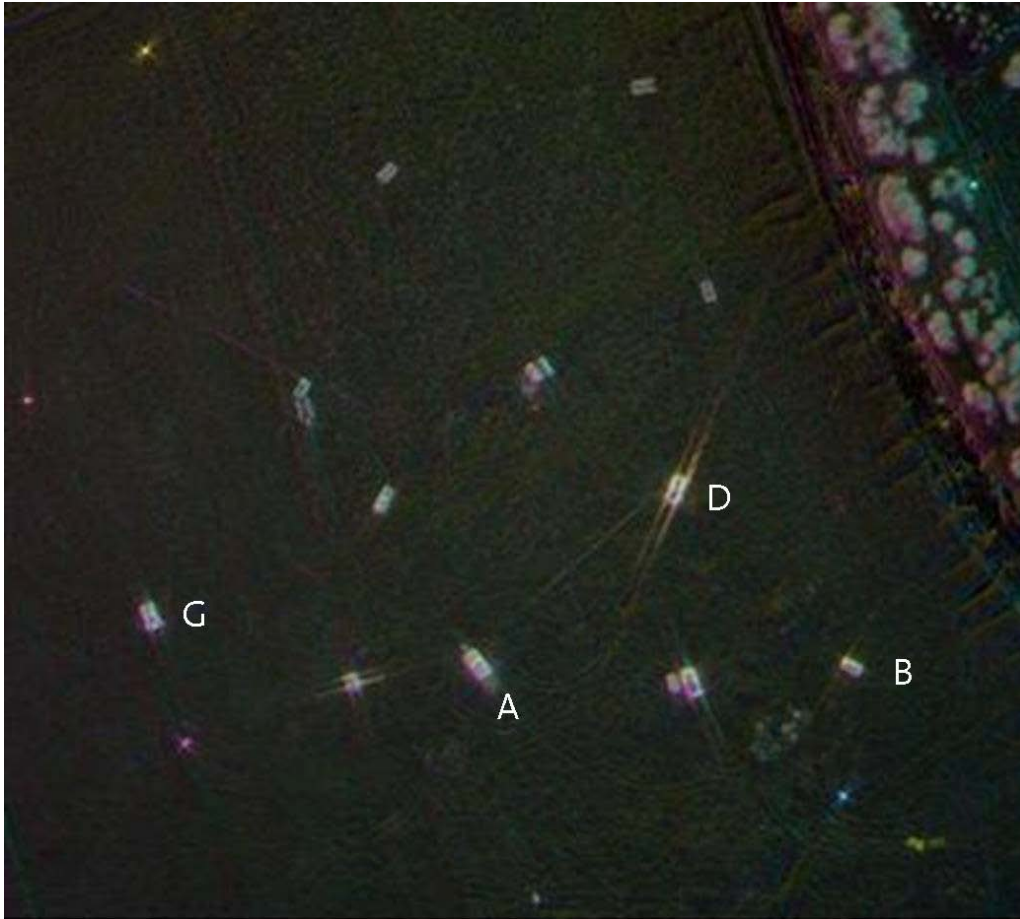
Figure 7-7 Sample SAR Image with annotated targets and reflectors

Figure 7-8 shows examples of target signatures at different orientations to the radar. The middle three images of target B in the upper row have more clutter visible due to lower reflections from the target at that angle and the nature of the image scaling.
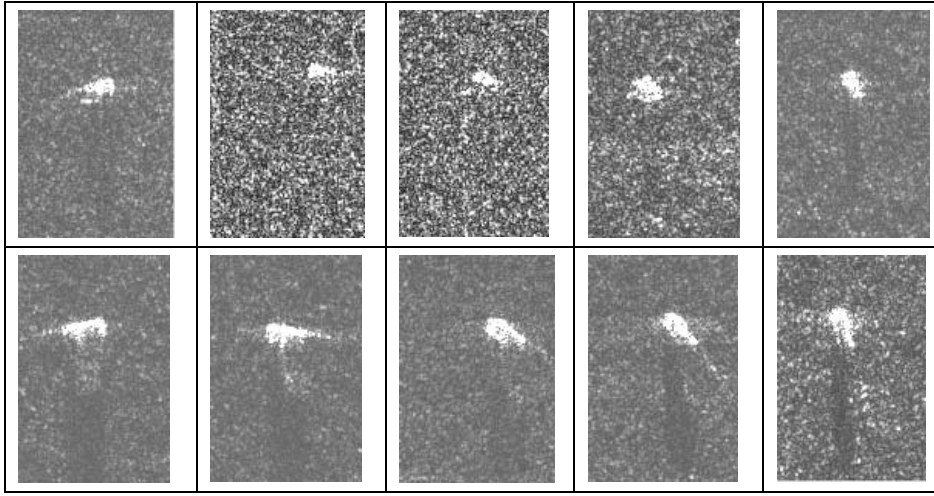
Figure 7-8 Target B (above) and Target A (below) image chips at 5 degree intervals

The four channels of this data have been polarimetrically calibrated through the use of dihedral and trihedral reflector targets. Calibration errors arise in the form of channel imbalance and cross-talk. For the ESR radar, the cross-polar isolation is better than 30 dB and therefore cross-talk can be considered to be a negligible effect. The significant effect is channel imbalance which can be measured using the responses from the dihedral and trihedral reflectors. Trihedral and dihedral scattering responses are given by

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \tag{7.5}$$

and

$$\begin{pmatrix} \cos(2\theta) & \sin(2\theta) \\ \sin(2\theta) & -\cos(2\theta) \end{pmatrix} \tag{7.6}$$

respectively, where θ is the seam angle of the dihedral.

Calibration of orthogonal linear polarisations can be performed by making use of a property of the dihedral reflector. If a linearly polarised wave is incident at an angle *θ* relative to the seam between the faces, the reflected wave is linearly polarised oriented at an equal but opposite angle to the seam, Figure 7-9.
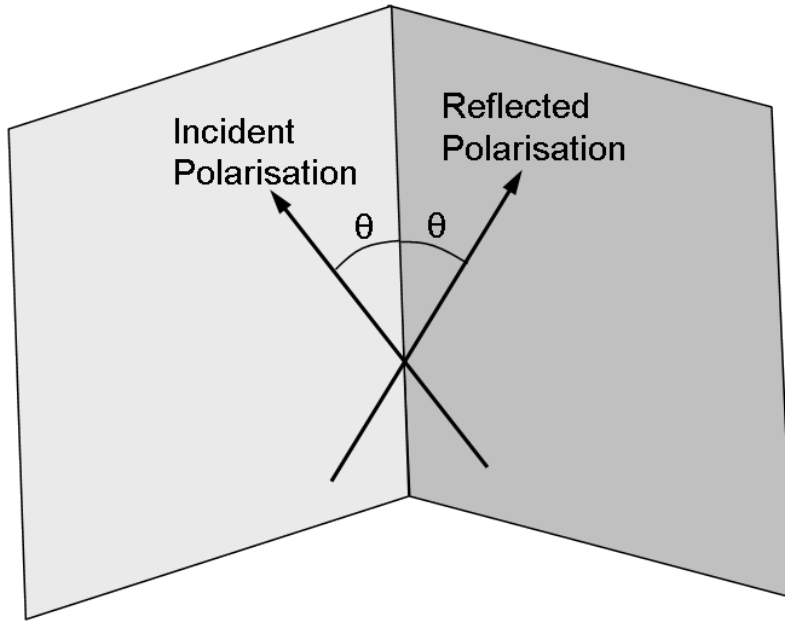
Figure 7-9 Polarisation reflection of a dihedral corner reflector

If the dihedral seam angle is set to 22.5 degrees, as for the trial, the dihedral response becomes

$$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad (7.7)$$

In this case, an incident linear polarisation will return a 45° polarisation relative to the incident wave, reflecting equal power in both the incident and orthogonal polarisations [27]. This response can be used to calculate a scaling factor for each channel to match that of the HH return.

A trihedral reflector produces strong returns in the HH and VV channels but no returns in either of the cross-polar channels.

The mean and standard deviation of the amplitude and phase of the HV, VH and VV channel balance factors relative to the HH channel were calculated with the results shown in Table 7-1. It can be seen that, with the mean amplitude close to unity but with a high standard deviation, it is acceptable to choose a scaling factor of unity for the amplitude. The phase values have relatively narrow spreads about the mean values. Thus, there are significant phase offsets from zero and it is these which embody the channel imbalances.

	amplitude mean	amplitude s.d.	phase (degrees) mean	Phase (degrees) s.d.
factor 2	1.16	0.29	116.7	3.5
factor 3	1.06	0.36	131.0	4.0
factor 4	1.03	0.12	-114.0	4.7

Table 7-1 Measured properties of channel imbalance factors

These data, therefore, appear to be ideal for analysis of the performance difference between linear and circular polarisations, although at a lower frequency than would be used by a missile seeker, resulting in slightly different scattering physics.

To facilitate this analysis, 335 target chips formed a training set for each vehicle, and 15 patches of rural clutter consisting of trees and grassland were used. Polarisation conversion was applied and the images stored in a file format suitable for input to the target detection process within the ATR software (Chapter 5.1).

7.3.4 Target sampling

In order to analyse the effects of varying polarisation on target detectability, it was necessary to create a database of target images which could then be processed. For RDM images such as that shown in Figure 7-10, where a number of targets are present in the scene, it is necessary to extract the target signature and an area of surrounding clutter, but to exclude returns from other targets. Typical extraction areas of the targets and decoys in the scene to produce a 'target chip' are highlighted by the white boxes.
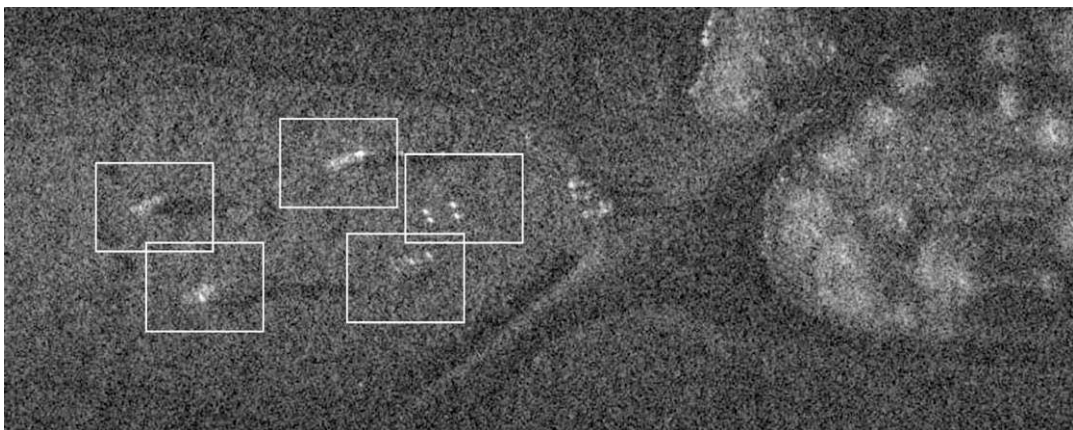
Figure 7-10 RDM Image with extracted target chip areas highlighted

For each of the three sensors, all instances of an object were extracted to form a database of image chips. Analysis of the radar signature of the extracted object

requires the ability to segment the object's returns from those of the background clutter (the snake segmentation technique described in chapter 5.1 was used).

Combining the image with the mask created by the segmentation allows calculation in both radar channels of the target to clutter ratio (TCR), which is a measure of the relative brightness of the target to the surrounding clutter, and one of the key factors in detectability. Having an accurate segmentation of pixels which belong to an object also allows the calculation of a number of size-based and statistical features that can be used to define the object.

7.4 Polarisation Conversion

7.4.1 Introduction

To compare the performance difference between linear and circular polarisations, the linearly polarised radar data needed to be converted into circularly polarised radar data in both the cross- and co-polar channels, also referred to as odd and even channels. This is done by means of a matrix multiplication described in the following section. To show that the conversion has been performed accurately, the circular response of trihedral and dihedral reflectors will be compared between ESR and MRMS data. Note, no calibration reflectors were identified in the MEMPHIS data set.

7.4.2 Theory

The MEMPHIS and ESR radars described in Section 7.3 operate with a linear polarisation instead of the circular polarisation used by MRMS. It is possible to convert fully polarimetric linear radar into odd and even circular polarisations.

The scattering characteristics of a point can be represented by a polarimetric time-independent scattering matrix. For a monostatic radar with linear polarisations H and V, this can be represented as

$$[S] = \begin{pmatrix} S_{hh} & S_{hv} \\ S_{vh} & S_{vv} \end{pmatrix} \tag{7.8}$$

Right and left-hand circularly polarised radar can be represented as a combination of H and V polarisations with a phase offset.

169

$$Right(R) = \frac{1}{\sqrt{2}}(1, i) \qquad (7.9)$$

$$Left(L) = \frac{1}{\sqrt{2}}(1, -i) \qquad (7.10)$$

where HH and VV are represented as

$$HH = (1, 0) \qquad (7.11)$$

The scattering matrix for right-hand circular polarisation is

$$[S'] = [R][S][R]^{T} \qquad (7.12)$$

where [R] is the right-circularly polarised component from equation (7.9). When multiplied through, this allows the calculation of the even (RR and LL) polarisations as well as the odd bounce (RL and LR). Expanding equation (7.12) for the even and odd bounce returns gives the power in each channel.

$$Even = \frac{1}{2}\left| S_{vv} - S_{hh} \right| + 2\left| S_{hv} \right|^{2} \qquad (7.13)$$

$$Odd = \frac{1}{2}\left| S_{hh} + S_{vv} \right|^{2} \qquad (7.14)$$

These formulae allow comparison of the linear and circular polarisations with respect to their ATR performance.

7.4.3 Testing

Examination of reflectors in the data allow some analysis to be made as to whether the conversion is working correctly. Dihedral reflectors should reflect strongly in the even channel, and trihedral reflectors in the odd channel. Reflector data were available for the MRMS and ESR radars.

ESR data

The ESR data contains reflections from both dihedral and trihedral reflectors. A reflector of each type was examined over a range of angles, using a TCR comparison based on segmentation of the reflector's signature against the background.

Table 7-2 shows the TCR values for a trihedral and dihedral reflector, for all available polarisations, and averaged over a number of viewing angles. The

trihedral is 16.5 dB brighter in the odd channel than the even channel, and equally, the dihedral is 13 .3 dB brighter in the even channel than in the odd channel.

	TCR (dB)				
	HH	VV	HV	Odd	Even
Trihedral	35.11	34.60	7.58	36.04	21.26
Dihedral	31.81	31.27	35.04	19.52	34.57

Table 7-2 TCR values for Trihedral and Dihedral Reflector

The linear polarisations also exhibit the expected scattering characteristics from the reflector types. A trihedral reflector produces a scattering matrix response shown by equation (7.5). Hence, we would expect strong returns in HH and VV and a weak return in HV. The results agree with this, as HH and VV have a TCR of ~35 dB, with the TCR for HV 27 dB below this level. That the linear polarisations exhibit a larger difference between the channels than the circular polarisations indicates that the calibration for the polarisation conversion may not be perfect. However, a channel separation of approximately 15 dB as measured between the odd and even bounce channels is greater than that measured in airborne trials by the MRMS radar.

A dihedral reflector produces a scattering matrix response given by equation (7.6), where θ is the seam angle of the dihedral measured from the vertical plane. If a seam angle of 22.5° is used, then the dihedral scattering matrix contains the values shown in equation (7.7). An equal amplitude in the HH, VV and HV channels would be expected. The dihedral results show the amplitude to be very nearly equal, with the HV channel slightly brighter than HH and VV. The increase in brightness is most likely caused by the dihedral seam angle (the angle between the seam at the back of the reflector and the vertical) being slightly greater than 22.5°.

MRMS data

The second dataset with reflectors present was the MRMS data. No dihedrals were present in any of the imagery, but a number of trihedrals were present, allowing testing of the channel separation between the odd and even channels.

During the data gathering sorties, few reflectors were imaged while the radar operated in Search mode (odd and even polarisations), whereas many reflectors were imaged in Track mode which uses odd polarisation only.

However, an array of small trihedral reflectors was present in the middle of the target array, and this was used to measure the difference in TCR between the odd and even channels. The reflector array is shown in Figure 7-11.
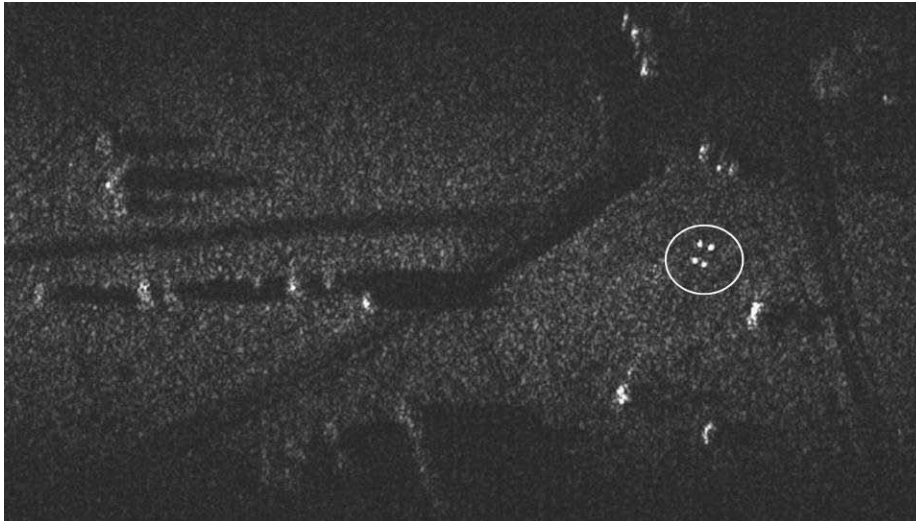
Figure 7-11 Four trihedral reflectors (highlighted)

The TCR values from these reflectors for the odd and even polarisations are shown in Table 7-3. As expected, the odd channel is brighter than the even channel, but the average difference, 7dB is lower than expected. It was found that this was due to incorrect orientation of the reflectors for that imaging geometry. The bulk of the data was gathered with the aircraft flying on a different track. It is interesting that the odd channel response, which may come primarily from the tripod and reflector mountings, is noticeably brighter than the even channel.

	TCR (dB)		
Ref	Odd	Even	Even / Odd
1	20.0	13.4	-6.6
2	19.9	11.8	-8.1
3	19.9	12.8	-7.1
4	16.6	10.3	-6.3
Mean	19.1	12.1	-7.0

Table 7-3 TCR values for reflectors at the Copehill plantation

Imagery of much larger reflectors in both polarisations was gathered during a flight. An image containing an array of trihedral reflectors, Figure 7-12, was analysed using the same process as above.

Table 7-4 shows the results for these reflectors. Due to the reflectors' greater size (500 m^2 compared to 14 m^2), the TCR values are much higher than for the set of four reflectors from Figure 7-11. The average E / O ratio increases to a mean value of -11.9 dB, which is likely to reflect the true performance of the system more closely as the effects of background clutter become less dominant.
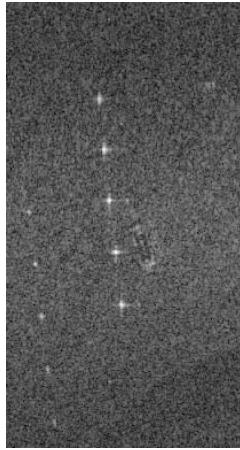
TCR (dB)			
Ref	Odd	Even	E / O
1	32.5	24.8	-7.7
2	36.5	25.0	-11.5
3	34.8	19.7	-15.1
4	32.8	19.6	-13.2
Mean	34.2	22.3	-11.9

Figure 7-12 Trihedral reflectors at Boscombe   *Table 7-4 TCR values for reflectors at Boscombe*

The difference in TCR between the reflectors in Table 7-3 and Table 7-4 is related to the Radar Cross Section (RCS) of the reflectors which form each array.

The difference in RCS should result in the 500 m^2 reflectors having a TCR 15.5 dB greater than that of the 14m^2 reflectors. The difference between the measured average in the odd channel is 15.1 dB, very close to the predicted figure.

The MRMS radar has 4.6 dB less channel separation than the ESR radar. This may be due to more cross-talk between radar channels in the MRMS radar than was present in the (more sophisticated) ESR radar.

7.5 FGAN Analysis

7.5.1 Odd/Even comparison

The FGAN data sets, over three imaging angles, had all images converted from linear to circular polarisation. Objects were then segmented to allow TCR measurements to be performed over the range of polarisations, Table 7-5.

Type	Samples	TCR (dB)		
		Odd	Even	Even/Odd
MBT	20	9.50	10.41	0.91
APC	20	9.20	10.51	1.31
Bus	21	10.44	9.93	-0.51
ref	19	10.70	9.88	-0.82
D1	21	7.57	9.80	2.22
D2	19	7.05	8.36	1.31
D3	21	10.86	11.42	0.57
D4	20	12.26	12.11	-0.14

Table 7-5 Target to clutter ratio for objects imaged by MEMPHIS

The two military vehicles which are examined in this data show that the contrast in the even channel is, on average, 1.1 dB brighter relative to the background than the odd channel. This trend was visible across all 20 samples and the three viewing angles.

The decoy targets also exhibit increased brightness in the even channel. This is to be expected due to the design of the decoys which are multiple cylinders and small trihedrals upon a flat metal plate. This will produce mainly a dihedral response from the plate / cylinder interface (as well as a smaller trihedral reflection). Decoy 4 has a similar response in odd and even channels, which is not expected for a trihedral reflector where the odd polarisation should dominate. It is believed that decoy 4 is a gridded variant of a trihedral. Figure 7-13 displays photographs of decoy 1 and decoy 4.
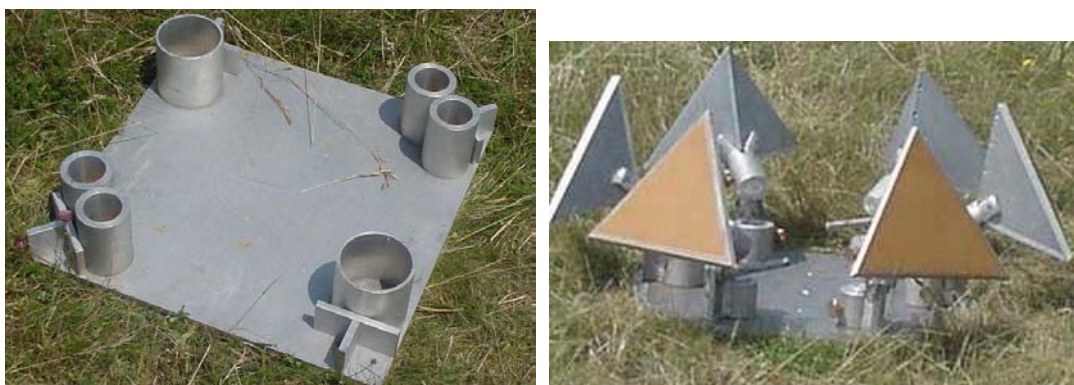
Figure 7-13 Decoy 1 (left) and Decoy 4 (right)

7.5.2 Linear polarisations

Target to clutter ratio measurements of objects imaged by MEMPHIS, including the linear polarisation results, are shown in Table 7-6.

Type	TCR (dB)				
	HH	VV	HV	Odd	Even
MBT	10.37	9.46	8.89	9.50	10.41
APC	9.73	9.90	9.46	9.20	10.51
Bus	10.70	10.41	6.29	10.44	9.93
Tower	12.08	8.89	5.74	10.70	9.88
D1	9.51	9.27	3.97	7.57	9.80
D2	8.74	7.90	3.66	7.05	8.36
D3	12.70	10.58	5.67	10.86	11.42
D4	15.37	5.30	5.50	12.26	12.11

Table 7-6 Complete target to clutter ratio for objects imaged by MEMPHIS

Splitting the data into three groups, armoured vehicles, complex decoy reflectors (D1 and D2) and complex trihedral reflectors (D4) yields Table 7-7.

Type	TCR (dB)				
	HH	VV	HV	Odd	Even
Armour	10.05	9.68	9.18	9.35	10.46
D1&2	9.13	8.59	3.82	7.31	9.08
D4	15.37	5.30	5.50	12.26	12.11

Table 7-7 Summary of TCR by target type and polarisation

It can be seen that, for the military vehicles, the HH polarisation has slightly higher TCR compared to VV. However, the average difference is less than 0.4 dB. Of all the polarisation channels (HH, VV, even, odd), even polarisation produces the best TCR and odd the worst. The results from the HH and VV linear polarisations fall between the circular polarisations, and the HV linear polarisation performs worst.

The improvement in TCR from choosing to use an even circular polarisation is likely to be reflected in lower false object densities since a higher constraint can be applied for target brightness when discriminating targets from clutter.

Decoys D1 and D2 exhibit a strong response in the HH and VV channels and a much weaker response in the HV channel, with a TCR of less than 4dB. The low cross-polar response is to be expected due to the decoy's construction from dihedral and trihedral reflectors.

The trihedrals should produce a response in HH and VV and no response in the cross-polarisations. Since the seam angle of the dihedral reflectors is 90 degrees, there should again be no component of reflection in the cross-polar channels and strong reflection in HH and VV.

The positive contrast which remains in the HV channel could be due to multiple reflections between the components of the complex reflector.

Decoy 4 is believed to be a gridded trihedral, in which a grid of vertical wires covers the open face. The effect of this surface is to reflect the vertical component of the incoming wave, thereby acting as a flat plate, while allowing the horizontal component to pass through and be reflected back by the trihedral itself. Therefore, the reflector will appear as a trihedral in HH and as a flat plate in the other polarisations.

The measurements from this reflector agree with the theory – the average HH TCR is the largest for all the objects at 15.4 dB, while the returns in the other polarisations are 10 dB below this value.

Given a strong signal in HH polarisation, and a weak signal in both VV and HV, it is expected that the returns in odd and even circular polarisations should be roughly equal, as shown in Table 7-5.

7.6 MRMS Analysis

7.6.1 TCR analysis

Unlike the FGAN radar described in Sections 7.3.2 and 7.5, the MRMS radar did not record linear polarisations, only the odd and even circular polarisations. This removes the need to perform any conversions to achieve circular polarisation, and allows direct comparison between the odd and even channels.

TCR measurements of objects imaged by MRMS are shown in Table 7-8.

Type	TCR		
	Odd (dB)	Even (dB)	Even/Odd (dB)
MBT	14.78	16.26	1.49
SPG	17.51	19.64	2.13
ADU	14.37	16.81	2.44
Land Rover	19.29	20.42	1.12
Big Tractor	11.57	13.23	1.66
Small Tractor	9.76	14.32	4.56
Tree	8.58	8.01	-0.57

Table 7-8 Target to clutter values for objects at Copehill Plantation

As with the MEMPHIS radar, the targets exhibit stronger returns in the even channel compared with the odd channel. The three military vehicles are, on average, 2 dB brighter compared with the background in the even channel.

The civilian vehicles also exhibit stronger returns in the even channel and these are especially prominent on the small tractor. However, only a few samples were available of these vehicles, hence, the results must be taken as indicative rather than established values.

A number of clutter objects consisting of trees within the scene were also extracted and the signal to clutter ratio measured. It was found that the ratio in the even channel was ~0.5 dB lower than in the odd channel. This gives a 2.5 dB increase in separation between targets and clutter objects when switching from the odd to even channel.

A greater target to clutter separation will help to reduce false alarms, although the residual false object density (RFOD) of rural clutter is already low using the odd channel and the latest implementation of the ATR algorithm.

7.6.2 TCR effects on the Residual False Object Density

The TCR of an object is an important feature used to help separate segmented targets from clutter. The feature is one of a number of constraints which are applied to detected objects to aid in the target / clutter decision process. Use of a channel with superior target TCR is also likely to help at the CFAR stage of the detection process.

Since vehicles appear to have a higher TCR in the even channel than the odd channel, the TCR constraint on detection can be raised by a suitable amount to make use of this difference. Since the average increase in TCR was 2 dB in Table 7-8, the constraint was raised by this level and a number of rural and urban clutter scenes passed through the ATR algorithm.

The false alarm rate for rural clutter is already extremely low in the odd polarisation, and, given the limited number of polarimetric datasets, no difference was found between the odd and even channels when examining multiple flights over rural clutter on Salisbury Plain.

Urban clutter was taken from imagery of the Defence Academy at Shrivenham, which contained a number of buildings. This provided a large, varied urban area. It was found that urban areas appeared to have a number of strong dihedral scatterers, and the immediate benefits of using the even channel with a higher TCR were not significant. Table 7-9 shows the measured false object count for five flights over Shrivenham using constraints developed on the odd polarisation data and then applied to the even channel data with the TCR threshold set at the same level, and then increased by 1 and then 2 dB.

Run	Odd	Even		
	0 dB	0 dB	1 dB	2 dB
3	10	14	14	14
6	18	15	9	9
7	9	12	11	10
9	7	14	14	12
10	16	18	15	12
Average	12	14.6	12.6	11.4

Table 7-9 False object counts for Shrivenham with varying TCR threshold

For the same levels of threshold, the even channel has a slightly higher urban false alarm rate than the odd channel, and only when the 2 dB increase in threshold is applied does the false alarm count drop below that for the odd channel. Therefore, in this respect, the odd and even channels have similar performance.

The constraints applied were optimised for odd channel data. Since there is a much larger set of target signatures in the odd channel, it is probable that the constraints are not optimised for the even channel. The feature study in section 7.7 shows that, using 10 GHz SAR data, the even channel provides better clutter separation on certain key features than the odd channel. Given the ability to train the algorithm on a set of even polarisation target signatures, it should be possible to reduce the number of false targets detected further.

7.7 ESR Analysis

7.7.1 TCR analysis

The four military vehicles identified in Section 7.3.3 were processed through the segmentation code. Images which were well focused and where the segmentation performed correctly were selected for analysis, and the TCR values over a range of aspect angles are displayed in Table 7-10. There are a far higher number of samples in the ESR data analysis, allowing outputs to be plotted smoothly.

	TCR linear			TCR Circular			
Target	HH	HV	VV	Odd	Even	E/O	E/O Std Dev
D	19.81	16.34	19.17	18.47	19.43	0.96	1.18
G	19.39	15.62	19.19	17.70	19.47	1.77	1.14
A	19.56	16.75	19.54	17.97	19.76	1.79	2.47
B	19.68	16.57	19.38	18.55	19.00	0.44	2.29
Mean	19.61	16.32	19.32	18.17	19.41	1.24	1.77

Table 7-10 TCR measurements for ESR trial vehicles

As with the previous two radars, it is again noted that the returns in the even channel are brighter than those from the odd channel, with an average difference of 1.24 dB.

There is sufficient variability of the data that the even channel is not always brighter than the odd channel, as can be seen by the standard deviation of the even/odd TCR ratio. However, given the overall higher performance of the even channel, it would be the better choice of the circular polarisations.

Unlike the results from the FGAN radar, at 10 GHz, the HH polarisation has slightly higher TCR compared to the even polarisation, although the average difference over all the target data is only 0.27 dB.

7.7.2 Feature analysis

Since a large dataset was available for the ESR sensor, it was possible to look at the effects of polarisation on statistical features of the data as well as the detectability.

For each of the polarisations, the targets and sections of rural clutter (trees and grassland) were used to form a dataset. On average, 265 images of each target

were used, and 340 detections from rural clutter formed the potential false object set. The data sets from MRMS and FGAN had approximately 10 and 21 images of each target in both polarisations. Hence, the data yield is too low to perform comparable analysis.

Since the task of detecting and extracting the targets from the background was not challenging in this imagery, the tests were based on measured features that aid in clutter rejection and target discrimination.

The features can be divided into size-based and statistical groups. To measure size, detected objects in the image are subject to a segmentation technique and the length and width of the segmentation have been measured here. The statistical features examined are the normalised standard deviation of power (NSDP), which measures the variation of power within the target signature, and the target to clutter ratio (TCR), which compares the relative brightness of the target's signature to the surrounding area. Where x_i are n returns from within the target outline, and x_j are m returns from the localised background clutter, the NSDP and TCR are

$$NSDP = \frac{\frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})^2}{(\frac{1}{n}\sum_{i=1}^{n}x_i)^2} \tag{7.15}$$

$$TCR = 10 * \log 10 \frac{\frac{1}{n}\sum_{i=1}^{n}(x_i)^2}{\frac{1}{m}\sum_{j=1}^{m}(x_j)^2} \tag{7.16}$$

These features were calculated over the full data set by passing the target and clutter sets through the ATR algorithm.

Target B

The NSDP for Target B is shown in Figure 7-14. This plot is of ROC (Receiver Operating Characteristic) curves, and shows the quality of separation of target to clutter using this feature. Targets tend to have a higher NSDP value than clutter objects. By setting an NSDP threshold, above which an object would be considered as a target, it is possible to see how the probability of classification as a target (P_d) varies compared with the probability that a clutter object is declared a target (P_{fa}). A perfect result would give a P_d of 1 and a P_{fa} of 0. The closer the curve lies to the left of the graph the more effective the feature is in discriminating between target and clutter.

The P_{fa} axis on these plots is only suitable as an indication of the relative performance of the four polarisations examined. Since only a single feature is examined at a time, the P_{fa} is not representative of what would be expected in a full system. The useful information from the graphs, however, comes in the relative positions of the plots for the different polarisations and their ranking in terms of performance.

The blue line represents the results for the even polarisation, the red line the results for the odd polarisation. It can be seen that the even polarisation gives improved target to clutter separation over the odd polarisation for this clutter type. The black and yellow lines represent HH and VV respectively. There is little difference between the performances of the two linear polarisations and, when compared with the circular polarisations, the quality of results is closer to that of the even polarisation than the odd.
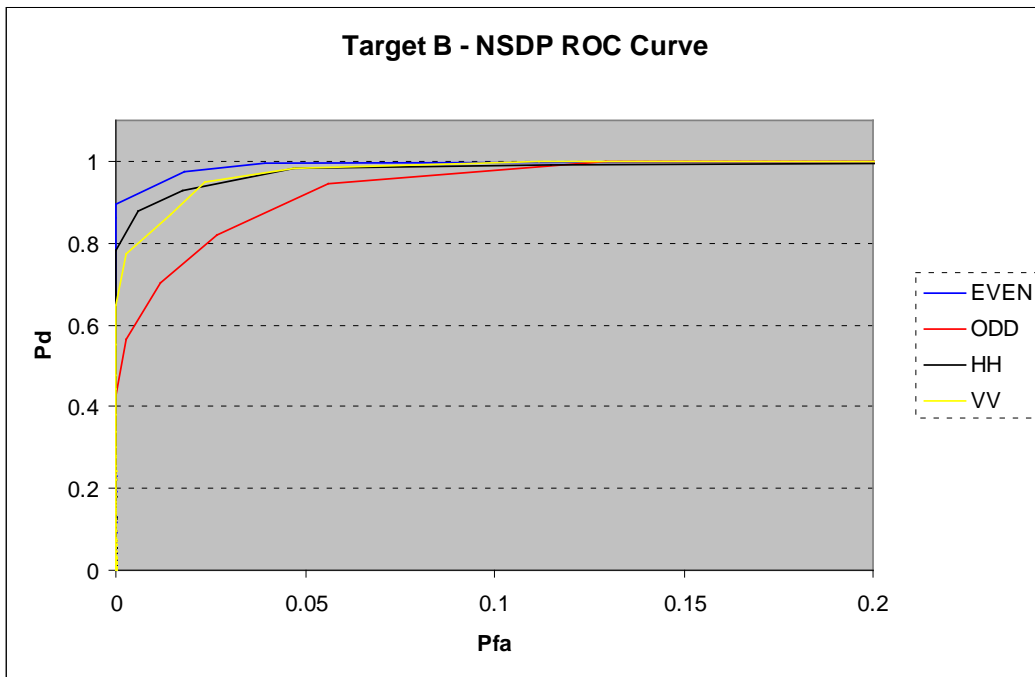
Figure 7-14 NSDP for Target B in linear and circular polarisations

The separability in the even channel may be due to the nature of dihedral reflections which produce the even channel response. If the target signature consists of a few bright dihedral responses, it will score more highly in the NSDP feature than it would for a more even spread of flat-panel or trihedral responses.

TCR has been analysed as a measure of target detectability against the background, but is also one of the key features in separating targets from clutter. While rural clutter such as trees may have similar length and width to a vehicle, the relative brightness against the background is normally much lower. Figure 7-15 shows ROC curves for the Target B in linear and circular polarisations. Here the difference between the polarisations is even more pronounced than for the NSDP. Again, the even circular polarisation has the best performance, with a P_d of near 1 reached for a very low rate of false alarms. The odd polarisation performs worst and the performance of the linear polarisations lies between the two, with VV slightly outperforming HH. The clutter objects used do not pose a challenge when separating them from clutter in the even channel response, hence the step-nature of the plot.
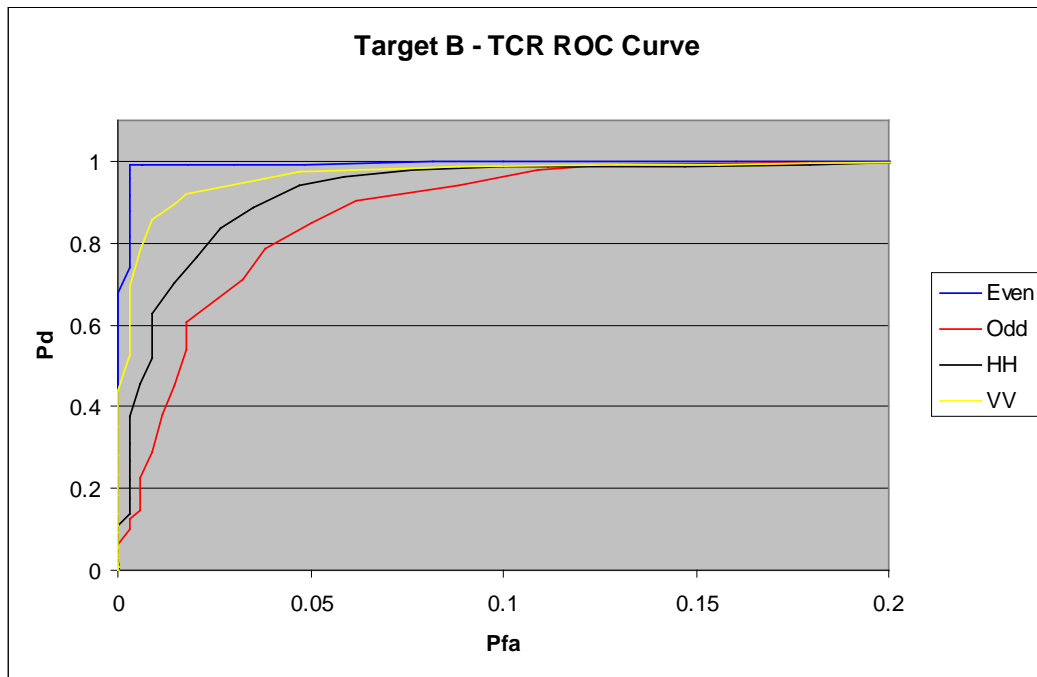
Figure 7-15 TCR for Target B in linear and circular polarisations

Target A

The same internal statistics were calculated for the Target A as for Target B.

Target A showed greater separation from the clutter than Target B, and to make the ROC curves more interpretable, the P_{fa} is plotted on a log scale, accentuating the separation between the curves.

Figure 7-16 shows the NSDP for the Target A in circular and linear polarisations. As before, the even channel outperforms the odd channel, and the performance of the linear polarisations lies between the two, although with the HH polarisation performing marginally better than VV.

The ROC curve for the TCR feature is shown in Figure 7-17. Again the even channel has the best performance, the odd channel the worst, and in this example the VV polarisation appears to outperform HH.
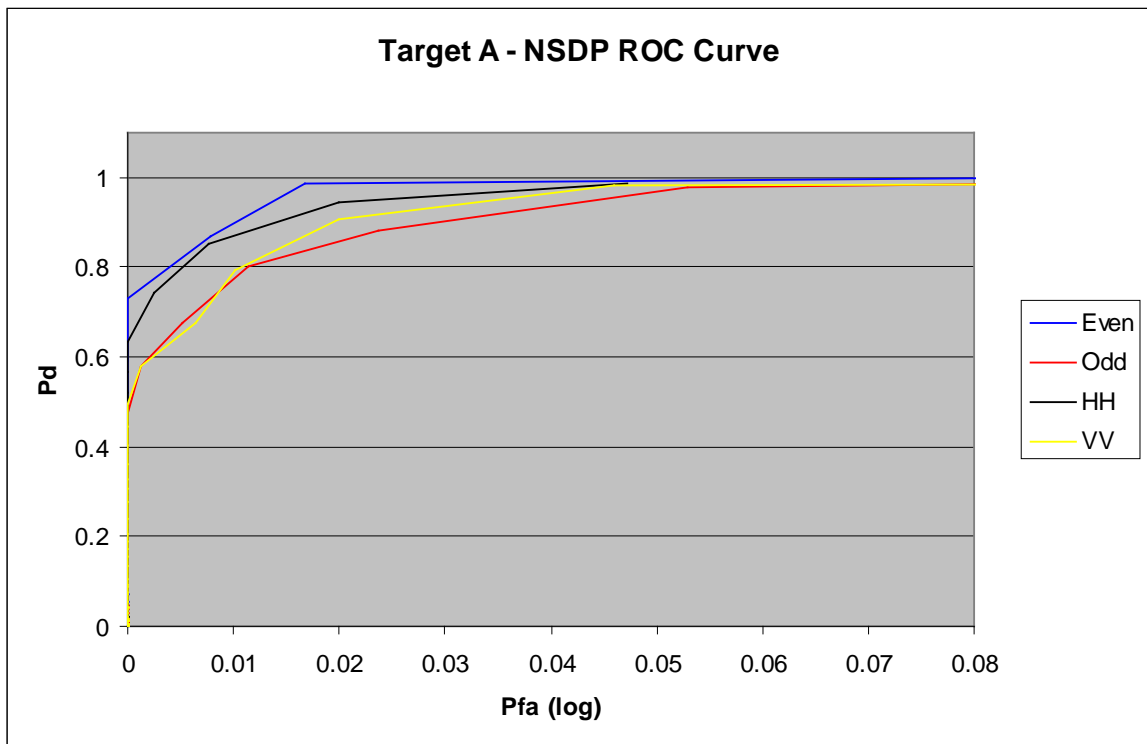
Figure 7-16 NSDP for Target A in linear and circular polarisations
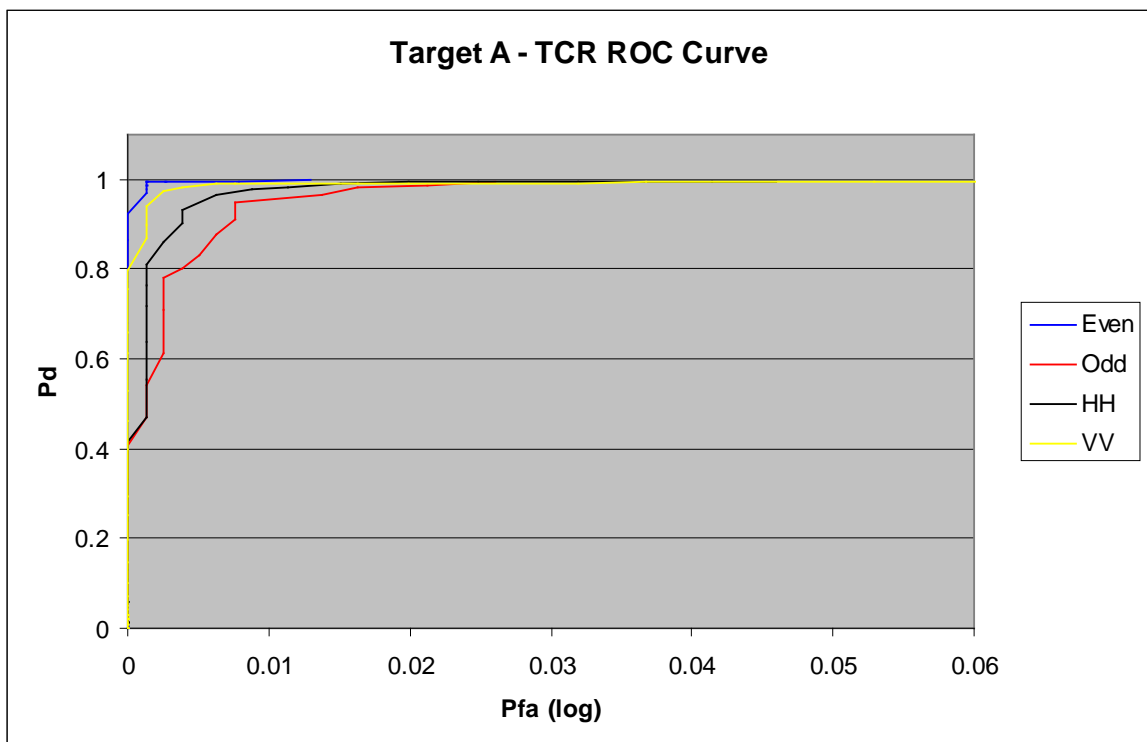
Figure 7-17 TCR for Target A in linear and circular polarisations

Object Sizing

The results from size measurements of the full data sets of two of the targets are shown in Table 7-11 for Target B and Table 7-12 for Target A.

The measured size of Target B is slightly larger in both length and width than its true size. There is little difference in the size estimates from the different polarisations. The standard deviation of measurements is lowest for the even polarisation, potentially due to the higher TCR allowing more consistent segmentation.

	Target B			
Polarisation	Length (m)	Std Dev (m)	Width (m)	Std Dev (m)
Even	7.57	1.05	4.56	0.59
Odd	7.43	1.14	4.27	0.64
HH	7.50	1.07	4.49	0.67
VV	7.50	1.06	4.44	0.64

Table 7-11 Sizing measurements for Target B

The Target A measurements are shorter in length than reality and roughly correct for width. Structure on Target A can cause significant shadowing of the rest of the target, and it is likely that this is causing the measurements to be smaller than expected. Images of Target A side-on to the radar, at 45 degrees, and head-on are shown in Figure 7-18, where the increasing self-shadowing becomes evident. Table 7-12 shows that the even polarisation has the smallest deviation on the results, although the small variation between the polarisations indicates that size measurement is not an area where performance will be affected since all polarisations have adequate TCR for good segmentation.

	Target A			
Polarisation	Length (m)	Std Dev (m)	Width (m)	Std Dev (m)
Even	6.30	0.64	3.41	0.54
Odd	6.34	0.73	3.44	0.68
HH	6.35	0.64	3.54	0.70
VV	6.33	0.78	3.42	0.67

Table 7-12 Sizing measurements for Target A
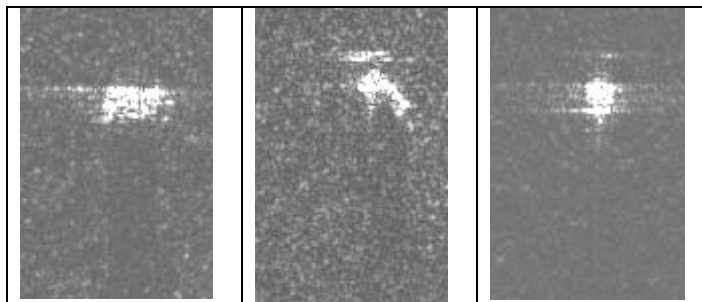
Figure 7-18 Images of Target A: side-on, at 45 degrees and head-on

7.8 Results Summary

7.8.1 Even and odd polarisation

The results from all three data sets indicate that the even polarisation performs better than the odd polarisation for target detection and discrimination.

Against a grassland background, for targets imaged at 35 GHz, the MRMS radar indicates an even/odd ratio for the TCR of 2 dB and for the FGAN radar 1.1 dB. For targets imaged at 10 GHz by the ESR radar, the TCR in the even channel is 1.24 dB higher than that measured in the odd channel, albeit at a different wavelength.

On average the MRMS radar has a TCR for military vehicles that is 6.7 dB greater than that of the FGAN radar. This difference is likely to stem from the much steeper imaging angle used by FGAN, i.e. 20 degrees compared with 10 degrees for MRMS. At these grazing angles, the clutter reflectivity increases approximately with the sine of the incidence angle, and should result in a 6 dB increase in clutter reflectivity for FGAN, accounting for the bulk of the difference between the radars. Residual differences may also be caused by the backgrounds (grass in all cases) or the target types imaged by the radars, since no targets are common between any of the datasets.

Sufficient 10 GHz data were available to examine the relative clutter performance of size and statistical features in the circular polarisations. Size-based measurements showed very little variation with polarisation and all data sets had sufficient contrast to allow the segmentation routines to work effectively. Statistical features showed increased separation between data sets, with features from the even channel demonstrating more potential to be used as clutter filters. However, the performance of features in both channels against rural clutter was good and the differences in system performance would be slight.

7.8.2 Circular and linear polarisation

For the 35 GHz FGAN and 10 GHz ESR radars, it was possible to compare circular and linear polarisations directly.

At 35 GHz, the even polarisation produced the highest target to clutter ratio, and the odd the lowest. Of the linear polarisations, HH performed better with slightly lower values than the even channel and VV was slightly below the even channel. Both were higher than the values for the odd channel.

At 10 GHz, the HH channel performed marginally better than the even channel. The VV channel was slightly above the odd channel which, again, gave the lowest TCR values.

Analysis of feature distributions at 10 GHz indicated that, for the two key features examined, the performance of the linear polarisations was between that of the even and odd channels.

7.8.3 Effects on performance

Using data gathered in odd and even polarisations by the 35 GHz MRMS radar, it was possible to analyse false target detections in rural and urban clutter.

The performance difference between even and odd channels in rural areas was sufficiently small that no difference was shown in the data examined. However, the RFOD in rural areas for the latest version of the ATR algorithm was already very low. Improvements in TCR due to switching to the even channel may improve this further but extended rural datasets would be needed to analyse the effects.

Performance in urban areas is much more challenging for radar seekers. Dual-polarisation MRMS data of the Defence Academy at Shrivenham showed that urban areas have a very similar number of false target detections in the odd and even channels when making use of the improved TCR in the even channel. Since

the only change in the algorithm when switching from its developed configuration for odd channel data was to raise the TCR threshold slightly, it is possible that an algorithm developed from a target set imaged in the even polarisation would perform better. However, any performance difference is unlikely to be large, and would not produce a solution to the problem of false target detections in urban areas.

7.9 Conclusions

This chapter has investigated the use of linear and circularly polarised data from three radars to assess the potential effects on target detection and attribute measurement performance for different seeker antenna types.

Circular polarisation is conventionally used by radar seekers, with a preference for the even polarisation due to improved target contrast and a reduction in rain backscatter. All three data sets indicated that the even channel produced a higher target to clutter ratio than the odd channel and is, therefore, preferred for target detection and segmentation. However, the contrast in the odd channel, which is between 1 and 2 dB below that of the even channel is still well suited for target detection and high quality segmentation.

The target contrast in linear polarisations was higher than in the odd circular polarisation and close to that of the even circular polarisation. The HH polarisation produced, on average, slightly higher contrast values compared with the VV, and in both channels the contrast was sufficient for detection and accurate segmentation.

A study of clutter rejection features was performed on a large data set from the 10 GHz ESR sensor. It was found that, for the purposes of separating target objects from clutter using two key statistical features, the even channel has the best performance, the odd channel the worst, with results from the linear polarisations lying between the two.

The effects of higher target contrast in the even channel was assessed by processing data of Salisbury plain and the Defence Academy at Shrivenham through the ATR algorithm. The odd channel performance in rural clutter is already very good, and no detectable improvement in performance in the rural areas processed was found. The challenging urban scene produced many bright reflections in the even channel, possibly due to bright dihedral scattering between building walls and the ground. Use of the extra 2 dB target contrast in the even channel was sufficient to allow a false target detection rate which was slightly

improved from that of the odd channel. Configuring the algorithm for even channel detection, which was not possible due to the low number of targets imaged in the even channel, could improve this performance further.

Given the good performance of the current antenna in rural areas, and the nature of urban clutter, it is unlikely that a change to even polarisation will significantly affect the capabilities of the radar seeker under good conditions. However, there may be cases where an additional few dB in clutter reflectivity (e.g. for shadow analysis) or target contrast could make a critical difference in system performance.

Should only one polarisation be available for static target ATR, the even polarisation channel would be the first choice for a seeker system. If an antenna is required to have a linear polarisation, HH appears to offer marginal improvement over VV.

Given the current system performance, marginal differences between the polarisation options, and the likely possible scenarios a seeker will face, the probability of mission success will be depend more upon the quality of image formation and subsequent processing than the choice of polarisation.

7.10 Future Work

A more extensive set of data, including a broad range of clutter backgrounds, should be tested at the likely frequencies that could be used for a radar seeker, ideally 35 and 94 GHz.

One consideration which needs to be investigated, should a missile seeker using a linear polarisation be considered, is the increased backscatter from rain, and the effects this may have on system performance. Backscatter from rain may be a more significant factor in linear and odd circular polarisations than for even circular polarisation [59]. The effects of rain should be further considered if it is felt that polarisation-sensitive backscatter could be a limiting factor in performance. This was discussed in Chapter 5.4.

8   Real-Beam ATR

8.1   Introduction

The bulk of the work presented, thus far, has focused on radar sensors which use forward squinted SAR to produce imagery with a high cross-range as well as down-range resolution.

Sensors have also been considered in which SAR imaging may not be possible, but a requirement to detect targets exists, and there is scope to examine new techniques to look for areas in which capabilities could be enhanced.

8.2   Sensor Fused Munition outline

A study was undertaken to evaluate the potential for new sensor and algorithm technology to enhance the performance of the next generation of Sensor Fused Munitions / Weapons (SFM / SFW). An SFM is a relatively simple use for a radar seeker, as it follows an unpowered trajectory, and is potentially fired by artillery to a target location. The warhead is fused directly by the detection and / or classification process from the seeker. An illustration of the SADARM (Sense And Destroy ARMor) SFM is shown in Figure 8-1.
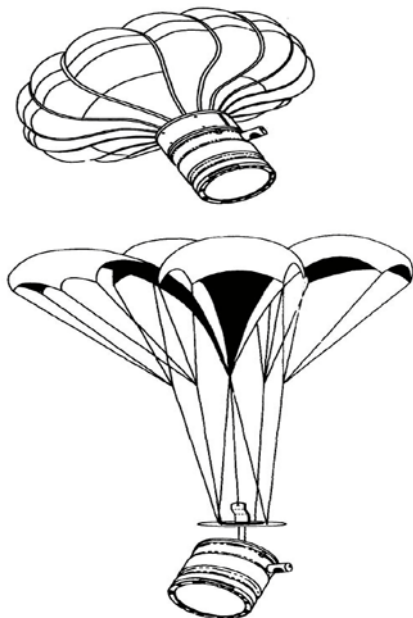
Figure 8-1 Illustration of SADARM during descent and search. Image courtesy of US Department of Defense

This study concentrates on sensor technologies which are low-cost but, when compared to the understood capabilities of existing weapons, may substantially improve object detection performance and reduce the likelihood of clutter or collateral objects being selected. The military benefits could include improved operation in complex environments, reductions in collateral damage, and increased overall system cost-effectiveness due to an increased probability of detection.

The study draws upon existing SFM systems and related research in this area. The approach taken is to step back from specific SFM applications to address generic low cost sensor technologies that could be used and identify where they might offer performance benefits.

The baseline assumptions used in the study are drawn from the open literature.

The value of the technique is assessed using a simple model of an SFM dropping at constant rate and performing a circular scan of an area as outlined in section 8.3. This representation is derived from known current SFM parameters. Modelling of wind and other complex environmental effects are not included in this study.

8.3   Understanding the system

8.3.1 Sensor Fused Munition Review

For the purposes of clarity and to place the work in context, a brief review of current SFM technologies was undertaken, with information largely taken from [12].

Three artillery launched SFMs have been developed and are in use: BONUS (Sweden and France), SADARM (US) and SMArt (Germany, also used by Switzerland and Greece). Also in use by the US is the air dropped CBU-97 which goes through several stages of deployment, resulting in 40 sensor fused munitions (per weapon) known as Skeets being spread over the target area. Table 8-1 gives an overview of the sensors used in these SFMs.

Weapon	BONUS	SADARM	SMArt	CBU-97
IR	Line Scan	Line Scan	Line Scan	2 colour, Single Pixel (2-4 micron, 4-5 micron)
Laser	Profile Detector	✘	✘	Profile Detector
RF	✘	Active mmW Passive radiometer	Active mmW Passive radiometer	✘

Table 8-1 Sensor systems employed by current SFMs

Two of the SFM, SADARM and SMArt, include active MMW sensors for the detection of targets.

Prior to any discussions of specific sensor technologies or upgrades, their operation should be better understood. Using freely available information on current SFMs, an overview of their design was formed and, from this, several important parameters could be extracted and compared. With these parameters in place some non-sensor specific analysis may be performed. This may then be taken through into specific sensor discussions.

Table 8-2 shows the variation in key parameters of the various munitions investigated.

Weapon	BONUS	SADARM	SMArt	CBU-97
Deceleration/Spin	Samara Wings	Ram air/vortex Ring parachute	Drag/auto-rotating parachute + de-spin flaps	Launched horizontally spinning from carrier
Descent speed	45 ms^{-1}	17 ms^{-1}	13 ms^{-1}	Freefall
Spin rate	15 Hz	7.6 Hz	3 Hz	30 Hz
Vertical metres/Rotation	3	2.2	4.3	n/a
Munition offset from vertical	30°	30°	30°	22°
Detonation height	<150 m	<150m	<150 m	<150 m

Table 8-2 Comparison of key attributes of current SFMs from open literature

It is clear that offset angles and detonation heights are similar between munitions. Much larger variation occurs in the descent speeds and spin rates; however, they are roughly in proportion between the munitions. This translates to a similar pattern in terms of ground coverage.

Lower descent speeds and higher spin rates are likely to make detecting (and tracking) the target easier as they would provide more opportunities to view the target. However, these will be at the expense of stability and robustness to wind effects (due to the larger parachute/wings required). Spin rates must be linked to sensor rates as they will affect the frequency of target viewing and the ability to view the required search area.

Assuming a nominal target extent of 4 m diameter (the smallest targets represent the most challenging cases) and a munition offset at 30° to the vertical, the munition has only a 7 m section of its descent in which it may fire and achieve a hit on the target. The BONUS munition would only view the target twice, SADARM three times and SMArt possibly only once within this distance. It is, therefore, critical that the sensors can acquire the target such that an opportunity to fire can be taken. In the case where there is only one image of the target, there is less opportunity to confirm acquisition and refine the positional measurement.

The small window of opportunity for firing, coupled with target positional uncertainty, is likely to lead to a low probability of a successful engagement. As an increase in firing opportunities is unlikely (without steerable warhead technology), the best mechanism for increasing the effectiveness is, therefore, an improvement in sensor technology in order to provide more accurate target estimates. This will then allow for optimisation of the hit point across the target.

8.4   Millimetre Wave Radar

Of the SFMs produced to date, SADARM and SMArt use MMW radars. It is understood that the radars used in these weapons are basic frequency modulated continuous wave (FMCW) designs operating in conjunction with an IR sensor. The low range resolution of the existing radars indicates that they are used as simple detectors of objects that have a higher RCS than the background clutter.

This section of the study investigates radar technologies which are low cost but can potentially improve target discrimination and aim-point selection performance.

A review of radar technology was carried out, and indicated that sensors with the following properties could be produced relatively cheaply:

- 94 GHz frequency
- 1.7° beamwidth
- 0.3 m range resolution

These figures were used as a basis for the systems modelled, with changes made to reflect the capabilities of recent advances in FMCW hardware.

The work in this chapter is split into the following sections. Section 8.4.1 contains information on the SFM radar, geometry, hardware and modelling. Section 8.5 outlines the scenarios studied, and section 8.5.1 shows the application of 1D and 2D ATR techniques to the data with a single-beam antenna. Section 8.5.5 demonstrates the utility of monopulse to aid detection through height measurement when an SFM with a multiple beam antenna is considered, and section 8.5.10 contains the conclusions.

8.4.1 SFM Radar Geometry

The radar in an SFM operates in a real-beam mode as opposed to the SAR imagery discussed in previous chapters. The radar measures the returned signal from the ground within the radar footprint. The footprint size is determined by the height, depression and beamwidth of the radar. A narrow beam, and thus a small footprint, is desirable since it allows higher antenna gain while reducing the target location uncertainty, as long as sufficient ground coverage is maintained.

To simulate the motion of an SFM, the trajectory of the sensor was set to remain stationary in the horizontal plane while descending vertically at a uniform velocity, and with a fixed scan rate. This is an idealised case for the weapon's descent, since the effects of drift and pendulation were not within the scope of the study.

Given the size constraints of the submunition if launched by artillery, the antenna diameter can be expected to be at most 135 mm. To give a smaller beam and, therefore, more accuracy, a 94 GHz radar was modelled, resulting in a 3 dB beamwidth of 1.7°.

The imaging geometry for an SFM is shown in Figure 8-2. The steep imaging angle, Ø, of 60° from horizontal results in down-range measurements that are a combination of ground range and height. The outer dashed lines represent the 3 dB width of the beam, θ_B, the resultant footprint width, d, and its projection to the ground is d_G.
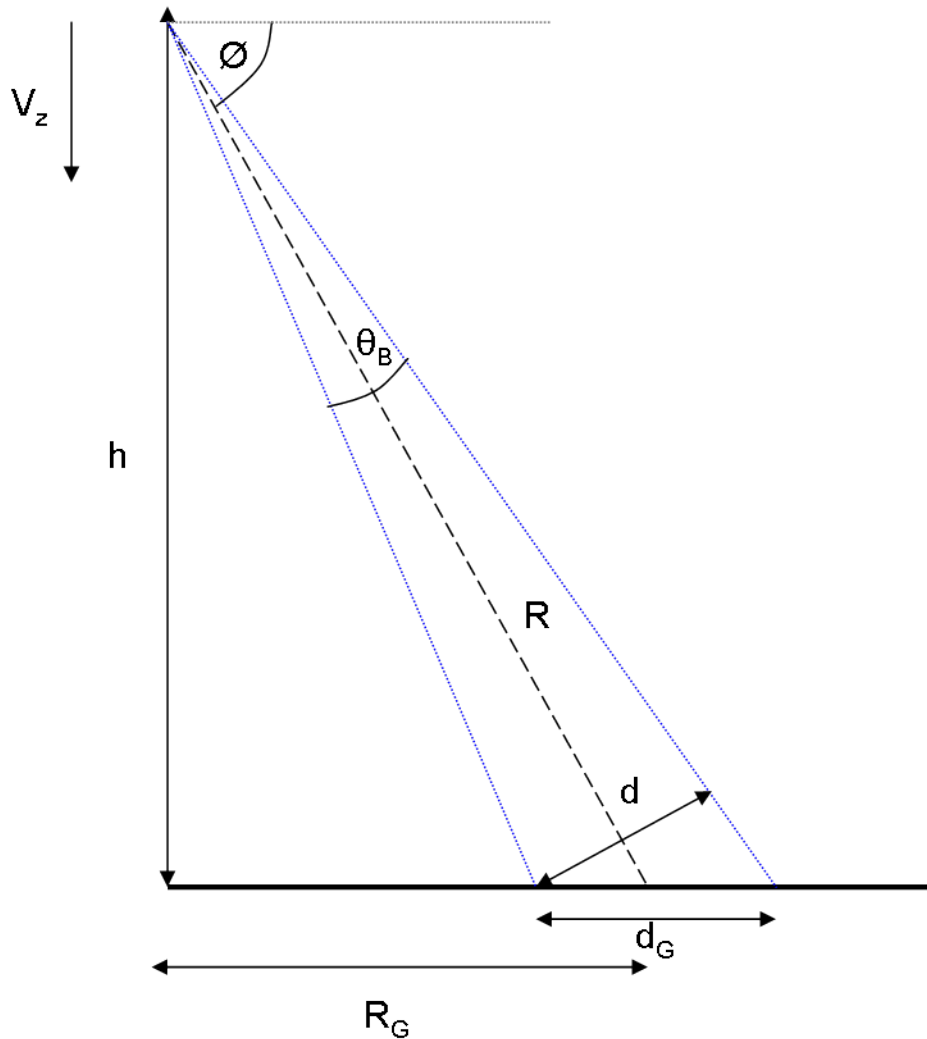
Figure 8-2 Imaging geometry for an SFM

Since the expected start of the search for an improved weapon could be at a height of 300 m, given the 60° lookdown angle, the range to the target is ~350 m. This results in a patch on the ground of 10 m in azimuth for the 1.7° beam. As the target may be viewed head-on and, given a typical main battle tank width of 3 m, significant contamination of the target signature by the surrounding clutter would be expected.

At a height of approximately 200 m, where existing weapons start their detection process, this patch is 7 m wide, and does not fall to within the 3 m target width until a height of 87 m, by which time the search area is significantly reduced.

While it is not possible to improve the beamwidth through use of a larger antenna due to weapon size constraints, adding a monopulse capability to the radar will improve the angle estimation while staying within size constraints.

A second effect of the requirement for a narrow beam is the ability to image the search area adequately. While IR sensors in SFMs may operate with a relatively wide Field of View (FOV), the radar FOV must be smaller to maintain accuracy. This has implications for the number of times a target can be imaged.

The rotation rate requirement, W, for a radar with height, h, descent speed, V_z, and desired number of SFM rotations with the target in the beam, n, may then be calculated using (8.1).

$$W = \frac{nVz\sin\phi\cos\phi}{h\sin\theta_B}$$

(8.1)

Given a descent rate of 20 ms^{-1}, the necessary minimum rotation rates to ensure complete ground coverage are shown in Figure 8-3. Any rotation rate below that shown for a given altitude will result in less than 100% coverage of the ground. As the weapon descends, the beam's centre position on the ground moves inwards at a constant rate. However, the size of the imaged area on the ground decreases. Hence, as the SFM descends, a higher rate is required to maintain coverage.
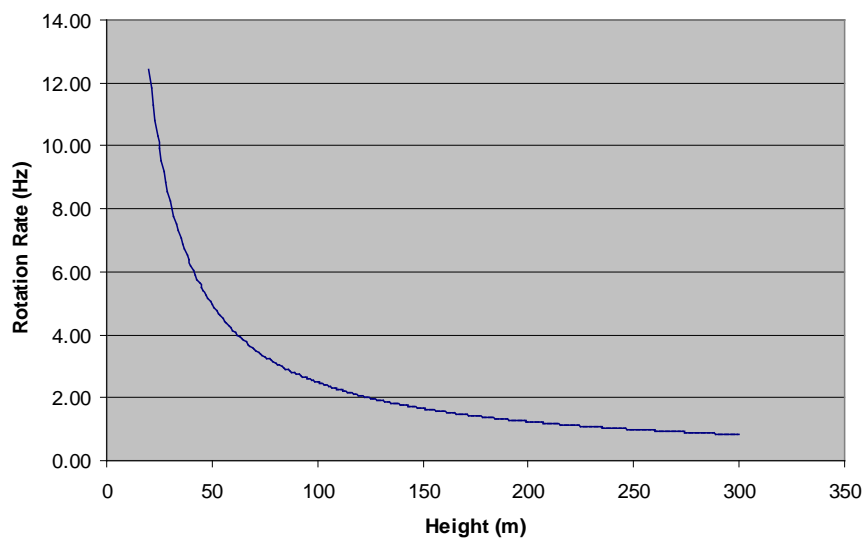
Figure 8-3 Height and required rotation rate to provide 100 % coverage for a 2° beam at 20 ms^{-1} descent

For full coverage of the ground within one rotation of the munition, the waveform repetition frequency (WRF) needs to be sufficient that areas of the ground are not missed. The minimum required WRF can be calculated from the munition spin rate, W, and beamwidth, θ_B, is given by equation (8.2), and illustrated in Figure 8-4. This ensures that a point target is sampled at least once within the 3 dB beam during a rotation.

$$WRF \geq 360\frac{W}{\theta_B}\cos\theta \quad \text{Hz} \tag{8.2}$$
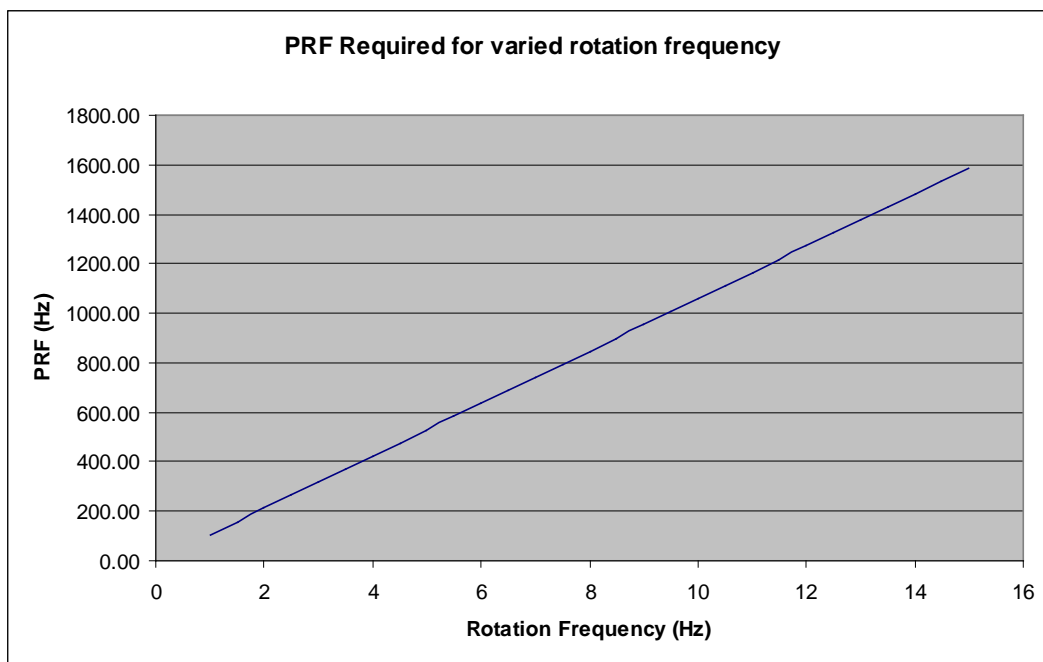
Figure 8-4 Required minimum PRF for varied rotation frequency for a 1.7° beam

Ideally, coverage of greater than 100 % would be obtained, allowing repeated looks at a target before firing; however, this may not be possible given the physical constraints on the sensor.

For the radar simulation in this report, a spin rate of 7.5 Hz was selected. This is close to that of the SADARM weapon which uses a MMW sensor. Changes to the spin rate do not significantly affect the nature of the radar processing, assuming a sufficient number of FM ramps fall on the target. Therefore, techniques that are shown to offer potential improvements will remain effective through the range of possible spin rates.

It is also necessary to ensure that the radar PRF is such that targets will be adequately sampled in the azimuth direction.

The key radar specifications used throughout this study are as shown in Table 8-3.

The waveform repetition frequency of 5 kHz is sufficient to allow six ramps on target within the 3 dB beamwidth as the sensor rotates. This requires a rapid ramp rate, but as described in section 8.4.2, should be possible with modern hardware.

Beamwidth (3dB)	0.03 rad (1.7°)
Monopulse Beam Separation	0.03 rad (1.7°) at 3 dB Power overlap
Rotation Rate	7.5 Hz
Range	231 m (200m height)
Waveform Bandwidth	600 MHz
Waveform Repetition Frequency	5 kHz

Table 8-3 Suggested FMCW radar parameters

8.4.2 Radar Hardware

Most conventional radars are based on a pulsed technique, deriving target range from the time of flight of the pulse. However, increasingly, radars are employing continuous wave (CW) based techniques that achieve the same performance with much lower peak transmit powers. This is because radar detection performance is related to average power, not peak power. The most common of these waveforms is FMCW, in which the transmit signal is ramped linearly in frequency, returns to the start frequency and then repeats. The target range is derived by 'mixing' the signal being transmitted with the signal received. The resultant signal has a frequency component that is proportional to range. Whilst this technique has been known since the 1940s, its practical implementation has only really been possible in the last two decades due to developments in signal processing and solid state microwave components.

The advantages of FMCW radar over a simple pulsed radar, which relies on short-pulses rather than bandwidth for resolution, are [71]:

▪ It is easier to achieve the bandwidth required for high range resolution with an FMCW radar than a simple pulsed radar [72]. High resolution allows better target discrimination and characterisation, together with good rejection of surface and rain clutter ;

- The sensitivity of a radar depends upon the average transmit power, not the peak power. Therefore, a 1 mW FMCW radar achieves the same sensitivity as a 10 W pulsed radar with 0.01 % duty cycle;

- The low transmit power facilitates the use of solid state components with low power supply voltages and high reliability;

- FMCW radar can be implemented with a simple, low cost homodyne architecture.

The disadvantages of FMCW radar are:

- Because the radar is transmitting and receiving simultaneously the radar can be desensitised by leakage of transmit power into the receiver. This can be overcome by having separate transmit and receive antennas (which is not an option for an SFM), or by incorporating a circulator and designing the RF paths to cancel dominant breakthrough components;

- The phase noise performance of the transmit waveform is critical and can desensitise the receiver.

FMCW radars are used for a diverse range of applications at frequencies throughout the microwave and millimetre wave bands. They are particularly well suited to short range applications requiring high range resolution such as in missile seekers (at 94 GHz) and for automotive cruise control (at 77 GHz).

The use of automotive radars as a commercial off-the-shelf (COTS) component was considered for applicability in this work. The automotive application is driven by cost, not performance, and is a highly integrated product designed specifically for that application. While the units are low-cost, it would not be feasible to try to adapt this type of radar for use in SFMs in a cost effective manner. This is because the radar sensitivity is poor and the antenna configuration, designed for integration within a car rather than an SFM, is inappropriate. The level of changes required would be equivalent to a bespoke solution.

Current FMCW architectures usually generate the swept linear frequency by modulating an oscillator at the final RF frequency, rather than generating the sweep at low frequency and multiplying, or mixing, up. This does not produce the best phase noise performance but was, until recently, the simplest implementation. In the case of millimetre wave frequencies, the oscillator of choice is a Gunn oscillator which has a very non-linear frequency tuning characteristic. To achieve good range

resolution, it is essential to transmit a highly linear frequency sweep [73]. A simple illustration of a frequency sweep and the returned signal from a point target is shown in Figure 8-5.
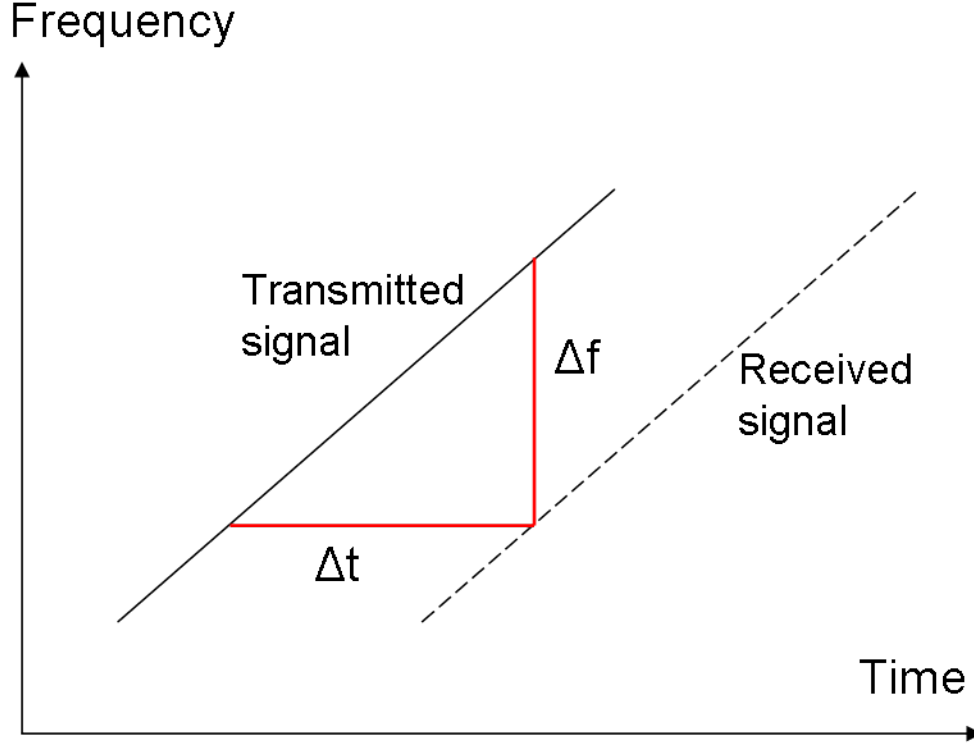
Figure 8-5 Illustration of FM frequency ramp and returned signal

The difference in frequency, f_r, between the transmitted and returned signal to a point target at range R is given by

$$f_r = \frac{2R}{c} \frac{\Delta f}{\Delta t} \tag{8.3}$$

where $\dfrac{\Delta f}{\Delta t}$ is the ramp-rate of the waveform and c is the speed of light.

For the parameters suggested for this system, defined in Table 8-3, and assuming a saw-tooth waveform, the ramp rate is 3 MHz/µs and, given a range of 231 m, the expected frequency shift is 4.62 MHz.

Since a range resolution of 0.25 m is possible with a bandwidth of 600 MHz, the change in frequency due to the ramp rate between the transmitted and received signal should be less than the change in frequency over the dimensions of a resolution cell. To have an accuracy of one quarter of a resolution cell, this implies a dimension of 6.25 cm and a round-trip frequency shift of 1250 Hz.

The required accuracy over the sweep is given by

$$SweepAccuracy = \frac{\Delta f_{cell}}{\Delta f_{trip}}$$ (8.4)

and in this case requires a sweep accuracy of 0.03%. For longer ranges, this accuracy will need to be even finer. This is a simplification, and the rate of variation can have an effect on the error magnitude. This is further explored in [74] and [75].

To achieve this level of accuracy, these systems have to include a feedback-based linearisation circuit which, at millimetre wave frequencies, can typically only be achieved up to a bandwidth of about 600 MHz, corresponding to 25 cm range resolution. In addition, the frequency sweep rate is limited by the feedback loop bandwidth and it can be difficult to sweep faster than 600 MHz in 1 ms.

Recent developments in microwave oscillators and the availability of low cost Direct Digital Synthesis (DDS) Integrated Circuits (ICs) now facilitate an improved FMCW radar architecture (that is currently being patented by QinetiQ). The new architecture uses a Dielectric Resonant Oscillator (DRO) at a microwave frequency which is multiplied up to the final, millimetre wave frequency. The DRO after multiplication has a significantly better phase noise performance than a Gunn oscillator and, hence, the radar has improved sensitivity.

The DRO acts as a fixed frequency Stable Local Oscillator (STALO). The frequency sweep is generated at low frequency (UHF) using DDS. DDS has the advantage that it can generate a perfectly linear sweep, so there is no need to provide a linearization circuit, and it has ultra-low phase noise. The swept waveform is up-converted with the DRO output to provide the final RF transmit waveform. The architecture is capable of producing swept bandwidths of greater than 1.5 GHz, corresponding to 10 cm resolution. This, in turn, equates to better rain and surface clutter rejection. An additional advantage of using DDS is that very high modulation rates can be achieved. A sweep bandwidth of 600 MHz at 10 kHz repetition frequency using DDS is realistic but would be impractical using conventional analogue techniques.

As DDS is a digital technique, there are disadvantages to its use. The waveform is sampled, therefore, due to quantisation the output signal is not spectrally pure and has spurious returns [76].

The proposed system is shown in Figure 8-6.
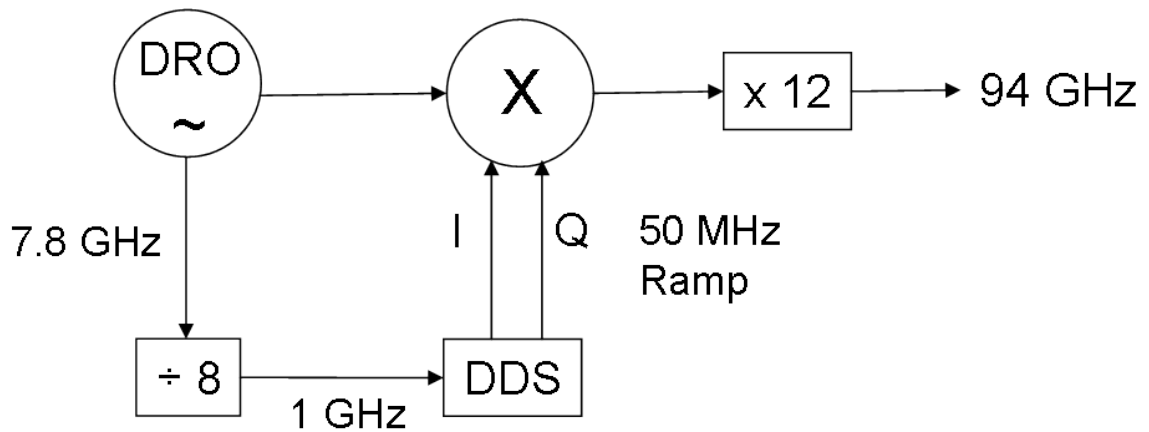
Figure 8-6 DDS Radar

In this new FMCW architecture, the DDS and STALO, together with the signal sampling system are locked together. This provides coherency from frequency sweep to frequency sweep enabling simultaneous range and Doppler measurement of the target and enhanced sensitivity. Whilst this is computationally intensive, there are now Field Programmable Gate Arrays (FPGAs) available that can process data at a sufficient rate. Also, improvements in recent years of Analogue to Digital Converters (ADC), in terms of the number of bits and the sampling rate, provide enhanced receiver dynamic range. 14 bit ADCs with 150 MHz sample rate are available off-the-shelf, providing the radar with an instantaneous dynamic range of about 80 dB.

This use of FMCW radars is ideal for short-range high-resolution applications. The step-change in performance should allow improved sensitivity, a faster update rate, coherent processing to provide sub-clutter visibility and lower cost.

This technology offers a low-cost, high performance radar option for future SFMs, and is preferable to the use of existing automotive radars, the low cost of which, driven by economies of scale, would not be possible for the custom design, and would offer far more limited performance.

8.4.3 Radar Modelling

The CLTG software (Section 4.4) was used for this study to produce simulated radar range profiles. The active MMW radar in an SFM will almost certainly be an FMCW design. However, the CLTG is currently unable to produce this waveform. Since the outputs from both a pulsed and FMCW radar are range profiles, it is possible to configure the CLTG pulsed radar mode to give the same output as would be expected from an FMCW radar.

In order to verify that the pulsed waveform used was giving broadly comparable results, a configuration was used to generate synthetic data that could be compared with real data from an FMCW. The FMCW radar chosen for comparison was the Thales Missile Electronics Dual-Mode Seeker (DMS) radar, which is a 94 GHz FMCW seeker. Real data from air-carry trials was available which included grassland, military vehicles, trees and urban areas. Example real-beam imagery mapped to the ground is shown in Figure 8-7.
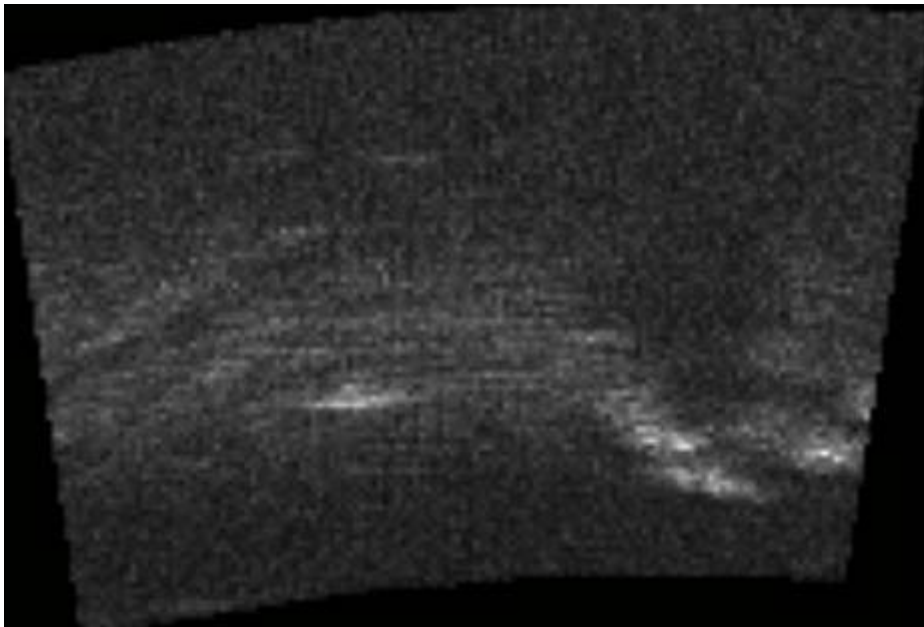
Figure 8-7 DMS Radar imagery mapped to the ground

Measuring the signal levels from these background types for both the DMS radar and the CLTG simulation allowed the configuration to be validated (Table 8-4).

Signal Type	Ratio (dB)		Difference (dB)
	DMS Radar	CLTG	
Clutter to Noise	6	9	3
Building to Clutter	8.6	9.7	1.1
Tree to Clutter	6.5	7	0.5
Target to Clutter	16	15.6	-0.4

Table 8-4 Signal ratios for DMS radar and CLTG simulation

The clutter versus noise ratios for the DMS measurements are lower than the CLTG measurements. Due to the nature of DMS image generation, it is possible that cells chosen as noise still contain some weak returns from the scene, effectively lowering the ratio.

However, for the ratio of signals from other objects to clutter, the levels are broadly comparable. In order to make the CLTG output match that of the DMS sensor, the reflectivity of the clutter types and target scatter point model were adjusted to produce the best fit.

The comparison was performed at a 10° look-down angle since that was the only geometry for which trials data were available. When moving to the 60° SFM geometry, reliance is placed upon the clutter calculations within the CLTG to adjust clutter reflectivity to suit the steeper angle.

These calculations are based on the expected change in the reflectivity of rough clutter as the elevation angle varies from a 10° default value. The correction applied to the clutter RCS is

$$rcs = rcs * \frac{f(i)}{f(10)}$$ (8.5)

where

$$f(i) = \sin(i)$$ (8.6)

for rough clutter [77].

This produces a 7 dB increase in clutter reflectivity, in-line with that expected for grasses and short vegetation listed in [68].

8.5   Test scenarios

Two scenarios were examined for this work:

- Scenario 1: Flat grass terrain with static MBT;
- Scenario 2: Rural scene (grassland and trees) with static MBT.

The rural scene file is shown in Figure 8-8. The top image shows the object type over the scene (e.g. grass, tree, road) and the lower image is a height map with darker red colours representing the tallest objects (trees), and changes in the background shade showing terrain slope.
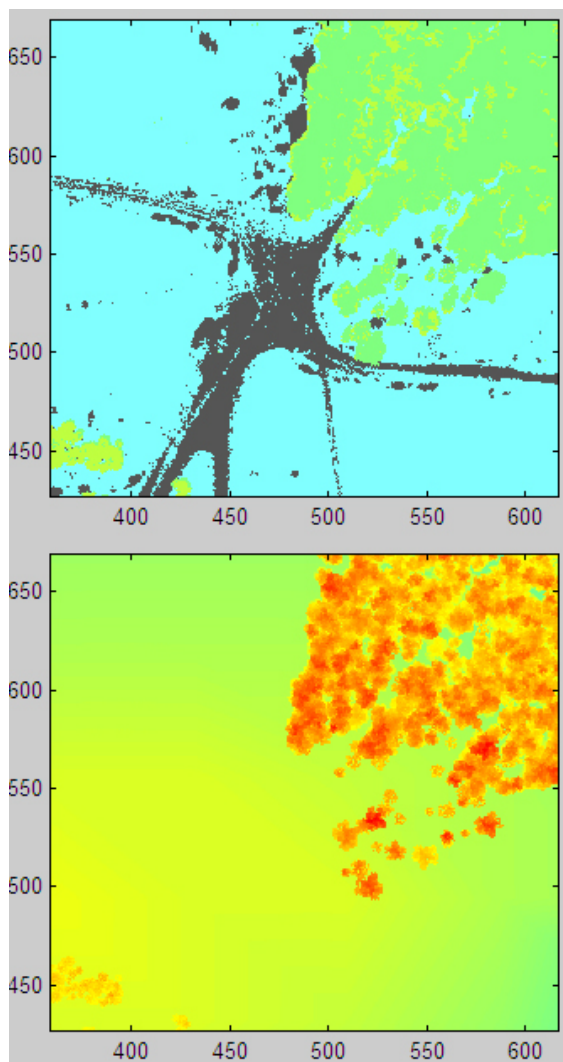
Figure 8-8 CLTG rural scene file. Axes units in metres

Each scene was imaged at three different heights to represent the descent of a submunition. 300, 200 and 100 m were chosen, the larger heights representative of what might be achieved by an improved submunition, allowing the imaging of a greater area during descent.

8.5.1 Single Channel Radar Techniques

A brief examination of ATR techniques was undertaken on the simulated SFM data. The aim was to identify techniques that held potential for future investigation rather than to develop an ATR process. Previous work had found that high fidelity data required sophisticated algorithms and a significant amount of tuning to cope with changes in the sensor performance. A more robust approach would increase the versatility of the sensor.

Two approaches to enhance the ATR capability from a real-beam radar have been considered; high range resolution profiling for 1D data and 2D algorithms for real-beam imagery.

The ability to use target height to aid detection by looking for the associated range-shortening is then investigated.

8.5.2 1D High Range Resolution (HRR) imaging

Section 8.4.2 suggests that the implementation of an HRR radar in an SFM would be practical. Existing weapons, which use a range resolution of approximately 5 m, are limited in the ATR techniques they can employ since the target dimensions are comparable with the range gate size. The use of fine range resolutions results in multiple returns from a target, which provide information on target structure that can be used in an ATR process.

High resolution range profiling is an established detection technique, used in existing radars such as the Tarsier radar for debris detection [78]. Radars for target acquisition typically operate at a shallow depression angle, hence the radar returns measure target brightness along the target's down-range extent. The imaging geometry for an SFM produces fewer returns from the target since the steep imaging angle mainly measures height, which is often the smallest dimension. The mixture of height and range information from the 60° angle makes the returns difficult to interpret. However, there is still the potential for implementation of specific processing to improve detection and classification performance.

An example image, covering 40 ramps as the radar scans over the target is shown in Figure 8-9. Down-range is in the *x*-direction, and number of ramps (effectively cross-range) is the *y*-direction. The target is visible as a series of bright returns, with the first return at near-range relating to the top of the target. The sensor height is

200 m, and the depression angle of 60° results in a 100 m ground range, shown on the *x*-axis.

A range profile over the target in the grassland scene is shown as the red line in the right of Figure 8-9. The same scene, without the target present, is shown as the blue line. The elevated area in the centre indicates where the beam illuminates the ground, with the lower returns to either side due to noise in the system and environment.
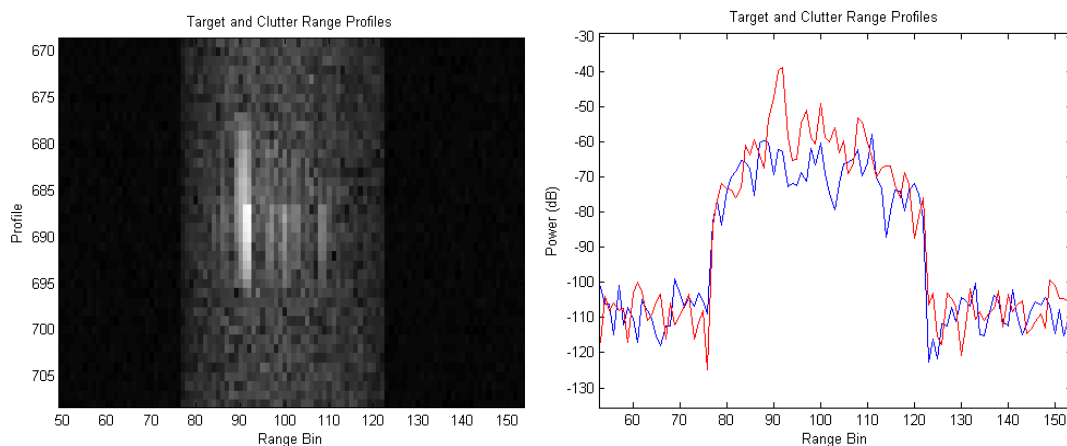
Figure 8-9 Target Image(left), down-range from left to right (0.25 m bins), cross-range in the vertical direction (individual range profiles) and range-profile plots (right) through the target (red) and clutter (blue)

The type of processing used in range profile analysis starts with a simple detector to identify bright areas and locate a section of profile that is of interest. This extracted profile can then be analysed using two techniques. The profile can be correlated with a database of templates stored on the weapon that represent returns expected from the target. If the match is sufficiently close, a detection can be declared and the weapon fired. Alternatively, statistical features can be calculated from the profile and, should the features satisfy criteria based on expected values for the chosen target type, detection can be declared.

Both types of detector described above are very sensitive to the radar data, and need to be tuned specifically for the chosen target type. While the CLTG model produces representative radar imagery, and has been baselined against air-carry trials data, it is not expected to be able to create data with sufficient fidelity to research detailed detection algorithms. Development of these algorithms would require access to trials data from high-resolution sensors of military vehicles measured in SFM geometries, or from detailed CAD models passed through complex signature prediction code.

8.5.3 2D Adaptive Thresholding

The scenario created using the rural scene file is a more challenging cluttered environment for target detection than the flat grassland that was also modelled.

As targets are expected to have a slightly higher radar reflectivity than the wooded background, an adaptive thresholding scheme was tested to detect the target returns. Initially, a 1D CFAR detector was tested. However, since this radar configuration measures the target in a sequence of range profiles multiple times during each scan, this was changed to a 2D detector applied to the 2D real beam image.

The imaging geometry of an SFM poses challenges for the commonly used ATR techniques. Thresholding each range profile to detect a target is problematic at lower ranges since the clutter patch shrinks to a comparable size to the target. Without background clutter in the profile from which statistics can be measured, it is difficult to set a threshold for the target. As the background noise is much lower than the clutter or target returns, any threshold set on the noise statistics is likely to pass both target and clutter.

While this may be a problem in 1D profiling, using the time dimension it is possible to make use of the high ramp rate of the radar to gather clutter statistics from consecutive profiles rather than from within a single profile. The technique applied measures returns at a constant range and uses the statistics from a window of preceding and following profiles at that range, separated by a few 'guard' profiles to avoid contamination from the test cells, to form an adaptive threshold, illustrated in Figure 8-10.

For an image such as that in the left of Figure 8-11, this works well. However, the image on the right, which shows the rural scene with slope, tree clutter and a target, imaged with the radar from a height of 200 m, is more challenging.
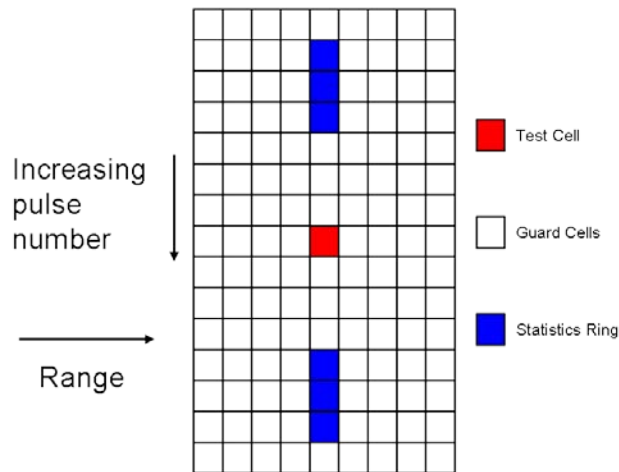
Figure 8-10 Adaptive threshold based on pulse history for each range
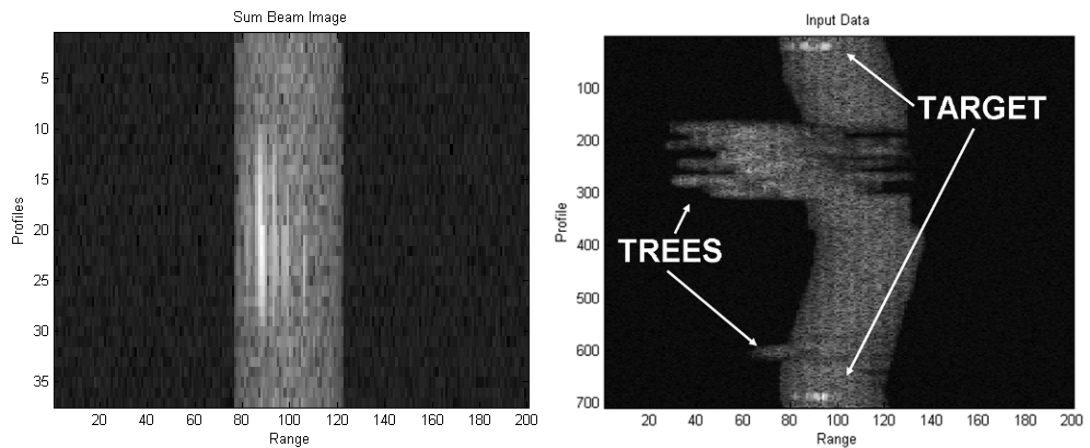
Figure 8-11 MBT imaged in flat grassland over 20° (left) and Copehill plantation over 380° (right). Range from left to right in 0.25 m bins, increasing range profile number from top to bottom

While the targets (at the top and bottom) of the right-hand image are in the main clutter band, wooded areas stand clear of the ground clutter because of their height. When running a threshold over these areas, the interface between noise and clutter is prone to false alarms.

Figure 8-12 shows the image after the 1D filter has been applied. Much of the noise has been rejected and the targets detected. However, there are alarms within the clutter region, scattered single pixel returns and larger areas from the edges of the trees.
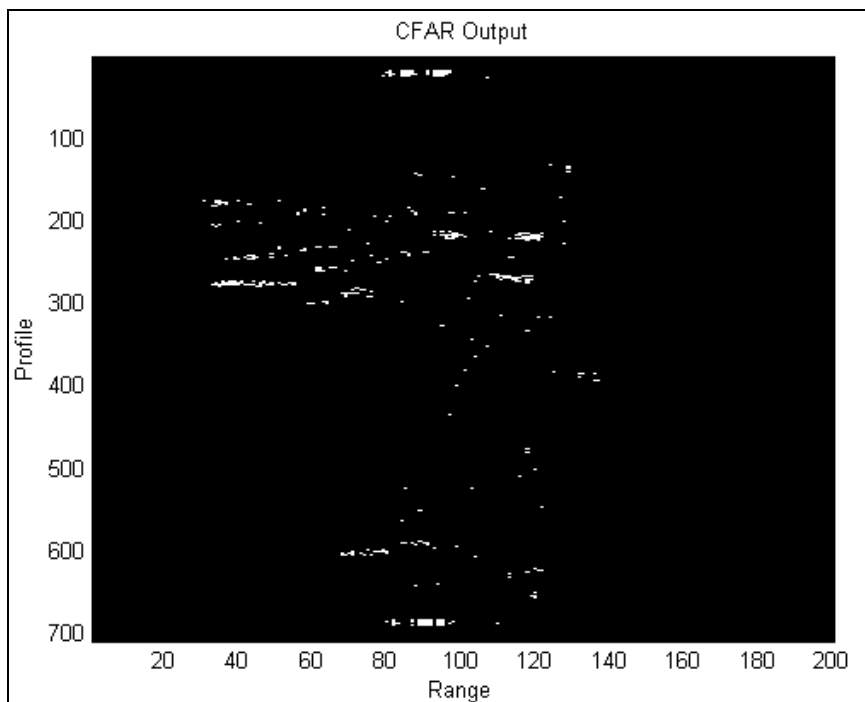
Figure 8-12 Rural image after 1D CFAR. Range from left to right in 0.25 m bins, increasing range profile number from top to bottom

Applying filtering in both time and range direction removes many of the smaller false alarms while keeping target detections. However, returns from the trees are still present (Figure 8-13).
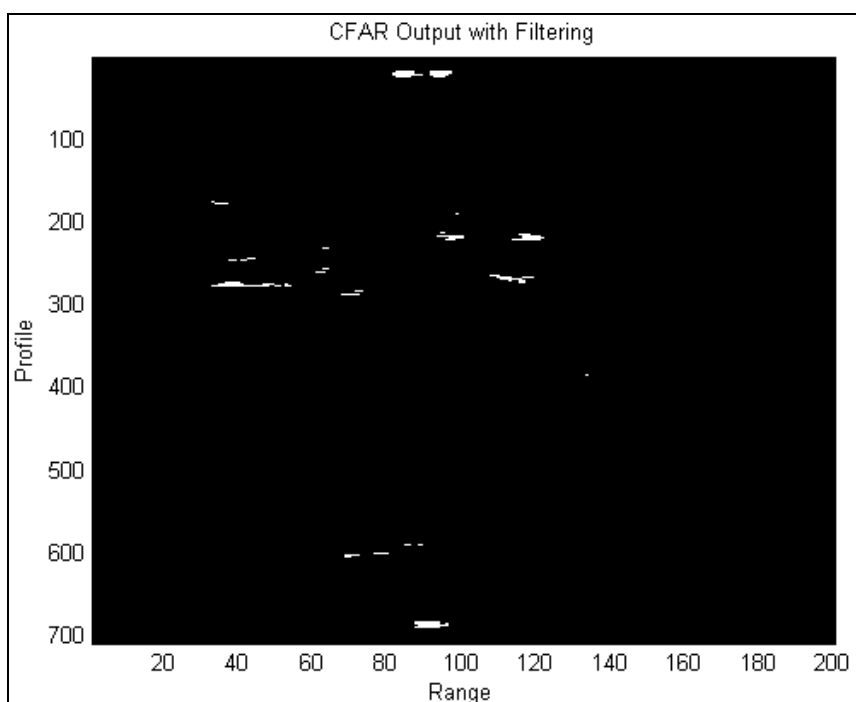
Figure 8-13 Rural image after CFAR and Filtering. Range from left to right in 0.25 m bins, increasing range profile number from top to bottom

This highlights the potential for 2D CFAR thresholding as a detector in an SFM. However, the scenes used here are relatively simple, and creating a detection algorithm that would be robust in realistic scenarios requires further research.

8.5.4 Height detection

Due to the SFM trajectory and the steep lookdown angle of the sensor, the down-range direction of the radar records a mixture of range (away from the point on the ground towards which the sensor is descending) and height. If a target is imaged towards the leading edge of the radar beam, it will appear clear of the clutter band in range since its height moves it closer to the radar. This offers the possibility of gathering information about a target's height as a potential detector and discriminant.

The image on the left of Figure 8-14 shows a series of CLTG range profiles which include the MBT target. Down-range (height) is from left to right in the *x* direction, and time / profiles descend in the *y* direction. The target is visible as bright returns, closer in range due to its height relative to the grassland background. The image in the right of Figure 8-14 shows the same target imaged twice in a 380° scan.

The bright return towards the top of the image is the target in the centre of the clutter patch, but standing out as closer in range due to its height. The second return is of the target imaged a second time as the sensor rotates.
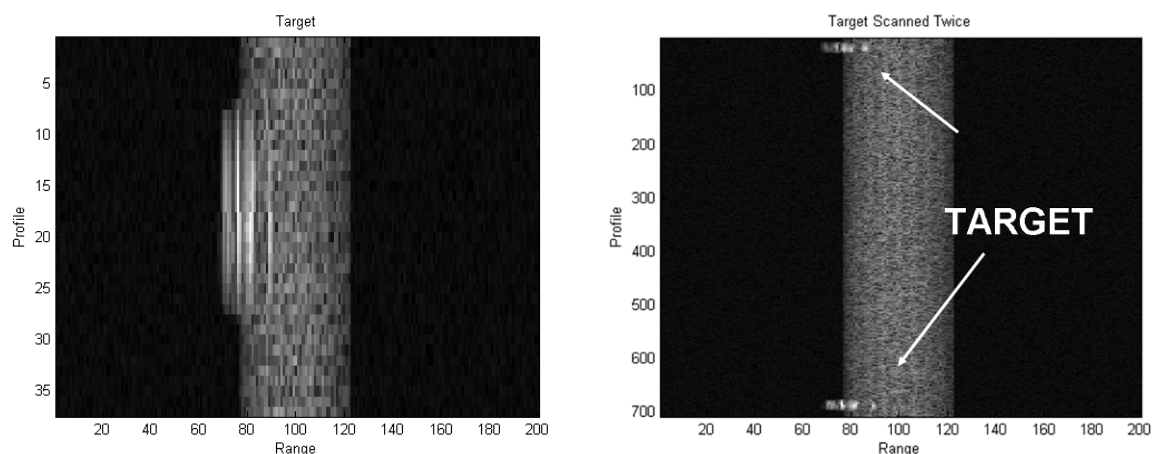
MBT, 20° scan, 7.5 Hz Scan Rate   MBT, 380° scan, 7.5 Hz Scan Rate

Figure 8-14 Simulated SFM Imagery from 200 m height. Range from left to right in 0.25 m bins, increasing range profile number from top to bottom

The lead pixel is defined to be the first non-noise pixel encountered in each profile when in the down-range direction, i.e. from the left edge of the image. Since the clutter to noise ratio is high, detection of this transition is simple. The output from a

simple lead-pixel detector is shown in Figure 8-15. The image to the left shows two peak returns closer to the radar corresponding to the two viewings of the target in the 380° rotation. The quantisation of range into 1 m bins is evident, with the returns coming from a range of 76 or 75 m.

A number of pixels down-range from that position are extracted, the power within the pixels summed and the results shown on the right of Figure 8-15. Returns from both viewings of the vehicle are clearly visible above the clutter band.
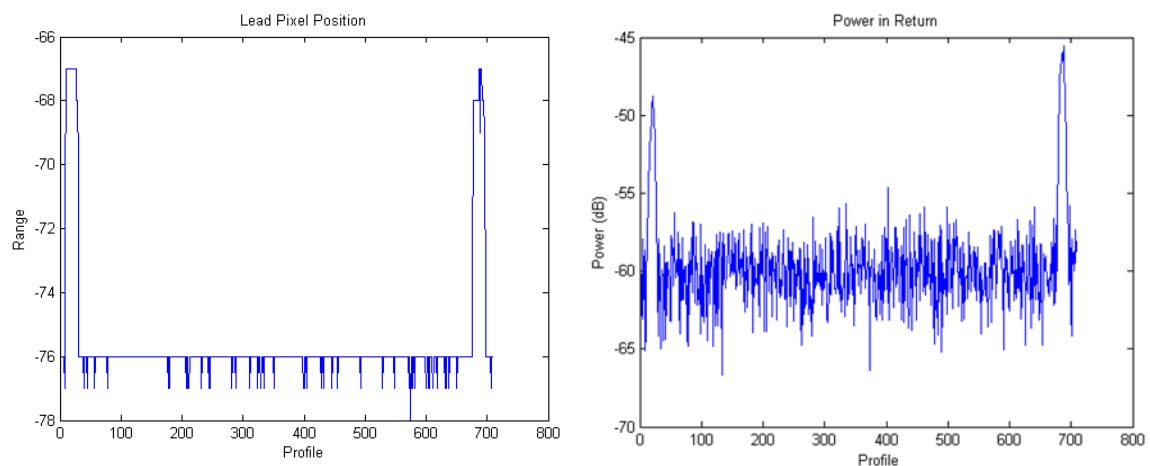
Lead pixel position   Power behind lead pixel

Figure 8-15 Lead pixel position (left), with range in 0.25 m bins on the y-axis, and lead pixel power (right) for two images of a target.

As is shown in Figure 8-15, looking for returns from outside the clutter band has the potential to work well when targets are positioned towards the leading edge of the beam, however it will be of little use when targets are further back in the beam and appear in the middle of the clutter band.

Since the beam on the ground is wide, especially when the SFM is at greater height, the target is more likely to appear within the clutter band as shown in Figure 8-16 rather than outside it as in Figure 8-14. With the target in this position, it is not possible to deduce any information about target height using this technique since it does not appear closer in range than the clutter band.
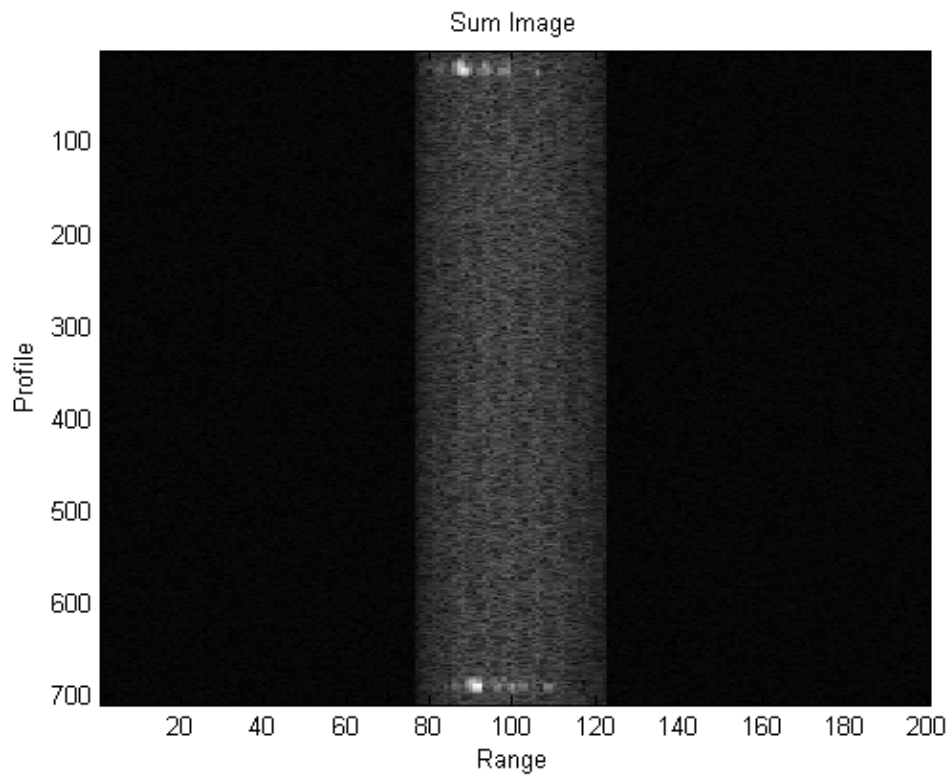
Figure 8-16 Vehicle imaged twice by SFM, scan rate 7.5 Hz. Range from left to right in 0.25 m bins, increasing range profile number from top to bottom

These results are for a very clean scenario with a high signal to clutter ratio and no other objects present. When a more realistic scene is used containing numerous clutter objects of varied height, the approach is less successful. Figure 8-17 shows a target in the Copehill scene, imaged twice, that has been positioned towards the leading edge of the beam such that it stands clear of the clutter. The plots in Figure 8-18 show the target still detectable in power, but less clear in height against the varied background. While the target may still be detectable, it is likely that the clutter objects such as trees could produce false alarms in a height detection algorithm.
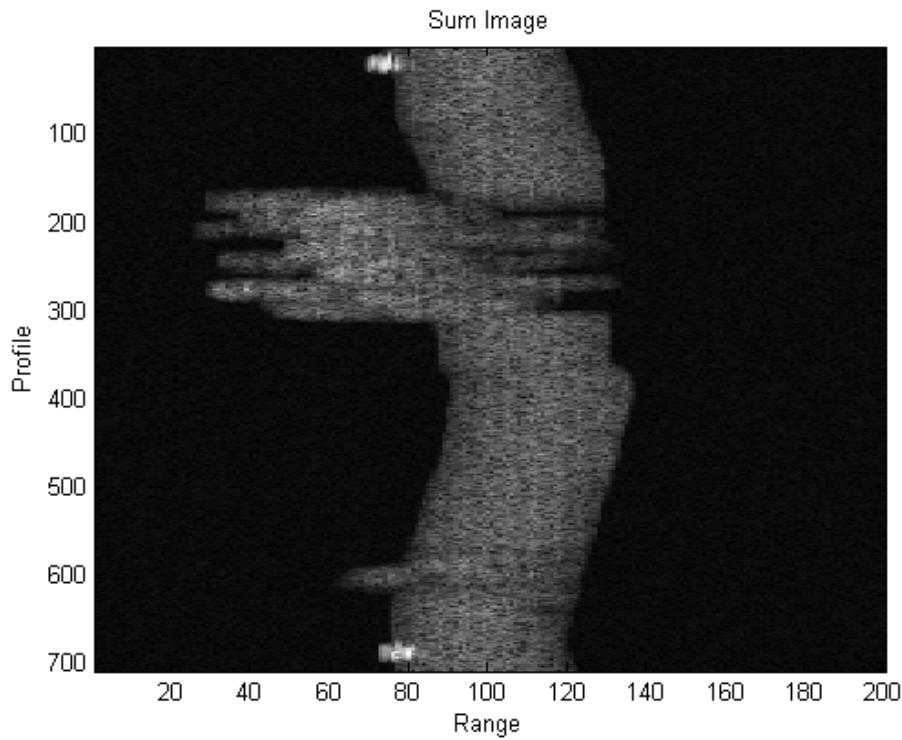
Figure 8-17 Copehill scene with two targets present at 200 m height. Range from left to right in 0.25 m bins, increasing range profile number from top to bottom
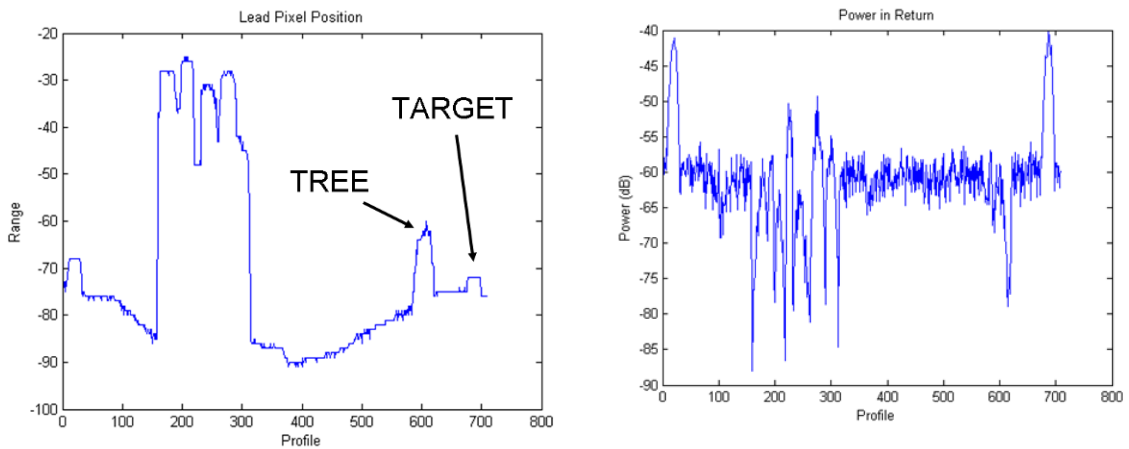
Lead pixel position   Power behind lead pixel

Figure 8-18 Lead pixel position in 0.25 m bins against range profile number (left) and lead pixel power (right) for two images of a target in the Copehill scene

Should the scene be imaged later in the trajectory when the radar has descended further, the target appears within the clutter band and the single tree now appears with a height similar to that of the first target in Figure 8-18.

Sum Image — Copehill with target present

Lead Pixel Position
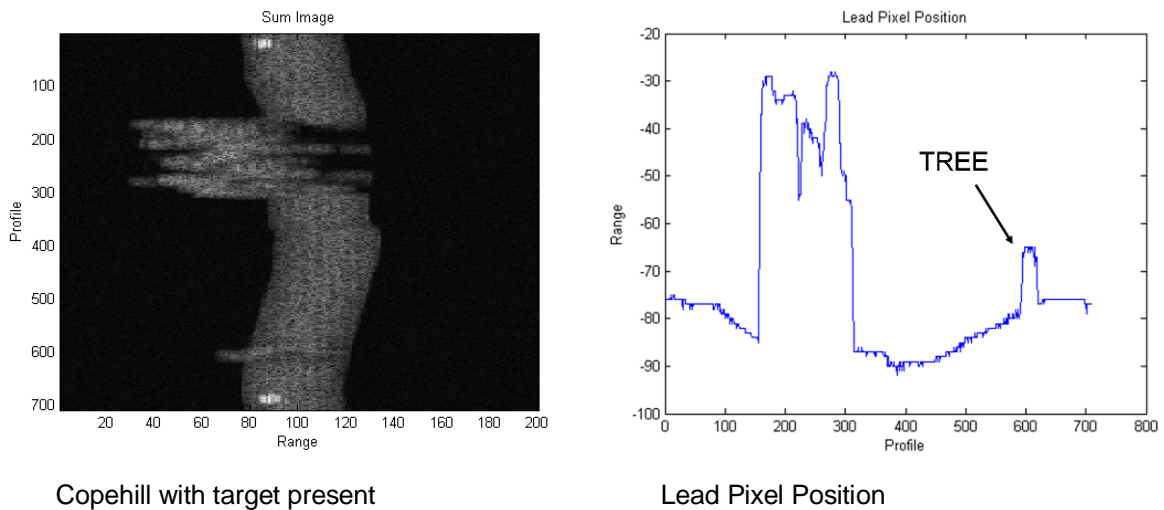
Figure 8-19 Copehill scene viewed from 190 m height. Intensity of returns against range in 0.25 m bins (left), and lead pixel position in metres against profile number (right)

While lead pixel detection is not a robust solution to aiding detection through height measurement, monopulse, described in section 8.5.5, appears to be a more promising technique.

8.5.5 Multi-channel ATR

Monopulse is an established radar technique that uses multiple beams, which are offset in azimuth and elevation angle from boresight, to provide the ability to achieve angular resolution within the radar beam. Historically, monopulse has been used to improve the angle estimation and aim-point selection for anti-air weapons. Monopulse has also been used to improve elevation angle estimation for terrain mapping SAR, for example [79].

The CLTG was configured to run in an amplitude comparison monopulse mode, simulating four simultaneous beams, allowing use of azimuth monopulse and elevation monopulse techniques. Figure 8-20 shows the relative position of the four beams in angle space. Because the target beams point in slightly different directions, if a target is offset from the antenna boresight, the amplitude of the return in one lobe will differ from the amplitude received through the others, and this difference is proportional to the angular error.

Beam A

Beam B

Beam D

Beam C
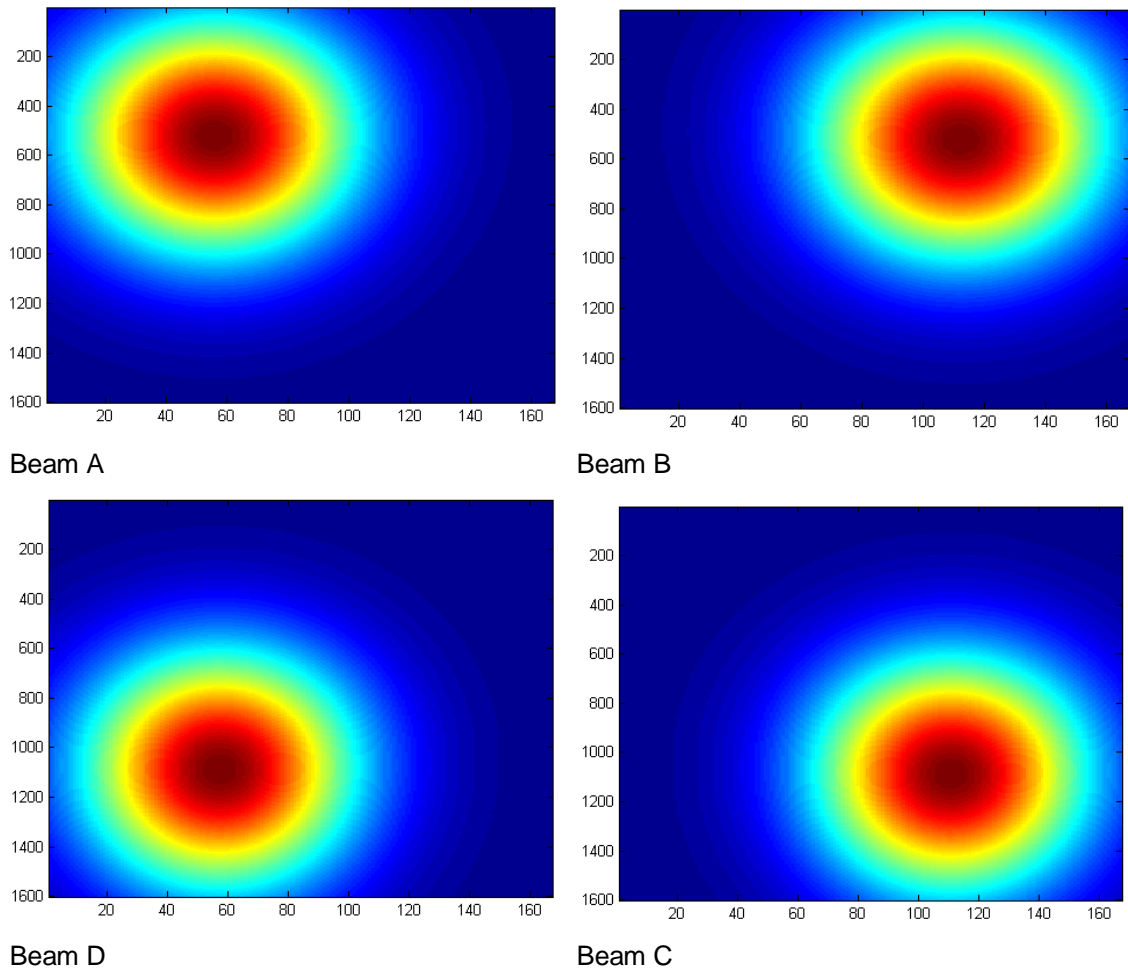
Figure 8-20 Beamshapes for the four monopulse beams

Summing all four beams produces a monopulse sum beam, equivalent to a single-beam antenna but with a broader beamwidth, equation (8.7). Summing the two beams on the left (A, D) and subtracting the sum of the two beams on the right (B, C) produces a monopulse azimuth difference beam, equation (8.8). In a similar manner, summing the bottom beams and subtracting the top beams produces an elevation difference beam equation (8.9). These beams, and the beamshape of a single-beam antenna are shown in Figure 8-21 for comparison.

$$Sum = (A + D + B + C) \qquad (8.7)$$

$$Azimuth\ Difference = (A + D) - (B + C) \qquad (8.8)$$

$$Elevation\ Difference = (A + B) - (C + D) \qquad (8.9)$$

Monopulse Sum

Single Beam

Azimuth Difference

Elevation Difference
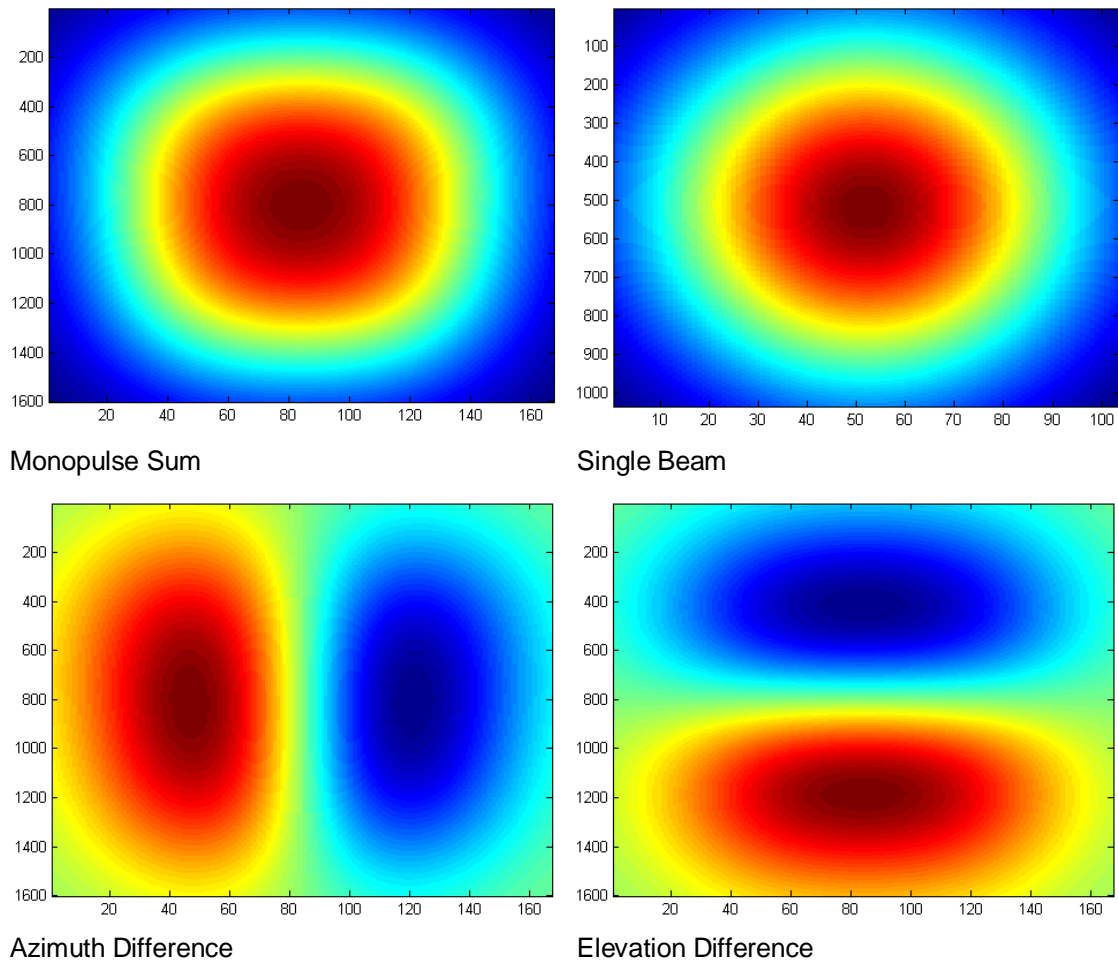
Figure 8-21 Monopulse beamshapes and single beam

The separation of the beams was chosen such that the beams are crossing at their 3 dB point in power. For a more detailed design, the optimum choice of squint angle could be calculated [80]. The left image of Figure 8-22 shows a cut through the centre of Beam A (left) and Beam B (right), and the separation between the beams at the 3 dB point. An azimuth monopulse slope example, equation (8.10), created from Beams A and B is shown in the right image of Figure 8-22.

$$Slope = \frac{A - B}{A + B} \qquad (8.10)$$

When an object is measured by the radar, a monopulse value can be calculated for each return, which can then be compared to the expected value from the slope to deduce the object's location to greater accuracy. This capability can be harnessed to improve SFM performance.

Beam A and B Power   Azimuth Monopulse slope
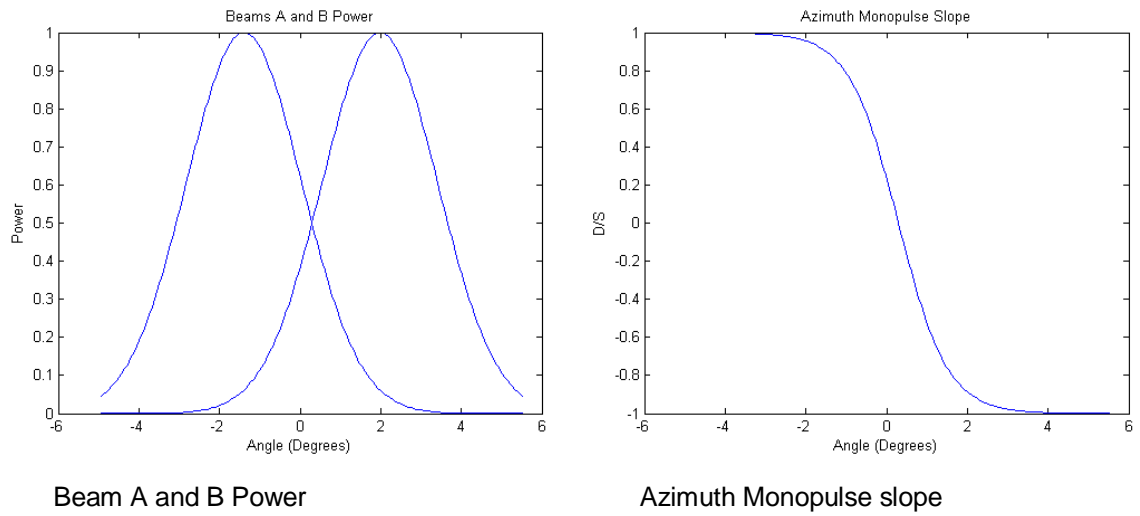
Figure 8-22 Beams A and B power profile and azimuth monopulse slope

8.5.6 Monopulse Tests

This section demonstrates the use of monopulse techniques against realistic targets in simulated SFM data created with the CLTG.

An example of the effect of using simulated monopulse radar in the SFM imaging geometry is shown in Figure 8-23.
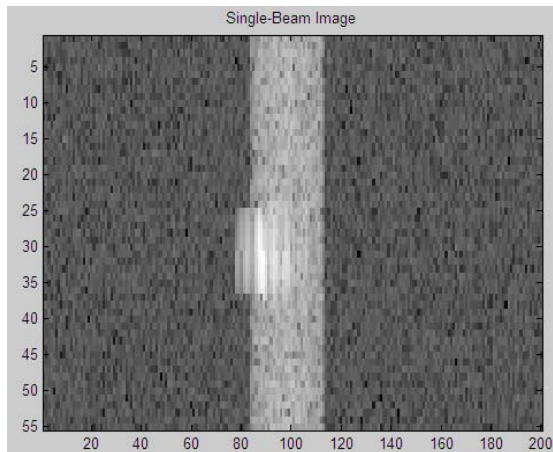
Single beam MBT from 200 m height
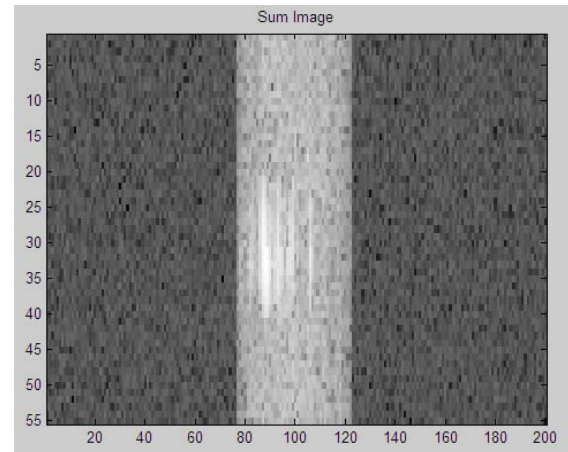
Sum monopulse beam, MBT, 200 m height
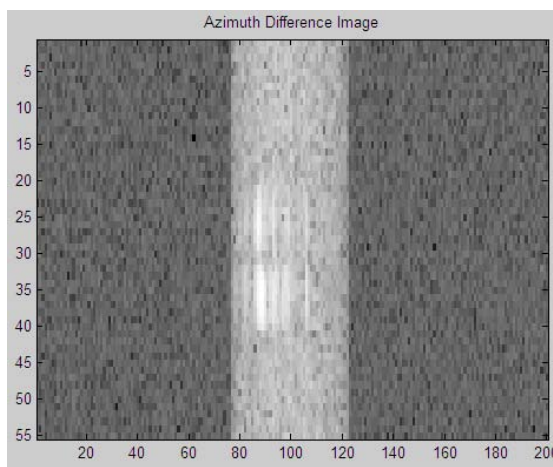
Azimuth Difference, MBT, 200 m height
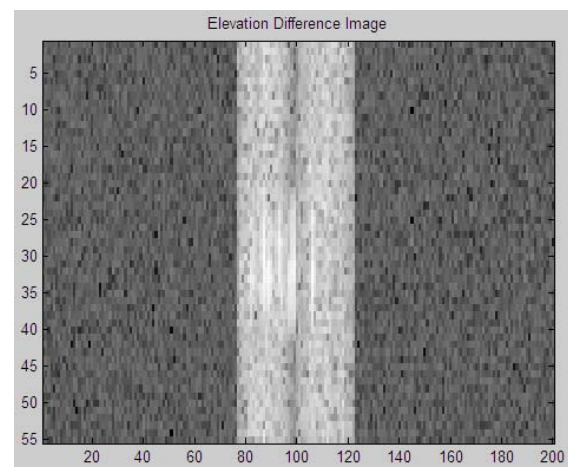
Elevation Difference, MBT, 200 m height

Figure 8-23 Single scans over an MBT for single and monopulse beams. Down-range from left to right in 0.25 m bins, increasing pulse number (time) down the image

The top left image in Figure 8-23 is that produced by a single beam antenna, and the top right that from the same antenna in monopulse mode.

The change from single beam to monopulse sum beam broadens the resulting beam in both elevation and azimuth by a little over 50%. The effect of this is to widen the clutter band in the down-range direction, now enveloping the target in the monopulse image, and to increase the number of ramps which include the target, shown as a stretch in the target returns in the *y*-axis of the plot.

While these changes would not be ideal for aiding target detection, the benefit of switching to a monopulse mode is the production of azimuth and elevation sum and difference beams.

The bottom left image in Figure 8-23 is the azimuth difference image, and a null can be seen when the target reaches boresight. The position of this null can fluctuate

slightly in each range cell depending on the distribution of scatterers within the target. The bottom right image is the elevation difference image, and here a null can be seen in the centre of the clutter patch, where the higher and lower beams cancel each other.

The three monopulse patterns can be used to calculate the azimuth or elevation monopulse value of any point in the image. At points where the target is within the beams, the monopulse value will estimate the position of the target within the beam.

The simple case of a bright reflector against a grass background (Figure 8-24) allows the potential of these monopulse techniques to be investigated for an SFM sensor. Here the target is clearly distinguishable against the clutter profile and noise background.

Sum Beam Image
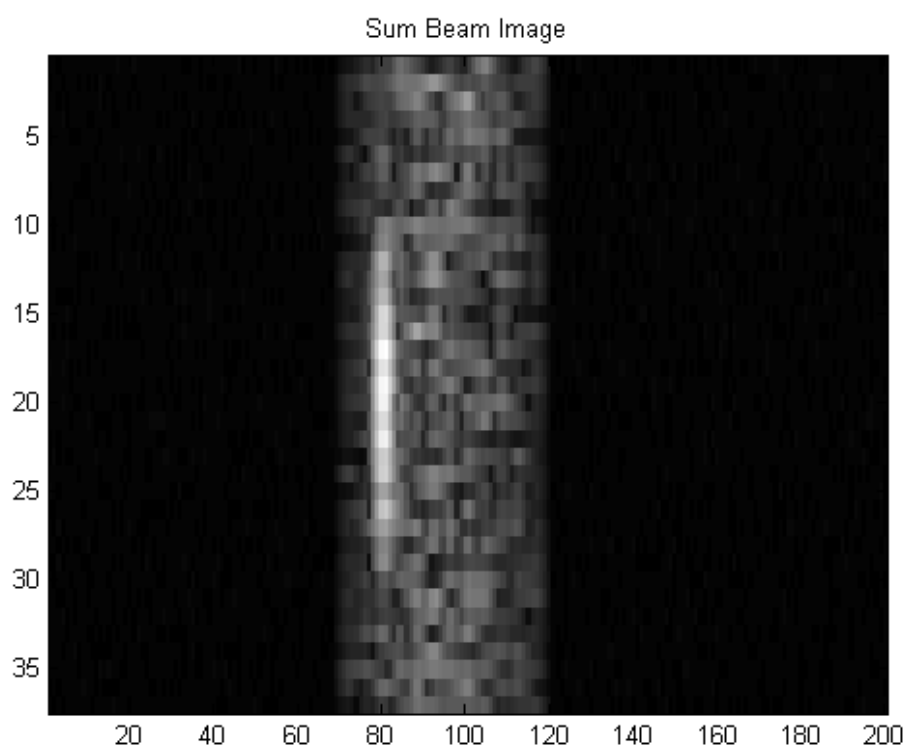
Figure 8-24 Bright point reflector against clutter background. Range from left to right in 0.25 m bins, increasing range profile number from top to bottom

Dividing the azimuth difference image by the sum image will produce a new image with values ranging from -1 to 1. Knowledge of the beam geometry can then be used to relate these values, which form the azimuth difference / sum (D/S) slope, to an angle within the beam.

Figure 8-25 shows the image produced when the azimuth D/S is calculated (left), and a cut through the target in the azimuth direction (right). The D/S value for a

single range bin (80) over a series of frequency ramps has been plotted and the slope is clearly visible. This slope, which exists over the target's extent in the beam (profiles 15 to 25) represents the angular correction between where the antenna is pointing, and where the target exists within the beam. Hence, as the beam scans onto the target to the left of the slope, the difference is positive (i.e. the target is to the right of boresight), and when the beam has passed the target the angular correction is negative. Noise in the system affects the monopulse calculation and is visible in the plot outside the target extent, but the slope is still close to a straight line as would be expected for this simple case.
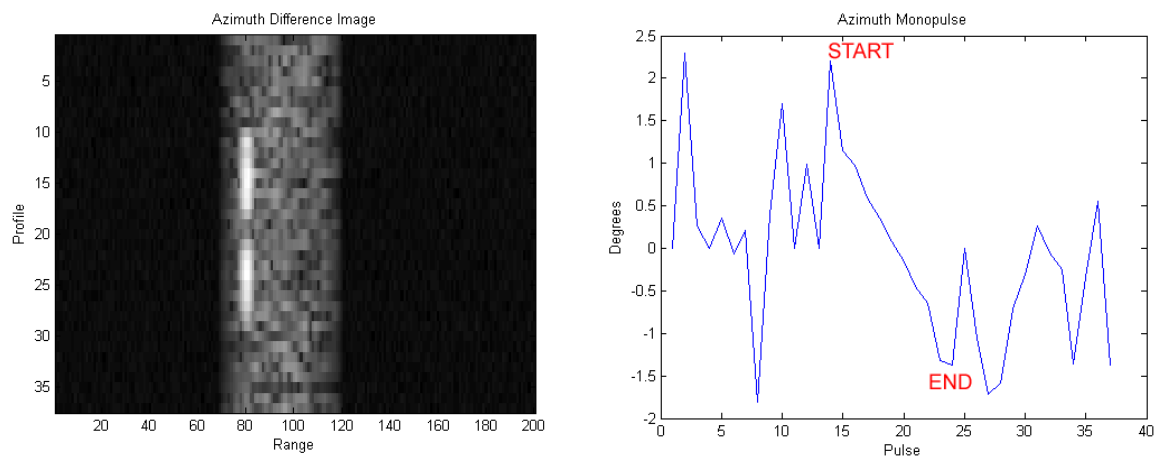
Point reflector azimuth difference image

Point reflector azimuth monopulse slope (°)

Figure 8-25 Azimuth monopulse for a bright point reflector. Intensity image on the left with range from left to right in 0.25 m bins, increasing range profile number from top to bottom. Azimuth monopulse slope for 40 pulses on the right

This extra angular information offers the possibility of significantly enhancing aim-point selection on the centre of the target at ranges where the beam footprint on the ground is significantly larger than the target's dimensions, as is used in systems such as Brimstone.

Figure 8-26 shows the same principles applied to the elevation monopulse beam for the point reflector. The elevation D/S has been calculated for an image without the reflector present and is shown on the top right. The lower images are the D/S slope for multiple ramps averaged, and for a single ramp.

Outside the clutter band, the D/S value is random and tends towards zero when summed. Over the clutter band, the D/S value changes from the near-range to far-range extent of the beam as shown by the plot in the top left. When a target is introduced, the target height makes the reflector appear closer in range than its true position on the ground. However, the elevation monopulse value for the target is for

its position within the beam, rather than range. This results in a step in the monopulse slope as shown by the lower images in Figure 8-26.

Point reflector elevation difference image   Clutter Elevation Monopulse profile (no target)

Point reflector elevation monopulse slope (averaged)

Point reflector elevation monopulse slope (single pulse)

Figure 8-26 Elevation monopulse for a bright point reflector and clutter, average and single pulse. Range on x-axis in 0.25 m bins

This step serves multiple purposes:

- Objects in the scene with a height extent can be detected using the variation from the expected monopulse profile;

- The size of this variation relates directly to the height of the target, and hence, a height filter can be used to discriminate targets from clutter;

- The monopulse value gives the true position of the target in the beam, reducing aim-point error when the beam is larger than the target.

Taking the difference between the expected ground elevation values and those returned by the scene, including the target, shows the effect of the point target (Figure 8-27). In this simple scene, the target stands out clearly as appearing at a

different angle to that predicted, and the size of the difference, 1.2° relates to a target height of 6.2 m, close to the true 5 m height of the reflector.
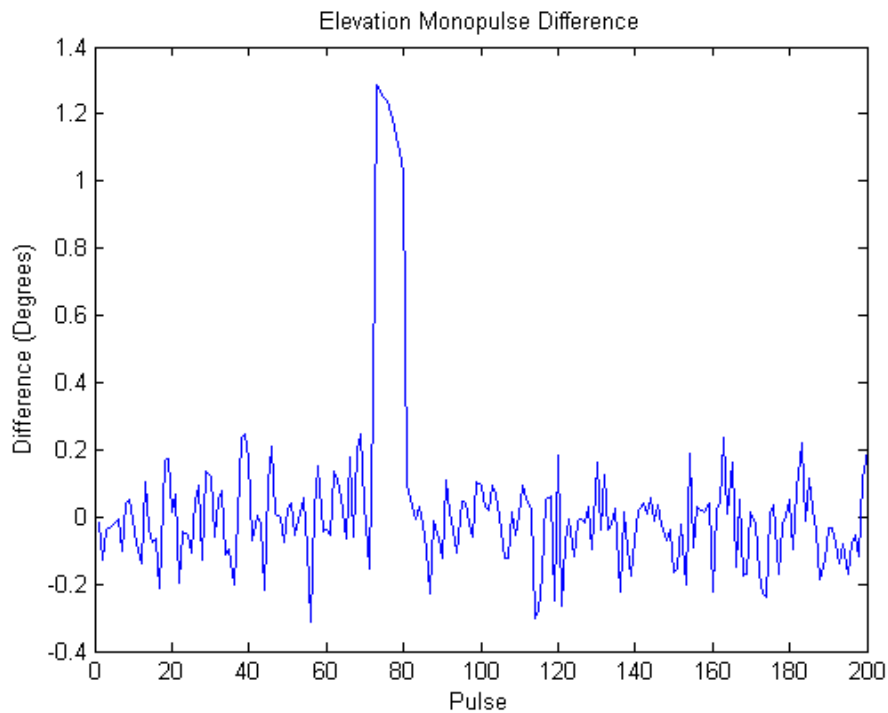
Figure 8-27 Difference between expected value and measured value. Monopulse difference signal on y-axis, pulse number on x-axis

Applying these techniques to a military target, the simulated MBT, still produces good results. With the vehicle against a grass clutter background, the azimuth monopulse slope for the target at range bin 88 is shown in Figure 8-28.
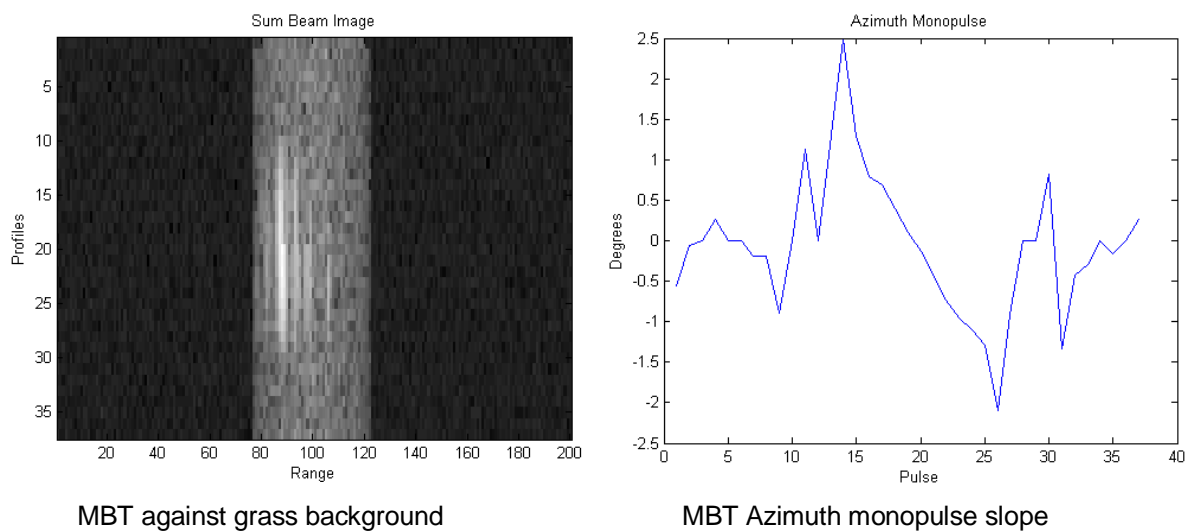
MBT against grass background   MBT Azimuth monopulse slope

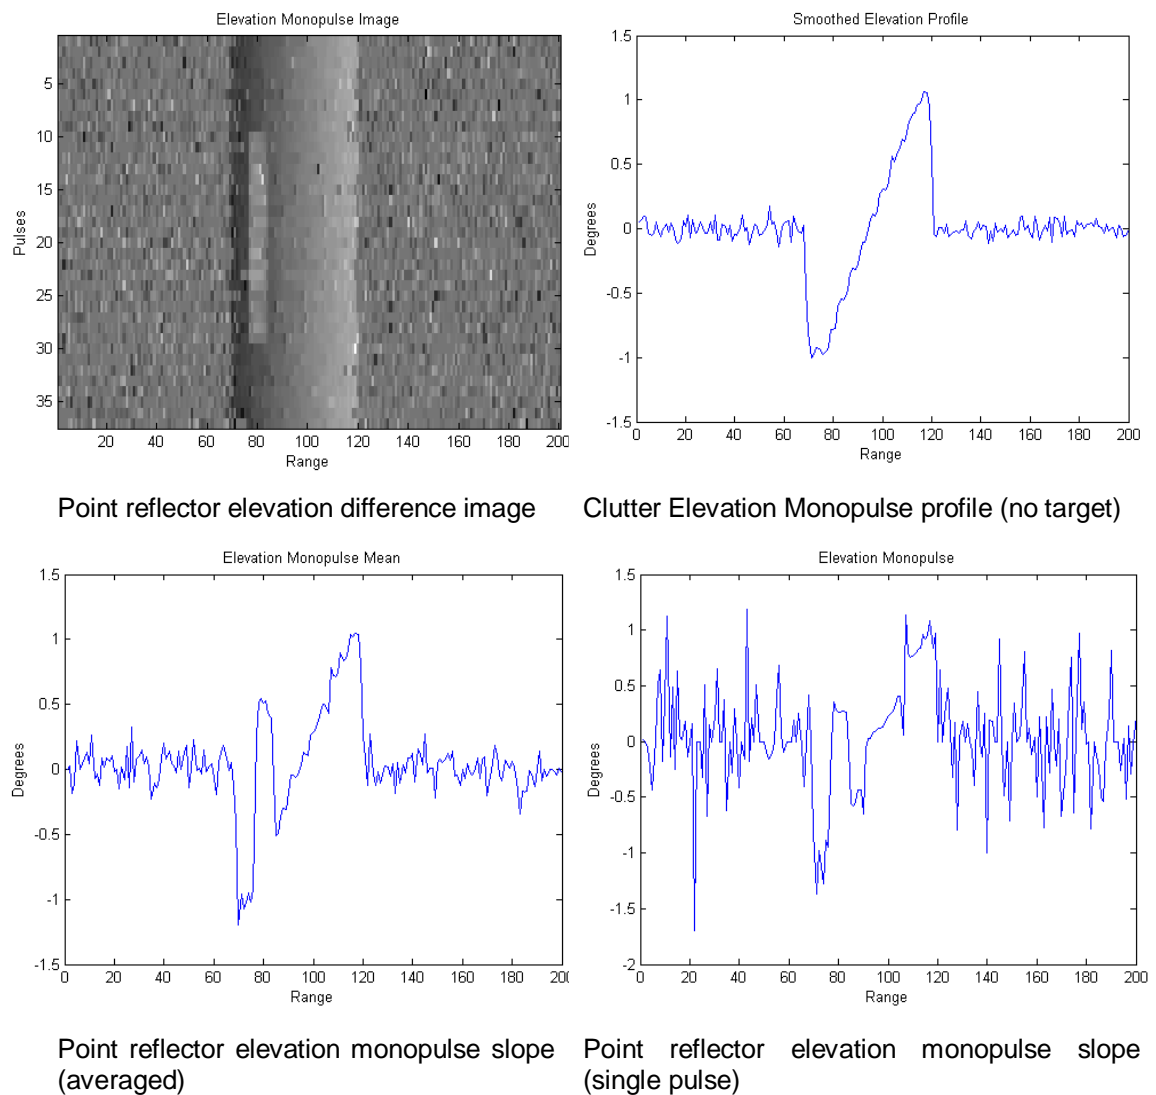*Figure 8-28 MBT Image and azimuth monopulse slope. Intensity image on the left with range from left to right in 0.25 m bins, increasing range profile number from top to bottom. Azimuth monopulse slope for 40 pulses on the right*

The bottom right plot in Figure 8-29 shows the elevation monopulse slope for the MBT. The target can be seen as a number of spikes in the slope, which relate to the height of the scattering areas of the target.

MBT elevation difference image

Clutter Elevation Monopulse profile (no target)

Elevation monopulse slope (single profile)

Elevation monopulse slope (averaged profiles)
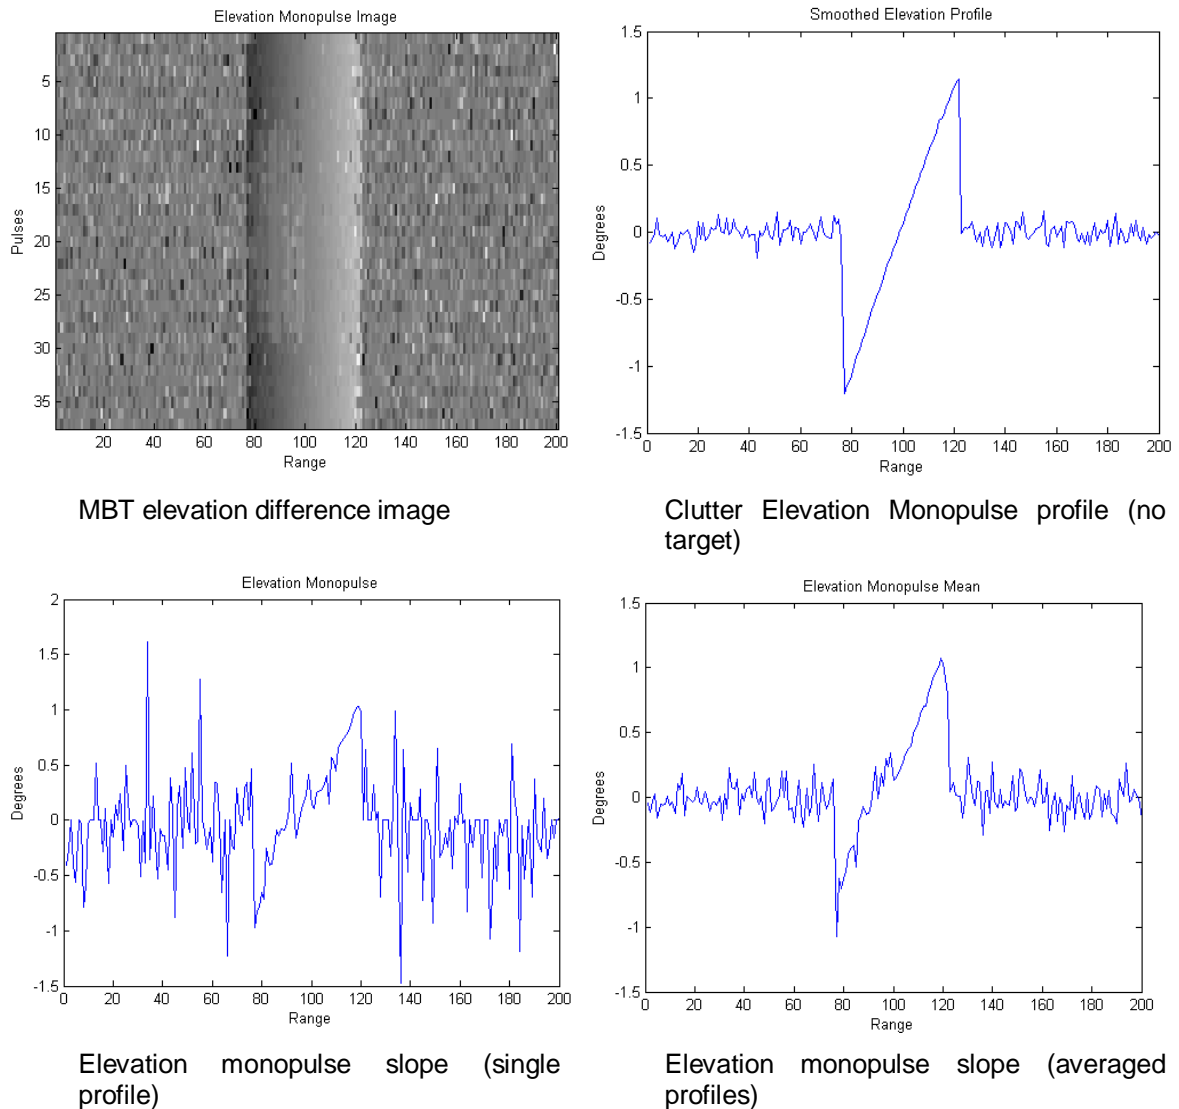
Figure 8-29 Monopulse slope and difference for MBT. Range on x-axis in 0.25 m bins

The difference plot in Figure 8-30 shows the target with a mean offset of 0.41° above the clutter. This offset relates to a target height of 2.24 m, close to the true height of an MBT.
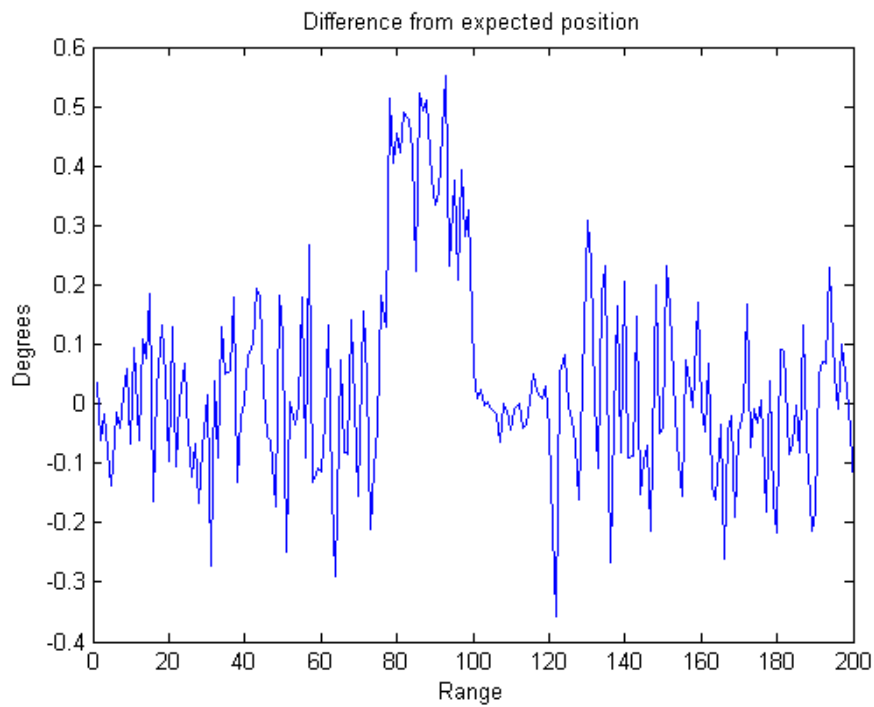
Figure 8-30 Elevation angle difference in degrees between measured and expected values. Range on x-axis in 0.25 m bins

8.5.7 Monopulse Height Detector

In order to test the utility of the monopulse measurements as an ATR technique, the outputs from the 2D ATR software developed in section 8.5.3 were combined with the monopulse measurements. This allowed a two-stage process to be employed for the detection of targets.

For the MBT model in a grassland scene, the input image and results of the ATR are shown in Figure 8-31.
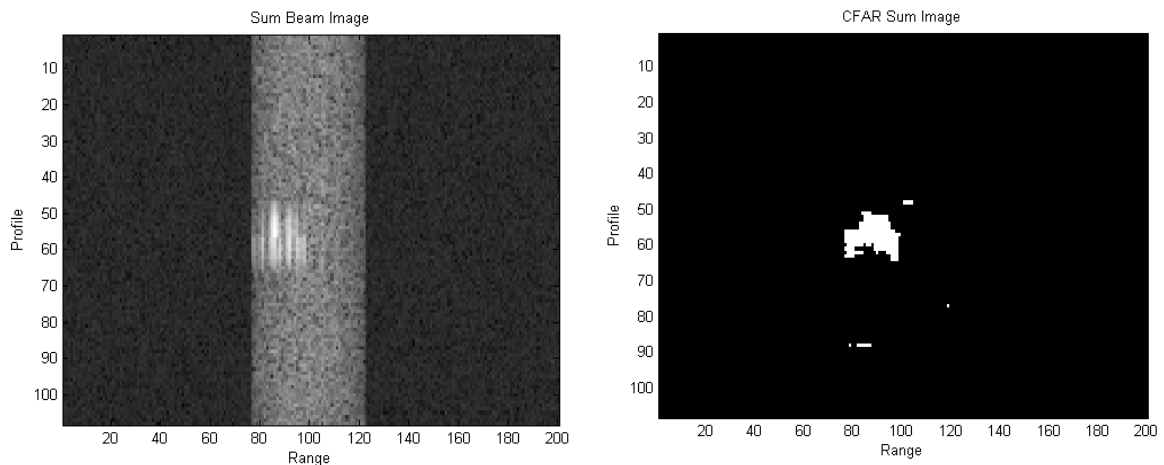
MBT, 200 m height   Thresholded Image

Figure 8-31 MBT in grassland (left), after thresholding (right), range on x-axis in 0.25 m bins

The monopulse slope through this target compared to the slope without the target present, and the difference from expected values are shown in Figure 8-32.
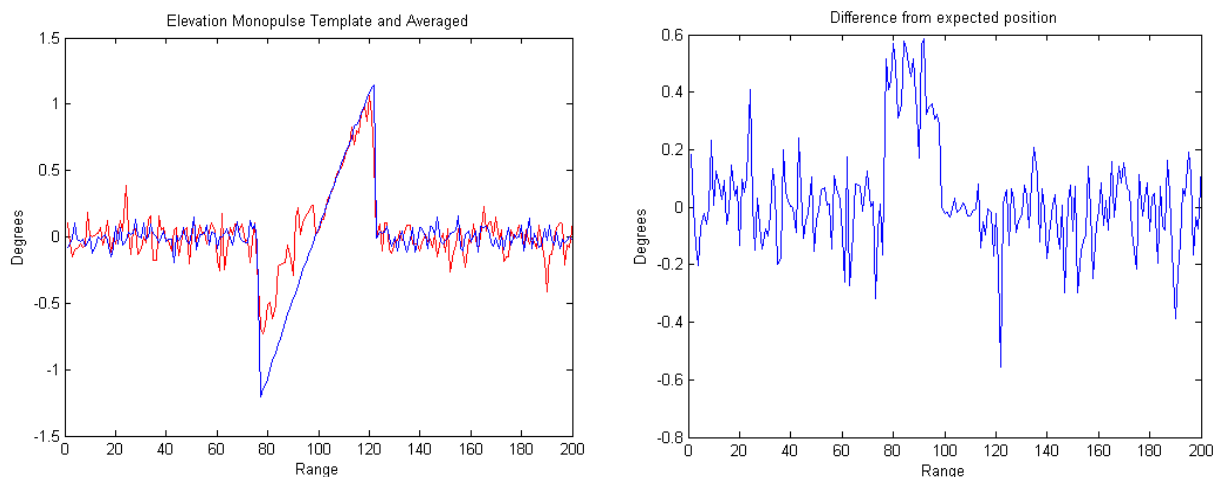
Elevation monopulse (red) against clean slope (blue)

Elevation difference from expected values

Figure 8-32 MBT Monopulse slope and difference, range on x-axis in 0.25 m bins

The elevation monopulse difference values can be used in two ways for aiding the target detection decision. Firstly, a threshold can be applied to the image based on the known target height and the monopulse error this would cause. The output of that process is shown to the left of Figure 8-33. Secondly, the monopulse difference image can be combined in an 'and' process with the 2D detector output, shown on the right of Figure 8-33.
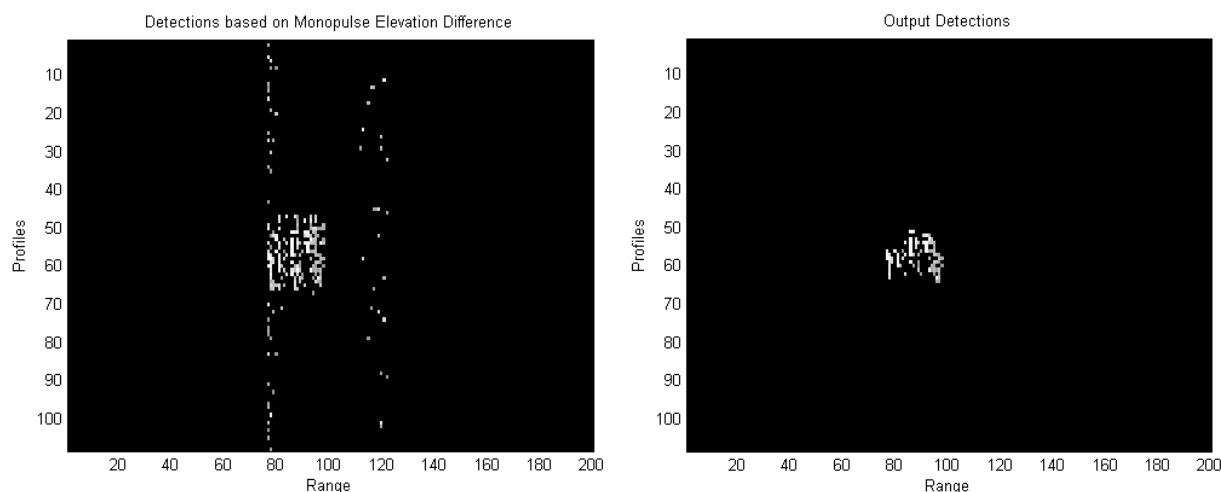
Figure 8-33 Thresholded monopulse values (left), and combined with 2D detector (right). Range on x-axis in 0.25 m bins

The monopulse threshold, based on expected target height, correctly selects the bulk of the pixels which make up the target return, and has scattered false alarms along the transition between clutter and noise, which could then be removed through a clustering algorithm which groups connected areas of pixels and rejects single-pixel detections. The combination of ATR and the monopulse image to the right of Figure 8-33 has no false alarm detections, and the mean monopulse difference value of the pixels selected by the thresholding is 0.41°, which corresponds closely to the correct target height.

From these results, a suitable approach to use of the elevation monopulse information would be to run an ATR scheme over the data and then cluster the results into objects. Elevation monopulse could then be used to calculate a height estimate for each object, and that information used to improve the targeting decision.

To test this, false targets were generated and run through the processing. A building 6 m wide, 10 m long and 3 m high was added to the scene shown in Figure 8-34. This object was similar in size and appearance to the MBT used previously but with less internal structure. The output of the initial detector, on the right of Figure 8-34, is very similar to that resulting from the MBT in Figure 8-31.
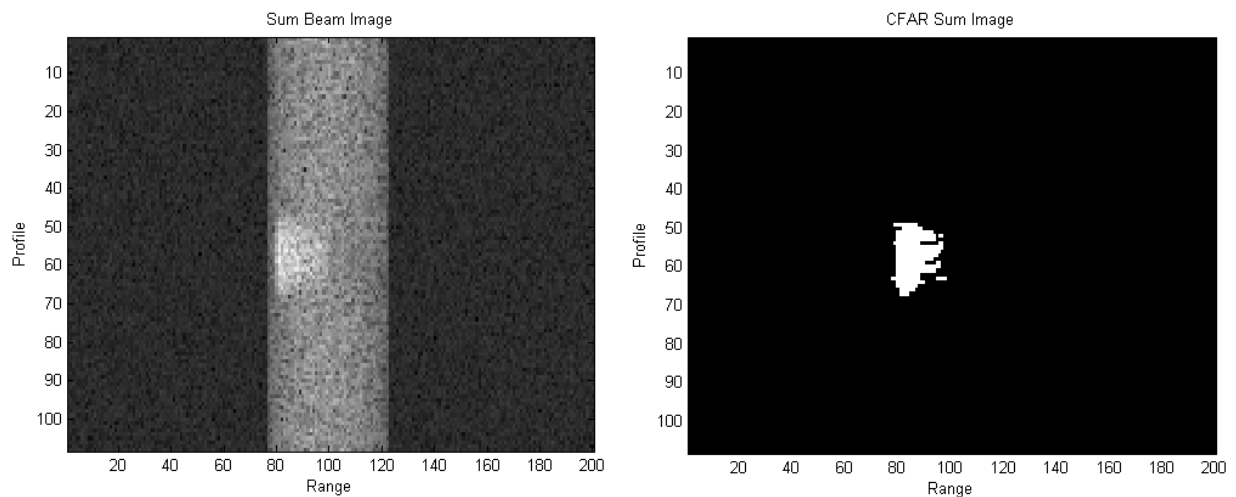
Figure 8-34 Building at 200 m and the thresholded output, range on x-axis in 0.25 m bins

When the monopulse information is utilised, a large fraction of the pixels within the cluster are rejected due to their height measurements exceeding the expected values for a target. Figure 8-35 shows the elevation monopulse slope plotted against the expected slope, and the pixels which remain after a threshold based on expected target height is applied. The number of pixels passed has dramatically reduced, offering added potential to reject the cluster. If the mean monopulse value of the pixels within the entire thresholded cluster in Figure 8-34 is considered, the value of 0.62 is significantly higher than the 0.41 measured from the target and could be used to reject the cluster.
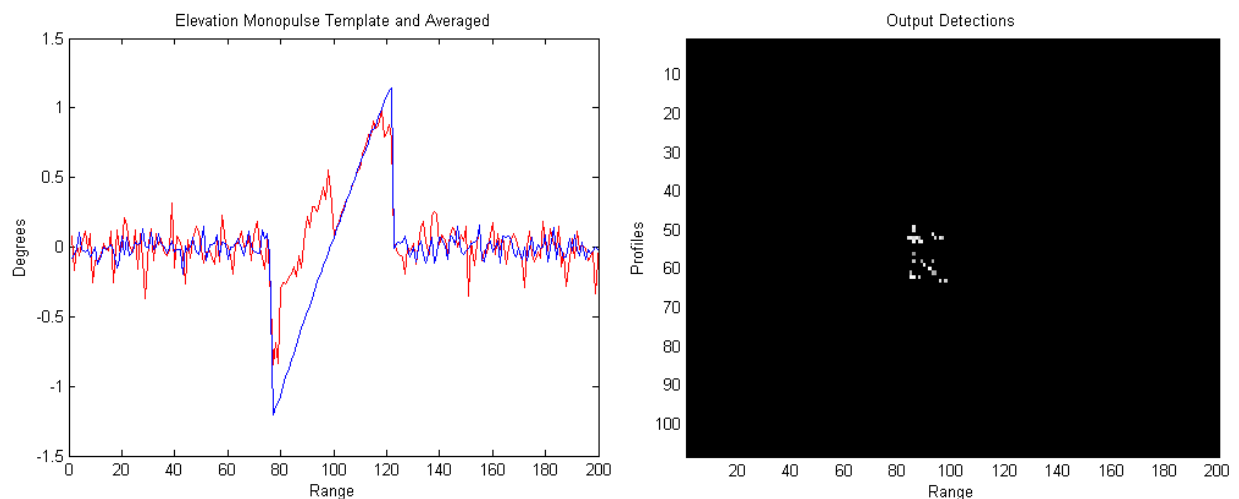
Figure 8-35 Building at 200 m, Monopulse elevation slope and output image, range on x-axis in 0.25 m bins

The capability of this height estimation to reject objects that do not fit into the expected range of target heights, when combined with improved ATR techniques,

should result in a sensor with relatively sophisticated detection processes giving the potential for a significant improvement in performance.

A similar structure, 6 m wide, 10 m long, but only 1 m high was then added to the scene, Figure 8-36.
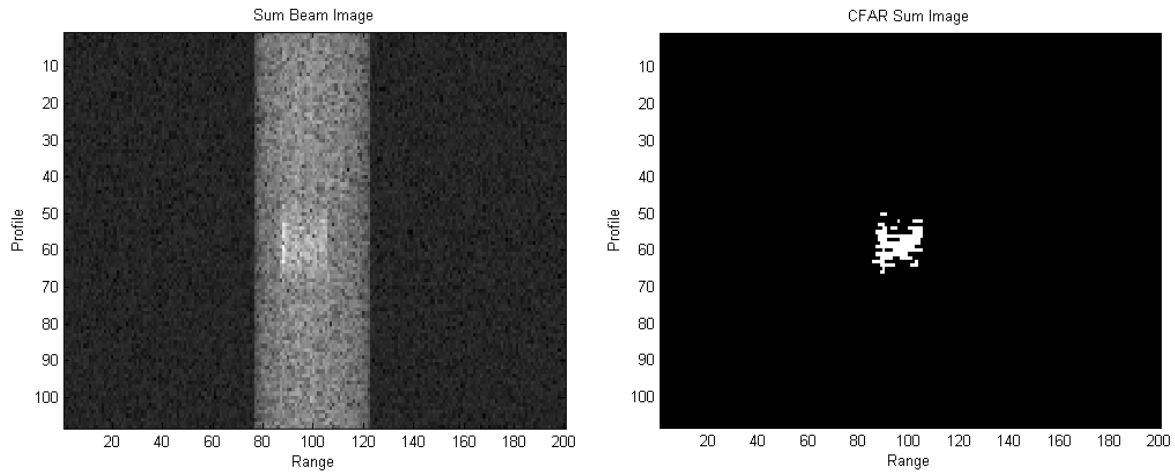
Figure 8-36 Object, 1 m high, image from 200 m and the thresholded output. , Range on x-axis in 0.25 m bins

Figure 8-37 shows the elevation monopulse slope plotted against the expected slope, and the pixels which remain after a threshold based on expected target height is applied. The number of pixels passed has dramatically reduced, offering added potential to reject the cluster. If the mean monopulse value of the pixels within the thresholded cluster in Figure 8-34 is considered, the value of 0.23 is significantly lower than the 0.41 expected from the target and could be used to reject the entire cluster.
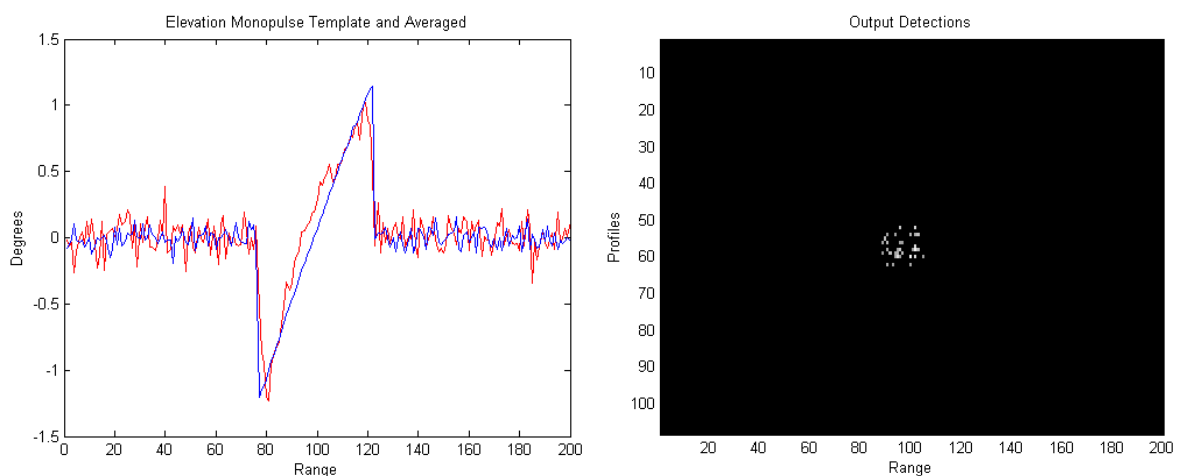
Figure 8-37 Object, 1 m high at 200 m, elevation slopes and output image. Range on x-axis in 0.25 m bins

The final test in the simple scene was a house-sized object, 15 m long, 10 m wide and 5 m high, shown in Figure 8-38
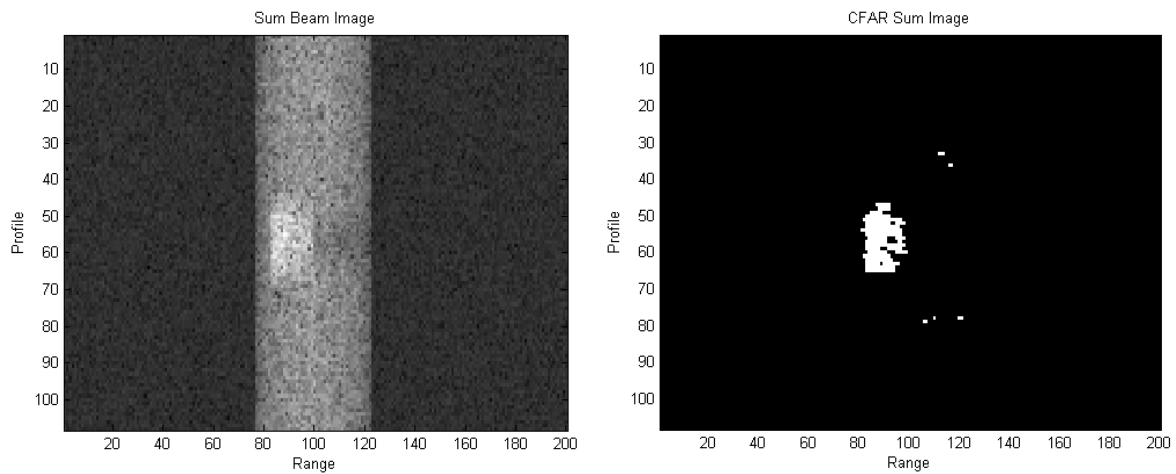
Figure 8-38 Building, 5 m high at 200 m and the thresholded output. Range on x-axis in 0.25 m bins

Figure 8-39 shows the elevation monopulse slope plotted against the expected slope, and the pixels which remain after a threshold based on expected target height is applied. The number of pixels passed through has dramatically reduced, offering added potential to reject the cluster. If the mean monopulse value of the pixels within the thresholded cluster in Figure 8-34 is considered, the value of 0.80°, twice the 0.41° measured from the target, means that a few pixels pass through the height filter, and the cluster can be rejected easily.
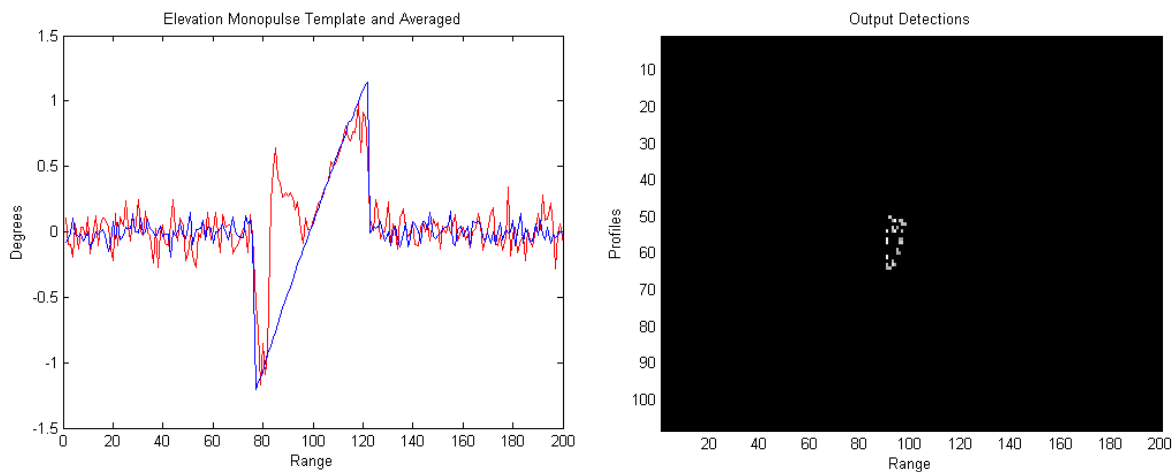
Figure 8-39 Building, 5 m high at 200 m range, Elevation monopulse slope (left) and output image (right). Range on x-axis in 0.25 m bins

8.5.8 Scene with trees and target

To examine the ability to detect targets in a more cluttered environment, a scene was created in which the MBT model was placed between two trees, each positioned approximately 7 m from the target in the azimuth direction to avoid obscuration. The CLTG object editor displaying this scene is shown in Figure 8-40. The upper image is a type map which defines reflectivities, where white is background, green is tree, and brown is vehicle. The lower image is a height map, in which colours transition from blue to light green as height increases.
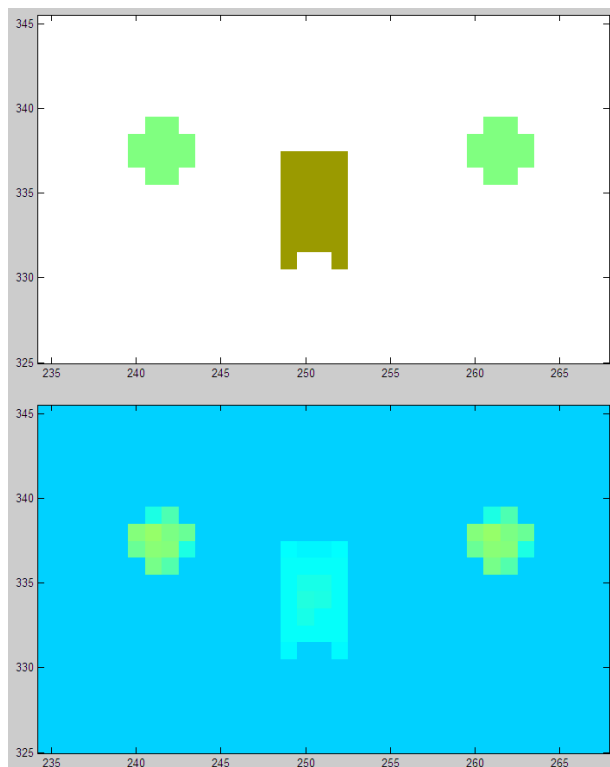
Figure 8-40 CLTG object editor containing an MBT and two trees
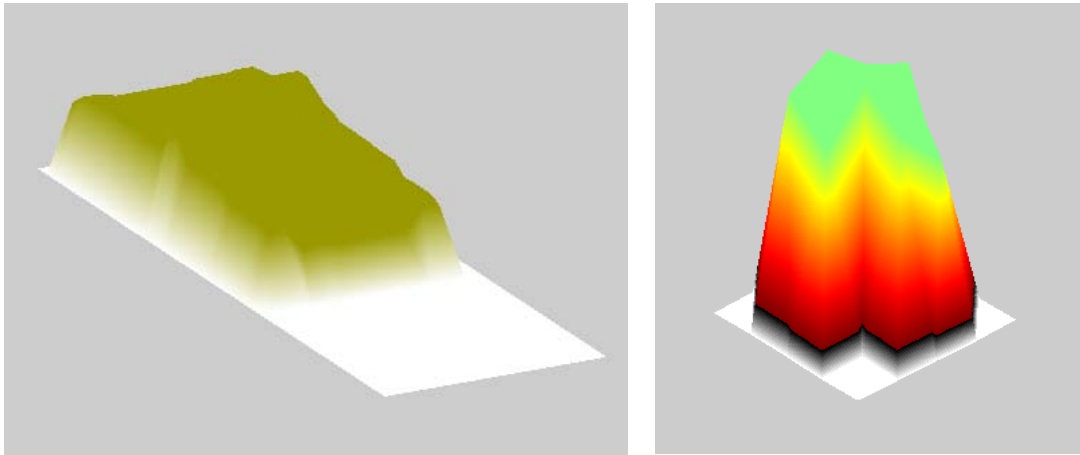
Figure 8-41 Simple MBT (left) and tree (right) object files

The scene was imaged from a height of 200 m. Figure 8-42 shows the image of the MBT in the centre, and fainter returns from the two trees on either side of the target.
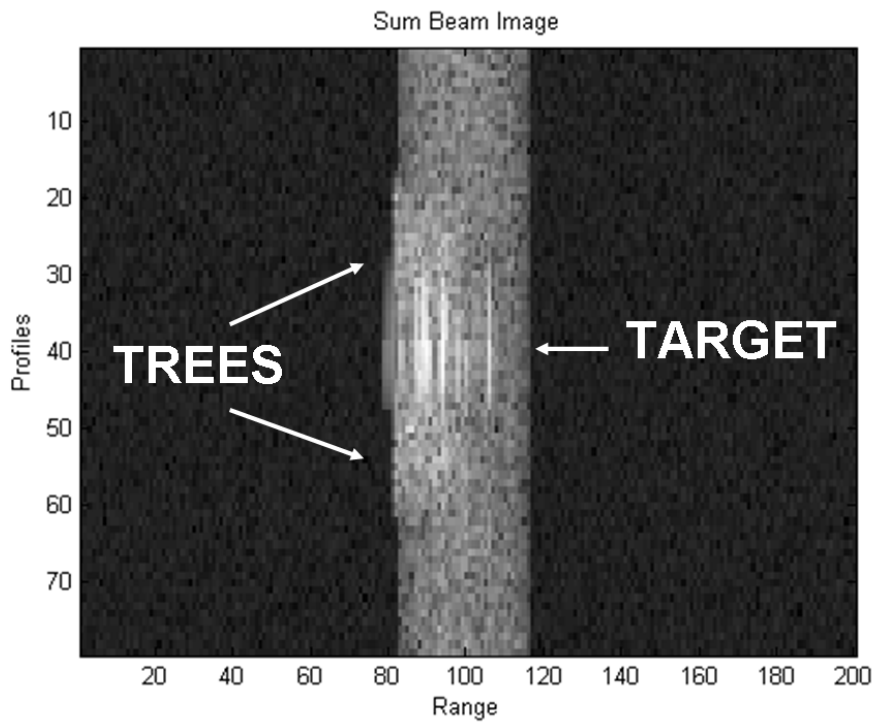
Figure 8-42 MBT and two trees imaged from 150 m. Range on x-axis in 0.25 m bins

The monopulse elevation difference image for this scene is shown in Figure 8-43. Here the two trees stand out as having a large difference from the expected slope, and the target returns in the centre are noticeable but less significant. Figure 8-44 highlights these differences, with the MBT (left) showing a small deviation from the expected slope.
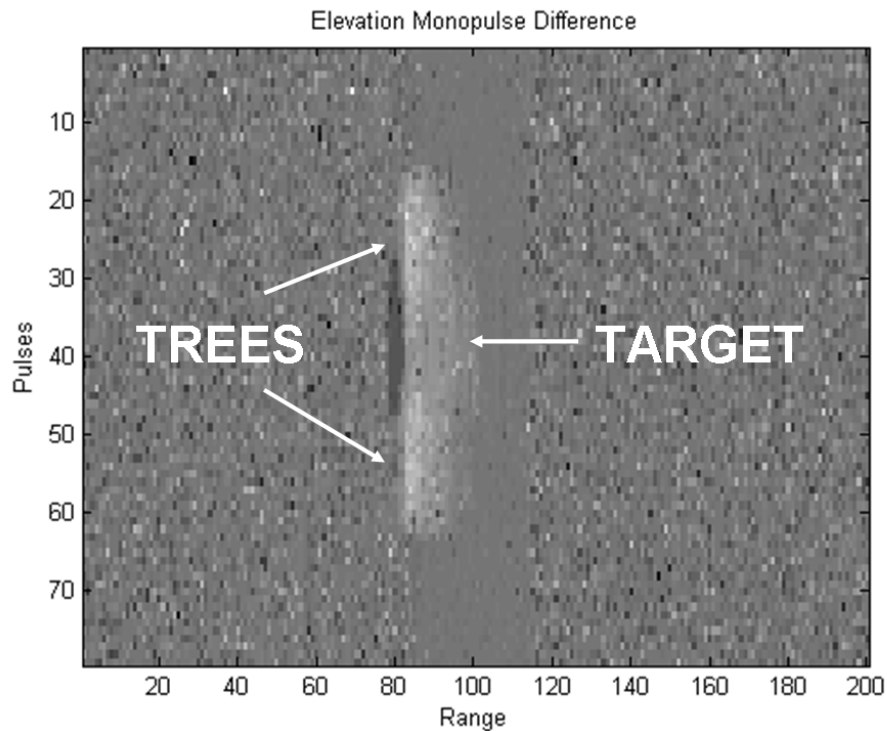
Figure 8-43 Monopulse elevation difference image. Range on x-axis in 0.25 m bins
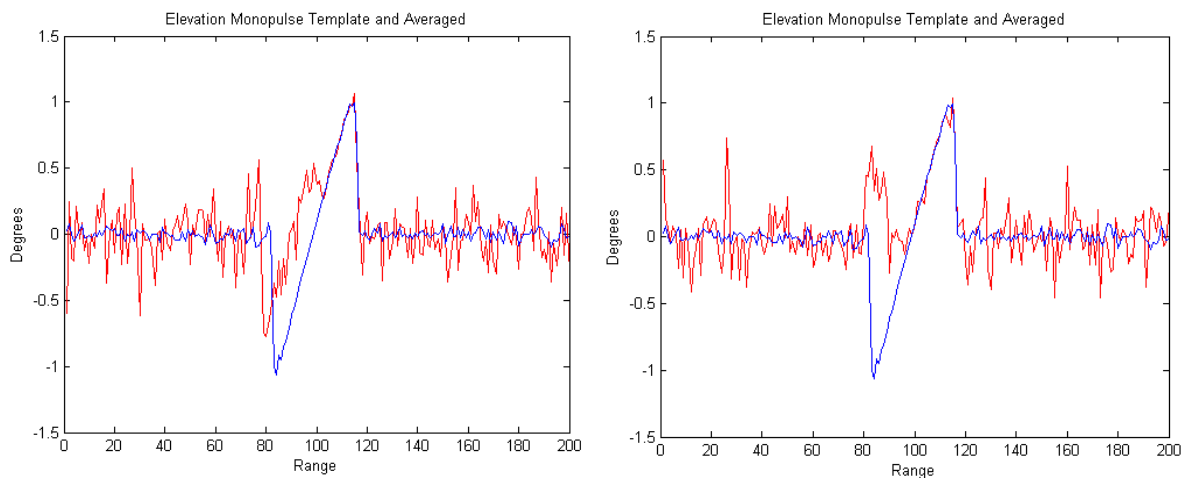
Figure 8-44 MBT (left) and tree (right) elevation monopulse difference. Range on x-axis in 0.25 m bins

Thresholding can be applied to the radar image to detect bright objects in the scene, and limits can be applied to the elevation monopulse image to detect only objects that fall within a certain height range. As shown in Figure 8-45, the detector, after simple clustering has been applied, rejects the returns from the trees, passing only pixels from the target. The monopulse threshold detects the target pixels, rejects the bulk of the returns from the trees, only detecting pixels at the down-range extreme of the trees. The application of clustering to this image should be able to fit the target extent.

234
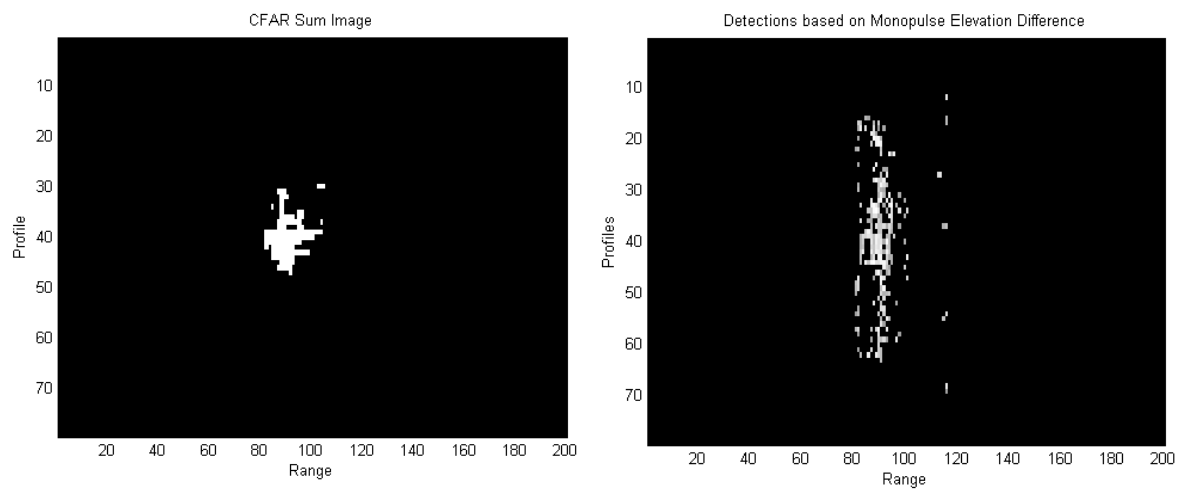
Figure 8-45 CFAR and clustering (left) and Monopulse (right) detections. Range on x-axis in 0.25 m bins

When these two techniques are combined, a cluster of pixels are output as shown in Figure 8-46. No pixels from the trees are passed though, but the bulk of the returns from the target remain.
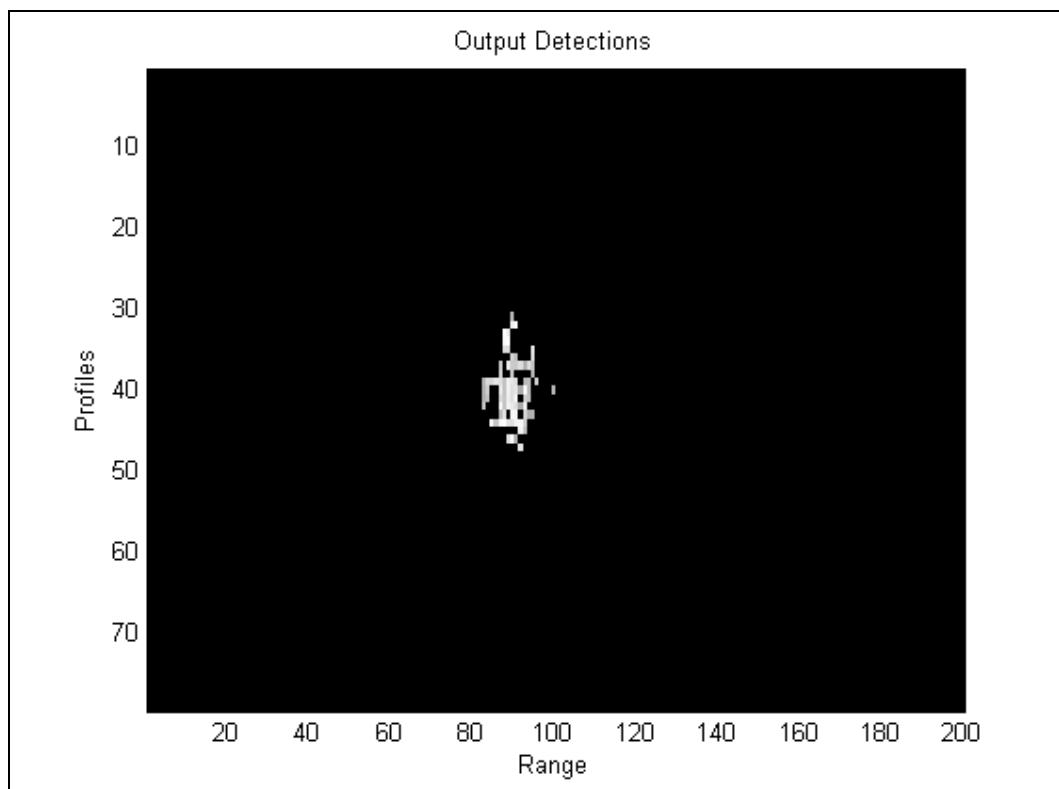
Figure 8-46 Output detections on MBT. Range on x-axis in 0.25 m bins

This shows the efficiency of the application of thresholding combined with elevation monopulse when applied to more cluttered scenes. Further work to investigate this would require the acquisition of real trials data to verify the models used.

8.5.9 Monopulse Accuracy

The monopulse corrections derived in the previous section indicate that good performance can be expected from the system. While the modelling has been made as accurate as possible, there are still differences between the simulated data and the information gathered by a real radar.

Since an accurate ground tracking filter was not developed, the monopulse slopes used to calculate the difference between measured and expected angle were based on imagery of the scene with no targets present. This process could be replaced with a tracking filter to establish the height variation of the ground within the scan, and automated generation of the expected profile based on knowledge of the sensor height and beam patterns.

The scatterers used to create the synthetic target models have an RCS that changes as the illumination angle changes, the value of the RCS being derived from ISAR turntable data. Since the SFM always illuminates the scatterers from the same angle, this does not change during the simulated descent and the target will remain correlated. However, the modelling produces the clutter which is decorrelated from ramp to ramp (apart from discrete clutter features). Coherent modelling suggests that, for the ramp rate and SFM descent speed used, the returns from the target should remain coherent for approximately 30 ramps. This is longer than the target is imaged, so the CLTG coherence is acceptable.

Clutter should decorrelate over approximately 10 ramps; hence, the CLTG method of using completely decorrelated clutter is less accurate. The effects of having uncorrelated clutter should not significantly affect the utility of the analysis.

Estimates of the monopulse angle measured from a target which allow height and accurate azimuth angle to be measured are affected by both noise in the system and clutter. For an amplitude monopulse system, the monopulse error due to noise can be estimated by equation (8.11) [66].

$$\sigma = \frac{\theta_{Bsum}}{k_m \sqrt{2n\,S/N}}[1 + D/S]^{0.5} \tag{8.11}$$

where:

- σ is the standard deviation of the angular error in degrees;
- θ_{Bsum} is the sum beamwidth in degrees;
- k_m is the D/S slope normalised to the sum beamwidth;
- n is the number of ramps used for the integration (usually 1);
- S/N is the signal to noise ratio in linear units;
- *D/S* represents the normalised ratio of the target's returns.

This equation applies whether a scatterer's return is noise or clutter limited. If the return is clutter limited, then SNR is replaced by SCR.

The SNR for an SFM in a rural area should be in the region of 15 to 25 dB depending on the background type and available power.

For a 15 dB SNR, the accuracy is expected to be 0.8°, and at 25 dB SNR, this error has fallen to 0.24°. Both of these represent a significant improvement over the finest possible beamwidth of 1.7° achievable at 94 GHz with the largest antenna that can be fitted into an SFM.

In the imagery shown, these errors are revealed by the variation in the height estimation over a target, and it has been shown that the magnitude of the errors does not appear to prevent height information being used.

8.5.10   Conclusions

A number of approaches have been shown that offer potential for improvement to the relatively basic radar system employed on the current generation of Sensor Fused Munitions.

Developments in FMCW radar hardware allow for fine range resolution and rapid ramp rates while retaining the low cost, robustness and low power requirements of solid-state electronics.

Improved radar resolution allows more sophisticated target detection algorithms to be used in place of the existing approach as a simple bright object detector. Detailed range-profiling and 2D image analysis should reduce the false alarm rate and allow high detection probabilities.

It has been shown that a multi-channel radar employing monopulse techniques has the potential to provide useful information about the height and position of objects in the scene. The height information could aid both the detection process and target classification decision, while azimuth monopulse can reduce the aim-point error of the weapon, significantly improving effectiveness.

9 Conclusions & Future Work

9.1 Conclusions

The aim of this thesis has been to investigate the sensitivity of ATR performance in missile seekers to variations in radar system parameters, and to investigate techniques that can provide useful additional information. The problem studied is the ability to select targets of a given class whilst minimising selection of false objects from a cluttered background. The techniques developed are robust to working within the constraints of a missile system. For each object (targets and false targets), it is desirable to pass high-quality attribute information to a target selection system.

In Chapter 3, the theory of SAR resolution from a squinted and scanning DBS missile seeker was developed. Chapter 4 described data gathered using an airborne system representative of a missile seeker. The data is of a high quality, with resolution suitable to support target recognition, but not identification. This data has been used to explore the sensitivities of ATR performance, and has been accompanied by synthetic data from a radar simulation. The synthetic radar data has, where possible, been validated against real data and provides a useful resource to expand the conditions available for analysis beyond those in the limited set of trials data.

Chapter 5 explored the size and shape information which could be extracted from a target's signature, introducing a robust model-matching approach before looking at segmentation performance in conditions of reduced sensitivity. The chapter began with an outline of the ATR process, and identified how model-matching could be applied to reduce the number of false targets caused by clutter objects. An approach was developed which used simple target models in the CLTG seeker simulation over a 360° rotation to produce radar images of that target from which the segmented outline was stored as a template. The intensity variation within the target signature is not used, resulting in an approach that is robust to changes in the target configuration. As these changes can rapidly degrade the performance of a template match that relies on intensity variation. The approach can also be used where insufficient data is available on a target to build a model of the fidelity required for a detailed match. Correlation tests were applied using template sets from two synthetic target models, and a range of target and clutter objects which

had been passed as potential targets by the ATR algorithm. The ability to reject 50% of false alarms while maintaining a high P_d was reported, indicating that the algorithm made additional use of the object shape which was not being captured by the spatial features already in the classifier. Given the relatively small overhead in generating shape-based templates and the speed of the classification process, it appears to offer good benefits as an additional processing stage.

The ability to detect targets and measure their dimensions under conditions of reduced sensitivity was tested. This is an area of particular concern for missile seekers, where the high frequency transmissions can undergo severe attenuation in heavy rain. While increased attenuation reduces background clutter contrast, the ability to detect armoured targets remains high, although slightly less so for unarmoured vehicles. Attribute measurement quality falls away before detection capability, but is still retained to levels where the imagery produced would otherwise be considered of low quality. These findings help to understand the range of conditions under which radars can be used for target selection, and indicate that a robust capability may exist at lower levels of sensitivity than expected.

Chapter 6 considered the use of information from shadows in high resolution radar imagery, primarily for the case of moving targets, where information in the bright signature is lost. The ability to detect shadows, with the option of cueing from the bright signature, was demonstrated, and air-carry data was analysed which showed that, where the bright signature was too distorted to give useful dimensions, separability between an armoured vehicle and a 4x4 could be achieved on the basis of shadow size information. These tests were expanded through the use of simulated data, which confirmed results from air-carry data, and added an additional large target class, again showing good separability.

The ability to measure velocity from changes in the shadow centroid position between images was tested over short time intervals. Whereas velocity measures from the target's bright signature were subject to errors in the ability to reposition the target in the cross-range direction using monopulse, the shadow's position could be measured to a high level of accuracy, and hence a good velocity measure produced. This was tested for the simulated data of three different vehicle types moving at different speeds and on different headings.

The challenges involved in using shadow information were explored. A range of velocities were identified, for a given seeker image formation time, above which the target's motion may result in a loss of shadow contrast that prevents reliable

segmentation. The contrast of the local background, and knowledge of the terrain slope were both identified as factors which would affect performance. Contrast being a requirement for segmentation of the shadow, and terrain slope affecting the shadow length, and therefore target height estimate.

Chapter 7 investigated the effects that the choice of polarisation made on target detectability, and the attribute information used to separate targets from clutter objects. Using data from three different radars, comparisons were made between HH, VV and HV linear polarisations and co- and cross-polar circular polarisations. The co-polar circular channel appeared to offer the best contrast between target and background, and the best separability between target and clutter objects when using statistical feature discriminants. However, the performance differences between the options were not so large that use of any of the channels would be ruled out.

Chapter 8 investigated ATR techniques for a different class of seeker radar, with a downward-pointing antenna, which could be used in a future sensor fused munition. A study of existing technologies identified the parameter space in which the sensor would need to operate. Updated hardware, which could support high-resolution monopulse imaging, was described. An approach was developed to make use of monopulse information to measure the height of objects, which was tested against simulated data and indicated that it gave a much improved capability to separate targets from clutter.

In summary, each section examines a technique through which seeker performance could be improved and tests it against representative data. These studies can be used to help inform seeker hardware designers and the processing used on the imagery.

9.2   Future work

In chapter 5 it was suggested that the model matching approach be tested against a broad set of real data, for example the MSTAR dataset. The range of vehicles present would be sufficient to test whether, as the research undertaken indicated, a limited set of simple models was sufficient to aid recognition of a target class. Representative clutter detections from the MSTAR dataset should be compared to the template sets to measure the capability to reject false targets that have passed detection and segmentation stages.

The effects of reduced sensitivity should be also tested against a broader set of targets, and a greater numbers of features examined to judge the degradation of capability. This could be combined with an analysis of the impact of sensitivity on other uses for the radar imagery, such as scene-matching with prior information. This can provide the capability to remove navigation errors, but will become challenging as clutter contrast is reduced.

For the shadow analysis and detection work of chapter 6, the analysis approach could be integrated as part of the existing ATR process, although additional air-carry data would be required for further testing as only a limited quantity of suitable data were available. The current results are from a single moving vehicle against rural clutter. The process would benefit from application to more challenging scenarios, with multiple moving targets and greater levels of clutter to provide potential confusers.

The polarisation investigation of chapter 7 considered contrast and size measures for a limited set of target and clutter types, using a single approach to ATR. An expansion of the target and clutter sets would be valuable to further test the findings. The implications of polarisation choice should be tested on a broader range of algorithms, as well a study into the benefit of polarimetric features compared to a single polarisation.

While the elevation monopulse approach developed in chapter 8 is very specific to that type of weapon, the technique could be applied to any radar sensor imaging from a steep geometry. The approach could be further tested through a modelling approach which includes drift and pendulation in the descent of the sensor, and a study into compensation for these effects. Real data could be gathered for model validation given a suitable radar and trials location with a steep depression angle.

References

[1]   Licata, W.H., 'Missile Seekers for Strike Warfare Beyond the Year 2000', RTO Lecture Series, Technologies for Future Precision Strike Missile Systems, RTO-EN-018, June 2001

[2] http://www.globalsecurity.org/military/systems/munitions/aim-7.htm

[3] http://www.raf.mod.uk/equipment/skyflash.cfm

[4] http://www.fas.org/man/dod-101/sys/missile/aim-54.htm

[5] http://www.globalsecurity.org/military/systems/munitions/aim-120.htm

[6] http://www.fas.org/man/dod-101/sys/missile/row/exocet.htm

[7] http://www.fas.org/man/dod-101/sys/smart/agm-84.htm

[8] http://www.fas.org/man/dod-101/sys/missile/row/as-20.htm

[9] http://www.army-technology.com/projects/brimstone/

[10] http://www.dtic.mil/ndia/2006psa_apr/mulholland.pdf

[11] http://www.fas.org/man/dod-101/sys/missile/agm-114.htm

[12] http://www.janes.com

[13] http://www.boeing.com/defense-space/missiles/jdam/index.htm

[14] http://guidebook.dcma.mil/38/dpas/12DavisPres.pdf

[15] http://www.boeing.com/news/releases/2005/q3/nr_050725m.html

[16] Pittman, D.N., Roberts, C.E., 'Determining GPS anti-jamming performance on tactical missiles', IEEE Position, Location and Navigation Symposium, ISBN: 0-7803-1435-2, pp641-648, April 1994

[17] Gebhard, L.A., *Evolution of naval radio-electronics and contributions of the Naval Research Laboratory*, NRL Report 8300, pp170, Naval Research Laboratory, Washington DC

[18] Papovic, V., Horvat, R. and Nikolic, N., 'Nikola Tesla – Lectures, Patents, Articles', Nikola Tesla Museum, Beograd, pp109-152, 1956

[19] Swords, S.S., *Technical history of the beginnings of RADAR*, Peter Peregrinus Ltd., ISBN: 086341043X, 1986

[20] http://www.cdvandt.org/huelsmeyer_rx.htm

[21] Marconi, G., 'Radio Telegraphy', Proc. Inst. Radio Engrs, pp215-238, 1922

[22] Lovell B., *Echoes of War*, Taylor & Francies, ISBN: 0852743173, 1991

[23] Oliver, C., Quegan, S., *Understanding Synthetic Aperture Radar Images*, Artech House, 1998

[24] Stimson, G.W., Introduction to Airborne Radar, SciTech Publishing Ltd, ISBN: 1-891121-01-4, 1988

[25] http://cartome.org/min-rez.htm, 'Air Standard 80/16 Mimimum Resolved Object Sizes for Imagery Interpretation', Air Standardization Coordinating Committee, March 1992

[26] Novak, L.M., 'A comparison of 1D and 2D algorithms for radar target classification', IEEE International Conference on Systems Engineering 1991, 10.1109/ICSYSE.1991.161069, pp6-12, 1991

[27] Currie, N.C., Brown, C.E., *Principles and Applications of Millimeter-Wave Radar*, Artech House, ISBN: 0-89006-202-1, 1987

[28] Dudgeon, E., Lacoss, T., 'An Overview of Automatic Target Recognition', Lincoln Laboratory Journal, Volume 6, Number 1, 1993

[29] Hansen, V.G., 'Constant false alarm processing in search radars', Int. Conf. on Radar – present and future, IEE Publ. No. 105, pp325–332, October 1973

[30] Kass, M., Witkin, A., Terzopoulos D., 'Snake: Active Countour Models', International Journal of Computer Vision, Vol 1, pp 321-331, 1988

[31] Klein, L.A., *Milimeter-Wave and Infrared Multisensor Design and Signal Processing*, Artech House, ISBN: 0-89006-764-3, 1997

[32] Pearson, K., 'On Lines and Planes of Closest Fit to Systems of Points in Space', Philosophical Magazine 2, pp559–572, 1901

[33] Fisher, R.A., 'The Use of Multiple Measurements in Taxonomic Problems', Annals of Eugenics, v.7, pp179-188, 1936

[34] http://demonstrations.wolfram.com/PatternRecognitionPrimerII/

[35] Novak, L.M., 'State-of-the-art of SAR automatic target recognition', IEEE International Radar Conference, ISBN: 0-7803-5776-0, pp836-843, 2000

[36] Bell, J., Petillot, Y., Reed, S., 'Automated approach to classification of mine-like objects in sidescan sonar using highlight and shadow information', IEE Proceedings – Radar, Sonar and Navigation, ISSN 1350-2395, pp48-56, Fenruary 2004

[37] Jahangir, M., Blacknell, D., Moate, C., Hill, R.D., 'Extracting information from shadows in SAR imagery', International Conference on Machine Vision 2007, 978-1-4244-1624-0, pp107-112, December 2007

[38] Jahangir, M., Rollason, M.P., 'Detecting moving targets in multiframe SAR imagery without using direct backscatter', IMA Mathematics in Defence 2009, Farnborough, UK, November 2009

[39] Sparr, T., Hansen, R.E., Callow, H.J., Groen, J., 'Enhancing target shadows in SAR images''', Electronics Letters, Vol. 43 No. 5, March 2007

[40] Stewart, N.A., 'Use of crosspolar returns to enhance target detectability', IEE Proc., Vol. 129, Pt. F, No. 2, April 1982

[41] Giuli, D., 'Polarization Diversity in Radars', Proceedings of the IEEE, Volume 74, Issue 2, 1986

[42] Peebles, P.Z., 'Radar Rain Clutter Cancellation Bounds Using Circular Polarisation', IEEE, International Radar Conference, pp210-214, 1975

[43] Gillespie, A.R., 'What does the law mean for autonomous air systems' technology', Dstl/CP38381, July 2010

[44] http://www.raf.mod.uk/equipment/brimstone.cfm

[45] Salmond, D.J., Rollason, M.P., Gregory, I.N., 'Target selection with communicating observers' Information Fusion 2003, Proceedings of the Sixth International Conference, Volume 1, pp665-671, 2003

[46] Krason, H., Randig, G., 'Terrain backscattering characteristics at low grazing angles for X- and S-band', Proceedings of the IEE, Volume 54, Issue 12, pp1964-1965, 1966

[47] Skolnik, M.I., *Introduction to Radar Systems*, McGraw-Hill, ISBN: 0-07-057909-1, 1981

[48] Ward, K.D., 'Compound representation of high resolution sea clutter', Electronics Letters, Volume 17 Issue 22, July 1981

[49] Blake, A.P., Blacknell, D., Oliver C.J., 'SAR clutter analysis and its resolution dependence', Radar 97 (Conf. Publ. No. 449), pp124-128, October 1997

[50] May, A.W., 'Cross Range Resolution of a DBS Radar', QinetiQ Internal Technical Note, March 2004

[51] http://www.mbda.co.uk

[52] Novak, L.M., Halverson, S.D., Owirka, G., Hiett, M., 'Effects of polarization and resolution on SAR ATR', IEEE Transactions on Aerospace and Electronic Systems, Volume 33, Issue 1, ISSN: 0018-9251, pp102-116, January 1997

[53] Chesnaud, C., Refregier, P., Vlady, B., 'Statistical Region Snake-Based Segmentation Adapted to Different Physical Noise Models', IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol 21, No. 11, November 1999

[54] Moate, C.P., Denton, J., 'SAR image delineation of multiple targets in close proximity', EUSAR 2006, 6th European Conference on Synthetic Aperture Radar, May 2006

[55] Kirscht, M., 'Detection and imaging of arbitrarily moving targets with single-channel SAR', IEE Proceedings – Radar, Sonar and Navigation, Vol 150, Issue 1, ISSN 1350-2395, April 2003

[56] Jao, J.K., 'Theory of synthetic aperture radar imaging of a moving target', IEEE Transactions on Geoscience and Remote Sensing, Vol 39, Issue 9, ISSN 0196-2892, August 2002

[57] Reddy, B.S., Chatterji, B.N., 'An FFT-based technique for translation, rotation, and scale-invariant image registration', IEEE Transactions on Image Processing, Volume 5, Issue 8, ISSN: 1057-7149, pp1266-1271, August 1996

[58] Lewis, J.P., 'Fast Normalized Cross-Correlation', Vision Interface 95, Canadian Image Processing and Pattern Recognition Society, pp120-123, May 1995

[59] Richard, V.W., Kammerer, J.E., Wallace, H.B., 'Rain Backscatter Measurements at Millimeter Wavelengths', IEE Transactions on Geoscience and Remote Sensing, Vol 26, No. 3, May 1988

[60] Ulaby, F.T., Elachi, C., *Radar Polarimetry for Geoscience Applications*, Artech House, ISBN: 0-89006-406-7, p11, 1990

[61] Schneider, A.B., Williams, P.D.L., 'Circular Polarization in radars – An assessment of rain clutter reduction and likely loss of performance', Radio and Electronics Engineer, vol. 47, pp11-29, 1977

[62] Ruegg, M., Hagelen, M., Meier, E., Neusch, D., 'Moving target indication with dual frequency millimieter wave SAR', 2006 IEEE Conference on Radar, ISBN: 0-7803-9496-8, May 2006

[63] Lombardo, P., Sciotti, M., Kaplan, L.M., 'SAR prescreening using both target and shadow information', Proceedings of the IEEE Radar Conference 2001, ISBN: 0-7803-6707-3, pp147-152, May 2001

[64] Jingjing Cui, Gudnason, J., Brookes, M., 'Automatic recognition of MSTAR targets using radar shadow and superresolution features', IEEE International Conference on Acoustics, Speech and Signal Processing 2005, ISBN: 0-7803-8874-7, Volume 5, pp589-592, March 2005

[65] Kingsley, S., Quegan, S., *Understanding Radar Systems*, McGraw-Hill, ISBN: 0-07-707426-2, p54, 1992

[66] Barton, D.K., *Modern Radar System Analysis*, Artech House, ISBN: 0-89006-170-X, p402, 1988

[67] Hill, R.D., Moate, C.P., Blacknell, D., 'Estimating building dimensions from synthetic aperture radar image sequences', IET Radar, Sonar and Navigation, Vol. 2, No. 3, pp189-199, 2008

[68] Ulaby, F.T., Dobson, M.C., *Handbook of Radar Scattering Statistics for Terrain*, Artech House, ISBN: 0-89006-336-2, 1989

[69] IEEE, 'IEEE Standard Test Procedures for Antennas', IEEE Std 149, ISBN: 0-471-08032-2, p77, 1979

[70] Schimpf, H., Essen, H., Boehmsdorff, S., Brehm, T., 'MEMPHIS – a Fully Polarimetric Experimental Radar', Geoscience and Remote Sensing Symposium 2002. IGARSS '02, IEEE International, 2002

[71] Stove, A.G., 'Linear FMCW radar techniques', Radar and Signal Processing, IEE Proceedings F, ISSN 0956-375X, pp343-350, October 1992

[72] Griffiths, H.D., 'New ideas in FM radar', Electronics and Communication Engineering Journal, Volume 2, Issue 5, ISSN 0954-0695, pp185-194, October 1990

[73] Griffiths, H.D., Bradford W.J., 'Digital generation of high time-bandwidth product linear FM waveforms for radar altimeters', Radar and Signal Processing, IEE Proceedings F, Volume 139, Issue 2, ISSN: 0956-375X, pp160-169, April 1992

[74] Griffiths, H.D., 'The effect of phase and amplitude errors in FM radar', IEE Colloquium on High Time-Bandwidth Product Waveforms in Radar and Sonar, INSPEC: 3937524, pp9/1-9/5, May 1991

[75] Middleton, R.J.C, Robertson, D.A., 'Predicting range point response from chirp non-linearity', International Conference on Radar 2008, ISBN: 10366726, pp128-132, September 2008

[76] Kroupa, V.F., Cizek, V., Stursa, J., Svandova, H., 'Spurious signals in direct digital frequency synthesizers due to the phase truncation', IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control, Volume 47, Issue 5, ISSN: 0885-3010, September 2000

[77] Skolnik, M.I., *Radar Handbook*, McGraw-Hill, ISBN: 0-07-057913-X, 1990

[78] Beasley, P.D.L., Binns, G. Hodges, R.D., Badley, R.J., 'Tarsier®, a millimetre wave radar for airport runway debris detection', Radar Conference 2004, EURAD. pp261-264, 2004

[79] Zink, M., Ottl, H., Freeman, A., 'Design of a Monopulse SAR system for the determination of elevation angles', IEEE Transactions on Geoscience and Remote Sensing, ISSN: 0196-2892, pp1198-1201, August 2002

[80] Berger, H., 'On the Optimum Squint Angles of Amplitude Monopulse Radar and Beacon Tracking Systems', IEEE Transactions on Aerospace and Electronic Systems, Volume AES-8, Issue:4, July 1972

[81] Beard, G.S., 'Multi-Spectral Target Detection Fusion', London Communications Symposium 2005, http://www.ee.ucl.ac.uk/lens/previous/LCS2005/lcs2005/68.pdf, September 2010

A   Published papers

A.1   **Multi-Spectral Target Detection Fusion**

G S Beard [81]

QinetiQ – University College London

Abstract: Scenarios for modern air to surface missile seekers require a high probability of target detection with a minimal false alarm rate to avoid collateral damage. Active and passive countermeasures are increasingly expected to appear on enemy targets, designed to counteract a seeker in a narrow or wide band of frequencies. A technique to improve on the performance of a single sensor is the use of two separate sensors such as radar and infrared. The detection and false-alarm performance of each individual sensor on data of military targets is examined, and compared to that achieved through the use of detection level fusion.

1 Introduction

The expected model of battlefield engagements has moved from the attack of massed armour on the battlefield to one comprising small-scale engagements of small groups of targets in urban areas. The emphasis is focused on the attack of smaller groups of targets while minimising collateral damage. Defensive aid suites in the form of active and passive countermeasures such as radar netting, infrared (IR) sheeting and decoy targets are increasingly expected to appear. These can be designed to counteract a seeker in one waveband or over a range of frequencies. The operational use of the seeker sensor for this work is envisaged as being a Terminally Guided Sub-Munition (TGSM) at the end of a larger chain of targeting sensors.

An important question for future seekers is whether a single-band sensor can deliver the desired level of performance against countermeasured threats, or if sensors operating in multiple wavebands are required. Taking advantage of the different physical nature of microwave and optical sensor systems, the fusion of the data can result in an enhancement in system performance.

In this paper the probability of target detection and rate of false alarm are calculated for both radar and IR sensors. Detection level fusion is then applied to the sensor imagery and the effects assessed.

Section 2 describes the data available and how the imagery was formed. Section 3 covers the radar detection processing and performance. Section 4 describes the IR processing and performance. Section 5 introduces the fusion options and the benefits these bring. Conclusions are drawn in section 6.

2. Data gathering and Image formation

Data for this work was gathered using a seeker sensor in both the radar and IR domains. The seeker consists of a real-beam radar and an uncooled imaging infrared sensor. A range of military vehicles, some with camouflage deployed, were imaged in a rural environment by the seeker mounted on a helicopter. The weather was clear and sunny resulting in a high solar loading of the scene with consequent high contrast in the IR images. Most of the vehicles were deliberately heated by running their engines prior to the start of measurement to further aid the IR sensor.

The radar sensor is typical of the type that could be utilised in low-cost TGSM roles. The down-range resolution of the radar is limited by bandwidth, and is far better than the cross-range resolution which is limited by the size of the antenna. Hence the sensor can resolve fine detail down-range but only coarse structure in cross-range. The radar data consists of a series of range-profiles for each swath (side-to-side scan) recorded with information on where the antenna was pointing. From this data, a ground map image can be formed for each swath as shown in Figure 0-1.

The IR sensor is co-located with the RF sensor; as the radar antenna scans across the field of view, the IR sensor looks in the same direction. Each IR swath comprises the data gathered by the middle column of a 64 by 64 pixel detector array, and so is comprised of a sequence of vertical strips 64 pixels high. By placing these strips side-by-side as the sensor scans, a 2D image of the scene is formed.
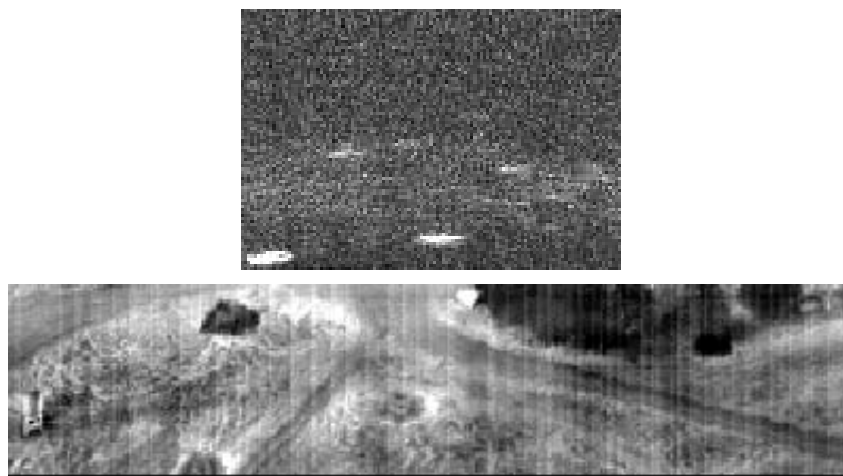

Figure 0-1 Radar (top) and IR (below) imagery of the same scene

3. Radar Processing

The radar target detection processing operates in a number of stages before a target is declared.

An adaptive threshold is used to detect pixels with intensities brighter than their local background. Every pixel in the image is processed sequentially. The pixel under test is surrounded by a square of pixels which are ignored – the guard area. This prevents returns from the target affecting an outer-ring of pixels – the statistics ring. The size of the guard is set such that the distance from the test pixel to the statistics ring exceeds the expected target dimensions. In order to be declared a possible target, the pixel value needs to exceed a threshold calculated from the statistics ring. The approach used was to calculate an ordered statistic threshold. This procedure orders the pixels in the statistics ring by brightness then sets the threshold such that k% of pixels are below this threshold. This procedure does not assume a particular clutter model so has robust performance across boundaries, but is computationally expensive.

The output from the thresholding is an image with scattered pixels that have exceeded the ordered statistic threshold. From this, little information is available about regions of interest and some clustering of pixels is needed. A technique known as snake delineation [2] is used. This process finds the region of pixels that bound a potential target after it has been initially detected. The snake consists of point nodes which are joined together to form a polygon that is initially seeded around a detection. The algorithm moves each node by a small amount and tests to find whether the new snake is a better delineation. If the delineation is better the node is left where it is, otherwise the node is moved back. The program keeps trying to move all the nodes one by one in random directions. When no improvement is found the delineation is complete and an outline is formed as shown in Figure 0-2.
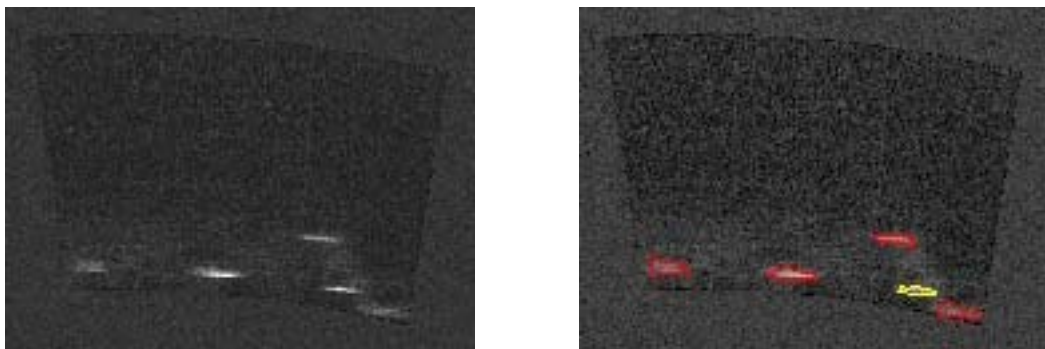
Figure 0-2 Radar image (left) and with snake segmentation (right)

The snake segmentation can be seen to produce an outline that matches a human visual interpretation of the object's extent in this example scene. Each segmented object has a number of features calculated. These are spatial features such as length and width as well as statistical features such as target to clutter ratio and fractal dimension. A number of data sets were examined to provide 'training data'. These allowed selection of certain features, and for constraints to be applied on the range of values within which a target should fall. These features were fairly successful at removing false alarms caused by clutter, and artefacts due to saturation in the radar imagery. The yellow outline of the second target from the

right in Figure 0-2 indicates an object that has been labelled as clutter by the feature discriminator.

Overall PD	0.95
Open PD	1.00
Treeline PD	0.67
FA/km^2	14.55

Table 0-1 Pd and Pfa Even channel

The detection results shown in Table 0-1 are split into three categories, detection of all targets, detection of those targets in the open, and detection of targets in the treeline. The rates of false alarm shown are normalised for a square kilometre, an area approximately 25 times larger than the sample search area.

4. IR Processing

The IR detection algorithm is designed to identify objects that are unusually bright or dark relative to the background and whose sizes and shapes are consistent with those expected from military targets. Statistical models are used to characterise the background and unusual objects are then identified as statistical outliers. The process follows similar stages to the radar technique, although using slightly different methods at each. Pixels with unusually strong edges, i.e. a steep gradient with respect to the whole image are found, clusters of pixels are formed and features are extracted from these clusters.

Figure 0-3 shows an IR swath with the detections enclosed by an ellipse. Any pixels within the ellipse are considered to be target returns for the purposes of calculating statistical measures.

Figure 0-3 IR image with detections highlighted

Overall PD	0.97
Open PD	0.96
Treeline PD	1.00
FA/km^2	99.86

Table 0-2 Pd and Pfa IR sensor

IR detections and false alarms were counted in the same manner as for the radar sensor. The IR had fewer problems detecting targets in the treeline but was more susceptible to the effects of camouflage. Targets with IR sheeting were detected via patches of heat escaping from the edges, and should such a patch be selected as the target location the seeker could well miss the vehicle. Should the target have been stationary without the engines running, IR detection of the targets would be

very difficult. The IR also suffered from a significantly higher false alarm rate than the radar sensor, as shown in Table 0-2.

While the IR sensor has a high probability of detection for the targets, it also has a large number of false alarms present. These are caused by man-made objects in the scene, but also by trees and 'bright' patches of ground.

5. Detection level fusion

Fusion using the two sensors aims at the integration of disparate and complementary data to enhance the information present in the images. This should lead to more accurate data and increased utility.

Fusion is a possibility at three levels: pixel, feature and detection. Pixel level fusion involves the data from each sensor being mapped onto an image with the same spatial dimensions such that each pixel in the new image represents a combination of the two sensors at that point. For sensors that have dissimilar geometries such as RF and IR this combination is a difficult problem. In feature level fusion, features are calculated from regions of interest and are combined, the aim being that the combined features produce better performance in a detection algorithm than the features from each sensor independently. This scheme would be practical using the data gathered, however for this work detection level fusion was employed since this is the simplest fusion scheme to implement. Detection algorithms are performed separately on data from each sensor to produce regions of interest and the locations of the regions are then combined. This gives an ability to reinforce valid target decisions and reject false alarms.
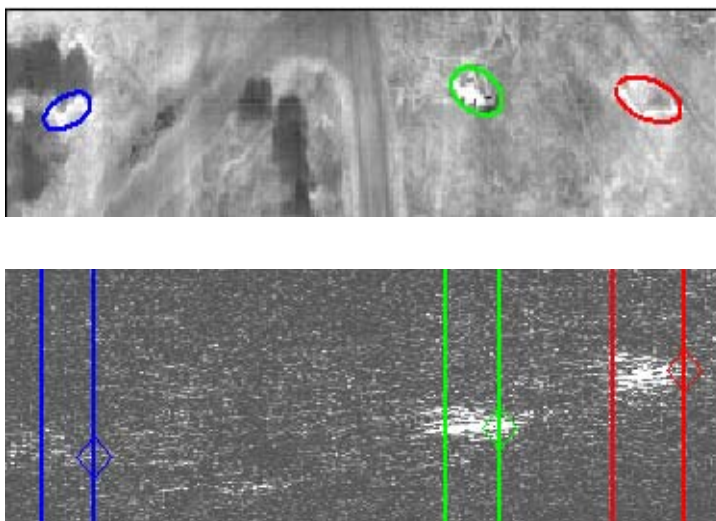
Figure 0-4 IR detections with azimuth mapped to RF image

Due to the co-location of the sensors and the antenna azimuth information available it is a simple process to correctly register the direction of detections from one sensor to the other. Figure 0-4 shows an IR swath with detected areas, and the azimuth extent of the detections mapped onto the RF image. While the azimuth maps across accurately, the different geometries of the sensors - azimuth and elevation angle for the IR and azimuth and range for the RF - makes the mapping in the vertical direction less trivial. For the results shown here the fusion has

effectively been performed manually, however for future work, the automatic registration of the detections should be possible.

Overall PD	0.94
Open PD	0.96
Treeline PD	0.80
FA/km2	12.02

Table 0-3 Pd and Pfa – AND Fusion

Each sensor's detection algorithms are currently configured for single-sensor detection performance which may not be the optimal setting for the best performance when their results are combined. Table 0-3 shows the effects of applying 'AND' detection fusion to the outputs of both sensors. Corresponding IR results were available for 4 of the 5 radar scenes and hence the fused treeline PD is constrained by a radar PD of 0.8 instead of 0.67.

None of the tree-based false alarms from the IR sensor are passed through since the tree areas do not give strong returns in the radar domain. The only source of false alarms that passed through the fusion process were man-made metallic objects in the scene.

6. Conclusions

Applying detection level fusion to the sensor system resulted in improved performance over that from each single sensor based on the reduction in false alarm rate. There was a slight drop in detection probability, however the reduction in the false alarm rate compensated for this.

In order for fusion to produce the best results it would be ideal for each individual sensor's detection scheme to be optimised for a fusion role. This could involve raising the detection probability slightly for each sensor, relying on clutter objects not producing a false alarm in both sensors.

The system would also offer flexibility for different scenarios. Scenes where targets are expected to have IR camouflage or in misty/rainy weather conditions could place emphasis on the radar sensor. Scenes where targets are expected to be operational and producing heat would favour the IR sensor.

References

[1] G.T. Maskall and S.K. Watson, "DMS Multi-Spectral Signatures and Data Fusion Options". QinetiQ/S&E/APC/CR032156. 24 Nov 2003. Unpublished QinetiQ report

[2] D. Pedlar, "Target Delineation and Classification using a Region-Based Active Contour and a support Machine classifier on SAR Imagery". EUSAR Conference, Proceedings Volume 1, Ulm, May 2004

Blank page