SHIP DETECTION METHODS FOR MARITIME DOMAIN AWARENESS USING SAR SATELLITE DATA

by

Colin Peter Schwegmann

Submitted in partial fulfillment of the requirements for the degree

Master of Engineering (Computer Engineering)

in the

Department of Electrical, Electronic and Computer Engineering
Faculty of Engineering, Built Environment and Information Technology
UNIVERSITY OF PRETORIA

April 2014

SUMMARY

SHIP DETECTION METHODS FOR MARITIME DOMAIN AWARENESS USING SAR SATELLITE DATA

by

Colin Peter Schwegmann

Supervisor(s):	Dr. W. Kleynhans
Co-supervisor:	Dr. B. P. Salmon
Department:	Electrical, Electronic and Computer Engineering
University:	University of Pretoria
Degree:	Master of Engineering (Computer Engineering)
Keywords:	Maritime Domain Awareness, Illegal, Unreported and Unregulated fishing, Synthetic Aperture Radar, Constant False Alarm Rate, Mean-shift.

Maritime Domain Awareness is the understanding of all aspects relating to the maritime domain that may have an effect on the security, economy or environment of a country bordering the sea. Large ships in South Africa have historically been monitored using ship-based transponder systems such as Automatic Identification System. These systems transmit geographical ship coordinates which allow operators to track ships. The greatest disadvantage of monitoring ships in this manner is that the transponders need to be installed and switched on in order to track ships.

The monitoring of ships can be done in another manner by taking advantage of Synthetic Aperture Radar imagery which allows for the monitoring of large portions of the Earth. This imagery is generated using an active sensor which makes use of radar pulses to observe areas under any weather condition, day or night. In this dissertation, various ship detection systems configurations were tested using Synthetic Aperture Radar imagery and experimental conclusions were drawn for each configuration. The system was tested against simulated Synthetic Aperture Radar imagery and actual Synthetic Aperture Radar imagery located within the South African Exclusive Economic Zone. Tests were performed in order to evaluate

which noise distribution model best models the Synthetic Aperture Radar imagery used in this study over South African coastal waters. The parameters and the effects of varying them within each system configuration were evaluated.

Experimental results found that the K-distribution is the noise distribution model that best describes the noise of sea-water found within the Synthetic Aperture Radar imagery used in this study. Furthermore, it was determined that ship detection system configurations using the Constant False Alarm Rate prescreening method had the highest average detection accuracies and lowest false alarm rates amongst all of the configurations tested. The various detection system configurations had similar detection accuracies at low thresholds but varied significantly in terms of the number of false alarms.

OPSOMMING

SKIPOPSPORINGSMETODES VIR MARITIEME GEBIEDBEWUSTHEID DEUR DIE GEBRUIK VAN SINTETIESE GLEUFANTENNE-RADARWAARNEMINGSATELLIETDATA

deur

Colin Peter Schwegmann

Studieleier(s):	Dr. W. Kleynhans
Medestudieleier:	Dr. B. P. Salmon
Departement:	Elektriese, Elektroniese en Rekenaar-Ingenieurswese
Universiteit:	Universiteit van Pretoria
Graad:	Magister in Ingenieurswese (Rekenaaringenieurswese)
Sleutelwoorde:	Maritieme Gebiedbewustheid, Onwettig, Ongeraporteerde en Ongereguleerde Visvang, Sintetiese Gleufantenne-radar, Konstantevalsalarm-koers, Gemiddelde Verskuiwing.

Maritieme gebiedbewustheid is die begrip van alle aspekte met betrekking tot die maritieme gebied wat moontlik 'n uitwerking op die sekuriteit, ekonomie of omgewing van 'n kusland kan hê. Maritieme gebiedbewustheid omvat die bewaking en begrip van 'n land se eie eksklusiewe ekonomiese sone. Groot skepe in Suid-Afrika is voorheen gemoniteer deur die gebruik van skipgebaseerde transpondeerdersisteme, soos die outomatiese identifiseringsisteem. Hierdie sisteme versend geografiese skipkoördinate wat operateurs in staat stel om skepe op te spoor. Die grootste nadeel van die nasporing van skepe op dié manier is dat die transpondeerders geïnstalleer en aangeskakel moet word sodat die skepe gemonitor kan word.

Die nasporing van skepe kan op 'n ander manier gedoen word, naamlik deur die gebruik van sintetiese gleufantenne-radarbeelde wat die regulering van groot oppervlakke van die aarde toelaat. Hierdie radarbeelde word opgewek deur die gebruik van 'n aktiewe sensor wat gevorderde radarbeginsels aanwend om areas onder enige weertoestande, dag of nag, waar te neem. In hierdie verhandeling is verskeie skipopsporingsisteemkonfigurasies getoets deur die gebruik van sintetiese gleufantenne-radarbeelde en eksperimentele gevolgtrekkings is

gemaak vir elke konfigurasie. Die sisteme is getoets teen gesimuleerde sintetiese gleufantenne-radarbeelde asook werklike sintetiese gleufantenne-radarbeelde, wat in die Suid-Afrikaanse eksklusiewe ekonomiese sone geleë is. Toetse is uitgevoer om te evalueer watter steuringsverspreidingsmodel die gevorderde sintetiese gleufantenne-radarbeelde in Suid-Afrikaanse kuswaters die beste modelleer. Die parameters en die gevolge van die afwisseling binne elke sisteemkonfigurasie is geëvalueer.

Eksperimentele resultate het bevind dat die K-verdeling die steuringsverspreidingsmodel is wat die steuring van seewater wat binne die sintetiese gleufantenne-radarbeelde in Suid-Afrika se eksklusiewe ekonomiese sone voorkom, die beste beskryf. Verder is bepaal dat skipopsporingstelselkonfigurasies wat die voorafsiftingsmetode van konstantevalsalarmkoerse bevat, die hoogste gemiddelde opsporingsakkuraatheid en laagste valsalarmkoerse het van al die konfigurasies wat getoets is. Die verskeie opsporingsisteemkonfigurasies het soortgelyke opsporingsakkuraathede by lae drempels getoon, maar het aansienlik gevarieer ten opsigte van die aantal vals alarms.

LIST OF ABBREVIATIONS

AIS	Automatic Identification System
APR	Automatic Position Reporting
ASAR	Advanced Synthetic Aperture Radar
ASM	Application Specific Message
ASP	Application Service Provider
CA-CFAR	Cell-Averaged Constant False Alarm Rate
CDC	Cooperative Data Centre
CFAR	Constant False Alarm Rate
CSP	Content Service Provider
DA	Detection Accuracy
DDP	Data Distribution Plan
EEZ	Exclusive Economic Zone
EM	Electro-Magnetic
ESA	European Space Agency
FAR	False Alarm Rate
GPS	Global Positioning System
IDE	International Data Exchange
IMO	International Maritime Organization
IUU	Illegal, Unreported and Unregulated (fishing)
LRIT	Long Range Identification and Tracking
MDA	Maritime Domain Awareness
NDC	National Data Centre
PRF	Pulse Repetition Frequency
RCS	Radar Cross Section
RDC	Regional Data Centre
RF	Radio Frequency
SAR	Synthetic Aperture Radar
Sat-AIS	Satellite Automatic Identification System
SUMO	Search for Unidentified Maritime Objects
TDMA	Time Division Multiple Access

UHF	Ultra High Frequency
VMS	Vessel Monitoring System
VHF	Very High Frequency
WSM	Wide Swath Mode

TABLE OF CONTENTS

CHAPTER 1

INTRODUCTION

1.1 PROBLEM STATEMENT

1.1.1 Context of the problem

Maritime Domain Awareness (MDA) is a term used to describe environmental, economic and security aspects related to the maritime domain. One of the main focus points of MDA is the observation of illegal, unreported and unregulated (IUU) fishing and employs methods to prevent it. It has been estimated that worldwide losses due to IUU fishing are approximately R235 billion annually [1]. While IUU fishing is a global issue, it has been found that developing countries are most at risk [1]. It is estimated that South Africa loses approximately R1 billion annually due to IUU and that some western African countries have reported losing up to 40% of their legal catches due to IUU fishing [1, 2].

Parties concerned with their own MDA typically employ a wide variety of techniques to monitor aspects of the maritime domain. For instance, sometimes land crews are employed to monitor specific maritime areas. Another common method used to improve MDA is that of airborne campaigns. Airborne campaigns use aerial-based surveying to monitor a region to describe a number of maritime related changes in ship traffic, environmental effects and for security reasons. These campaigns monitor a region using techniques such as optical imagery, radar and even human-based surveillance.

The installation of ship-based transponders, such as Automatic Identification System (AIS), has become law for many ships. This allows countries to improve their MDA as transpon-

der systems allow for tracking of ships at sea. Transponders are devices that transmit the ship's geographical coordinates at specific time intervals. IUU fishing can be monitored by correlating actual on-the-sea observations with the reported (or lack thereof) transponder data.

An alternative method of monitoring ships at sea is to use Synthetic Aperture Radar (SAR) satellites. These satellites generate SAR imagery which can be processed to allow for observation of large tracts of the Earth in almost any weather condition [3]. One of the benefits of using SAR satellites to observe areas is that by using active sensors, areas and specifically ships can be observed without the any active participation of the ship. This allows for SAR imagery to highlight "dark" ships which represent ships that have no associated transponder data [4]. Thus, SAR imagery and transponder data can be used in conjunction to obtain a better overall picture of an area for the purpose of improved MDA.

SAR imagery can be used to detected ships and to do so it is processed using a wide variety of methods. Each of these methods have various parameters that may differ depending on the region of observation, the SAR imagery composition and the physical properties of the ships within the area [4]. It is for this purpose that a thorough investigation of a number of aspects relating to the detection of ships within SAR imagery around South Africa's oceans is necessary.

1.1.2 Research motivation

A significant gap in the usage of SAR image ship detection techniques for South African waters has been identified. The application of ship detection methods using SAR images has become a feasible endeavour within the South African context. Specifically, a flexible system that can observe ocean waters using SAR imagery has, to the author's knowledge, not been researched and implemented in South Africa. Related to this, the various parameters and system nuances of the ship detection systems processing SAR imagery are not well understood for the specific context of South African coastal waters.

1.2 RESEARCH OBJECTIVE AND QUESTIONS

The objective of the research in this dissertation is to identify, evaluate and suggest possible optimal ship detection methods for SAR satellite imagery from ENVISAT's Advanced SAR (ASAR) sensor off of the coast of South Africa. Furthermore, the study aims to investigate aspects related to the detection of ships at seas such as sea-noise modelling and expected ship sizes. By focusing on a single region, the study aims to investigate these aspects for a South African context with the hopes of improving South Africa's MDA. The fundamental research questions this dissertation aims to address are:

- Can conventional sea-noise models be used to model South African waters in ASAR images?
- Within each step of the ship detection process, does a single method exist which significantly improves the effectiveness of the detection of ships for ASAR imagery focused on South Africa?
- Does a combination of ship detection methods exist that will provide an optimal ship detection effectiveness compared to any other combination for ASAR imagery focused on South Africa?

1.3 HYPOTHESIS AND APPROACH

The hypothesis for ASAR imagery acquired off of the South African coast is: *It is hypothesised that sea-noise within ASAR imagery can be modelled using the K-distribution. ASAR images that have the a) assumed K-distribution and b) are processed using CFAR-based system configurations will provide higher levels of ship detection accuracy and robustness. It is further hypothesised that there exists a ship detection system configuration that will outperform all other configurations.*

To evaluate the hypothesis a thorough investigation is needed. First, an understanding of the fundamental concepts and methods behind ship detection needs to be presented. From this a selection is made and the characteristics and mechanics of these methods are explored in be detail. Thereafter, an extensive set of experiments will be conducted in order to thoroughly

test each of the chosen methods. These experimental results will either support or reject the hypothesis postulated in this dissertation.

1.4 RESEARCH GOALS

The main research goal of this study is to test and compare various ship detection methods applied to ASAR images within the South African region. The study will investigate if there is an optimal collection of methods that can be used to detect ships at sea. Another research goal is to obtain a deeper understanding of the the various system parameters and assumptions commonly used in ship detection methods and how these specifically relate to methods applied to synthetic imagery as well as geographically limited SAR imagery.

1.5 RESEARCH CONTRIBUTION

This dissertation aims to contribute to the current body of knowledge in a number of ways. Primarily, the study aims to investigate the most effective manner of ship detection in coarse resolution ASAR satellite imagery for South Africa's coastal waters. It also aims to determine if assumptions such as sea-noise models used in other studies can be applied to sea-noise in SAR imagery of the coast of South Africa. The dissertation provides in-depth detail to better understand of how various aspects of the entire ship detection process may be used to improve the overall effectiveness of ship detection as well as improve the understanding of the parameters of the system. By providing a deeper understanding on a local scale, the study may emphasise factors about the parameters that can be used to improve ship detection effectiveness on a more general scale.

1.6 OVERVIEW OF STUDY

Five chapters are dedicated to this study and are: Chapter 2 provides a general overview of the important aspects relating to ship detection and concepts such as MDA and Marine Ship Monitoring. Chapter 3 delves into the various selected methods for ship detection and provides a mathematical description of each. This chapter provides a description of how these methods are jointly used to form a ship detection system. Chapter 4 provides a concise description of the experimental design and associated results. Chapter 5 provides a discussion of the results and the observations made from these. Finally, Chapter 6 provides a conclusion

about the study, the results and how they relate to the research questions.

CHAPTER 2

LITERATURE STUDY

2.1 OVERVIEW

The purpose of this chapter is to investigate the fundamental principles upon which the detection of ships at sea rely. The literature study will address the most prominent literature for a number of topics related to the detection of ships. This chapter aims to achieve the following:

1. Give a general overview of how each maritime aspect is related and how each aspect relates to the awareness of a country's own maritime area.

2. Describe the operations and differences of the most prominent transponder based methods for detecting ships.

3. Describe and explain the concept of SAR imagery and how it can be used as the basis for ship detection.

4. Provide the most up to date methods used to detect ships in SAR imagery. This includes a summary of the typical system process that is used for ship detection in SAR as well as the most prevalent methods used in each step of the process.

The rest of the chapter is arranged into four sections each detailing the aspects laid out in these objectives.

2.2 MARITIME DOMAIN AWARENESS

2.2.1 Introduction to Maritime Domain Awareness

MDA is concerned with all the actions and aspects related to the maritime domain. A concise description of a maritime domain is "All areas and things of, on, under, relating to, adjacent to, or bordering on a sea, ocean, or other navigable waterway, including all maritime-related activities, infrastructure, people, cargo, and ships and other conveyances" [5].

In terms of marine safety, MDA means keeping a country aware of any threats that may affect the country's security or natural resources either by natural or human forces. MDA is a global concern and many countries are actively involved in the enhancement of MDA through a number of common goals, some of which include:

- Rapid identification of potential and actual threats to a country's maritime domain.
- Planning and execution of countermeasures to the threats via informed decisions.
- Integration of intergovernmental and other departments to improve the overall awareness of maritime activities.
- Sharing of information between local, national and international partners.

MDA is a major concern for a number of countries to ensure surety of as safe passage to and from a country. In addition, commercial and environmental concerns also need to be addressed and monitored. The purpose of MDA is to seek out collaborating partners such that each country involved can better police its own maritime domain. A single country cannot maintain absolute awareness of maritime activities related to it in both a global and local sense and as such MDA is a global, cooperative effort. Sharing relevant information freely between cooperative parties will improve situational awareness.

2.2.2 Major interest areas of MDA

There are a number of important areas in which MDA is concerned. The three biggest are the Environmental Monitoring, Commercial Ship Monitoring and Security Monitoring.

Environmental Monitoring: The MDA initiative seeks to monitor a wide variety of environmental ocean-based effects such as ocean current and Arctic sea-ice behaviour as well as adverse effects such as oil pollution and refuse disposal. This is done to maintain the well-being of the ocean environment in and around a country.

Commercial Ship Monitoring: Commercial Ship Monitoring is where MDA tries to gather information from various sources to accurately track ships. Having access to this information can provide an cooperative party with many benefits including the tracking of assets on commercial voyages. Due to the fact that many countries rely heavily on their maritime commercial commerce, having up-to-date and reliable information concerning each commercial ship is highly beneficial to a country.

Security Monitoring: Effectively monitoring a country's surrounding Exclusive Economic Zone (EEZ) and beyond is of interest to a country's military. Using multiple sources of information such as SAR satellites, conventional radar and ship transponders, a country can infer intent and take appropriate action to ensure its own security. By collaborating with other allied countries, a network of information can be utilised to further improve each country's maritime domain knowledge.

2.2.3 Exclusive Economic Zone

A country's EEZ is the body of water surrounding a coastal nation for which the coastal nation has jurisdictional rights. An EEZ includes the zone of ocean waters from five to 322 kilometres (three to 200 nautical miles). The coastal country is required to sign the United Nations Convention of the Law of the Sea (UNCLOSE) to establish its territorial and EEZ waters. The country that borders the coast is then able to use the area within the EEZ as an economic resource as well as being able to explore, conserve and manage it.

2.2.4 Key players in the MDA initiative

The United States of America first completed a national security presidential directive named NSPD 41/HSPD 13 in December of 2004 and was used as the basis for the "National Plan to Achieve Maritime Domain Awareness" [5]. In this directive a number of important goals were made in order to protect the USA and corroborate with international partners to further this

goal.

The directive established a "Maritime Security Policy Coordinating Committee" that was tasked with the development and overseeing of a "National Strategy for Maritime Security" as well as additional, supplemental plans to help achieve national maritime domain awareness of the maritime environment in and around the USA. Essentially, plans for plausible threat detection, intelligence gathering/integrations, recovery of commerce/infrastructure were all proposed to help provide improved security of the USA's maritime environment. An important point of the directive was to create an "outreach" program whereby the USA government would solicit support from foreign governments to enhance its own maritime security. By being more aware of their own maritime environments and sharing this information with the USA, foreign governments themselves benefit as their own maritime environment would enjoy increased security. Many of the governments interested in MDA have used this plan as a basis for their own directives to achieve maritime domain awareness within their own region.

The Candian government was another influential government in the shaping of a global MDA effort when it released an update to its national security policy in 2004. The first of its kind for its country, the policy is a comprehensive list of all strategies intended to address threats to Canada [6]. One of the focuses of the policy was an increase in maritime security as well as additional accountability for marine security in and around Canada by establishing Marine Security Operation Centre (MSOC). MSOCs are intended to help Canada improve its responsiveness and awareness relating to maritime security. Much like the USA security policy, the Canadian government recognised that to effectively monitor all domains related to maritime security an improved collaboration between itself and other countries was necessary. Within the document the Canadian government discusses that it will seek to collaborate with other governments to identify and implement new maritime related technologies.

In much the same way other key maritime players have followed suite. Norway, while not actually part of the European Union (EU) heavily influenced the EU's decisions regarding the broad maritime security/awareness policy put in place for the EU and surrounding nations [7]. Other groups such as the Arctic Council have all signed and agreed to similar maritime domain awareness policies.

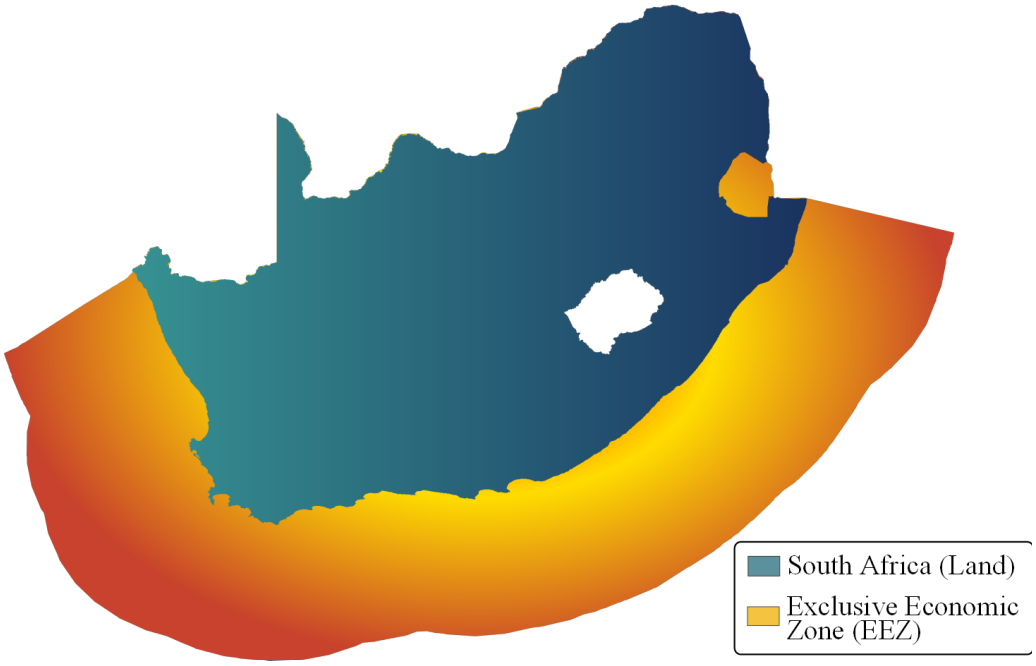

Figure 2.1: South Africa and its EEZ. South Africa covers approximately 1.2 million square kilometers yet its EEZ covers more than 1.5 million square kilometers.

2.2.5 South Africa as an MDA participant

South Africa has a coastline of approximately 3900 km in length and is the 39th longest coastline in the world. Its EEZ occupies an area of approximately 1.5 million square kilometres (Figure 2.1). This means that South Africa's EEZ covers a larger area than South Africa's land (1.2 million square kilometres). The Cape Sea Route is one of the most prominent routes between the East and West. Approximately 30% of European- and Americas-bound oil from the Middle East comes through the Cape Sea Route [8]. Furthermore, South Africa is located at a maritime choke point in that it is surrounded by three oceans - the Indian, South Atlantic and Southern Oceans.

South Africa has a number of well developed ports and related maritime infrastructures. South Africa operates six major ports namely: Cape Town, Durban, East London, Port Elizabeth, Richards Bay and Saldanha Bay which account for an excess of 90% of South Africa's imports and exports (in terms of value) carried by sea. In addition, the commercial fishing industry of South Africa employs in excess of 25000 people and creates a gross income of more than R4.8 billion annually [8].

All of this leads to the conclusion that South Africa is a major maritime nation. It is dependent on the sea for its trade, economic and security well-being. It is to this end that many South African internal parties collaborate to improve its MDA. Parties involved within South Africa include the Department of Transport, South African Maritime Safety Authority (SAMSA), Department of Defence: Maritime Defence/SA Navy and the South African Air Force (SAAF). South Africa also seeks to have local and global partners to extend its own MDA. South Africa is committed to the Standing Maritime Committee which comprises of all the Southern African Development Community (SADC) countries. The Standing Maritime Committee allows for South Africa to participate with regional countries on military maritime activities. This is in addition to other international partners such as India, Brazil and the United Nations.

Another important related point is that due to the positioning of South Africa, there is a denser concentration of polar orbiting satellites and thus the possibility of a greater amount of remote sensing data can be obtained.

2.2.6 MDA Monitoring

2.2.6.1 Traditional MDA Monitoring

MDA is geared to use every available source of information regarding maritime related activities and environments. Traditionally, ships have been monitored with ship monitoring transponder systems such as Automatic Identification System (AIS), Satellite AIS (Sat-AIS), Long Range Identification and Tracking (LRIT) and the Vessel Monitoring System (VMS).

Transponder systems are a cooperative system as the ship has to ensure its transponder is turned on before it can be monitored. If a ship's transponder is sabotaged or damaged then the ship cannot be tracked until the transponder is fixed. This means that the traditional method of ship tracking has a major disadvantage and therefore the monitoring of issues such as IUU has become a complex problem for countries bordering on a sea. To accompany traditional MDA monitoring, a newer MDA monitoring method that uses active sensors within a satellite can be used to improve a country's own MDA.

2.2.6.2 SAR satellites as an MDA monitoring method

SAR satellites use relative movement and radar principles to capture images of areas of the Earth which allow for wide area observation under any weather condition. SAR satellites provides a wealth of benefits as a source of maritime domain related information. Some of the major benefits of using SAR satellites include:

- Large swath widths that allow surveillance of thousands of square kilometres in a single pass of a SAR satellite. The cost to monitor such an area using other, conventional monitoring systems (such as air-borne optical imaging) would be infeasible and as such SAR satellite surveillance is able to monitor larger areas for a reduced cost per square kilometre. To understand how useful SAR is at observing large areas two SAR images with South Africa's coastline superimposed over the images is shown in Figure 2.2. To cover the same length of South African coastline using other surveillance methods would be costly.

- Large areas can be monitored under almost all weather conditions/cloud cover.

- Depending on image resolution, highly reflective objects such as large sea vessels are clearly visible to the naked eye within SAR images.

- The way in which SAR images are captured allow pollutants such as crude-oil and drilling fluids to be easier to identify at sea.

- SAR is a means of observing an area remotely. It therefore allows imaging of areas that would usually be restricted due to either natural (storms) or other reasons (restricted air space).

- The increase in the volume of SAR satellite imagery has lead to a reduction of the time required to monitor an area. This allows authorities with sufficiently advanced SAR monitoring systems to more rapidly identify and respond to illegal activities.

- By using historical SAR imagery, a profile of maritime activities for a region can be built up. This allows a number of benefits including the detection of unusual behaviour by ships.

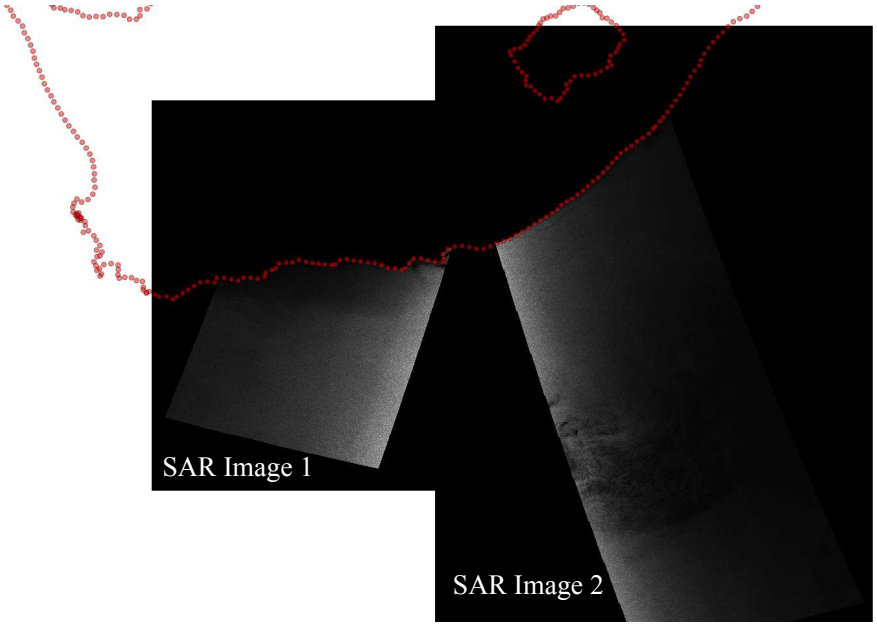

Figure 2.2: Two SAR images with South Africa's coast superimposed in red of the images. Notice that with just two SAR images a large percentage of South Africa's coast can be monitored.

One of the greatest benefits of using SAR satellites to improve MDA is that it is a remote-based surveillance method. SAR does not have the same issue of cooperation as traditional transponder-based ship tracking methods as it does not require the active participation of the ship. This allows SAR to highlight ships which do not have their transponders switched on. In some cases this can lead to identifying ships with malicious intent by identifying those ships that regularly turn off their transponders (Figure 2.3).

By combining SAR imagery with the ship positions garnered from transponder-based ship detection systems, a more complete picture of a country's MDA can be garnered.

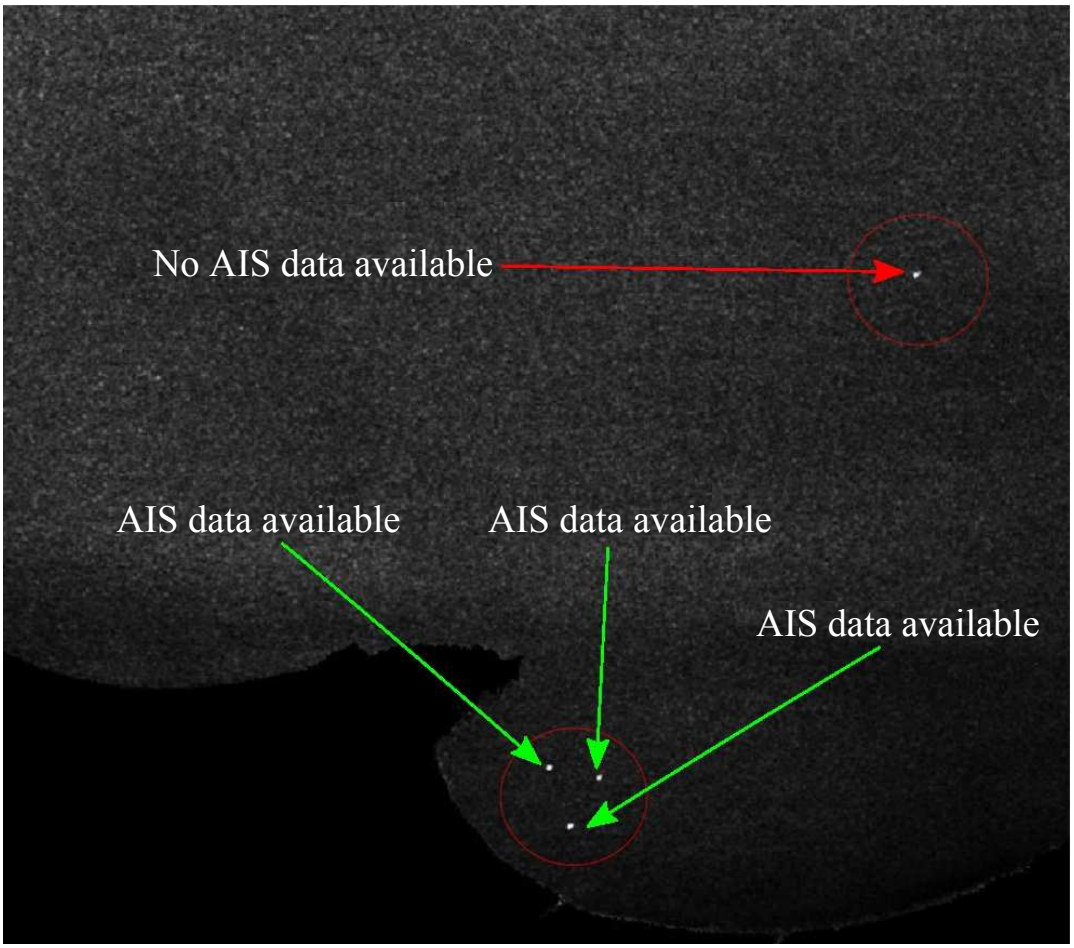

Figure 2.3: Four ships that have been detected using a SAR image off the coast of Cape Town, South Africa. Of the four ships only three of them have corresponding AIS positional data. Using AIS only (instead of SAR imagery) would lead to one of the four ships not being detected.

2.3 MARITIME SHIP MONITORING

Maritime Ship Monitoring is an MDA term that seeks to monitor ships at sea using a transponder-based system. A transponder is a device that is installed on a ship and transmits information such as the ship's position and destination. Most systems send the information directly to a ground station whilst others use relay satellites to improve coverage. In all cases, the transponder based system is known as a cooperative tracking system. This means that to track a ship a transmitting transponder is required. If the transponder is switched off, damaged or sabotaged then the ship cannot be effectively tracked and is known as a "dark"

target.

The following section will detail some of the prominent transponder systems, namely: AIS, LRIT and VMS.

2.3.1 Automatic Identification System

AIS is a ship monitoring system concept that was first introduced in Sweden in 1993 and later developed into the ship tracking system. In 2000, the International Maritime Organisation (IMO) mandated that by January 2002, all ships with a gross tonnage of 300 or more, all cargo ships of 500 gross tonnage and all passenger ships (irrespective of size) must carry an AIS transponder [9]. This was a major move towards more effective ship monitoring on a global scale.

2.3.1.1 How AIS works

AIS is a complex system that broadcasts messages between ships, aids to navigation and base stations [9]. AIS is built on very stringent technology standards which are governed by the IEC [9]. These standards ensure that the messages and performance standards of AIS are guaranteed across multiple devices and ships. AIS's main components include a broadcast scheduling procedure, specific broadcast frequencies and AIS-specific message composition. AIS has two types of transponders: Type A and B. Type A includes a high power (12.5W) Very High Frequency (VHF) transmitter for the maximum AIS range available, a dual channel receiver with either a built-in or external Global Positioning System (GPS) device. This was the originally specified AIS type specification and can received and transmit all the required AIS information. Type B provides a specification for less expensive AIS functionality. It makes provision for a less powerful VHF transmitter (2W) as well as a dual channel receiver. It specifies a built in GPS receiver and can only transmit a subset of AIS information.

2.3.1.2 AIS Broadcast Scheduling/Slot Map

AIS is built upon a technology called Time Division Multiple Access (TDMA). This time-multiplexing technique allows AIS to efficiently allocate limited frequency bandwidth. AIS has two broadcast frequencies (87B which is 161.975 MHz and 88B which is 162.025 MHz)

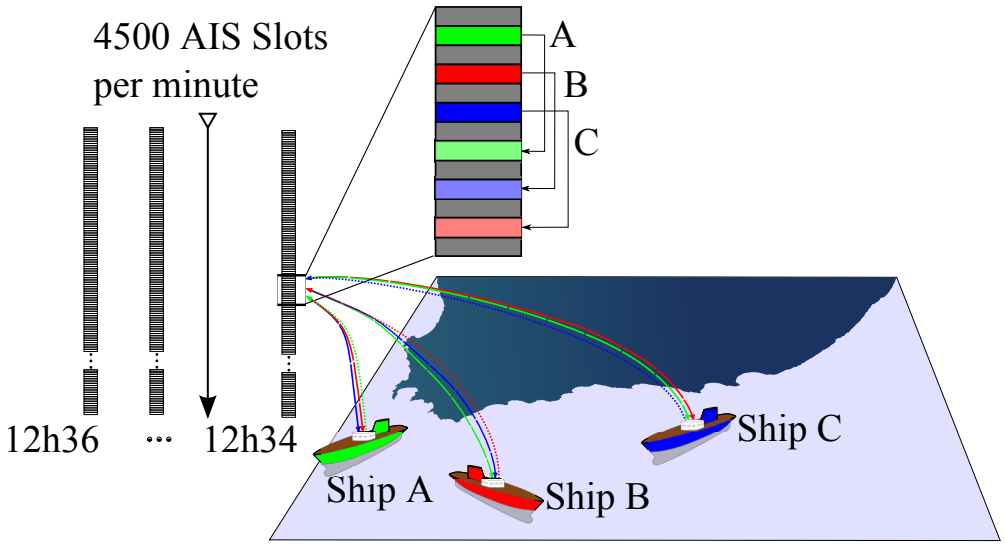

Figure 2.4: Typical representation of the slot map used in AIS Type A messaging system. A ship transmits its information at a certain time slot and then reserves a slot in advance.

and allows a fixed number of communication channels of up to 4,500 different channels every 60 seconds (which amounts to half that every 60 seconds per channel). Communication is bidirectional but to avoid transmission interference both channels are used when sending/receiving data.

Ships automatically update their details a number of times every 60 seconds and it is mandatory for ships to register their next time slot during their current slot update. Due to the manner in which the AIS system was built, it is able to intelligently adapt to the number of ships within a specific area. Message processing preference occurs for closer ships than ones that are further away. For an example of how the system operates please refer to Figure 2.4.

2.3.1.3 Satellite AIS

Satellite AIS (Sat-AIS) is an extension to the AIS specification as it allows for near-global ocean coverage. Normal AIS is a terrestrial-based system whereby messages are sent directly to terrestrial stations. Sat-AIS takes this a step further by allowing messages to be sent to a satellite which then relays it to the satellite's related ground station. The vertical height of transmission means that the overall coverage of Sat-AIS is far greater than the direct, horizontal coverage of normal AIS. One difference between Sat-AIS and terrestrial

AIS is that Sat-AIS has a greater transmission delay than normal AIS - depending on the application this may be an important point to note. Sat-AIS can transmit almost the full complement of Type A AIS information. This means that both types of AIS can be processed in much the same way and, depending on application, how the AIS data is received may be irrelevant. One of the downsides of using Sat-AIS is that the acquisition of positional data is not always available due to satellite revisit times whereas terrestrial AIS is.

2.3.1.4 AIS applications

AIS can be used in a number of ways such as ship traffic and collision avoidance, navigation aid, accident investigation ("black box") as well as for search and rescue.

2.3.2 Long Range Identification and Tracking

In May of 2006, the IMO introduced plans into the Safety of Life At Sea resolution which detail the implementation and enforcement of a LRIT system [9]. LRIT is a MDA-related initiative that became mandatory for all ships from 2008 onwards and is a mandatory installation for all passenger ships, high speed ships, cargo ships (over 300 tonnes) and offshore drilling units.

LRIT allows Member States (governments/search and rescue units) to request reports to obtain ship identity/location information in a timely manner. Requests may be made from ships operating under the Member State's flag and ships that will fall within the Member States coastline proximity. The system was initially designed to allow contracted governments to assess the risk of a ship off its coast and effectively respond to any possible threat. It has since been extended to allow for Search and Rescue operations, as well as to improve the safety (and to protect) the marine environment.

2.3.2.1 How LRIT works

The LRIT system is divided into a request/response system. Ships on international voyages need to send reports with their identity, position and time/date stamp at least 4 times a day (every 6 hours). Member states can request/purchase reports that inform the state when a ship is within 1000 nautical miles from its coastline. Similarly states with ports can request

reports for ships that wish to enter the port at a specific, predetermined time. LRIT therefore allows Member States to effectively monitor its coastline.

2.3.2.2 LRIT Reporting System

LRIT has a request/response system that requires various components linked together to allow for successful communication. LRIT contains two report operating modes [10]. On-demand Reporting which allows an LRIT-enabled ship to have the ability to provide a positional report wherever it may be and do so without human interaction. The LRIT system should also be able to respond to polling requests for this information. On-demand reporting affords states the ability to easily and quickly gauge a ship's whereabouts/heading direction. Pre-scheduled Reporting allows for a ship to automatically report its position at a certain reporting frequency. This type of reporting is also referred to as Automatic Position Reporting (APR). Message reporting frequencies can range from a maximum frequency of one report every 15 minutes (96 reports every 24 hours) to a minimum of 1 report every 6 hours (4 reports every 24 hours).

2.3.2.3 LRIT Broadcasting and Receiving Procedure

There is a basic system by which all LRIT-enabled ships adhere to (Figure 2.5). The procedure is as follows: the Member State will transfer and receive information from specific Data Centres (DC). Three types of Data Centres are present in the LRIT specification: National Data Centres (NDC) service a single Contracting Government/Member State to provide security of its own national region; Regional Data Centres (RDC) which service a group of Contracting Governments/Member States to cover a specific, nearby-to-each-state geographic region; and finally Cooperative Data Centres (CDC) which service two or more Contracting Governments/Members States to allow for cooperation between Member States to monitor areas/ships that are not necessarily geographically near to each other.

Communication between regions and ensuring that only the correct information is passed to specific Member States is enforced by what is known as a Data Distribution Plan (DDP) server. Communication of a APR or On-demand report occurs through what is known as a Application Service Providers (ASP) or Content Service Provider (CSP). If certain data are required by a different Member State then information is passed through what is known as

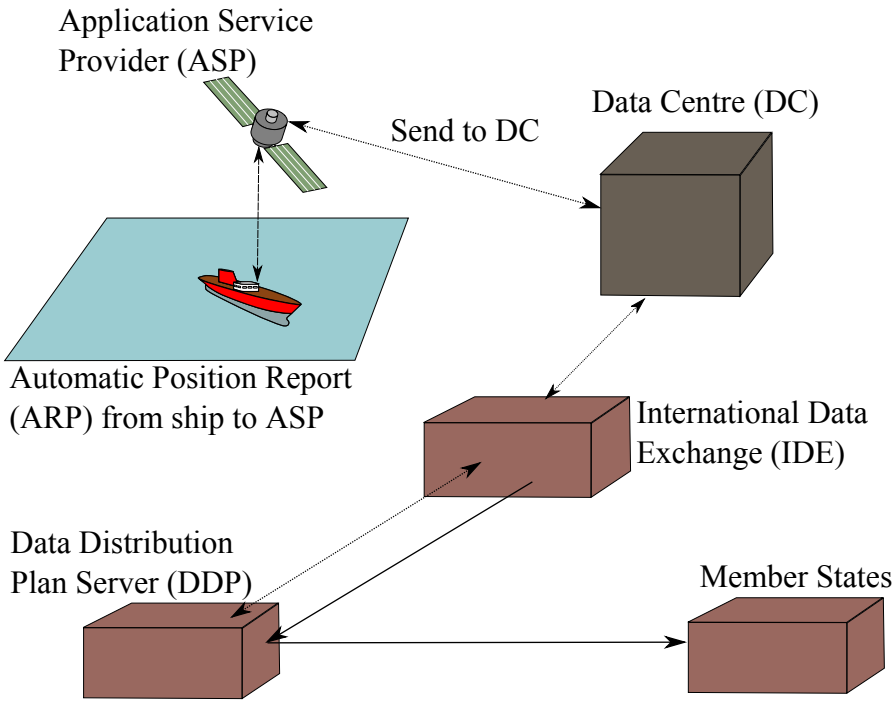

Figure 2.5: Typical representation of the data flow within an LRIT network. Ship locations are sent via an ASP to a DC (either national, regional or cooperational). Data are distributed via the IDE to Member states depending on the rules put forth in the DDP.

an International Data Exchange (IDE). The rules of how information can be routed through the IDE is governed from within the DDP.

2.3.2.4 LRIT Range and Receivers

Unlike terrestrial AIS, LRIT receivers only operate vertically as it is a space-borne ship monitoring system. Information is sent to a satellite (which is the ASP) which then forwards the information to the specific Data Centre associated with that specific ship. The space-borne nature of the system means that LRIT has a global coverage and ship-to-ASP coverage is approximately 400 kilometres.

2.3.2.5 LRIT applications

LRIT has a number of applications including monitoring ships for the purpose of security and especially as an aid to search and rescue due to its on-demand reporting.

2.3.3 Vessel Monitoring System

A vessel monitoring system (VMS) is a collection of standards that allow for ships to be tracked at sea. VMS is equipped on certain ships (usually fishing ships) and provides satellite-based positional information and tracking.

VMS's main advantage comes from the fact that it can greatly simplify the tracking and identification of commercial fishing ships. By requiring ships use VMS to fish in an area (usually the EEZ), would be incursions are avoided and illegal fishing activity and the like can quickly be identified.

Even though the preferred method of communication is a satellite-based system, VMS can also use ground based transponder stations. Reasons for using non-satellite based VMS include less expensive set-up time/costs but this comes at the expense of reduced coverage.

2.3.3.1 How VMS works

VMS works in much the same way as LRIT. Reports from each of the ships are sent to a satellite which then relays the information to a terrestrial base station. The base station validates the information, stores it and then notifies the monitoring agency that the data is available. Organisations such as a fishery monitoring centre then retrieves the data and stores its own database. The ships current position are then correlated to the current state of the marine area. If the ship is prohibited from being in the area (restricted EEZ boundary) then the fishery control can detect discrepancies.

VMS has one primary function: provide the location of a ship at any given time. This information typically is sent periodically to a base station via satellite transmission. Typically VMS systems are comprised of a GPS antenna/receiver, a computer processing unit and an antenna and transmitter to be used to send information to the satellite system. For an example of VMS, see Figure 2.6.

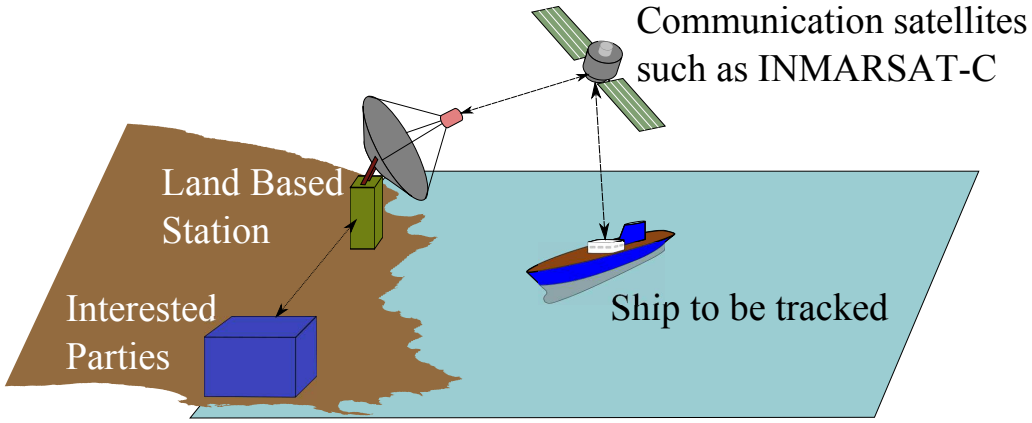

Figure 2.6: Typical representation of the data flow within a VMS network.

2.3.3.2 VMS applications

VMS has a number of applications including fishery control/catch register, secure communication channel provider and a search and rescue aid.

2.3.4 Marine Ship Monitoring Systems Comparison

This section provides a short comparison of the advantages and disadvantages of the various ship monitoring systems (see Table 2.1).

Which of the methods shown in Table 2.1 to use in a project is dependant on the project requirements. Greater coverage means LRIT or Sat-AIS will be chosen above AIS or if extensive logging is required VMS should be chosen. In the end the data garnered from each method differs slightly but can be used in much the same way as discussed next.

2.3.5 Ship Transponder Data Usage

Ship monitoring data can be used in a number of ways. The most prominent manner in which ship monitoring data is used is part of a simple ship tracking platform. Ship positions are reported and these are then referenced against a map to locate the ship. Monitoring ships using transponders is the most common way of monitoring ships in South Africa. The simplicity of the system allows for relatively cheap receivers to be set out and used to track ships off of the coast of South Africa.

Table 2.1: Comparison of the different four studied Marine Ship Monitoring Systems.

Ship Monitoring System	**Advantages**	**Disadvantages**
AIS	Simple, Effective within coverage area, real-time localisation.	Limited range, not fitted to all ships.
Sat-AIS	Extended AIS coverage, can map almost the entire ocean.	Longer delays for positions, requires satellite (expensive) .
LRIT	Global Coverage, uses existing comm. platform, excellent early warning system.	DDP is single point of failure, unencrypted transmission from ship to ASP, GPS spoofing, non-realtime localisation, opt-in.
VMS	Automatically handles fishing logging, global coverage, uses existing comm. platform.	Unencrypted from ship to FMC, GPS Spoofing, non-realtime localisation, opt-in.

Another possible use of ship monitoring data is for the validation of other ship detection methods. The other ship detection methods produce detections and associated ship positions. These positions are then compared to the positions reported by the ship transponder systems to determine the detection accuracy (if the ship transponder data and SAR data are time and location synchronised). Using transponder data in this manner can be problematic because a) the data can be sparse depending on which transponder system is used, b) ships need to be cooperative in order to get actual positional data and c) it is difficult to synchronise transponder data with SAR data.

One novel way of using ship transponder data is to build a ship distribution map. This map assigns probabilities to positions within a given geographical region based on how likely a ship will traverse that point. The likelihood of a ship traversing a given point is determined by the positional information garnered from transponder data and how many times that geographical point has seen activity from various ship transponders. These probability maps can be built up using sparse positional data over long periods and can help to identify

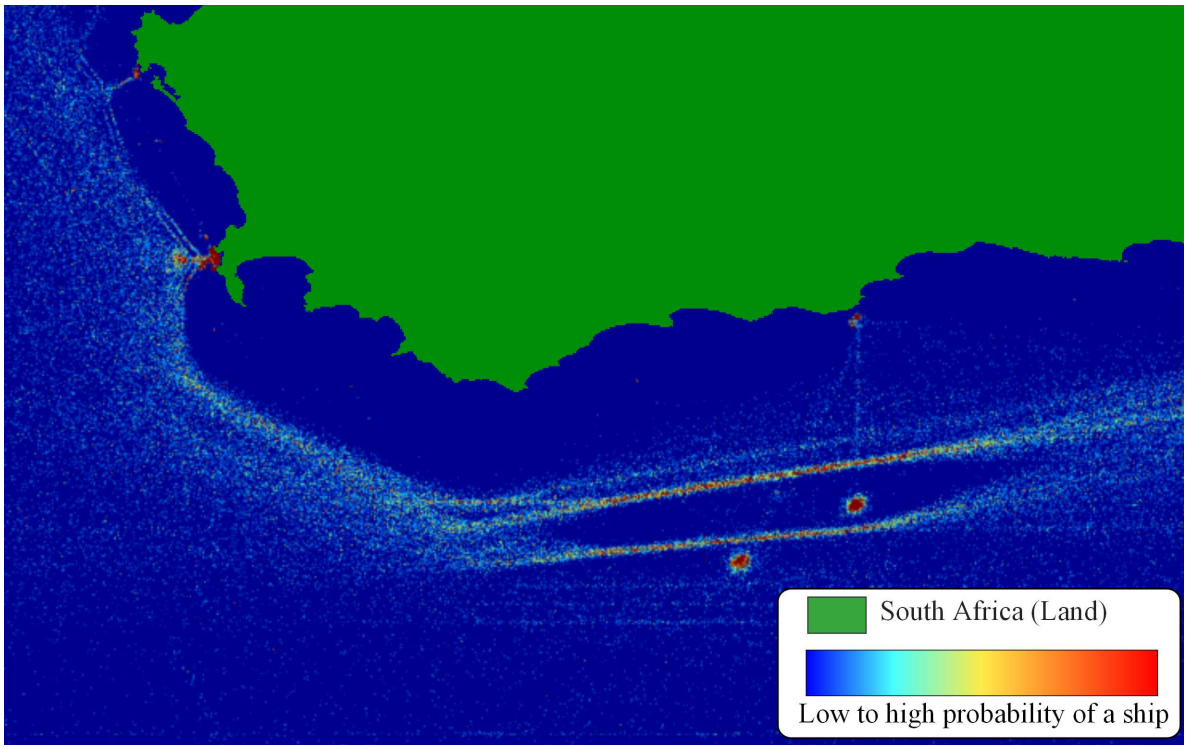

Figure 2.7: A ship distribution map of the west coast of South Africa. This map was generated using LRIT positional data and it highlights clear pathways that ships use to/from harbors.

interesting information around certain geographic areas such as typical shipping lanes used in that area. An example of such a ship distribution map can be seen in Figure 2.7.

2.4 SYNTHETIC APERTURE RADAR

2.4.1 Introduction to Synthetic Aperture Radar

Acquiring and analysing SAR imagery of the Earth has become more feasible recently [3]. There are numerous applications of SAR satellite imagery including: mapping and monitoring of vegetation, monitoring of sea and icy areas, target tracking and interferometry (InSAR). Satellite imagery from such satellites as ENVISAT and RADARSAT-2 provide a wealth of information and allow detailed observations of large areas in almost any weather condition, day or night. SAR satellite imagery is an excellent tool to help improve a country's MDA. A thorough understanding of SAR and its qualities will allow for a better understanding of the operation and analysis of SAR imagery which can, in turn, improve a country's understanding

of its own MDA. Furthermore, by understanding the past, current and future of SAR satellites (and imagery) one can design flexible systems that can be used across a broad range of applications related to MDA such as ship detection.

The next sections each deal with various aspects of SAR including the basic operation (Section 2.4.2), polarisation (Section 2.4.3), imaging phenomena (Section 2.4.4) and a discussion on past, present and future SAR satellites (Section 2.4.5).

2.4.2 Basic SAR Operation

Synthetic Aperture Radar (SAR) is a active (self-illuminating) remote sensing method and is an extension to phased arrays [3, 4, 11]. SAR satellites use an active, coherent microwave-based sensor in order to obtain high-spatial resolution of a an area. The sensor transmits an electromagnetic signal at some specified frequency (or more commonly note wavelength) which is then altered by the scene and recorded by the sensor. SAR is able to remotely sense areas in almost any weather condition, day or night, even through cloud cover. The fact that SAR is an active sensor improves the prevention of it being affected by radio interference but, it should be noted, that in most cases it is not immune to interference. In high-precision systems, cloud cover and other signal interference plays a large part in determining the return signal. For this dissertation, however, it is assumed that these effects are negligible at the level of observation and bandwidth (C-band) used.

In order to achieve a high spatial resolution, large bandwidth pulses are transmitted (in range channel) and a longer synthetic aperture antenna is synthesised due to the relative motion of the sensor platform along the flight path (azimuthal channel). SAR observes an area by the recording the transmission and reception of electromagnetic (EM) waves in the form of pulses. Reflected pulses (echos) are recorded from the objects and used to place objects relative to one another and the sensor. SAR can be mounted on either an airborne or a spaceborne transceiver. For the purpose of this dissertation the SAR will, in general, refer to the satellite based SAR systems.

The pulses are transmitted at a fixed rate known as the pulse repetition frequency (PRF). To reach targets at a very far distance the pulses need to modulated on a carrier frequency. These carrier frequencies determine the type and range of applications a SAR system can

be used for. Carrier frequencies vary with wavelength and different carriers have different properties that allow different applications. The pulses require other electromagnetic waves to ensure that the pulses are able to reach the targets. A list of prominent electromagnetic wavelengths, their frequencies and most common applications is presented in Table 2.2.

In most cases the pulses from a SAR satellite can be directed by varying the viewing geometry. The angle at which the system looks at the scene is called the antenna look angle and can be adjusted by steering the antenna or radar beam. The angle formed between the look direction and the normal to the ground is called the incidence angle. The line directly underneath the satellite is known as the nadir or suborbital track.

The radar observes a specific tract of the surface known as the antenna footprint. The size of the footprint perpendicular to the satellite is known as the swath width and is determined by the look/incidence angle. The resolution in this direction is known as the range or across-track resolution. The resolution of the footprint parallel to the flight of the satellite is therefore known as the azimuth resolution. Fig. 2.8 gives an example of the geometry encountered when using SAR side-looking satellites.

The azimuth resolution is inversely proportional to the parallel, along-track length of the radar antenna. For a fine azimuthal resolution, one would require an arbitrarily long antenna. Synthetic aperture radar overcomes this by *synthetically* extending the radar by coherently summing together the pulses received as the antenna moves along the track. This synthetically long antenna is the principle upon which SAR is based.

Each pulse from the satellite is modified by the target in some way recorded by the reflectivity of the target's backscatter function. A coherent imaging radar operates by recording the complex reflected values. If the data is only recorded using the magnitude then the instrument is operating in an incoherent manner. The magnitude P_r and phase φ expressions are related to the power scattered back from the target and can be expressed as

$$P_r = \frac{P_t G A_e}{(4\pi R^2)^2} \sigma_A^0, \tag{2.1}$$

$$\varphi = \frac{4\pi}{\lambda} R + \varphi_{scatter}, \tag{2.2}$$

where P_r is the received power, P_t is the transmitted power, R is the distance to the target, G is the antenna gain, A_e is the effective antenna area and σ_A^0 is the radar cross section per

Table 2.2: Common electromagnetic waves, their wavelengths and frequencies along with their most common applications.

Band	**Frequency (GHz)**	**Wavelength (cm)**	**Applications**
VHF	0.03 - 0.3	1000 - 100	TV & FM Radio
P	0.28 - 0.39	107 - 77	SAR
UHF	0.3 - 1	100 - 30	TV
L	1 - 2	30 - 15	GPS, GSM, DAB, SAR
S	2 - 4	15 - 7.5	Sat. Comms., 802.11b/g, SAR
C	4 - 8	7.5 - 3.75	SAR
X	8 - 12.5	3.75 - 2.40	SAR, Radar, Comms.
Ku	12.5 - 18	2.40 - 1.67	Sat. and Terrestrial Comms.
K	18 - 25	1.67 - 1.18	Radar, Sat. Comms.
Ka	25 - 40	1.18 - 0.75	Sat. Comms., close range radar

unit area (for non-point targets). The variable λ is the wavelength of the trailing wave and $\varphi_{scatter}$ is the function that models the phase modulation the target applies to the original pulse (depends on polarisation, volume scattering, etc).

The distance R can be determined by measuring the time delay between the transmission of the radio pulse from the satellite and the reception of the received backscatter. Another important point to note is that an image may be converted from a slant range to ground range by utilising the incidence angle. This is important as it determines exactly how the final image will be structured.

2.4.2.1 Range resolution

The resolution of the radar depends on a number of SAR attributes such as the sensor viewing geometry at acquisition, sensor properties and the form of the pulse that is transmitted to the targets. The spatial resolution of side-looking SAR system is far better than a nadir

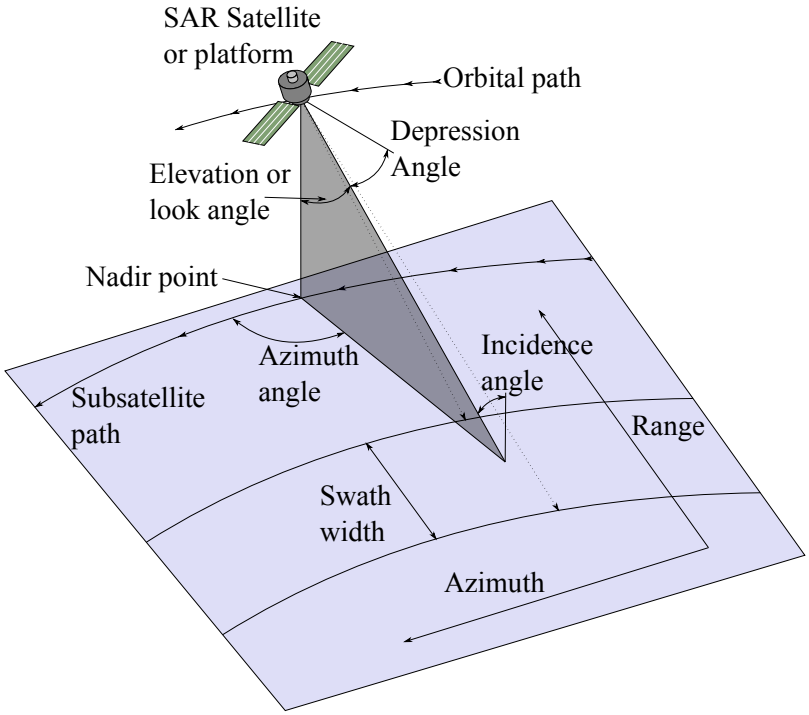

Figure 2.8: SAR Operational Basics including: Azimuth angle/resolution, Range resolution, Swath Width, Look/Incidence angle and Suborbital (nadir) track.

system because the side looking configuration avoids over-saturation of the signal. We can determine the spatial resolution between targets (in the slant range) by resolving two return signals from two targets. The spatial resolution, δ, is defined as

$$\delta = \frac{c}{2B}, \tag{2.3}$$

where c is the speed of light and B is the bandwidth of the transmitted pulse. To achieve a larger bandwidth the transmitted pulses' width needs to be very narrow. It is necessary to linearly modulate the frequency of the pulse. A pulse that is modulated in terms of frequency is known as a chirp. Essentially, chirping "compresses" the pulse in time allowing for a dual benefit of a narrower pulse (range resolution improvement) but still maintaining the benefit of having a wider pulse (more energy sent and thus more energy returned). To retrieve the chirp information digital pulse correlation is used to find the pulse return in the signal.

2.4.2.2 Doppler Frequency

A target on the ground has a Doppler frequency when it is observed at a specific position relative to the an orbiting platform. The shift in Doppler frequency (f_D) between a transmit-

ted pulse and received pulse within the image is exactly proportional to the relative velocity between the target and the satellite is

$$f_D = \frac{2v_{\mathrm{p}}x_{\mathrm{d}}}{R\lambda} = \frac{2v_{rel}}{\lambda}. \tag{2.4}$$

The velocity of the platform is v_{p}, v_{rel} is the target's velocity relative to the platform, λ is the wavelength, x_{d} is the distance from the plane perpendicular to the flight line whilst R is once again the slant range distance. Concentric cones represent surfaces of constant Doppler frequency relative to the satellite. When these cones intersect the Earth, a set of hyperbola are formed. These create zero-Doppler shift surfaces which are perpendicular to the direction of orbiting. By measuring the slant range at a specific instant and knowing the Doppler frequency at that time, the target can be located accurately on the ground. It is important to note that azimuth resolution information can be acquired using the matched filtering between the Doppler history of the targets as they pass through the antenna. Coherent radar is required for this as the Doppler frequency closely corresponds to the change in the phase differences between the transmission and reception of a transmitted pulse.

2.4.2.3 Azimuth resolution

One important aspect of SAR is that azimuth resolution is independent from range resolution. This is due to the fact that the Doppler profile (of a target) passing through the antenna will increase linearly with range. Since Doppler frequency resolution and Doppler profile are inversely proportional the angular resolution improves linearly with range also and thus the azimuthal resolution δ_{az} is independent of the range resolution. The seemingly paradoxically fact is that the resolution of a SAR improves as the along-track length of the antenna decreases,

$$\delta_{az} = \frac{L}{2}, \tag{2.5}$$

where L is the azimuthal length of the radar antenna.

2.4.3 SAR Polarisation

For a given transmitted polarisation pulse, the reflectivity of a target will have a number of different polarisations in the backscatter. To properly characterise the reflectance of a target two orthogonal polarisations must be transmitted as well. For most of the current lifetime

of SAR, satellites have not had the ability to transmit and receive in multiple polarisations (see [11]). The polarisation of a signal is written using two letters to indicate how the signal was transmitted and received. The first letter is the transmitted polarisation and the second is the received signal's polarisation. For instance, HH indicates that the signal was transmitted and received in the horizontal direction whereas HV indicates horizontal transmission and a vertically received signal. In general there are four polarimetric modes available to modern SAR satellites and are:

Single Polarisation: The most prominent form of polarisation before ENVISAT. The SAR satellites would transmit either an HH or VV and would capture only the same (linear) transmitted polarisation.

Dual Polarisation: This is the most common form of polarisation transmission. This can lead to transmission and reception of either like- or cross polarised components. For instance HH and VV or HH and HV. This form of polarisation does not retain the relative phase between polarisations and thus is devoid of phase information (between polarisations).

Full Polarisation: This is the richest option because it allows full characterisation of the complex matrix of the backscatter at all resolved points in the scene. Almost all modern SAR satellites being developed or planned have this as the polarisation system onboard. Fully polarimetric radar allows for a full description of the scattering properties of targets and essentially replaces the scalar value σ_A^0 (equation 2.2) with a complex vector.

Quadrature polarisation SAR is configured such that the transmitted pulse is sent in both (H and V) orthogonal polarisations. The satellite must then process the received data such that data represents all the possible combinations of transmit and received polarisations. This data will then provide the best quantitative scene description possible for SAR imaging.

2.4.4 SAR Imaging Phenomena

There are four main forms of phenomena that need to be considered when processing SAR images. The following presents a short overview of these phenomena and why they occur.

Foreshortening Foreshortening is a phenomena which occurs due to the side-looking nature

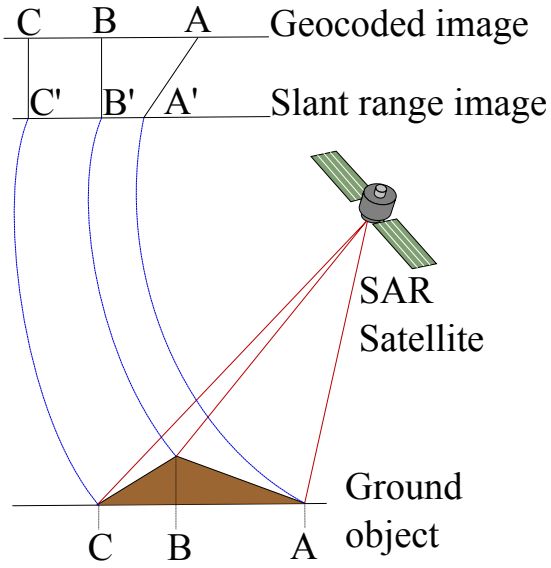

Figure 2.9: The effects of foreshortening. Notice that in the uncorrected slant image points A' and B' appear much closer together than their actual distance on the ground.

of SAR imagery. Foreshortening causes lengths of objects with differing heights to appear nearer to one another than they actually are on the ground. The effect of this is that the facing edges of an object (such as a mountain's facing side) appear to be shorter or 'leaning' in the resulting SAR images. Fig. 2.9 shows how the facing edges of an object appear shorter in SAR images before correction.

Layover Layover is another SAR image phenomenon. Layover appears as if physical elements have been 'swapped' in a SAR image as compared to the actual points on the ground. This occurs because taller physical elements will reflect the radar before lower elements and when these are recorded the taller elements is recorded as arriving before the shorter elements and thus 'first'. This means that the taller elements are recorded followed by the shorter elements irrespective of their actual positioning on the ground. Fig. 2.10 shows how the higher points of an object may cause an SAR image to record a 'swapping' of physical points.

Shadow Due to the fact that SAR satellites emit EM waves at an angle a shadow of one objects face will be cast upon one another. This shadow causes dark regions within an image and can make processing areas difficult because the shadow hides details of the object within the shadow. Fig. 2.11 shows the effects of shadow within a typical SAR

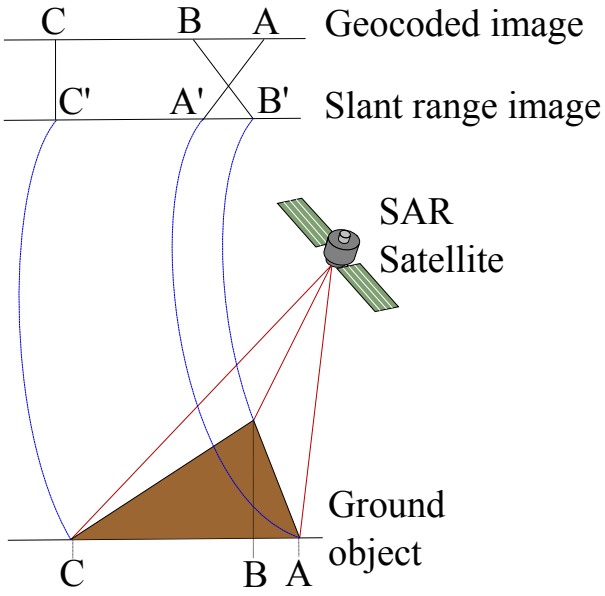

Figure 2.10: The effects of layover. Notice how points A' and B' are in reverse compared to how they are placed on the ground.

image.

Speckle Speckle is an SAR image phenomena that occurs when multiple wavefronts interferences occur (constructive and destructive speckle). When multiple signals from the radar are received and their phases coincide (either additively or subtractively) they produce an interference in a SAR image known as speckle. Speckle is a very common phenomenon in SAR imagery and is usually dealt with by either filtering the image before processing or designing the processing of the image to account for the speckle within the image. For an example of a SAR image with speckle see Fig. 2.12.

2.4.5 SAR Satellites

There have been a number of SAR-based satellites developed since the inception of the SAR concept in June 1951 at Goodyear (Later Lockheed Martin). The original concept grew from airborne systems into spaceborne systems as is found today. Table 2.3 details the most prominent historical SAR satellites since the first SAR satellite SEASAT, launched in 1978.

Three satellites of the current era in remote sensing will be described next. The first, EN-

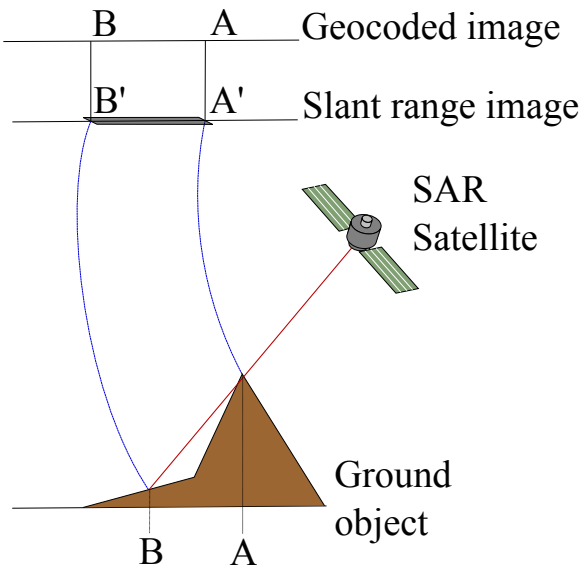

Figure 2.11: The effects of shadow. The height of the object causes a shadow that is visible on the SAR image.

VISAT, is a somewhat recent satellite that went offline in 2012. It provided a wealth of data and represents a good benchmark against which the newest satellite missions can be compared against. It also represents the most applicable historic satellite data that will be used in this dissertation. RADARSAT-2 is the only currently operational satellite of the three discussed mentioned here and is a very important satellite to a number of governments and companies due to the data it provides. Finally, a look at the Sentinel satellite mission will provide a look into the future of SAR-based remote sensing.

2.4.5.1 Environmental Satellite

The environmental satellite (ENVISAT) [12] is a multi-purpose satellite launched in March 2002 by the European Space Agency (ESA). It was, at the time of launch, one of the largest, most advanced civilian observational satellites. It contained nine advanced instruments including instruments for measuring sea-surface temperature, ozone monitoring and reflectance. For the purpose of SAR observation a special instrument was installed onboard known as Advanced SAR (ASAR). ENVISAT was the only orbiting polarimetric SAR satellite until RADARSAT-2 was launched in 2007. The initial ENVISAT mission was planned for 5 years but operated for nearly 10 and delivered over a petabyte of information. Unfortunately,

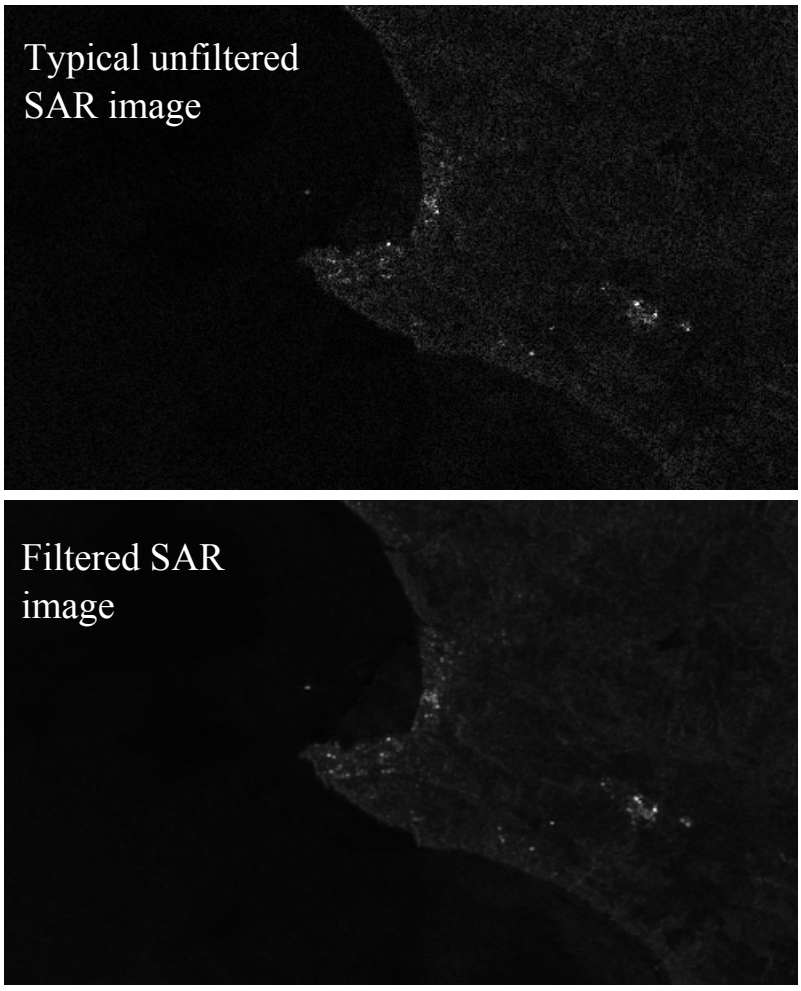

Figure 2.12: Speckle noise on a SAR image is characterised as bright and dark pixel values that are randomly distributed throughout the SAR image. The top image shows an unfiltered SAR image and the bottom shows the result of a spatial 9×9 median filter.

communication was lost to the satellite on the April 8, 2012.

The ENVISAT satellite had a sun-synchronous, polar orbit with a period of 101.6 minutes. The satellite had a revisit time of 35 days and orbited the earth at approximately 790km above ground. The ASAR instrumentation used the C-band with a frequency of 5.331 GHz, and was able to observe using a number of different data modes and polarisations. A summary of this information can be found in Table 2.4.

Table 2.3: Historic satellites and their operating characteristics.

Satellite	Country	Year	Operating Band (Frequency)	Applications
SEASAT	USA	1978	L (1.275 GHz)	Global surface wave field and polar sea ice monitoring
ERS-1/2	Europe	1991/5	C (5.25 GHz)	Surface height change detection,
JERS-1	Japan	1992	L (1.275 GHz)	Global forest mapping
RADARSAT-1	Canada	1995	C (5.3 GHz)	Change detection
SRTM (SIR-C/X-SAR)	USA/Germany	2000	C (5.25 GHz) / L (9.6 GHz)	Digital Elevation Modelling

2.4.5.2 RADAR Satellite 2

The RADAR satellite 2 (RADARSAT-2) [12] is a SAR satellite launched in December 2007 by the MacDonald Dettwiler and Associates (MDA) company built in association with the Canadian Government. It is part of a pair of satellites designated with the RADARSAT moniker. It has a number of advanced sensing device improvements over RADARSAT-1. It is the only fully polarimetric SAR satellite in operation today (this will change in the future with the launch of Sentinel 1).

The RADARSAT-2 satellite has a sun-synchronous, polar orbit with a period of 100.7 minutes. The satellite has a revisit time of 24 days and orbits the earth at approximately 798km above ground. The onboard SAR instrumentation uses the C-band with a frequency of 5.405 GHz and can map regions using a full polarimetric system. A summary of this information can be found in Table 2.4.

Table 2.4: Modern satellites in this dissertation and their operating characteristics.

Satellite	**Sun Synchronous**	**Altitude (km)**	**Orbital Period (mins.)**	**Repeat Cycle (days)**	**Operating Band (Frequency)**	**Polarisation**
ENVISAT	Yes, Polar	782 - 792	101.6	35	C - 5.331 GHz	Dual
RADARSAT-2	Yes, Polar	798	100.7	24	C - 5.405 GHz	Full/Quad
Sentinel 1	Yes, Polar	693	96	12 / 6	C - 5.405 GHz	Full/Quad

2.4.6 Sentinel Satellite Collection

The Sentinel collection of satellites is an extension of the ENVISAT mission by the ESA. The mission consists of (currently) 5 satellites but may be extended to more in the future. Sentinel 1 is a group of 2 satellites that is the planned to be launched. Sentinel 1-A launched early 2014 whilst Sentinel 1-B is expected to launch in 2015. When both satellites are in operation it will provide a significantly lower revisit time compared to ENVISAT and RADARSAT-2.

The Sentinel-1 satellite will have a sun-synchronous, polar orbit with a period approximately 96 minutes. The satellite will initially have a revisit time of 12 days which will become a 6 day repeat cycle when Sentinel 1-B is operational. The SAR satellite is expected to orbit at approximately 693km above ground. The onboard SAR instrumentation will use the C-band with a frequency of 5.405 GHz and will be able to map regions using full polarimetry. Another interesting attribute of the Sentinel mission is that it will have dedicated X-Band Radio Frequency (RF) channels for data transmission to allow receivers on the ground much faster data acquisition times. A summary of this information can be found in Table 2.4.

2.5 SHIP DETECTION

2.5.1 Introduction to Ship Detection

Maritime Domain Awareness may use SAR to monitor various aspects of the seas. From Oil-spill detection to sea-ice state information, SAR is an extremely useful tool in the surveillance of large tracts of the sea. One of the prominent usages of SAR is the detection and monitoring of ships at sea. This section intends to present the ubiquitous methods of ship detection as well as the general system typically used to process SAR images to detect ships at sea. The first part of this chapter details the specific process most ship detection systems adhere to. The second part entails giving specific details about some of the state-of-the-art methods of detecting ships at sea.

Ship detection using SAR satellite imagery is divided into methods that detect ships at sea using either the ship itself or its wake [4]. For the purpose of this disseration it is assumed that all ship detection is done using only the ship and its reflectance in the SAR image. This is done to avoid the inherent problems with wake detection such as unknown sea-state and stationary ships.

Ship detection using SAR satellite imagery can be split into several steps. These steps are found in one form or another throughout the literature and the sequence of steps has the same basic structure. The typical steps are shown in Fig. 2.13.

This section of the dissertation will discuss some of the different methods used in each step and provide a general overview of the methods that exist in literature today. In the section following this one an extensive, mathematical look at each of the methods used in this dissertation will be given.

2.5.1.1 Preprocessing - Land Masking and Image Filtering

SAR-based ship detection systems almost always have an initial preprocessing stage in which the areas of land around bodies of water (ocean, large lakes) are “masked” out [4]. In some instances this step is performed last depending on the requirements of the processing of the SAR imagery. This step ensures that results generated do not concern itself with points on

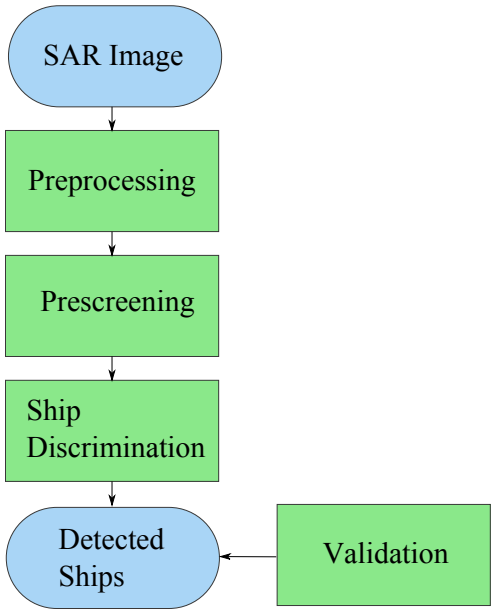

Figure 2.13: Typical ship detection system steps found in literature. Some systems forego a number of steps and/or combine two or more steps. The output is the detected ships superimposed on the image or a list of detected ship positions.

land that may be incorrectly identified as ships.

The most common approach to land masking is to geo-reference the image (register the image relative to the earth's surface) and then match the geographic shoreline to an area on the image. A simple method involves registering the image via the geodetic parameters from the SAR imagery [13]. In almost all land masking cases there exists a "buffer zone" which represents a level of inaccuracy with the registration. The method employed by [13] use an FFT based correlation technique to better register the land mask and achieve approximately 100 meters registration accuracy (difference between estimated and actual coastlines). A recent geospatial registration method makes use of a relatively new image processing technique known as Speeded Up Robust Features (SURF) to detect interesting features and register a SAR image accordingly [14]. There is no mention of actual registration error in terms of meters but an accuracy of between 1.1 and 1.2 pixels is expected on the SAR image (this translates to approximately 150 meters for the given ASAR sensor).

Alternative automatic shoreline detectors generally involve advanced image processing techniques to detect shoreline boundaries. Two methods seek to solve automatic shoreline de-

tection using in image processing principles such as Canny Edge detection and boundary-following algorithms [15, 16]. The results for both methods provide good-to-excellent results with the first stating a 1 pixel difference between detected and actual coastline and the second stating a 75% correct shoreline identification accuracy. The method employed almost universally today takes a georeferenced image and uses a shapefile with coastlines to remove the land from the image.

One of the most prominent artefacts visible on a SAR image is that of speckle noise. It is a well studied, multiplicative EM interference pattern formed when a coherent signal (such as radar) reflects off objects that are smaller than the minimum resolution cell of the image [3]. The various sub-scatterers can destructively or constructively interfere with one another thus causing either "dark" or "bright" speckle noise. The most basic method to deal with noise (from an image processing perspective) is to use some form of image filtering. Whilst the land masking step is almost unanimously used in some form or another to identify coastline boundaries in SAR satellite imagery, image filtering to remove speckle is a very situational step in the ship detection process. This is due to the fact that image filtering can possibly remove details that may be useful for certain types of ship detection systems. It is recommended to carefully consider the use of image filtering at the same time and in conjunction with the prescreening and ship discrimination algorithms to be used [4]. In most cases using multi-look processing and algorithms such as the Constant False Alarm Rate (CFAR) and its variants are sufficient enough to account for speckle noise and therefore no preprocessing filtering is needed.

Noise removal methods to remove speckle exist and include the adaptive Wiener Filter and the much more recent Bilateral Filter [17, 18]. Both try to remove speckle noise by using adaptive filters whilst trying to keep as much of the original image content intact. The adaptive filter used in the adaptive Wiener filter was compared to 3x3, 5x5 and median filters [17]. It was shown to have excellent results (using an edge map quality metric) with the adaptive filter providing the better results of all those tested. The Bilateral filter performs SAR de-speckling without removing too much detail from the image but has some errors when the number of looks is small [18]. Multilook processing is a form of spatial fitlering and can be used to process the image to improve the detail present in a SAR image [3].

2.5.2 Prescreening - Global, Adaptive and Other Thresholding Methods

Prescreening is the most prominent area of research in SAR ship detection. Most of the literature found regarding ship detection uses some form of prescreening to detect the ships within a SAR image. The various methods can be grouped into three broad categories that have various levels of performance depending on a number of important SAR parameters such as resolution, swath width, radar reflectivity (Radar Cross Section or RCS), polarisation and number of looks (multi-RCS images). Global methods create a single criterion and apply this criterion to all pixels within a SAR image to detect ships. Adaptive or local methods (typically) use a small, localised window to process each pixel based on its local neighbourhood. There is a final group of prescreening algorithms that do not necessarily fall into either of the above categories and is grouped as a separate class of algorithms even though they may share characteristics with some of the other methods. Short discussions of each of these groups will be given in the following sections.

2.5.2.1 Global Methods

Global prescreening methods operates by using a single threshold value against which each pixel within an image is compared. There are a number of ways in which the threshold can be selected. The simplest is to generate a histogram of the pixel values and select a value that will ensure a specified quartile of the image pixels (5% or 1% for instance). One of the more novel techniques selects a global threshold value using a K-distribution CFAR detector and then a two parameter CFAR adaptive method is applied to the global thresholded result to obtain ship positions [19]. Finally, a different approach to global thresholding involves normalising the radar cross section within the SAR image and then applying a global threshold to this normalised image [20].

2.5.2.2 Adaptive Methods

Adaptive methods or local methods base their algorithms and thresholding values on the individual pixels in SAR image and each of its surrounding neighbourhood pixels. Adaptive methods discriminate based on a value in much the same way as global methods but adjust this value based on the information provided by a pixel and its nearest neighbours. The main

point of an adaptive filter setup is to generate ship detections where the false alarm rate is constant. Constant False Alarm Rate (CFAR) detectors try to identify areas with values that differ significantly from the background. Many methods exist to do this. The two main types of methods involved either a) assuming a specific distribution for the pixel intensities while calculating the required threshold from the neighbouring pixels or b) estimating the distribution using nonparametric methods and then using this information to calculate the required threshold. By selecting a threshold value adaptively the idea is to actively adapt to local image statistics as the window is moved around which helps limit the number of false alarms.

The most common method used in CFAR is to assume a specific distribution of the noise in the image and detect all points that differ from that noise. The most common CFAR variation is called the Gaussian distribution or two-parameter CFAR detector [4], [19]. The two parameters refer to the background mean and background standard deviation. A more advanced (and equally ubiquitous) CFAR detector is known as the K-distribution CFAR detector [4], [19]. The K-distribution is mathematically intractable and thus the distribution needs to be solved numerically for each sub-window within the image [4]. The Cell-Averaging CFAR (CA-CFAR) algorithm is a special case of the K-distribution algorithm and simply uses the background pixels to estimate average background pixel values and then identify unusually bright pixels [3], [4], [21]. Recently some authors argue that the K-distribution is not a sufficient distribution model for background pixel intensities and other distributions such as the Alpha-stable distribution should be used [22]. In much the same vein a nonparametric Parzen-Window-Kernel-Based algorithm [23] uses a Parzen window technique to estimate the underlying distribution and then uses this information in the usual CFAR manner to build a threshold with which to determine whether a pixel is above or below a certain threshold.

2.5.2.3 Other Methods

CFAR may be the most common prescreening method found in literature, but it is not the only one available. Another method of ship detection that is not precisely either a local or global method is Wavelet Transform (WT) ship detection [24]. Wavelet transforms are a signal processing concept that allow for multiresolution data analysis. A detector can be built to discriminate between sea and ships by combining the WT of SAR images and *a priori*

knowledge of the sea and ship statistics. Another interesting method by which ships can be detected is by studying their physical backscattering properties within the SAR image (Rice Factor) and then exploiting this information in a CFAR-like manner by building a Rice Factor image and then adaptively thresholding this image [25]. Finally ship detection methods can use a learning based system to detect ships in SAR imagery. One well known method is the Genetic Algorithm (GA), Radial Basis Function (RBF) Neural Network [26]. This method uses a GA to train a RBF neural network (NN) to detect ships within a SAR image. While this method is certainly interesting, it requires a training stage and therefore more imagery to create a working prescreening system.

2.5.3 Ship Discrimination

After the prescreening step a secondary ship discrimination step is usually applied to SAR images [4]. The reason for this may be two fold - either the system designer wishes to group similar pixels together to form 'ships' or the designer wishes to apply even further filtering so that the overall likelihood of ship detection is increased. The latter is due to the fact that despite being extremely cautious with the design and selection of parameters for the prescreening stage, false positives may still be present and therefore additional post-processing is sometimes required to produce acceptable results. It should be noted that this step is sometimes combined with prescreening to produce a set of detected ship centre points or ship longitude and latitude locations.

The first and simplest method applied after the prescreening step is a simple filter such as a Butterworth filter [13]. The Butterworth filter must be used with great care. An improvement of Butterworth filtering uses a morphological filter and clustering method to group and detect ship pixels [20].

There exists further still complicated methods for ship discrimination. Two of the more popular are Resonance Agglomeration and Elimination of Local Noise [27]. Resonance Agglomeration runs through the resulting image and groups any two objects that are h pixels or less apart where h is chosen in a heuristic manner [27]. The Elimination of Local Noise techniques splits the image into 100x100 pixel sub-images and then applies a threshold to this area [27]. The 100x100 pixel area is further divided into four 50x50 areas and the same procedure is applied to these four sub-pixels. The standard deviations between the large

sub-image and small sub-images is compared with the idea that the standard deviation of the pixel area may be high when an ship lies outside yet close to one of these noise areas. It should also be noted that these two methods are used in one of the more popular ship detection systems known as the Search for Unidentified Maritime Objects (SUMO) [4, 27]. Two other alternatives for ship discrimination include the ever-present Neural Network and a secondary two-parameter CFAR detector as the post-processing step after initial prescreening of the input SAR image [4, 19].

2.5.4 Validation

Two main methods of ship detection validation are available to a ship detection system designer. In all cases it should be noted that the amount of imagery available is typically very sparse. That is there are usually only a few SAR images at any one time and within these SAR images there are usually at most a hundred or so ships within an image. This effectively means that the imagery that is available should be used in the most efficient manner. The next two section discuss the two different ways in which ship detection results can be verified.

2.5.4.1 Ground Truth

A SAR image may sometimes have a corresponding file which indicates where the actual ship positions are of the ships within the image. When a SAR image is accompanied by this information it is said that the SAR image has a corresponding Ground Truth (GT). Ground truths can be generated and presented in a number of ways. The most common is simply an georeferenced image with the ship pixels highlighted within the image. Other GT methods include a list of latitude and longitude points or special shape files that contain the ship position. Using the performance metrics discussed further on, one can empirically test the performance of a ship detection system if some form of GT is available.

2.5.4.2 Ship Transponder Data Comparison

As discussed in section 2.3, ships themselves have transponder data that they either emit to terrestrial or orbiting (Sat-AIS) receivers. As discussed previously, ship transponders typically broadcast information about the ship such as latitude/longitude position and heading.

Using historical transponder data allows for a sort of Ground Truth-like scenario where ship positions can be known relative to the position and time at that position. The difficulty with this method of ship position validation is acquisition rates. SAR satellites and ship transponders typically do not transmit and receive at the same times. This leaves a situation where an area is observed with SAR but little or no transponder data coincides with the SAR acquisition time. This requires interpolation in time of ship locations so that when a SAR image is processed at a certain time all points within a geographic region can be interpolated to determine roughly where the ships were at the time of the SAR acquisition. Examples of interpolation include linear velocity prediction, Kernel Density Estimation techniques as well as more advanced regression and prediction techniques such as Gaussian Processes [28].

CHAPTER 3

METHODS

3.1 OVERVIEW

This chapter will detail the various methods used in this dissertation to detect ships within SAR images. The system is split into four main components, namely: Data Synthesising, Prescreening, Ship Discrimination and Validation. An example of how the system is connected is shown in Figure 3.1.

The first part of this chapter cover some fundamental concepts related to SAR images. The concept of using synthetic versus real data will be discussed as well as how both of them are used.

The second part of this chapter will look at three prescreening methods: 1) Global, 2) CFAR and 3) Wavelet thresholding. Each method will be given a mathematical overview as well as a description each of the method's dependent parameters.

The third section of the chapter delves into the final step called ship discrimination and acts as an additional filtering step to improve the overall ship detection system's accuracy. Two methods are discussed: 1) Morphological Closing and Connected Component Analysis and 2) Mean Shift Clustering.

Validation plays an integral part of determining the viability of the SAR ship detection system in real-world situations and is discussed in Chapter 4.

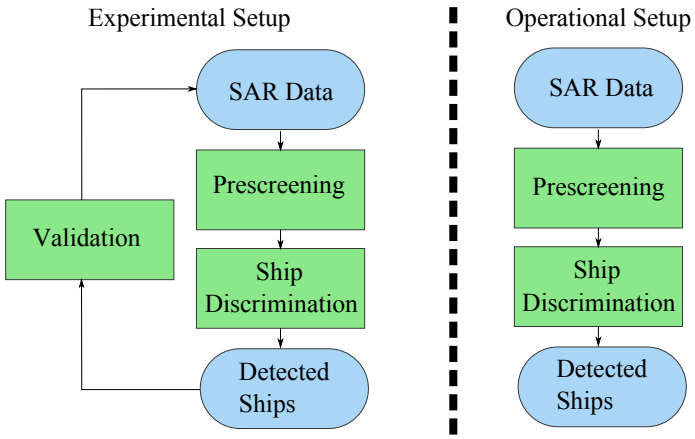

Figure 3.1: The experimental and operational block diagrams shown alongside one another. Notice that the operational block diagram does not use a validation step as it deals with as-of-yet unseen data. The experimental setup for the block diagram is the one used in this dissertation and the validation step is used to compare the results of the detected ships to their actual positions.

3.2 DATA GENERATION

3.2.1 Introduction

Testing a ship detection system using a large number of actual SAR images is the ideal scenario. Unfortunately, SAR imagery is expensive and in many cases difficult to acquire to test certain hypothesis. To properly evaluate a system, however, it is required that the system's robustness is tested. If only a small number of real world images are available then synthetic imagery can be used to test the robustness of the system.

Synthetic SAR images are generated by using an underlying distribution to model the sea noise. For the purposes of this dissertation two models were found that best describe the noise distribution function at sea: the Rayleigh distribution and K-distribution.

Three sections will detail how the synthetic SAR images are synthesised and processed by the ship detection system. The last three sections look into how sea clutter is modelled these two different noise models.

All of the notations found herein are derived from the seminal work by Gonzalez and Woods

[29].

3.2.2 Digital Image Fundamentals

3.2.2.1 SAR Digital Image Formation

In SAR satellite imagery the pixel values are backscatter values recorded by the SAR instrument and are a combination of both intensity and phase. These values are represented as complex values however, none of the methods used in this dissertation require phase information and so only intensity values are used.

Assuming a SAR input image $\mathbf{I}$ with dimensions $X \times Y$ where $x = \{0 \dots X-1\}$, $y = \{0 \dots Y-1\}$ and $x, y \in \mathbb{N}_0$ such that image $\mathbf{I}$ can be defined as

$$\mathbf{I} = \left\{ \left\{ I(x,y) \right\}_{x=0}^{x=X-1} \right\}_{y=0}^{y=Y-1} \tag{3.1}$$

$$= \begin{bmatrix} I(0,0) & \cdots & I(0,Y-1) \\ I(1,0) & \cdots & I(1,Y-1) \\ \vdots & \ddots & \vdots \\ I(X-1,0) & \cdots & I(X-1,Y-1) \end{bmatrix}. \tag{3.2}$$

The symbols x and y are spatial coordinates, $I(0,0)$ is the pixel intensity value at the top left point known as the origin. X refers to the number of columns in the image and is defined such that $0 \leq x < X$. Similarly, Y defines the number of rows in an image and is defined as $0 \leq y < Y$.

Each point in the matrix is known as a pixel and stores a level of intensity extracted from the sensor (in this case the processed SAR image). Each pixel value in a digital image is quantised into the region $[0, L-1]$, where L represents the number of levels of intensity and is defined as

$$L = 2^{\text{bits}}, \tag{3.3}$$

where bits refers to the number of bits used to store an image. The most popular choice for bits is usually bits $= 8$. The total number of bits to store an image (without compression) is therefore $\text{bits}_{\text{total}} = X \times Y \times \text{bits}$. An example of an 8-bit image can be seen in Figure 3.2.

	0	1	2	3	4	5	6
0	100	100	100	100	100	100	100
1	100	190	190	100	100	100	100
2	100	100	100	100	100	100	100
3	100	100	100	255	255	100	100
4	100	100	255	255	100	100	100
5	100	100	100	100	100	100	100
6	100	100	0	100	100	100	100
7	100	100	100	100	100	100	100

Figure 3.2: An example of an 8-bit digital image. This image has 256 levels of gray ranging from a pixel intensity of 0 which represents a black pixel and 255 which represents a white pixel.

3.2.2.2 Histograms

Digital images are composed of a finite number of discrete intensity levels. When these intensity levels are analysed they can provide information about the image for a variety of purposes. To analyse the intensity levels of an image a histogram of the image intensity values is created.

If we assume that a histogram that uses only the available intensity levels is created then the histogram will have L bins. The histogram is a discrete function such that $\text{hist}\,(r_l) = n_l$ where r_l is the lth intensity value and n_l is the number of pixels in the image with intensity r_l and $l \in \mathbb{N}_0$ with $l = \{0 \ldots L-1\}$. In many instances it is beneficial to normalise the $\text{hist}\,(r_l)$ so that the values sum to one. This is done by dividing n_k by the total number of pixels in the image $X \times Y$ [29]. This creates a normalised histogram or probability distribution function (PDF) as

$$p\,(r_l) = \frac{n_l}{X \times Y} \quad \text{with}\, l = \{0 \ldots L-1\}. \tag{3.4}$$

The other benefit of describing the image in these terms allows intensities to be modelled as a random variables. Eq. 3.4 represents the probability of $I(x, y)$ having the intensity level of

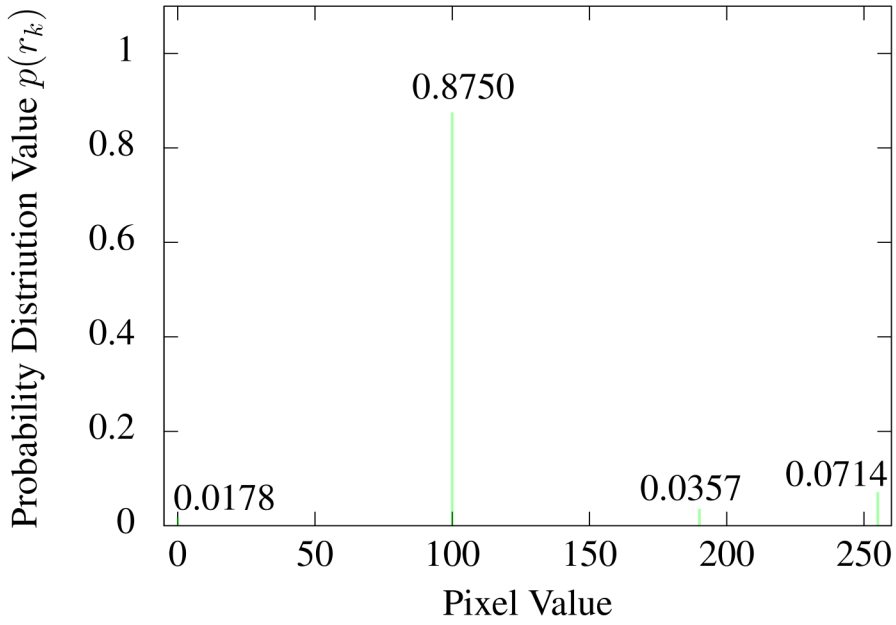

Figure 3.3: PDF of the sample image Figure 3.2. Notice that, as expected, the probability for 100 is the highest.

r_l. The PDF is drawn with r_l on the x axis and $p(r_l)$ on the y axis. This means that darker images tend to have a PDF concentrated to the left of the centre and lighter images have higher concentration of pixels to the right of the centre. Figure 3.3 shows what the PDF of Figure 3.2 would look like.

3.2.3 Sea Clutter Modelling

The purpose of an active SAR sensor is to self-illuminate a region using radar and to measure the return from various objects to build an image of the scene. In addition to the signals sent out from the radar, additional, unwanted signals may also be received. Noise refers to the unwanted alteration of the return signals by either an internal source (electronic components) or external source (thermal signal distortions, etc.). Clutter is a type of noise that is generated from additional signals from unwanted targets.

Clutter can come in a number of forms and includes returns from objects such as weather, buildings, birds, moving leaves and sea returns. Only clutter caused by sea returns is important as the focus is detecting ships on the sea. When observing objects on land most of

the returns are from stationary targets. The movement of the sea caused by waves causes a surface clutter within SAR images. Surface clutter itself is distorted by noise such as speckle in the image.

3.2.3.1 Why the CDF and PDF are discussed

The CDF and PDF are the two random value statistical functions which can be used to describe how a noise model describes a set of values matching a specific noise distribution. Specifically by having both the PDF and CDF it is possible to generate values that will fit the desired noise model. In the case of this dissertation the inverse of the CDF is used to generate the required noise distribution for the given CDF. The inverse CDF is evaluated using a randomly generated uniform distributed value and the output value from the inverse CDF will fit the PDF of the noise model. This technique is known as inverse transform sampling [30].

3.2.3.2 Rayleigh Distributed Sea Clutter

To understand the development of modern modelling techniques for sea clutter we need to understand the history of the field. In 1976, Goodman proposed the initial statistical model for clutter in single polarisation SAR imagery [31]. He stated that if the area illuminated in a single resolution cell within a SAR image was much larger than the wavelength of the signal, then clutter within each pixel could be represented using the super-position of random sub-scatterings within the resolution cell/pixel. If the number of sub-scatterings was sufficiently large and the amplitudes were all independent random variables then according to the central limit theorem the backscattering field within the resolution cell could be described using a Gaussian distribution. He went on to further show that the amplitude of the backscattering ends up being a Rayleigh distribution [31]. The Rayleigh distribution's PDF can be defined as

$$p\left(\mathbf{I}\,|\alpha_{\text{ray}}\right) = \frac{\mathbf{I}}{\alpha_{\text{ray}}} \exp\left(-\frac{\mathbf{I}^2}{2\alpha_{\text{ray}}}\right), \tag{3.5}$$

where $\mathbf{I} \geq 0$ and is the SAR image pixel intensities and $\alpha_{\text{ray}} > 0$ is the Rayleigh scale parameter. In much the same way, the Rayleigh distribution's cumulative distribution function (CDF) can be defined as

$$P\left(\mathbf{I}\,|\alpha_{\text{ray}}\right) = 1 - \exp\left(-\frac{\mathbf{I}^2}{2\alpha_{\text{ray}}}\right). \tag{3.6}$$

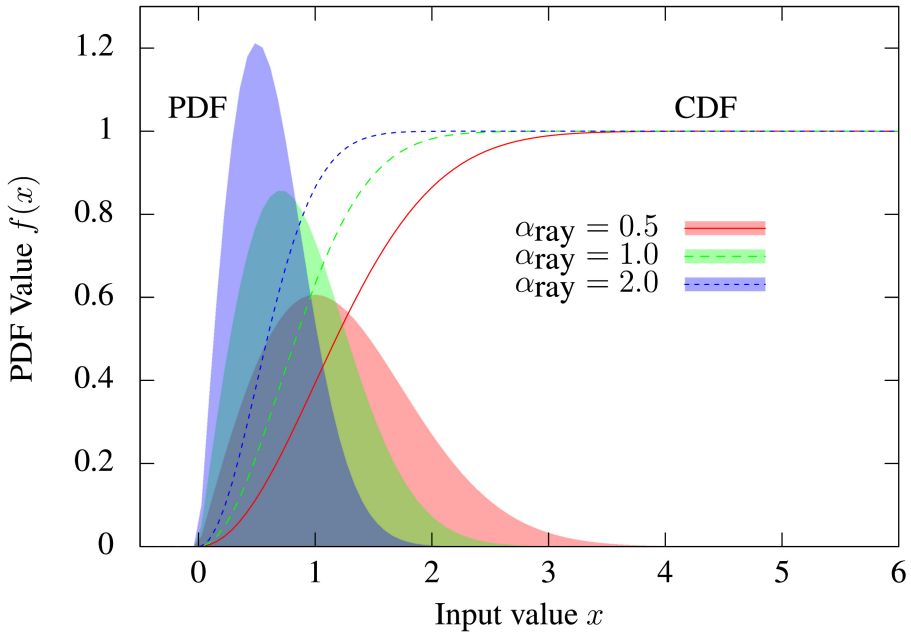

Figure 3.4: The cumulative and probability distribution functions for the Rayleigh distribution for various values of α_{ray}.

Figure 3.4 shows the effect of varying α_{ray} for both the CDF and PDF. By observing the graphs we can see that as α_{ray} is increased the PDF flattens out. This in turns causes the CDF rate of change to decrease and thus take longer to saturate. When α_{ray} is small it results in "spikey" noise (similar to salt-and-pepper noise) throughout the image whereas larger values of α_{ray} will result in a more distributed level of noise.

3.2.3.3 K-Distributed Sea Clutter

For a number of years it was believed that the Rayleigh distribution was a proper fit for SAR data. However, it has been shown that the distribution of clutter in high-resolution SAR imagery deviates from the the assumption of a Rayleigh distribution [32]. Statistical distribution models that have longer tails than the Rayleigh distribution have been proposed as better fits for SAR imaging [33].

The clutter found within SAR images relating to the sea has an intrinsic, rapidly varying speckle component as well as a more slowly varying intrinsic mean amplitude. The K-distribution can model both of these using a compound representation of the two compon-

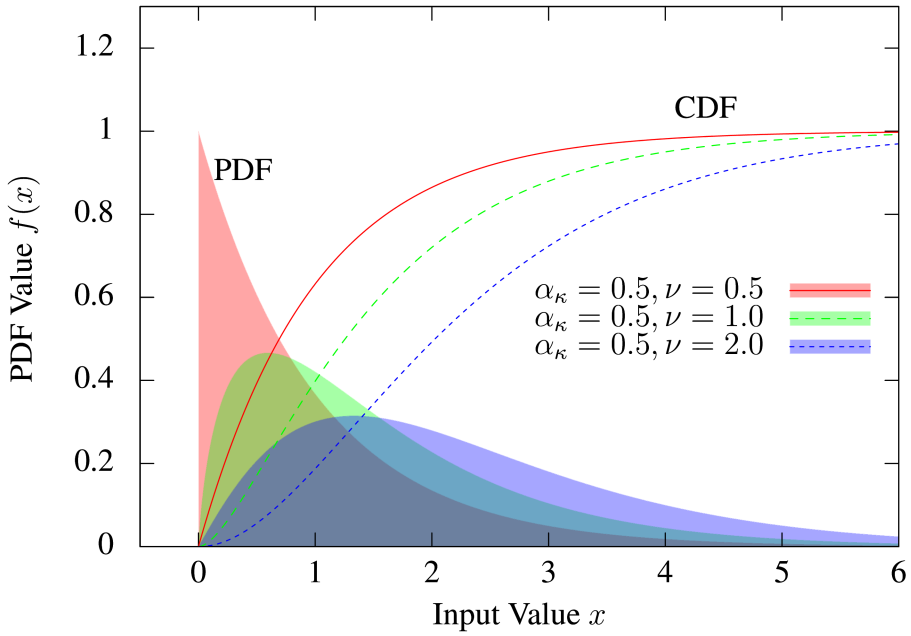

Figure 3.5: The cumulative and probability distribution functions for the K-distribution for $\alpha_\kappa = 0.5$ and various values of ν.

ents [34]. The K-distribution is the most frequently used distribution model for sea clutter in SAR imagery [32, 33, 35]. The PDF of the K-distribution is

$$p\left(\mathbf{I} \,|\, \alpha_\kappa, \nu\right) = \frac{2}{\alpha_\kappa \Gamma\left(\nu+1\right)} \left(\frac{\mathbf{I}}{2\alpha_\kappa}\right)^{\nu+1} \beta_\nu \left(\frac{\mathbf{I}}{\alpha_\kappa}\right), \tag{3.7}$$

where $\mathbf{I} \geq 0$ and is the SAR image pixel intensities, α_κ is the K-distribution scaling parameter, ν is the shape parameter with $\nu > -1$. Γ is the gamma function and β_ν is the modified Bessel function of the second kind with an order of ν. The K-distribution models the speckle using a Gaussian distribution which is modulated by the amplitude mean which itself is modelled using a gamma distribution.

The CDF of the K-distribution is given as

$$P\left(\mathbf{I} \,|\, \alpha_\kappa, \nu\right) = 1 - \frac{2\mathbf{I}^\nu \left(\frac{\nu}{\alpha_\kappa}\right)^{\frac{\nu}{2}} K_\nu \left(2\mathbf{I}\sqrt{\frac{\nu}{\alpha_\kappa}}\right)}{\Gamma\left(\nu\right)}. \tag{3.8}$$

Figures Figure 3.5, Figure 3.6 and Figure 3.7 all show the effects of varying both ν and α_κ on the CDF and PDF of the K-distribution. Between the three graphs $\alpha_\kappa = 0.5, 1.0$ and 2.0. Within each graph ν is varied as well between $\nu = 0.5, 1.0$ and 2.0. We can see that

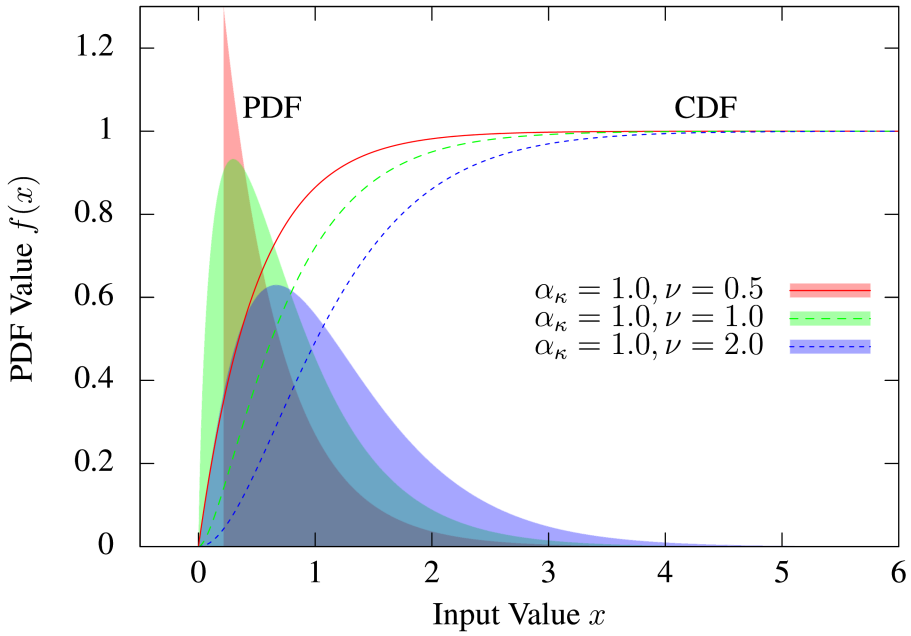

Figure 3.6: The cumulative and probability distribution functions for the K-distribution for $\alpha = 1.0$ and various values of ν.

in all three graphs ν has the property of extending the PDF tail of the distribution. Larger values of ν provide longer tails in the PDF which in turn decrease the rate of change of the CDF which causes the CDF to saturate to unity at higher values. Similar to the α_κ for the Rayleigh distribution, lower value of ν will cause more "spikey" noise. As α_κ is varied we can see a compression in the range of the PDF. For low values of α_κ, the PDF takes up a greater range of input values and is thus stretched out more than when α_κ is at a higher value. Increasing values of α_κ decrease the input range and thus cause the distribution maximum to increase. By increasing α_κ the rate of change of CDF increases. It is interesting to note that the parameters have inverse effects on the CDF. Increasing ν cause a decrease in the rate of change of the CDF whereas increasing α_κ causes an increase in the rate of change of the CDF.

To calculate the probability of detection in K-distribution clutter is either done by approximating the parameters or numerical integration [34]. These methods are required as both parameters have no closed form solution and require numerical, iterative methods to calculate. This means that, generally, either a Maximum-Likelihood (ML) or Method of Moments (MoM) approach have been used to estimate the parameters. Each of these methods have

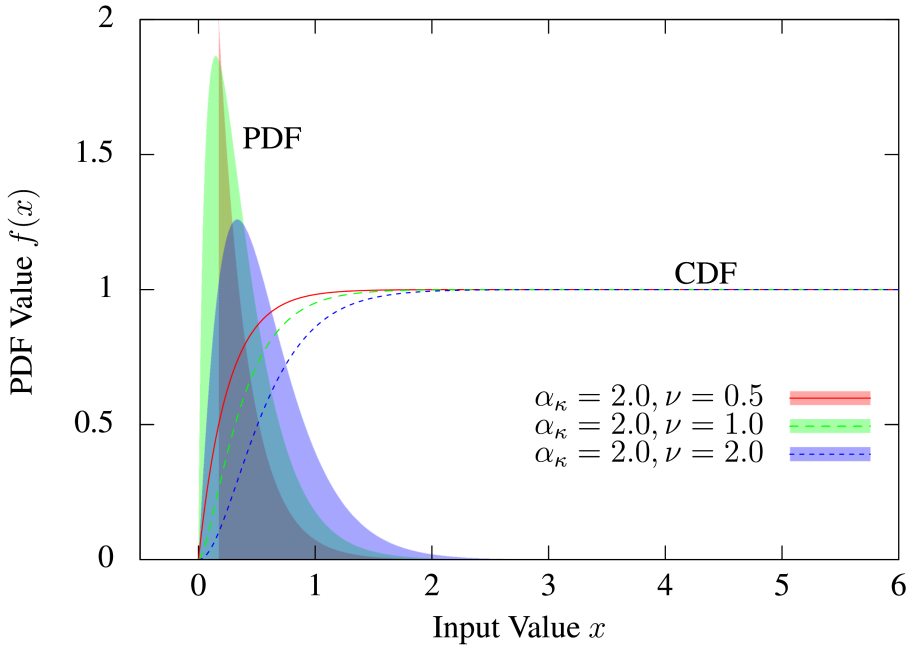

Figure 3.7: The cumulative and probability distribution functions for the K-distribution for $\alpha = 2.0$ and various values of ν.

their own advantages and disadvantages [36, 37, 38].

3.3 PRESCREENING METHODS

Ships within SAR images show up as pixels that are much brighter than the surrounding sea pixels. This is because, relative to the backscatter from the ocean clutter, the Radar Cross Section (RCS) of the ships is much higher and therefore ships reflect SAR signals much better than the ocean. The ship detection step of prescreening helps to further highlight possible ship pixels in SAR images. This section will detail three types of prescreening methods: Global, CFAR and Wavelet. The following sections will provide mathematical descriptions for each of these methods.

3.3.1 Global Image Thresholding

The global image thresholding prescreening method is a global, non-adaptive thresholding method that is applied to the entire SAR image [20]. To select pixels that are bright a pixel threshold value P_t is selected. All pixels that have intensities above this value are deemed as

ship pixels and all others are deemed non-ship pixels. The global thresholding method uses a SAR image as input and produces a binary image as output.

Assuming an input image $\mathbf{I}$ as defined before, the global thresholding method can be described as

$$\mathbf{J}(\mathbf{I}, P_t) = \left\{ \left\{ J(\mathbf{I}, x, y, P_t) \right\}_{x=0}^{x=X-1} \right\}_{y=0}^{y=Y-1}, \tag{3.9}$$

with

$$J(\mathbf{I}, x, y, P_t) = \begin{cases} \text{true} & \text{if } I(x,y) \geq P_t, \\ \text{false} & \text{otherwise.} \end{cases} \tag{3.10}$$

$\mathbf{J}$ is the global threshold binary result image for a given pixel threshold value P_t. $I(x,y)$ is the pixel under test with and its corresponding pixel ($J(x,y)$) is set to either true or false using equation 3.10.

3.3.2 Constant False Alarm Rate Prescreening

Global thresholding methods suffer from the drawback that a single threshold value applied to each pixel in the same way may not be sufficient to highlight pixels that are bright relative to its neighbours for all cases across a SAR image[4]. Due to geometry of SAR image acquisition, some areas will appear brighter than others within a SAR image. A global threshold method will highlight these areas as ship pixels if the threshold is set below the bright area's average pixel level. These are undesired effects of global thresholding methods and to combat this local, adaptive thresholding methods were created.

Constant False Alarm Rate (CFAR) prescreening is a local, adaptive thresholding-like method [39]. The CFAR prescreening method differs from global prescreening methods in that it uses a pixel's neighbours to determine if the pixel is bright instead of whether the pixel is bright relative to all other pixels in the image. The benefit of this is the method can highlight pixels that are bright relative to their neighbours but not necessarily bright relative to other pixels within the image. This also means that pixels that are ships and are not particularly bright may still be highlighted by the CFAR thresholding method because it is assumed that the ship pixels will be brighter than the sea clutter directly adjacent to it.

The Cell-Averaging CFAR (CA-CFAR) prescreening method is one of the versions of CFAR

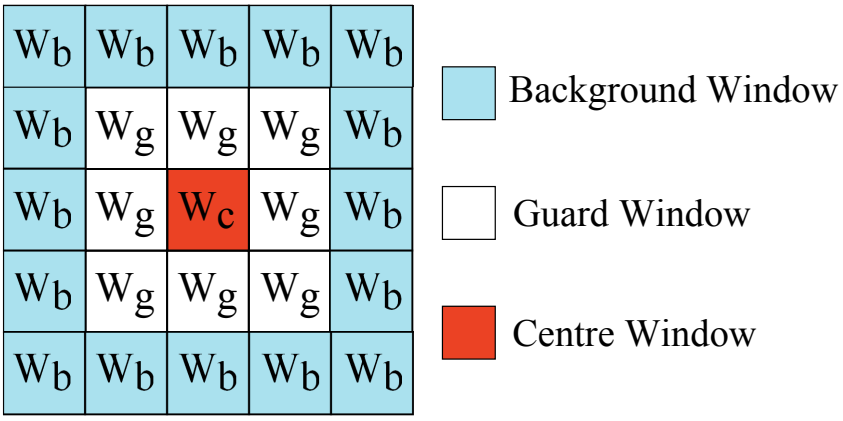

Figure 3.8: The three CFAR Windows used in the CA-CFAR method.

available and uses the mean value of the neighbouring pixels to determine if pixels are bright relative to their neighbours [4]. Central to the idea of CA-CFAR is what is known as "windows" relative to a centre window. Three windows are used throughout this study and the processing of SAR images: the centre, the guard and background windows. For an example of how the three types of windows please refer to Figure 3.8.

The centre window W_c is the pixel or pixels that are to be tested against the neighbours. The guard window W_g is the collection of pixels directly adjacent to the centre window. Pixels that surround the guard window are known as the background window W_b pixels. The background window is used to calculate a background mean and the guard window is used to ensure that this mean value is not affected by pixels that are closest to the pixel being tested (W_g removes the correlation between centre window pixels and background window pixels).

The CA-CFAR prescreening method uses a SAR image as input and produces a binary image $\mathbf{J}(\mathbf{I}, T)$ as output. Assuming an input image $\mathbf{I}$ with the sea clutter modelled using a K-distribution [4, 34, 38, 40] and image dimensions as defined before the CA-CFAR prescreening method can be defined as

$$\mathbf{J}(\mathbf{I}, T) = \left\{ \left\{ J(\mathbf{I}, x, y, T) \right\}_{x=0}^{x=X-1} \right\}_{y=0}^{y=Y-1}, \tag{3.11}$$

with

$$J(\mathbf{I}, x, y, T) = \begin{cases} \text{true} & \text{if } I(x,y) > \mu_b(x,y)\, T, \\ \text{false} & \text{otherwise.} \end{cases} \tag{3.12}$$

The output image $\mathbf{J}(\mathbf{I}, T)$ is evaluated at a threshold value T using input image $\mathbf{I}$. $I(x, y)$ is the pixel value located at pixel position (x, y) within $\mathbf{I}$. The variable $\mu_b(x, y)$ is the mean background value at pixel $I(x, y)$ and computed as

$$\mu_b(\mathbf{I}, x, y, S_b, S_g) = \frac{BG(\mathbf{I}, x, y, S_b) - GD(\mathbf{I}, x, y, S_g)}{S_b^2 - S_g^2}, \tag{3.13}$$

where S_b is known as the background square length such that $W_b = S_b \times S_b$ and similarly S_g is known as the guard square length such that $W_g = S_g \times S_g$ and $S_b > S_g$. BG is the background pixel value and GD is the guard pixel value and is defined as

$$\mathrm{BG}(\mathbf{I}, x, y, S_b) = \sum_{i=-\lfloor S_b/2 \rfloor}^{\lfloor S_b/2 \rfloor} \sum_{j=-\lfloor S_b/2 \rfloor}^{\lfloor S_b/2 \rfloor} I(x+i, y+j), \tag{3.14}$$

$$\mathrm{GD}(\mathbf{I}, x, y, S_b) = \sum_{i=-\lfloor S_g/2 \rfloor}^{\lfloor S_g/2 \rfloor} \sum_{j=-\lfloor S_g/2 \rfloor}^{\lfloor S_g/2 \rfloor} I(x+i, y+j), \tag{3.15}$$

with $S_g > 1$, $S_b > 1$ and both are odd and i and j are the horizontal and vertical indices respectively.

The CA-CFAR prescreening method has four main configurable variables namely: the centre window size, the background and guard window sizes and the CFAR threshold parameter T. A single pixel value for the window centre is assumed due to the coarseness ($75\,\mathrm{m} \times 75\,\mathrm{m}$) of the SAR imagery (ENVISAT ASAR Wide-Scan Mode imagery) that is available to test the system. The background and guard square lengths (window sizes) are chosen based on physical variables within the SAR image such as ground resolution of pixels and maximum expected ship size. The background window is usually chosen such that it is larger than the guard window and closely approximates the ocean's background statistics. The guard window size is selected such that it is larger than the maximum expected ship size which in turn is related to the SAR image resolution. As mentioned before, the selection of parameter T is the most complex aspect of using a CFAR thresholding method. The threshold parameter determines how many times brighter a pixel needs to be compared to its neighbours to be deemed a bright (ship) pixel. By using various threshold values it can be determined what is the best threshold for a given image.

3.3.3 Wavelet Thresholding

In digital image processing, a 2D image can be said to be composed of a number of signals that vary along the x- and y- axis. To correctly analyse these signals sometimes it is necessary to transform them from their basic form into another form. In many cases, when transforming images this way will reveal a great number of details that were previously unobservable.

Transformation allows for the capturing of both the low-resolution information (the repeating background signal) and the high resolution information (the individual, localised variations in the background). So how can a signal be converted and manipulated while keeping resolution across the entire signal and still be based in time? Wavelets are finite windows through which a signal is viewed. In order to move the window about the length of the signal, the wavelets can be translated about time in addition to being compressed and widened. The importance thereof is that the windows can be increased for a wide-eye view of the data (low resolution) as well as reduced down to capture details of the local changes (high-resolution). This allows SAR imagery to be processed to determine where details in local changes (ships) arise compared to low-resolution sea clutter (noise).

3.3.3.1 The Analysis Wavelet and the Continuous Wavelet Transform

Assuming a complex valued function, ψ called the mother wavelet is a well localised function. Furthermore, assume its Fourier transform is Ψ and it satisfies the following conditions [41]:

$$\int_{-\infty}^{\infty} \psi(t)\, dt = 0, \tag{3.16}$$

$$0 < c_\psi \equiv \int_{-\infty}^{\infty} |\Psi(t)|^2 \frac{d\omega}{|\omega|} < \infty, \tag{3.17}$$

where c_ψ is the admissibility constant. This mother wavelet ψ can be contracted, dilated and translated to form ψ_{cdt} by

$$\psi_{\text{cdt}} = \frac{1}{\sqrt{a}} \psi\left(\frac{t-b}{a}\right), \tag{3.18}$$

where a is known as the scale variable with $a \in \mathbb{R}_+$ and b is known as the time variable with $b \in \mathbb{R}$. This creates a continuous family of over-complete functions. This allows any time

signal to be represented as

$$s(t) = \frac{1}{c_\psi} \int_{-\infty}^{\infty} db \int_{0}^{\infty} S(b,a) \frac{1}{\sqrt{a}} \psi\left(\frac{t-b}{a}\right) \frac{da}{a^2}. \tag{3.19}$$

The function $S(b,a)$ is known as the coefficients of the wavelet transform signal and is defined as

$$S(b,a) = \int_{-\infty}^{\infty} \frac{1}{\sqrt{a}} \psi^*\left(\frac{t-b}{a}\right) s(t)\, dt, \tag{3.20}$$

where ψ^* is the complex conjugate of ψ. Equation 3.19 and equation 3.20 are known as the forward and reverse continuous wavelet transforms (CWT) for a continuous signal $s(t)$.

3.3.3.2 Discrete Wavelet Transform (DWT)

The above CWT can be converted into a Discrete Wavelet Transform (DWT) using the following conversion of the mother wavelet,

$$\psi_{m,n}(t) = a_0^{-m/2} \psi\left(\frac{t - nb_0}{a_0^m}\right), \tag{3.21}$$

where $m, n \in \mathbb{Z}$. The constants a_0 and b_0 are chosen such that the discretised a and b accurately represent the signal at the various scales and dilations [41]. Using the above equation the discrete wavelet transform coefficients $S(m,n)$ and its inverse transform are defined with

$$S(m,n) = \int_{-\infty}^{\infty} \psi^*_{m,n}(t)\, s(t)\, s(t)\, dt, \tag{3.22}$$

$$s(t) = o_\psi \sum_m \sum_n S(m,n)\, \psi_{m,n}(t), \tag{3.23}$$

where o_ψ is a constant value used for normalisation. The DWT converts the continuous time signal $s(t)$ into a discrete time signal and then extracts the DWT coefficients using equation 3.22. The original signal $s(t)$ is recovered using equation 3.23.

3.3.3.3 Haar Wavelets

There are a multitude of functions that can be used to define ψ to transform input signals. A common choice for ψ was proposed by A. Haar in 1909 [41], [29]. Assume φ defines a window

function such that

$$\varphi(t) = \begin{cases} 1 & \text{if } 0 \leq t \leq 1 \\ 0 & \text{otherwise} \end{cases}. \tag{3.24}$$

The mother wavelet function is defined in terms of the scaling function as $\psi(t) = \varphi(2t) - \varphi(2t-1)$ [41]. From this the Haar wavelet function can be described as

$$\psi(t) = \begin{cases} 1 & \text{if } 0 < t \leq \frac{1}{2} \\ -1 & \text{if } \frac{1}{2} < t \leq 1 \\ 0 & \text{otherwise} \end{cases}. \tag{3.25}$$

The combination of the Haar wavelet and Haar scaling functions is sufficient to perform multi-resolution analysis on a 1D signal. Finally the orthogonal wavelet basis is defined as $\psi_{m,n}(t) = 2^{-\frac{m}{2}}\psi(2^{-m}t - n)$ where $m, n \in \mathbb{Z}$.

3.3.3.4 Extending Haar Wavelets into Two Dimensions

To apply the above to SAR images, the 1D case (above) should be extended to 2D. To do so the scaling function must become a two-dimensional function and three, two-dimensional wavelet functions are required [29]. The scaling function given in equation 3.24 can be extended to two dimensions as

$$\varphi_{m,n}(x,y) = \varphi_{m,n}(x)\,\varphi_{m,n}(y). \tag{3.26}$$

The wavelet transformation can be similarly expanded to two dimensions using:

$$\psi^H_{m,n}(x,y) = \psi_{m,n}(x)\,\varphi_{m,n}(y), \tag{3.27}$$

$$\psi^V_{m,n}(x,y) = \varphi_{m,n}(x)\,\psi_{m,n}(y) \tag{3.28}$$

$$\psi^D_{m,n}(x,y) = \psi_{m,n}(x)\,\psi_{m,n}(y), \tag{3.29}$$

where a function dependant on x indicates the function works along the horizontal axis and a function dependent on y acts along the vertical axis. The H, V and D represent different wavelets for the Horizontal, Vertical and Diagonal directions respectively [29]. An important point to note is that the above equations are separable which allows the 2D DWT to be

implemented as the 1D DWT of the rows of an input image $\mathbf{I}$ and then perform another 1D DWT on the resulting columns [29].

Separability allows for translation and scaled basis functions to be defined as:

$$\varphi_{j,m,n}(x,y) = 2^{j/2}\varphi_{m,n}\left(2^j x - m, 2^j y - n\right), \tag{3.30}$$

$$\psi^H_{j,m,n}(x,y) = 2^{j/2}\psi_{m,n}\left(2^j x - m, 2^j y - n\right), \tag{3.31}$$

$$\psi^V_{j,m,n}(x,y) = 2^{j/2}\psi_{m,n}\left(2^j x - m, 2^j y - n\right), \tag{3.32}$$

$$\psi^D_{j,m,n}(x,y) = 2^{j/2}\varphi_{m,n}\left(2^j x - m, 2^j y - n\right), \tag{3.33}$$

where j is the chosen image scale and is selected such that $j \geq j_0$ where J_0 is an appropriately selected scale value. [29].

The 2D DWT of an input image $\mathbf{I}$ with dimensions $X \times Y$ is computed using

$$W_\varphi(j_0,m,n) = \mathbf{W}_{\varphi,j_0,m,n} = \frac{1}{\sqrt{XY}}\sum_{x=0}^{X-1}\sum_{y=0}^{Y-1}\mathbf{I}\cdot\varphi_{j_0,m,n}(x,y), \tag{3.34}$$

$$W^H_\psi(j,m,n) = \mathbf{W}^H_{\psi,j,m,n} = \frac{1}{\sqrt{XY}}\sum_{x=0}^{X-1}\sum_{y=0}^{Y-1}\mathbf{I}\cdot\psi^H_{j,m,n}(x,y), \tag{3.35}$$

$$W^V_\psi(j,m,n) = \mathbf{W}^V_{\psi,j,m,n} = \frac{1}{\sqrt{XY}}\sum_{x=0}^{X-1}\sum_{y=0}^{Y-1}\mathbf{I}\cdot\psi^V_{j,m,n}(x,y), \tag{3.36}$$

and

$$W^D_\psi(j,m,n) = \mathbf{W}^D_{\psi,j,m,n} = \frac{1}{\sqrt{XY}}\sum_{x=0}^{X-1}\sum_{y=0}^{Y-1}\mathbf{I}\cdot\psi^D_{j,m,n}(x,y), \tag{3.37}$$

where $n = m = 0, 1, 2, \ldots, 2^j - 1$ ([29]) and $\mathbf{W}_{\varphi,j_0,m,n}$ is a different representation of the coefficients $W_\varphi(j_0,m,n)$ chosen to match the other image notations used throughout this dissertation. Details regarding the reverse 2D DWT have been omitted because only the forward transform is required for the wavelet prescreening method but details about it can be found in [29].

3.3.3.5 Wavelet prescreening

Wavelet prescreening requires the SAR image to be processed at some scale $j+1$ (the original image) and then the four other images $\mathbf{W}_\varphi, \mathbf{W}^H_\psi, \mathbf{W}^V_\psi$ and $\mathbf{W}^D_\psi$ at scale j be spatially multiplied to obtain the processed wavelet image . This can be accomplished by taking the

dot product of each image

$$\mathbf{R} = \mathbf{W}_{\varphi} \cdot \mathbf{W}_{\psi}^{H} \cdot \mathbf{W}_{\psi}^{V} \cdot \mathbf{W}_{\psi}^{D}. \tag{3.38}$$

The resultant image $\mathbf{R}$ contains high values for points that are correlated between the images (and scales) whilst noise (which de-correlates differently for each of the functions) will have lower values. The ship pixels are enhanced whilst the Rayleigh sea noise (the chosen sea clutter distribution assumed in [24]) is reduced in $\mathbf{R}$. To detect the ship pixels, the prescreening procedure applied to $\mathbf{R}$ to generate a binary result image $\mathbf{J}$ using

$$\mathbf{J}(\mathbf{R}, c) = \left\{ \left\{ J(\mathbf{R}, x, y, c) \right\}_{x=0}^{x=X-1} \right\}_{y=0}^{y=Y-1}, \tag{3.39}$$

with

$$J(\mathbf{R}, x, y, c) = \begin{cases} \text{true} & \text{if } R\,(x, y) \geq \mu_{\mathbf{R}} + c \cdot \sigma_{\mathbf{R}}, \\ \text{false} & \text{otherwise.} \end{cases} \tag{3.40}$$

The two statistical values $\mu_{\mathbf{R}}$ and $\sigma_{\mathbf{R}}$ are the mean and standard deviation values calculated from $\mathbf{R}$. The value c is an adjustable parameter that determines when a pixel is deemed a "ship" pixel (similar to T for CA-CFAR).

3.4 SHIP DISCRIMINATION

The final step of the ship detection system is to generate a list of ship pixels positions SP by grouping pixels that are nearby one another in the resulting binary CA-CFAR image $\mathbf{J}$. To equally represent each ship so that comparisons to ground truth data are fair, each group of ship pixels is reduced to a single, centre pixel ("centre of mass"). There exists a number of methods to group pixels and find the corresponding centre pixel including K-means and Agglomerative Hierarchal Clustering [42], Resonance Agglomeration [27] or Morphological Closing and Connect Component Analysis [29], [43] and Mean Shift Clustering [44]. As the number of ships (and therefore clusters) are not known ahead of time, the clustering methods such as K-means and Agglomerative Clustering methods which require the number of clusters be set before processing the data, is not tested here.

The two methods for ship discrimination discussed in this section are the a) Morphological Closing and Connect Component Analysis and b) Mean Shift Clustering method for grouping pixels and finding ship centres. The Morphological Closing and Connect Component Ana-

lysis ship discrimination is grounded in digital image processing theory whereas Mean Shift Clustering takes a statistical, pattern recognition approach to clustering pixels.

3.4.0.6 Why finding the pixel centres is important

Finding the centres of the pixel groups is important for the evaluation of the ship detection system. The most important reason for applying an additional filtering step after the prescreening step is to reduce the number of false alarms and increase the detection accuracy. By grouping large groups of pixels into single pixel representations large numbers of falsely detected pixels are 'removed'. Furthermore, due to the fact that some ships are larger than others, reducing each ship detection to a single pixel ensures that all ships are weighted evenly in the validation stage. When compared to one another each ship therefore has the same representation and will equally affect the performance of the system. Chaining the prescreening and ship discrimination stages ensures that the number of false alarms is reduced and thus improves performance of the system.

3.4.1 Connected Component Ship Discrimination

The Morphological Closing and Connect component analysis is a form of Resonance Agglomeration and its usage requires the selection of a grouping size g_p. A grouping size g_p is the maximum allowable distance between two pixels to be considered for merging into a single group. If two pixels are greater than g_p pixels from one another then they are considered two separate groups. It can be assumed that due to the coarseness of the testing data if $g_p = S_g$ we ensure that only pixels that are less than the maximum expected ship size are grouped together.

A square structuring element D of size g_p is used in a morphological closing operation to group all pixels that are less than g_p pixels from each other. The closing operation on prescreening result image $\mathbf{J}$ can be described as

$$\mathbf{C}(\mathbf{J}) = (\mathbf{J} \oplus D) \ominus D, \tag{3.41}$$

where $\mathbf{C}(\mathbf{J})$ is the resulting morphologically closed image. The morphological closing operation connects the gaps between pixels that are g_p pixels from one another. This effectively

merges all pixels that are g_p pixels apart and forms Ω groups within the the image which denotes the number of ships detected. The operators $\oplus$ and $\ominus$ represent morphological dilation and erosion, respectively [29].

The next step of the ship discrimination is to create a list containing all the centres of the ships. If the list of connected components is denoted by V_ξ, where $\xi = 1, 2, 3, \ldots, \Omega$, then connected component analysis of $\mathbf{C}$ can be done iteratively as

$$V_\xi = (V_{\xi-1} \oplus D_8) \cap \mathbf{C}, \tag{3.42}$$

where D_8 is the 8-connectivity structuring element [29], V_ξ is a set of all the connected component regions within $\mathbf{C}$ [29]. The process of clustering the pixels terminates when equation 3.42 produces the same amount of groups ($V_\xi = V_{\xi-1}$).

Calculating the ship centroids requires identifying the location of each connected component in V_ξ. This is accomplished by computing the area of the ξth pixel group by incrementing a counter for each pixel [43]

$$A_\xi = \sum_{(x,y)\in V_\xi} 1, \tag{3.43}$$

where A_ξ is the area for the ξth connected component.

Once the area has been found it can be used to calculate the centroid coordinates of each pixel group. For every pixel in each connected component, the summation of the x and y coordinates respectively divided by the area will provide the centroid coordinate in each direction (discussed extensively in [43, 45]). This is expressed as

$$x_\xi^{\text{centre}} = \left\lfloor \frac{1}{A_\xi} \sum_{(x,y)\in V_\xi} x \right\rfloor, \tag{3.44}$$

$$y_\xi^{\text{centre}} = \left\lfloor \frac{1}{A_\xi} \sum_{(x,y)\in V_\xi} y \right\rfloor, \tag{3.45}$$

where x_ξ^{centre} and y_ξ^{centre} are the centroid x and y coordinates, respectively, for the ξth connected component. The above procedure generates a list of ship centre position coordinates, $\text{SP} = \left\{ \{x_\xi^{\text{centre}}, y_\xi^{\text{centre}}\} \right\}_{\xi=1}^{\xi=\Omega}$.

3.4.2 Mean Shift Clustering Ship Discrimination

3.4.2.1 Mean Shift Clustering Mathematical Basis

Mean Shift Clustering is a nonparametric estimation technique used to analyse multimodal data space [44]. The algorithm does not require a fixed number of clusters be known *a priori* and the shape of the cluster can be arbitrary.

Assume a single centre point or kernel window centre $\boldsymbol{\rho} = (\rho_1, \rho_2, \ldots, \rho_d)$ where d indicates the number of dimensions in a $\mathbb{R}^d$ space. In this case $d = 2$ for the prescreened result binary image $\mathbf{J}$. Furthermore, assume a collection of n data points $\boldsymbol{\rho}_i, i = 1, 2, \ldots, n$ such that $\boldsymbol{\rho}_1 = J(0,0), \boldsymbol{\rho}_2 = J(0,1), \ldots, \boldsymbol{\rho}_n = J(X-1, Y-1)$ with $n = X \times Y$. The multivariate kernel density estimation $\hat{f}(\boldsymbol{\rho})$ is

$$\hat{f}(\boldsymbol{\rho}) = \frac{1}{nh^d} \sum_{i=1}^{n} K\left(\frac{\boldsymbol{\rho} - \boldsymbol{\rho}_i}{h}\right). \tag{3.46}$$

$K(\boldsymbol{\rho})$ is the kernel window at some $\boldsymbol{\rho}$ and h is the bandwidth or window radius used to determine how far apart points are to be grouped together. For the purpose of Mean Shift Clustering a specific class of radially symmetric kernels is chosen such that

$$K^S(\boldsymbol{\rho}) = K(\boldsymbol{\rho}) = c_{k,d} k\left(\|\boldsymbol{\rho}\|^2\right), \tag{3.47}$$

where $k(\rho)$ is called the profile of the kernel when $\rho \geq 0$ and $c_{k,d}$ is assumed strictly positive and ensures $\int K(\boldsymbol{\rho})\, d\rho = 1$ [44].

The local minima or modes of the density are stationary points where $\nabla \hat{f}(\boldsymbol{\rho}) = 0$. Assuming a profile function $g(\rho) = -k'(\rho)$ the gradient of $\hat{f}(\boldsymbol{\rho})$ can be written as

$$\nabla \hat{f}(\boldsymbol{\rho}) = \frac{2c_{k,d}}{nh^{d+2}} \sum_{i=1}^{n} (\boldsymbol{\rho} - \boldsymbol{\rho}_i)\, g\left(\|\frac{\boldsymbol{\rho} - \boldsymbol{\rho}_i}{h}\|^2\right), \tag{3.48}$$

$$= \frac{2c_{k,d}}{nh^{d+2}} \left[\sum_{i=1}^{n} g\left(\|\frac{\boldsymbol{\rho} - \boldsymbol{\rho}_i}{h}\|^2\right)\right] \left[\frac{\sum_{i=1}^{n} \boldsymbol{\rho}_i g\left(\|\frac{\boldsymbol{\rho} - \boldsymbol{\rho}_i}{h}\|^2\right)}{\sum_{i=1}^{n} g\left(\|\frac{\boldsymbol{\rho} - \boldsymbol{\rho}_i}{h}\|^2\right)} - \boldsymbol{\rho}\right]. \tag{3.49}$$

The density estimate at $\boldsymbol{\rho}$ can be computed using a new kernel $G(\boldsymbol{\rho})$ derived from profile $g(\rho)$. This gives the first term in square brackets. The second term in square brackets is known as the mean shift $\mathbf{m}(\boldsymbol{\rho})$ and can be written as a function of h if

$$\mathbf{m}_h(\boldsymbol{\rho}) = \frac{\sum_{i=1}^{n} \boldsymbol{\rho}_i g\left(\|\frac{\boldsymbol{\rho} - \boldsymbol{\rho}_i}{h}\|^2\right)}{\sum_{i=1}^{n} g\left(\|\frac{\boldsymbol{\rho} - \boldsymbol{\rho}_i}{h}\|^2\right)} - \boldsymbol{\rho}. \tag{3.50}$$

The vector $\mathbf{m}_h(\boldsymbol{\rho})$ is the difference between the weighted mean of the points using kernel G for weights and the centre of the kernel window $\boldsymbol{\rho}$. This vector always points in the direction of which the density increase is maximum [44].

3.4.2.2 Mean Shift Clustering Algorithm

The procedure for finding the mean shift is guaranteed to converge to a point where the gradient of the density function is zero [44]. The procedure to do so for a single point $\boldsymbol{\rho}^t$ is

- Compute the mean shift vector $\mathbf{m}_h(\boldsymbol{\rho}^t)$.
- Shift the vector using the computed mean shift vector such that $\boldsymbol{\rho}^{t+1} = \boldsymbol{\rho}^t + \mathbf{m}_h(\boldsymbol{\rho}^t)$

Each mode will have a set of points that converge to that mode. This is known as the basin of attraction for that mode [44]. The points with the same basin of attraction are all associated with a single mode and thus each data point is associated with a single mode. This mode is the "centre" of mass for that cluster and all points with the same mode are in the same cluster.

The procedure above is repeated until the point $\boldsymbol{\rho}^t$ converges to a single mode. This is repeated for each of the n points in $\boldsymbol{\rho}_i$ so that each of the data points is associated with a single mode. In the case of applying the mean shift clustering procedure to $\mathbf{J}$, the collection of modes are the ship centres and all pixels in a single cluster represent a single ship. The bandwidth parameter h determines how nearby ship pixels need to be in order to be clustered together to form a single ship cluster and centre of mass (ship centre).

CHAPTER 4

RESULTS AND DISCUSSION

4.1 OVERVIEW

This chapter will present experiments that test the result of applying the selected methods (discussed in Chapter 3) to SAR imagery. The chapter is split into sections that each describe experiments to test the components of the ship detection system. The first section details the images in the experiments. Following this, the system components are first tested separately and then tested together to determine the performance of the system. Following each test, a short discussion about the results is given.

As an additional note it should be stated that when a method is referred to as a thresholding method it means that the method is used to select a thresholding value against which every pixel in the image is compared. So a prescreening method using CFAR as the threshold selection criterion is referred to as CFAR thresholding or CFAR prescreening for the sake of brevity.

Finally, the theory of detection is an important section of any detection system evaluation. For the sake of similarity to other works in the field, two main metrics were used: Detection Accuracy and False Alarm Rate. The theory of detection is well developed and is more than extensive enough to have its own section. Again for the sake of brevity it is assumed that the reader is familiar with the criterion used herein and the author refers the reader to [46] as the defining RADAR theory of detection paper. Furthermore, see [4] for a discussion on the connection between classic terms and those more commonly used in ship detection literature.

4.2 DATA DESCRIPTION

4.2.1 Synthetic SAR imagery

The synthetic imagery for this dissertation was generated in two steps. The first step generated a specified number of "ships" in the image. Ships appeared as randomly generated bright pixels, grouped together, with pixel values in the range of 33 to 142 with a mean pixel value of 88. Furthermore, each of the ships were represented in the SAR image as a randomly orientated collection of two to five pixels long. The motivation for the above assumed values is discussed in section. 4.3.2.1.

Once the "ships" had been generated, a certain level of noise was added to the image either specified by a required Signal-to-Noise (SNR) value or the parameters for the distribution. As discussed previously, two types of noise models for the sea were used in this study, namely: the Rayleigh distribution and K-distribution.

4.2.2 Real SAR imagery

Six SAR images off South Africa's coast were acquired for the experiments in this study. The six images were all acquired from the ENVISAT ASAR sensor in WSM sensor mode. Each image had a ground resolution of 75m $\times$75m. As previously discussed, only the intensity (real) part of the SAR data was used as a representation of each SAR image. Each of the six images had a VV polarisation. It should be noted that whilst VV is not the optimal polarisation for ship detection, it can be used as a baseline performance indicator for an operational system.

Each of the SAR images had an associated ship centre position image. Each centre position image was hand-created by scanning through each ASAR image and marking the centre position of each ship found. These binary positional images were used to compare the actual ship positions to the positions generated by the ship detection system and determine the system's performance. The six images and more details about each of them is provided in Table 4.1.

Table 4.1: Information about the real SAR images used in this dissertation.

Satellite	Acquisition date	Image size	Ground resolution (m)	Polarisation	Number of Ships
Image 1	2012/02/18	13360 × 19770	75 × 75	VV	17
Image 2	2012/02/22	9135 × 10796	75 × 75	VV	43
Image 3	2012/03/03	8769 × 9172	75 × 75	VV	21
Image 4	2012/03/08	12819 × 17695	75 × 75	VV	18
Image 5	2012/03/18	9818 × 9166	75 × 75	VV	2
Image 6	2012/03/27	13554 × 19771	75 × 75	VV	41

4.3 PARAMETER SELECTION MOTIVATION

A number of the methods used in ship detection requires a parameter to be set in order to optimise the Detection Accuracy (DA) and False Alarm Rate (FAR) of a system. Setting these parameters is one of the most complex aspects of working with SAR-based ship detection systems. This section will use a experiment to derive appropriate parameters to be used in this study.

DA is the total number of correctly detected ships divided by the total number of known ships in the image. The DA ranges from 0 to 1 where a DA of 0 represents no corresponding ships were detected and 1 represents perfect correspondence between the detected ships and actual ships in the image. The second metric, FAR, is the total number of falsely detected ships divided by the total number of pixels known to not contain any ships. The false alarm rate ranges from 0 to 1 with a FAR of 0 representing no false alarms and 1 representing only false alarms.

It is prudent to discuss the general philosophy of parameter selection used here. Parameters are selected such that they reduce the DA and FAR of a system and are selected on a per-situational basis. A single parameter will not work equally well for two different system configurations and two different data/image inputs [4]. For the sake of this study, a reasonable range of values is one which both encompass what the limits of the method are and is

a range that can typically be expected in practice. It is also assumed that the methods will produce results that behave monotonically, deterministically and linearly within this range. This is a reasonable assumption because none of the methods tested here have any stochastic elements and on a per-image basis the system behaviour is deterministic.

To understand what is meant exactly in this study by a reasonable assumption we can take a look at the prescreening using CA-CFAR. For this prescreening method it would make little practical sense to choose the threshold value $T = 0$ or $T > 255$ (for 8-bit intensity images). This is because the method computes the mean μ of a neighbourhood of pixels and then tests if the centre pixel is T times brighter than μ. Selecting $T = 0$ would mean that every pixel is automatically T times brighter than the mean pixel value (except in fringe cases where all pixels are zero) and similarly assuming $T > 255$ would ensure that none of the pixels are selected because the minimum pixel value is only 255 times less than the maximum and irrespective of μ. So clearly choosing T as either of these choices is not a reasonable choice for this parameter. The reasonable range in this case would be $0 < T < 255$. Similar arguments for other threshold selection methods can also be made. The monotonic, deterministic (non-stochastic) nature of the methods used means it is sufficient to test a few of the sample points from the range of possible threshold values to characterise the overall performance (and limits) of the system.

4.3.1 Using SAR images to determine synthetic imagery generation parameters

The synthetic SAR generator has a number of variables which account for various ship sizes, orientations, ship pixel intensities, noise levels and noise types. Varying these the researcher can generate synthetic SAR images which are statistically similar to actual ASAR images.

For this experiment a pixel index vs. pixel value graph was generated to help guide the selection of parameters. This graph was selected instead of a conventional PDF graph because the PDF graph discards the spatial information associated with an image and only describes the pixel value information. Using a pixel index vs. pixel value graph allowed the researcher to get a sense of the pixel intensity distribution over the image.

4.3.1.1 ASAR ship area statistics

4.3.1.2 Data description

The imagery used for this experiment included the six ASAR ENVISAT images and the six corresponding images indicating the various ship centre locations.

4.3.1.3 Experimental procedure

The experimental procedure was as follows:

1. A 25×25 sub-image is extracted for each of the ship centres in each of the six images. A total of 142 ships are identified across the six images.

2. Determine the mean μ and standard deviation σ of the 142 centre pixel values.

3. Add each of the 142 sub-images together and divide by 142 to get a "mean" sub-image of size 25×25.

4. Each pixel in the image will have an index $i = 0, 1, \cdots, (25 \times 25 - 1)$. Plot each pixel value versus its pixel index.

5. Pixels are visually subdivided into one of three groups based on pixel value: high, medium or low.

6. Threshold the mean image using the low, medium and high pixel value lines to create three 25×25 binary images.

7. Compute three SNRs by comparing the mean sub-image ("input image") to the three binary image(s) ("reference image").

4.3.1.4 Experimental results

Figure 4.1 presents a graph that shows the pixel indices, $i = 0, 1, 2 \ldots, (25 \times 25 - 1)$ for a 25×25 sub-image, versus the pixel value and three lines which indicate whether pixels fall in the low, medium or high pixel intensities. Figure 4.2 shows four images: the top left image

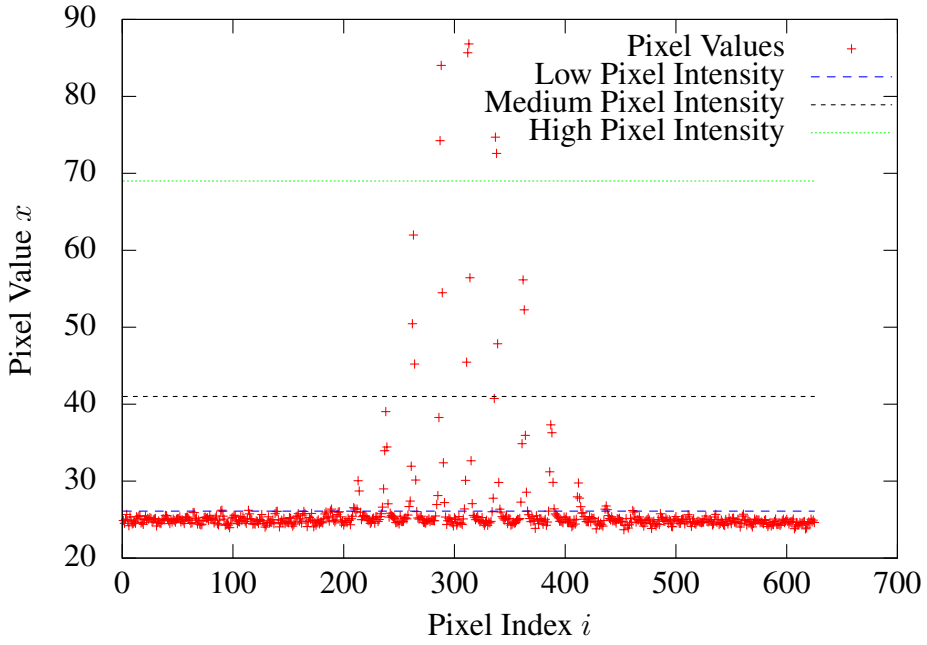

Figure 4.1: Graph of pixel index vs. pixel value in mean sub-image. Notice the three distinct groups of pixel values: low, medium and high. Pixels near the mean pixel value $\mu_{\text{sub-image}}$ are considered background pixels.

is the mean sub-image and the the three images from top right to bottom left indicate the low, medium and high pixel values respectively.

A pixel is a low value pixel if it falls above the low pixel intensity line. All low pixels are marked true and and all other false in the second image in Figure 4.2. The same is repeated for the other lines with all pixels above a line marked true and below it marked as false. This creates the three binary images shown in Figure 4.2. That is the pixels were manually marked using visual inspection of the pixel index vs. pixel value graph.

Looking at the sub-image statistics approximately 94.4% of the pixels are within 10% of the background mean of $\mu_{\text{sub-image}} = 25$ for the sub-image. Using only the centre pixel, the mean value is approximately $\mu_{\text{centre}} = 87.5$ with a standard deviation of +/- 54.57. Three noise-levels are when comparing the binary images to the mean sub-image: SNR $= -5.77\,\text{dB}, -6.54\,\text{dB}$ and $-8.56\,\text{dB}$ for the low, medium and high pixel value binary images respectively.

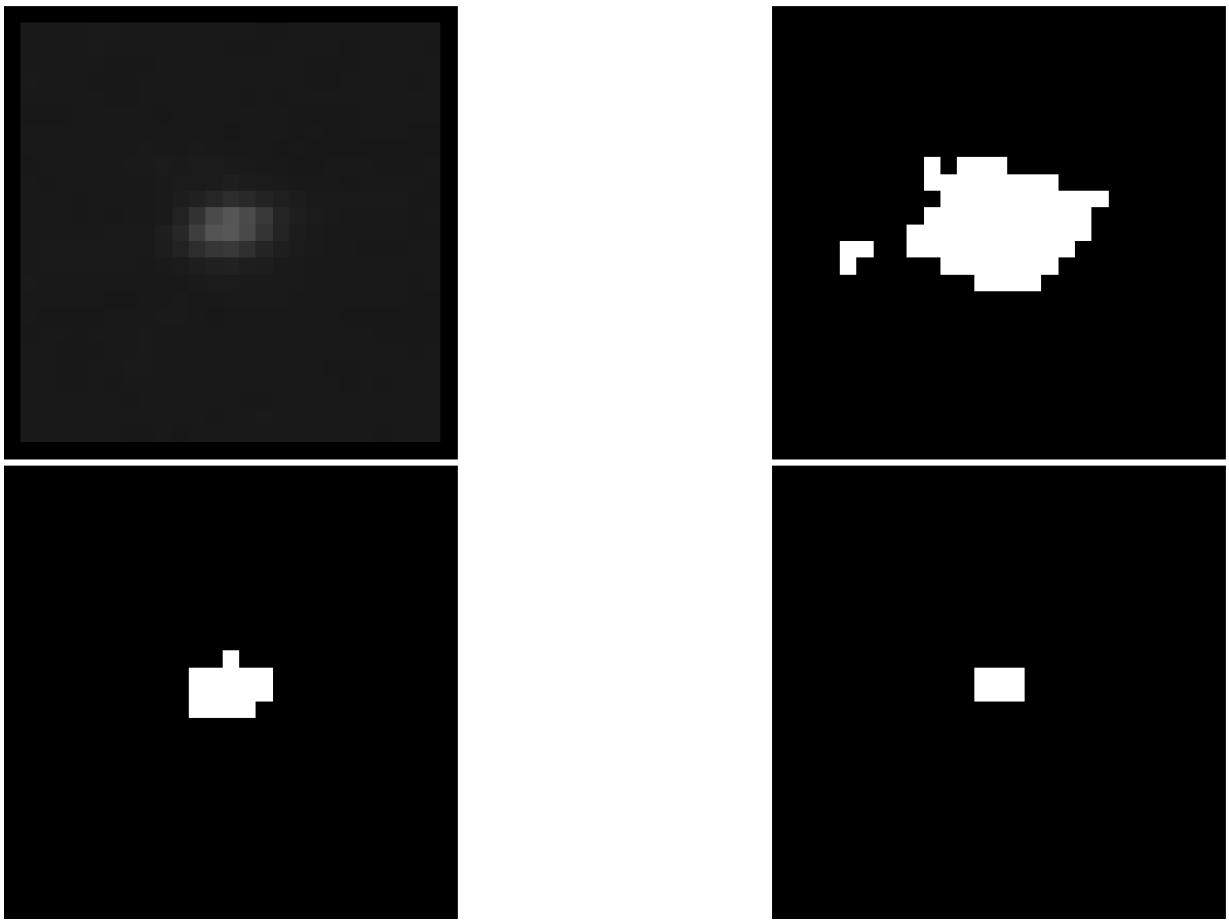

Figure 4.2: Moving from top left to bottom right: mean sub-image, low pixel thresholded image, medium pixel thresholded image and high pixel thresholded image. The black border in the mean sub-image is added for clarity.

4.3.1.5 Results discussion

To start of the discussion it was assumed that the sub-image size of 25×25 was sufficient to encapsulate any orientated ship and size present for the 142 ships in the six ASAR ENVISAT images. It was also assumed that each sub-image only had one ship present. Figure 4.1 indicates that there are three groups of pixel intensities: a low group (near $I(x, y) \approx 25$), a medium group ($I(x, y) > 41$) and the high pixel group ($I(x, y) > 69$). As expected, the centre pixel has the highest value $I(x, y) = 87.5$ because it is the pixel that is certain to be in all 142 sub-images and thus its average value among the 142 sub-images is the highest.

Examining Figure 4.2 reveals the pixels that are 1 pixel away from the centre are highly likely

to be high values because of the restricted number of orientations for ships and the implied ship size. Another point to note from the mean sub-image is that most of the high value pixels lie in the horizontal direction. This implies that for the six ENVISAT ASAR images, most of the ships lie horizontally or diagonally with a greater horizontal component. The top right image in Figure 4.2 indicates that all pixels in the low group and up are once again placed very much in the centre. Due to the fact that this image highlights pixels near to the mean it is insufficient to be used to make statements other than the fact that all of the pixels above the mean are located nearer to the centre pixel than the edges.

From the bottom left image in Figure 4.2 we can infer the expected dimensions and orientations of a ship in the imagery. The maximum expected length of a ship in any direction is five pixels. This is an agreeable result because the longest ship in service at the time when the ASAR images were taken was the 380m long "TI Africa" oil tanker [47]. This ship translates to approximately five ($380 \div 75 = 5.066$) pixels long in an ENVISAT ASAR WSM image. The assumption of ships being at least two pixels long is validated in final image in Figure 4.2. In this binary image the vertical direction has a two pixel length and three pixel in the horizontal direction indicating that the minimum ship size with very high correlations between the sub-images is two pixels long.

This experiment aimed to investigate what the expected variables found in ASAR imagery off the coast of South Africa might be. The mean image allowed for a reasonably clear picture of what ships look like in ENVISAT ASAR imagery. The experiment helped define synthetic SAR image generation parameters such as ship pixel values, orientations and sizes as well what level of noise surrounded a ship. A point to note is that if a historic collection of ship sub-images were to be gathered for a specific set of SAR imagery, one could effectively use this information as *a priori* knowledge about a typical ship found in that type of SAR imagery. This could allow for the automatic selection of parameters such as guard and background window sizes based on analysis of the mean sub-image. Over time the mean sub-image would better describe what is the expect ship shape and description as more and more samples were added to the collection. Furthermore, this would allow a system to dynamically adapt the parameters as the ships found in the input images vary over time.

4.3.2 Validation tests

In this section a test is performed in order to verify the validity of the assumptions made in the previous section using synthetic and actual SAR images.

Two synthetic images were generated for the Rayleigh and K-Distributions respectively. The parameters used for each were: image noise level of SNR $= -6.54\,\text{dB}$, ship pixel values in the range $33 \leq I(x, y) \leq 142$, maximum ship length of 5 pixels with a random orientation and image size of 1000×1000. A Receiver Operating Characteristic (ROC) curve analysis was drawn for the two synthetic images and six ASAR ENVISAT images using three prescreening methods namely: Global, CFAR and Wavelet thresholding. For more details about ROC curves please see appendix A.1. The guard and background window sizes for the CFAR prescreening method were 7×7 and 9×9 respectively. The results of the various ROC analyses can be found in Figure 4.3, 4.4 and 4.5.

Figure 4.3, 4.4 and 4.5 has increasing values of threshold for each prescreening method are top right to bottom left (threshold decrease from left to right). For each method in each graph, the ROC curve attains 100% DA and 100% FAR when the threshold is set to its lowest 0. This is an expected because such a low threshold highlights every pixel in the image regardless of pixel value. Similarly when the threshold value is 256 (bottom left) then all of the methods fail to register a single detection and both the DA and FAR are 0%.

Next it should be noted that while the shapes of the various ROC curves between the synthetic image and actual image results are similar they are not identical. This would require a more complex synthetic image generation system to achieve synthetic images that are almost identical to actual SAR images. However, the trends for the synthetic image results agree very well with the trends actual image results. This indicates that the synthetic imagery is at least sufficient for initial testing and is a good enough indicator of the expected performance of the system.

Across all three of the graphs, the CFAR prescreening method has the largest Area-Under-Curve (AUC). This denotes the method is empirically better than the other two even without any form of ship discrimination, irrespective of image input (synthetic or real). The K-distribution is handled better by the CFAR prescreening method because the method is built

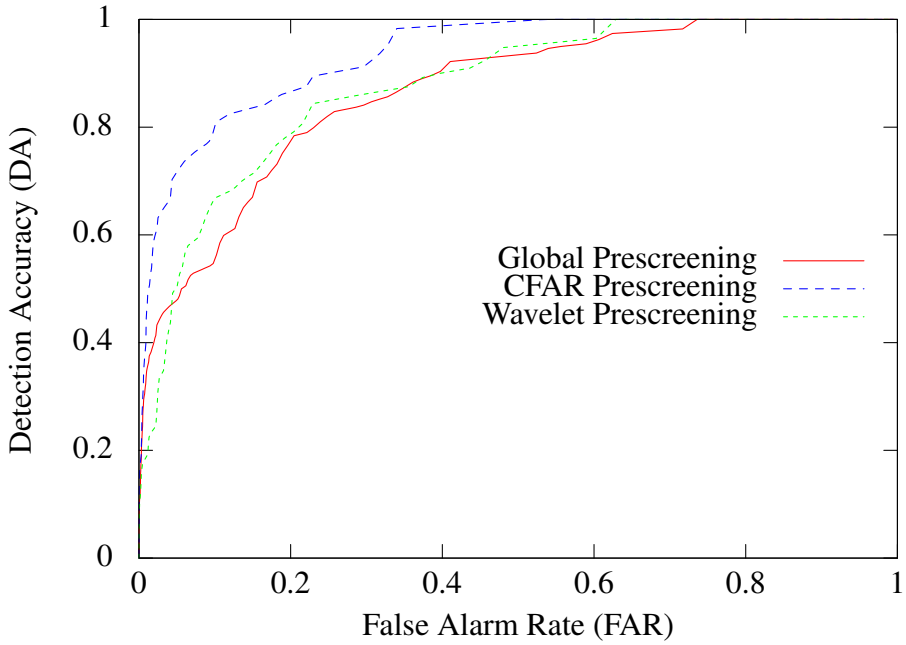

Figure 4.3: ROC curves for the Rayleigh Distribution SAR image.

upon the assumption that the underlying SAR image distribution is a K-distribution. The other methods take longer to reach a DA of 100% in the K-distribution than in the Rayleigh distribution ROC results.

4.3.2.1 Summary of assumed values

The results of the previous experiment and tests justify the following:

1. Synthetically generated ship pixel values should be in the range $33 \leq I(x, y) \leq 142$ with an approximate mean pixel value of $\mu = 88$.

2. An average noise level of SNR $= -6.54\,\text{dB}$ should be used to emulate the noise dispersion in synthetic SAR images.

3. The orientations of the ships are random but there should be a slightly higher concentration of horizontal ship components than vertical ship components.

4. The ship sizes should range from two pixels to five pixels long.

5. The guard size of the CFAR prescreening method should be at least 7×7 so that it is

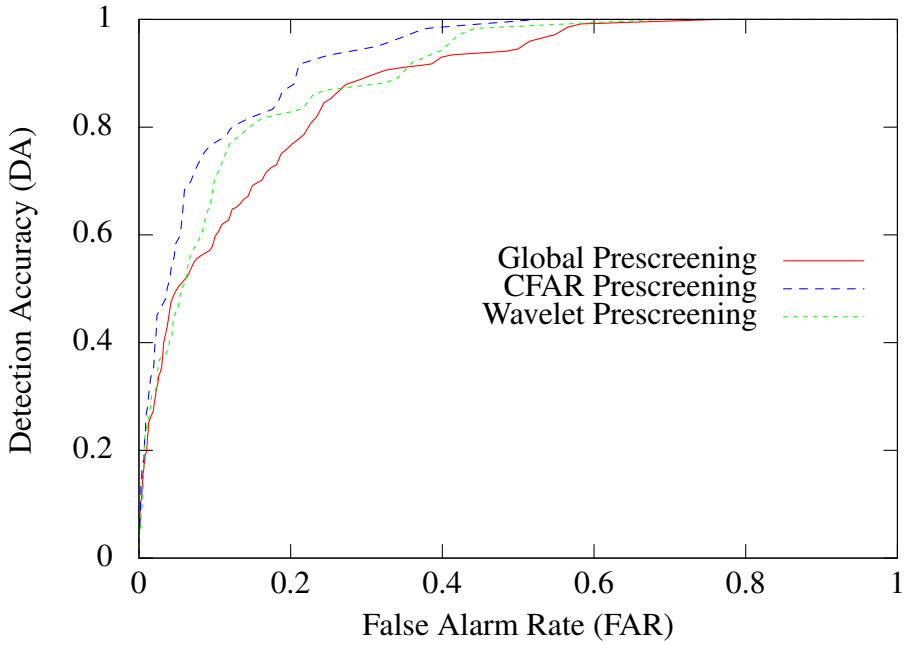

Figure 4.4: ROC curves for the K-Distribution SAR image.

slightly larger than the maximum possible ship length of approximately 5 pixels.

6. The background window size should be at least 9×9 to ensure a sufficient representation of a ship's background statistics.

7. The thresholds for three prescreening methods should be selected such that they cover the low, medium and high expected ship pixel values in the SAR imagery. Specifically these are $I(x, y) = \{33, 88, 142\}$ which equates to a) $P_t = 33, 88$ and 142 for the Global prescreening method, $T = 1.80, 2.90$ and 7.73 for the CFAR-based prescreening method and $c = 1.80, 2.90$ and 7.73 for the Wavelet-based prescreening method ($T_1 = \frac{255}{142} = 1.7977 \approx 1.80$ and similar calculations for the other values).

4.4 SAR IMAGE STATISTICS TESTING

4.4.1 Experiment description

This experiment was conducted to compare the generated synthetic SAR image statistics to the actual SAR images statistics. A set of synthetic SAR images were generated and then their respective PDF are compared to the mean PDF of the actual SAR images. To represent

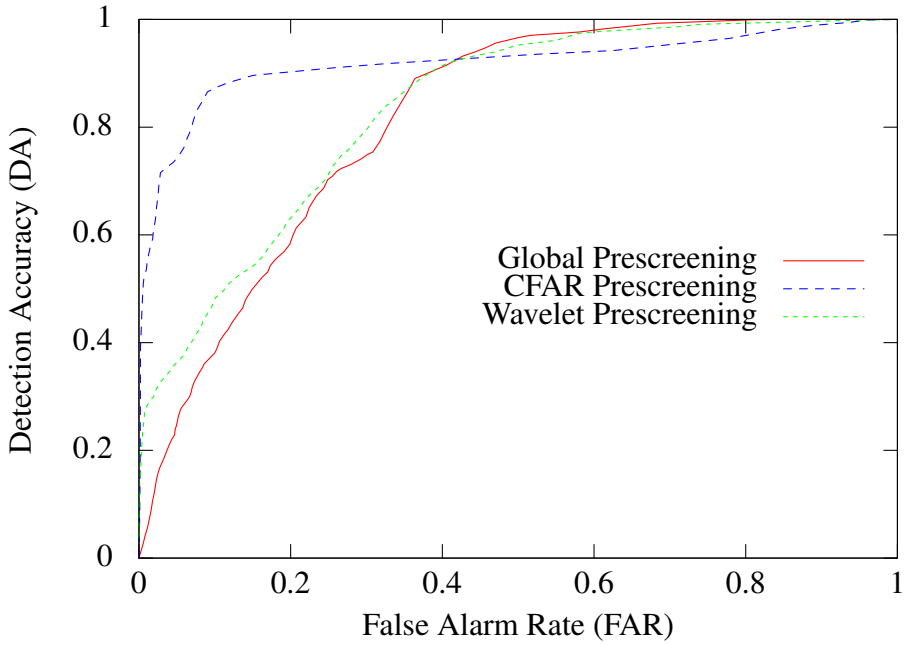

Figure 4.5: ROC curves for the mean of the six ASAR images.

the actual SAR images as a single PDF all six of the SAR PDFs are computed and then a mean PDF is computed from these six PDF. Two common PDF comparison functions are used, namely the Bhattacharyya and Quadratic-Chi distance metrics [48]. A full description of each of these metrics can be found in the appendix A.2. Note that these metrics are unitless and have a range of values $[0, \infty)$ with values closer to zero indicating higher correlation between two PDFs.

4.4.2 Data description

The data used for this is experiment was the six actual SAR images, ten synthetic Rayleigh distributed SAR images and ten synthetic K-distributed SAR images. The synthetic images each had a image size of 500×500 with ships randomly distributed as discussed before. The ten images were linearly spaced between SNR $= -30$ dB and SNR $= 10$ dB. The lower bound was selected because at this SNR ships are visually indistinguishable from noise and background pixels. The upper bound is selected because at this SNR ships are perfectly distinguishable (visually) from noise and background pixels. Ten images provide a sufficient difference between successive images to highlight the differences between SNR. This also ensures that a wide range of parameters values are tested for each distribution. Finally,

using ten images in this specific range ensures that the 7^{th} image lies approximately at the recommended SNR as discussed earlier.

4.4.3 Testing procedure

The testing procedure involved first calculating the mean PDF for the actual SAR images followed by the PDF for each of the synthetic SAR images. This was followed by a test to compare the synthetic SAR image PDFs to the mean SAR image PDF.

4.4.3.1 Comparison of Synthetic SAR PDFs to mean SAR PDF procedure

1. A PDF for each SAR image (10 Rayleigh, 10 K-distribution and 6 actual) was calculated.

2. The mean SAR PDF is calculated by adding the six actual SAR image PDF together and dividing by six.

3. Each of the ten Rayleigh distributed SAR image PDFs is compared to the mean SAR PDF. For each of PDFs, the Bhattacharyya and Quadratic-Chi distance between that PDF and the mean SAR PDF are computed.

4. This produces ten distance values for the Bhattacharyya distance metric and ten distance values for the Quadratic-Chi distance metric for the Rayleigh distributed SAR image PDFs.

5. A mean and standard deviation of the Bhattacharyya and Quadratic-Chi distances were calculated.

6. Steps 3, 4, 5 are repeated for the K-distribution SAR image PDFs.

4.4.4 Experimental results

Figure 4.6 shows the mean PDF for the six real SAR images. The mean SAR PDF is drawn in red. The Rayleigh distribution and K-distribution synthetic SAR image that produced the lowest Quadratic-Chi distance to the mean SAR PDF is shown in blue and green respectively.

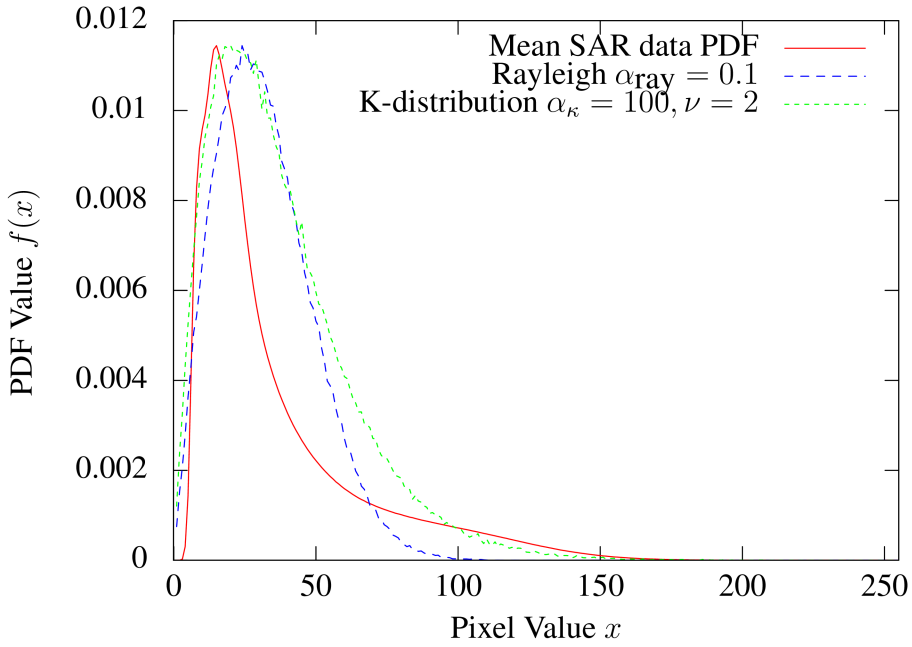

Figure 4.6: The mean SAR PDF along with the the Rayleigh and K-distribution PDFs that had the lowest Quadratic-Chi distance to the mean SAR PDF.

All three of the PDFs have approximately the same rate of increase for the left tail. The mean SAR PDF reaches its peak faster than the other two. It has a rapid descent for the first half of its right tail and and a more slowed descent from approximately $x = 50$. Finally, the mean SAR PDF reaches $f(x) = 0$ at just after $x = 150$. The Rayleigh distribution PDF decreases in a constant, more accelerated rate and reaches $f(x) = 0$ at approximately $x = 100$. The right tail of the K-distribution is similar to the mean SAR PDF and reaches $f(x) = 0$ at approximately $x = 150$.

Looking to the left side of Table 4.2, the Rayleigh PDF with the lowest Bhattacharyya distance is at SNR $= -12.22$ dB with a $D_{bhat} = 6.5718$. The Rayleigh PDF with the lowest Quadratic-Chi distance is at SNR $= -7.78$ dB with $D_{QC} = 5.0327$. Looking at the right side of Table 4.2, the K-distribution PDF with the lowest Bhattacharyya distance is at SNR $= 1.11$ dB with $D_{bhat} = 6.2616$. The K-distribution PDF with the lowest Quadratic-Chi distance is at SNR $= -16.67$ dB with $D_{QC} = 4.7119$.

Table 4.2: Bhattacharyya and Quadratic-Chi distances from the mean SAR PDF. The distances between each of the 10 synthetic SAR image histograms for each distribution is compared to the mean SAR PDF. A mean and standard deviation value is provided at the bottom of the table.

SNR(dB)	**Rayleigh distribution** D_{bhat}	**Rayleigh distribution** D_{QC}	**K-distribution** D_{bhat}	**K-distribution** D_{QC}
-30.00	7.2044	12.086	6.4450	5.7351
-25.56	6.7837	12.544	6.8972	5.1643
-21.11	6.9802	11.771	6.8221	5.3685
-16.67	7.0282	10.066	6.7746	5.5127
-12.22	6.5718	8.2099	6.7402	4.7119
-7.78	7.3546	5.0327	7.3360	10.273
-3.33	7.1979	5.9923	7.0431	11.912
1.11	7.5435	5.3295	6.2616	11.723
5.56	7.4670	5.1947	6.3862	9.5131
10.00	7.8958	6.9489	6.5519	7.6250
Mean	7.2027	8.3175	6.7258	7.7538
Std. dev.	0.3863	3.0516	0.3256	2.8529

4.4.5 Results discussion

This experiment compared two types of noise distributions to actual SAR image PDFs. By comparing the distributions to the mean SAR PDF the numerical results justifies two distributions best match the SAR imagery across all cases, on average, can be ascertained.

Looking at the results seen in Figure 4.6, we can see that each of the two distributions match the mean South African SAR image PDF well. Starting to the left of the head all three PDFs rise at approximately the same rate. The peak values for all three distributions are relatively similar and positioned near to one another meaning that most of the image content will look

similar because the background (lower) pixel values all have similar distributions.

The most significant difference between the three PDFs occurs at the right tails. The mean SAR image PDF has a rapid decrease in the first half of the right tail and then a much slower decrease in the last half or the right tail. This is because, in general, there are much fewer objects at sea that backscatter the SAR imagery to a high degree. This in turn means that more pixel values in the SAR image need to be darker and therefore more the head and most of the tail needs to occur before a pixel value of 50 (which is a dark grey). The remaining objects have increasing backscatter values and thus increasing pixel values. For darker ASAR images such as those at sea, the last half of the tail had an inverse relationship between pixel value and PDF value. That is, the higher the PDF value the lower the pixel value. The other two distributions have tails that either decrease in a manner that matches the mean SAR PDF but settle too quickly at zero (Rayleigh) or decrease slower than the mean SAR PDF but converge better at high pixel values (K-distribution).

Both visually and empirically (using the two distance metrics) the K-distribution presents a distribution that is closer to the mean SAR image PDF than the Rayleigh distribution. This allows for the conclusion that noise in ASAR imagery near South Africa's coast can be effectively modelled using the K-distribution.

4.5 PRESCREENING TESTING

4.5.1 Experiment description

This experiment investigated the performance of three prescreening methods: Global prescreening, CA-CFAR prescreening and Wavelet prescreening. The testing image was created to evaluate the limits of each of the three prescreening methods. The prescreening methods were tested against various levels of noise, at various thresholds, to ensure that the methods were performing correctly irrespective of the images being fed to the them.

The parameters for the three prescreening methods were selected using the experimental results from the previous experiment involving the mean sub-images (section 4.3.1.1).

4.5.2 Data description

A manually created, binary test image of size 100×100 was used in this experiment. The image consists of three single pixel groups and nine straight lines. Three straight lines form a pixel group and each pixel group had a pixel line in one of three directions: horizontal, vertical or diagonal. In the first group each pixel line is two pixels long, in the second five and in the final group seven. All other pixels were set to zero. Three levels of noise were added to the image: a) SNR $= -30$ dB (ships visually indistinguishable from noise), b) SNR $= -6.54$ dB (average noise dispersion around ships in ASAR WSM images) and c) SNR $= 10$ dB (ships perfectly distinguishable from noise). The three test images used in this experiment is shown in Figure 4.7.

4.5.3 Testing procedure

The experiment proceeded as follows:

1. The Global prescreening method was applied to each of the three test images. This was repeated for three global threshold values $P_t = 33, P_t = 88$ and $P_t = 142$.

2. The CA-CFAR prescreening method was applied to each of the three test images with $T = 1.80$. This step was repeated for three sets of background S_b and guard window sizes S_g. The three sizes were as follows: $[S_b = 5, S_g = 3]$, $[S_b = 9, S_g = 7]$ and $[S_b = 11, S_g = 9]$.

3. The above step was repeated with $T = 2.90$.

4. The Wavelet prescreening method was applied to each of the three test images. This was repeated for three threshold values $c = 2, 8$ and $c = 15$.

5. Each of the result prescreened images were compared to the original image (no noise) and from that DA and FAR values were computed.

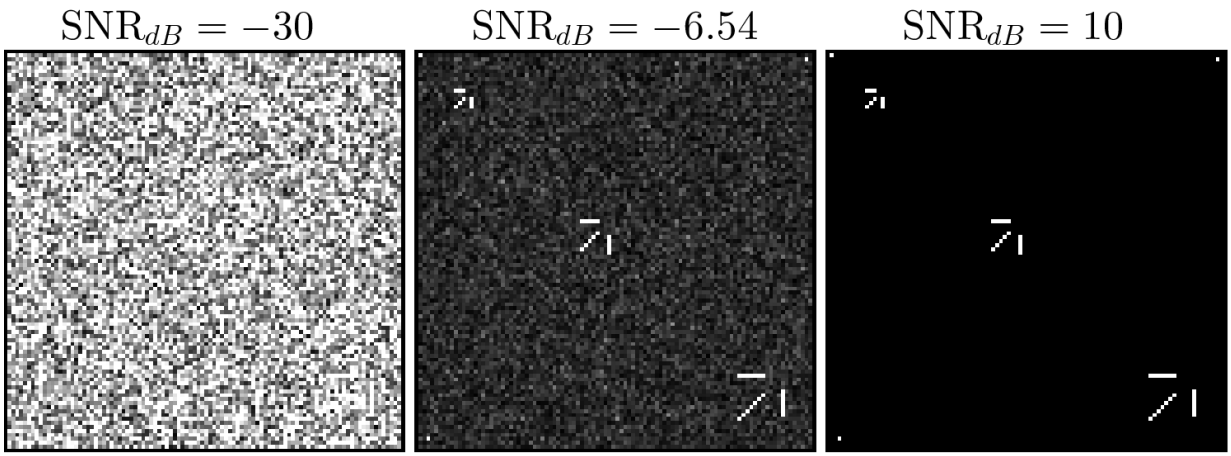

Figure 4.7: The three standard images used throughout the experiment. The SNR values are with respect to a perfect detection image. A black border was added to improve image clarity.

4.5.4 Experimental results

The results of the experiment are shown in Table 4.3. There are five groups of three with each group indicating a single prescreening method its various parameter permutations. There are three groups for the CA-CFAR method because the threshold and window sizes are varied separately.

The Global prescreening method, irrespective of threshold value, had the highest DA along with the highest FAR. The CA-CFAR prescreening method had the lowest overall FAR across all parameter permutations. When the threshold was set to a high value $T = 7.73$ the prescreening method did not detect any ships unless the ships were indistinguishable from the noise. As expected, the Wavelet prescreening method got a high DA and low FAR when the threshold was low and the noise dispersion was average. Increasing the noise and threshold resulted in a low DA and FAR.

Table 4.3: Results of applying the prescreening methods to the three test images in Figure 4.7. The methods are tested with varying parameters. The DA is shown with the FAR in brackets, both are in percentages.

Prescreening method	**SNR = -30 dB**	**SNR = -6.54 dB**	**SNR = 10 dB**
Global, $P_t = 33$	100.0 (94.0)	100.0 (82.0)	100.0 (0)
Global, $P_t = 88$	100.0 (90.0)	100.0 (66.0)	100.0 (0)
Global, $P_t = 142$	100.0 (78.0)	98 (8.0×10^{-2})	100 (0)
CA-CFAR, $T = 1.80$, $[S_b, S_g] = [5, 3]$	60 (6.27)	88 (0.5)	98 (0.01)
CA-CFAR, $T = 1.80$, $[S_b, S_g] = [9, 7]$	60 (3.30)	96 (1.05)	100 (0.0)
CA-CFAR, $T = 1.80$, $[S_b, S_g] = [11, 9]$	30 (1.12)	100 (5.10)	100 (0.0)
CA-CFAR, $T = 2.90$, $[S_b, S_g] = [5, 3]$	0.0 (8.0×10^{-2})	21 (4.0×10^{-2})	98 (4.0×10^{-2})
CA-CFAR, $T = 2.90$, $[S_b, S_g] = [9, 7]$	0 (0)	98 (1.0×10^{-2})	96 (4.0×10^{-2})
CA-CFAR, $T = 2.90$, $[S_b, S_g] = [11, 9]$	0 (0)	98 (0.5×10^{-2})	100.00 (0)
CA-CFAR, $T = 7.73$, $[S_b, S_g] = [5, 3]$	0 (0)	0 (0)	98 (0)
CA-CFAR, $T = 7.73$, $[S_b, S_g] = [9, 7]$	0 (0)	0 (0)	98 (0)
CA-CFAR, $T = 7.73$, $[S_b, S_g] = [11, 9]$	0 (0)	0 (0)	100 (0)
Wavelet, $c = 1.80$	69 (4.64)	100 (5.10)	100 (0.1)
Wavelet, $c = 2.90$	0 (1.10)	100 (0.35)	96 (4.0×10^{-2})
Wavelet, $c = 7.73$	0 (0)	42 (0)	96 (0.0)

4.5.5 Results discussion

The prescreening stage of the ship detection system is arguably the most important part of the whole ship detection system. A consistent and correct prescreening method makes the ship discrimination task trivial because only the correct ship pixels are highlighted and need to be grouped together. If the prescreening method is sub-optimal then the ship discrimination method plays an important role in reducing the FAR.

This experiment aimed to highlight attributes of the three prescreening methods by varying each one's parameters and the synthetic image SNRs. This allows for a relatively complete picture of each prescreening method's limits and performances.

As expected the Global prescreening method had a high FAR for the low SNR values because the method is indiscriminate with regards to pixel value. Pixels above the threshold will be selected irrespective of the number of pixels selected. There is a dramatic improvement in the FAR when the threshold is increased and the SNR of the synthetic image is above or equal to $-6.54\,\text{dB}$. This indicates that if the SNR of the input images are below the $-6.54\,\text{dB}$ then the Global prescreening method will have difficulties producing low FAR values consistently across a number of test images. This can be remedied by increasing P_t however this will come at a loss in FAR.

Due to the fact that the CA-CFAR prescreening method has three configurable variables (S_b, S_g and T) there are far more permutations of the method than the other prescreening methods. As expected, input images with no noise (SNR $= 10\,\text{dB}$) that are prescreened with the CA-CFAR method with a low or high threshold produce high a DA and FAR. Similar to the Global prescreening method, when the SNR is low far more ship pixels are detected. The method has the lowest FAR of all the methods and this makes sense as the method is designed to keep the FAR constant. Despite this result, there are still false detections when SNR $= -6.54\,\text{dB}$. Area statistics vary from image to image and thus for some images increasing the threshold will cause fewer bright ship pixels to be detected.

The varying of S_b and S_g shows that having a larger window sizes is preferable to smaller window sizes. A proper representation of the background and neighbourhood statistics is not possible when the window size is too small. This causes the CA-CFAR method to incorrectly

prescreen pixels that should be selected as ship pixels.

Finally, the Wavelet results are shown in Table 4.3. Once again, a low threshold and a low SNR image produce an image with a high FAR and moderate DA. The FAR at this low threshold and low SNR value is better than the Global prescreening method but worse than the CA-CFAR. The Wavelet prescreening method's results improve as the threshold is increase but ultimately do not perform as well as the other prescreening methods. The DA is the worst of the three methods but the FAR lies between the Global prescreening method and CA-CFAR.

The Wavelet prescreening method operates by building a correlation image by processing the original image at a) different scales and b) different orientations. The reason this works for noisy images is that because the noise is a random process adjacent noise pixels across the various images will have a lower correlation at different rotations. For instance, if two random noisy pixel values are correlated in the vertical direction it is very likely that they will not be correlated in the horizontal and diagonal directions simultaneously. Using this principle and multiplying different Wavelet images together will result in pixels across scales and orientations of high correlation increasing in value while those that do not correlate are decreased. The problem with this is that at the different scales and orientations noisy pixels directly adjacent to the actual ship pixels will correlated better than noise pixels that are not directly adjacent. Essentially, the ship pixels correlate well across scales and orientations and therefore the pixels directly adjacent to those will also correlate better than those farther away. This causes incorrect detections and increases the FAR slightly but can create problems for the DA. This effect can be reduced by increasing the Wavelet correlation multiplier value c. This will eliminate the some of the noisy, partially correlated pixels but keep the highly correlated ship pixels.

In summary, when prescreening SAR images that have a (possible) medium level of K-distributed noise, use the following guidelines:

- When using the Global thresholding method it is better to use a higher threshold value, above the expected pixel mean, to improve FAR results.
- When using the CA-CFAR thresholding method it is better to use slightly higher than the recommended guard and window sizes to improve ship pixel reconstruction in the

binary output image. In practice, a threshold value between the low and medium threshold values should be selected to ensure a sufficiently high DA and low FAR.

- When using the Wavelet thresholding method it is better to use a thresholding value closer to the mean threshold value to improve the reconstruction of ship pixels in the result image.

4.6 SHIP DISCRIMINATION TESTING

4.6.1 Experiment description

This experiment investigated the performance of two ship discrimination methods: Connected Component Analysis and Mean-Shift ship discrimination. The same test image used in testing the prescreening methods was used in this experiment (Figure 4.7) to investigate the grouping and centre finding capabilities of the ship discrimination methods. The centre pixel for each of the groups in the test image was marked on a centre location image. The results of each method with different parameters is compared to this centre location image to determine each methods DA and FAR. It should be noted that the test image purposefully contains single pixel ships in order to determine how the ship discrimination system will handle them. Even though it is assumed that ships in the available imagery cannot be smaller than two pixels, it is prudent to test the ship discrimination methods with ships that lengths of one pixel to ensure that the ship discrimination methods are performing as expected.

Three parameter values were selected from the assumed values present in section 4.3.2.1. Recall that the parameters of the two ship discrimination methods determine how far apart two pixels need to be in order to be grouped together. Knowing the maximum and minimum expected ship lengths in pixels allows for the selection of the following parameter values:

- $h = g_p = 1$ for a minimum ship size of two pixels (current pixel plus pixels that are one pixel away from it).

- $h = g_p = 4$ for a maximum ship size of five pixels (current pixel plus pixels that are four pixels away from it).

- $h = g_p = 7$ for a grouping size larger than the maximum ship size of five pixels.

Figure 4.8: The test image (left) and the 12 ship centres in the centre point image (right).

4.6.2 Data description

The test image used for this experiment was the same as the one created for the ship discrimination experiment (section 4.5.2). There are 12 groups of pixels and therefore 12 centres in the test image: one group per pixel point (three centres) and one group per pixel line for three lines for each of the three groups (nine centres). Figure 4.8 shows the test image and its associated binary ship centre location image (respectively).

4.6.3 Testing Procedure

The ship discrimination testing proceeded as follows:

1. Three grouping distances for the Connect Component ship discrimination method were used: $g_p = \{1, 4, 7\}$.

2. Three bandwidth values for the Mean-Shift ship discrimination method were used: $h = \{1, 4, 7\}$.

3. The six above method permutations were applied to the binary ship image present in Figure 4.8.

Table 4.4: Experimental results for the two ship discrimination methods with three different parameter values each.

Ship discrimination method	DA (FAR) in %
Connected Component, $g_p = 1$	100.0 (0)
Connected Component, $g_p = 4$	91.67 (0)
Connected Component, $g_p = 7$	50 (1×10^{-2})
Mean-Shift, $h = 1$	100 (7×10^{-2})
Mean-Shift, $h = 4$	91.67 (3×10^{-2})
Mean-Shift, $h = 7$	50 (0)

4. For the six images, the DA and FAR were computed by comparing the ship discrimination method result images to the centre location image.

4.6.4 Experimental results

The results for this experiment is shown in Table 4.4, Figure 4.9 and Figure 4.10. The table shows that the Connect Component ship discrimination method had slightly better DA and FAR across the three parameter values. Both methods performed best when $1 \leq g_p \leq 4$ and $1 \leq h \leq 4$.

Figure 4.9 and Figure 4.10 show the effects of varying the two parameters for the two ship discrimination methods. For all three parameter values, both the Connect Component and Mean-Shift ship discrimination method were able to correctly identify the single ship pixels indicating baseline correct performance. The Connect Component ship discrimination method correctly detected the ship groups when $g_p = 1$. Increasing the grouping size increased the range of pixels to be grouped and thus incorrectly groups the top left pixel lines as one group instead of three. The same result occurs when $h = 4$ for the Mean-Shift ship discrimination method. When the grouping size or bandwidth parameter is large as the maximum expect ship size (seven pixels long) then all the groups in the image are misidentified as being a single group each.

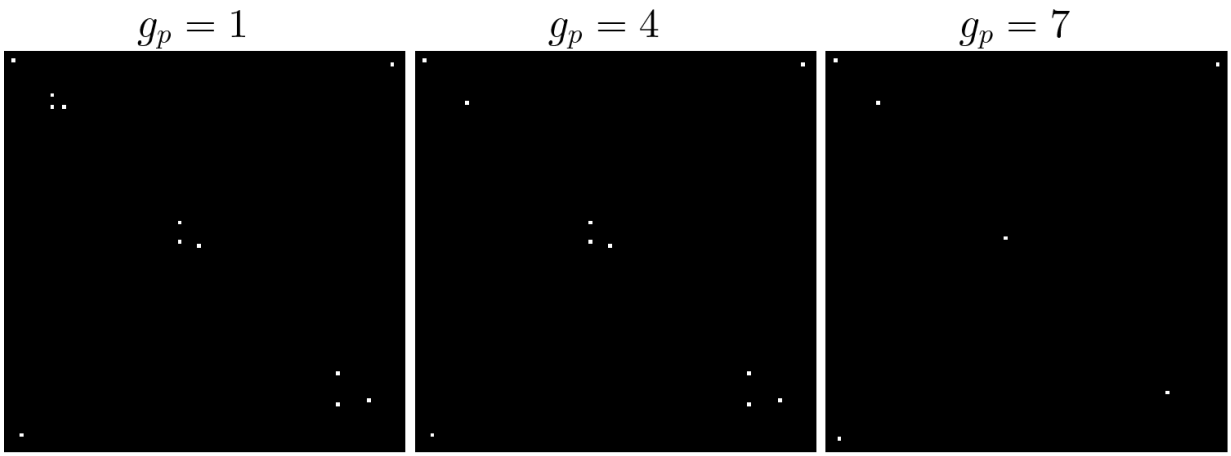

Figure 4.9: The results of applying the Connect Component ship discrimination to the test image. As g_p is increased the number of groups decrease because the acceptable distance between two pixels to form a group is increased.

4.6.5 Results discussion

The ship discrimination stage has the important function of reducing the number of false detections. By grouping pixels together as a single group and assigning a centre pixel allows each pixel group (ship) to be weighted evenly when compared to the other ships in the "ground truth" image. This experiment tested the ability of each of the two ship discrimination methods to correctly group and find the centre of mass of the various groups. The methods were tested against synthetic ships of varying size, from one pixel in length to seven, to ensure that the methods had a sufficient level of generality so that they processed whatever was received from the prescreening methods correctly.

The Connect Component method had a 100% DA and 0% FAR (within the bounds of synthetic, well-defined input images). As expected, increasing the grouping size causes larger groups to be grouped as a single centre pixel. The Mean-Shift procedure performed worse when $h = 1$ indicating that the same value for the two parameters do not produce the same results. Although when $h = 1$ the DA is slightly better, the FAR is 2.3 times worse. This may not seem significant enough but this is due to the small size of the image - larger images will cause the FAR to be exponentially worse. When $h = 1$, every pixel is isolated as its own group and thus every true pixel is highlighted in the ship discrimination binary result image. When the bandwidth parameter is increased, it expands the region of grouping (similar to

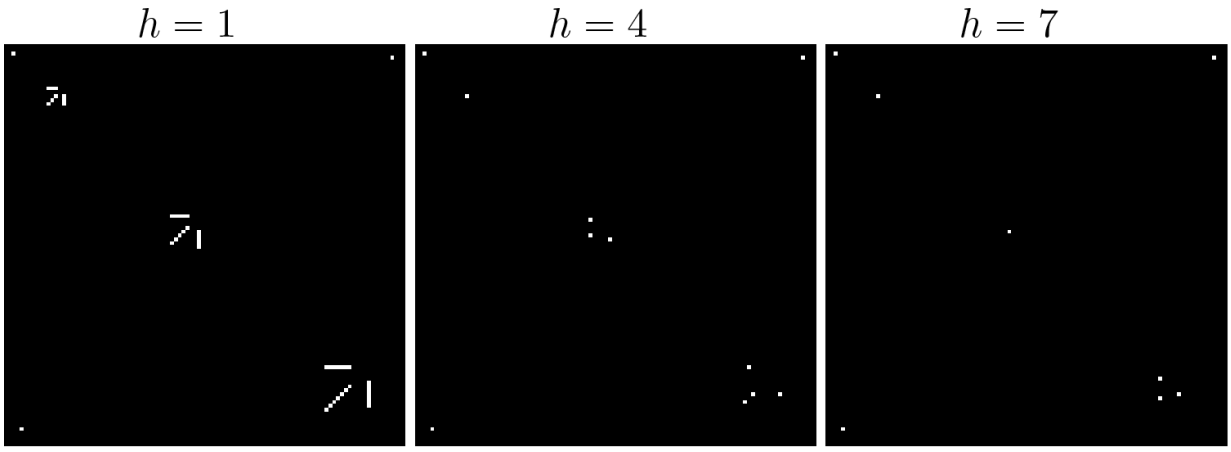

Figure 4.10: The results of applying the Mean-Shift ship discrimination to the test image. As h is increased the number of groups decrease because the acceptable distance between two pixels to form a group is increased.

g_p but in a multidimensional manner) which allows the Mean-Shift method to have sub-pixel accuracy.

In summary, the recommended grouping size would be $g_p = 1$ and bandwidth would be $h > 1$. These deductions are valid only for the ASAR imagery used in this dissertation.

4.7 SHIP DETECTION SYSTEM TESTING

4.7.1 Experiment description

This experiment tested different combinations methods within a ship detection system using both real and synthetic SAR images. The experiment used various parameters for the methods and various levels of noise (synthetic images) and actual SAR images to test the generality of the system.

All of the values used in these experiments use a subset of the possible parameter values as discussed before. For a list of the assumed values and a discussion of how they were obtained please see section 4.3. The two performance metrics used to describe the system performance are the DA and FAR.

4.7.2 Data description

Six synthetically generated K-distribution SAR images and the six actual ASAR images were used. The six K-distributed images were generated to span an SNR range of $-30\,\text{dB}$ (ship pixels indistinguishable from noise) to $10\,\text{dB}$ (ship pixels perfectly distinguishable from noise). The 4^{th} image has an approximate mean of SNR $= -6\,\text{dB}$ which is similar to the SNR dispersion found in the mean sub-image experiment.

4.7.3 Testing procedure

The testing procedure proceeded as follows:

1. Generate six K-distributed synthetic SAR images of size 500×500 with noise levels SNR $= \{-30, -22, -14, -6, 2, 10\}$.

2. 15 randomly located and orientated ships are placed in each image. The centre positions for each ship is recorded in a centre location image.

3. Process each of the six K-distribution images as follows:

   (a) Three prescreening methods at three different thresholds each applied to each image.

   (b) For each of the nine above prescreening result images apply the two different ship discrimination methods at their optimal parameter values. Specifically $g_p = 1$ and $h = 4.0$ (section 4.6.1).

   (c) Calculate the DA and FAR for each of the result images from the previous step.

   (d) Repeat a), b) and c) three times and record the mean DA and FAR results.

4. Calculate the mean and standard deviation for the DA and FAR.

5. Repeat 3. and 4. for the six actual SAR images.

4.7.4 Experimental results

The results of testing the entire system is split into two sections: the synthetic SAR ship detection results and real SAR ship detection results.

4.7.4.1 Synthetic Imagery System Test Results

Table 4.5, Table 4.6 and Table 4.7 represent the results of the ship detection processing of the six K-distribution SAR images for the Global, CFAR and Wavelet prescreening methods respectively and two ship discrimination methods. The format of the results in all tables are DA and FAR in parenthesis. Note that for the results in Table 4.6 were obtained using the recommended guard and background sizes of $S_g = 5$ and $S_b = 7$ respectively.

The results of the Global prescreening method (Table 4.5) shows that this system configuration had had a high DA and high FAR across the three thresholds irrespective of ship discrimination method used. It was noted that around the calculated SNR $= -6\,\text{dB}$, the Global prescreening method obtains a minimum FAR of 2.052×10^{-5} with a 100% DA. This system configuration had the highest mean DA but lowest FAR across all of the system configurations tested.

The CFAR prescreening method results (Table 4.6) shows a system configuration that had an overall acceptable DA and a low FAR across all images tested. The differences in DA between the two ship discrimination methods are negligible with both methods performing almost the same groupings of the correctly detected pixels. The DA found across all thresholds for the CFAR prescreening method is second only to the Global prescreening method. The CFAR prescreening method had the overall best FAR with the lowest standard deviation amongst all of the system configurations tested. It should also be noted that the Mean-Shift ship discrimination system provides slightly better FAR results compared to the Connect Component method for medium to high thresholds.

The Wavelet prescreening method results (Table. 4.7) shows a system configuration that provides DA below that of both the Global and CFAR prescreening methods but with FAR results better than that of the Global prescreening method. The system performs best when the threshold is set to the low value and has a relatively stable DA across all noise levels

except at the most noisy (SNR $= -30$ dB). The mean DA decreases linearly as the threshold value is increased indicating that in all likelihood a threshold lower than the low threshold would provide higher DA but increase FAR as well (a property all of the system configurations appear to have due to the system's deterministic and monotonic nature nature).

4.7.4.2 Real Imagery System Test Results

Table 4.8, Table 4.9 and Table 4.10 represent the results of the ship detection processing of the six ENVISAT ASAR WSM images for the Global, CFAR and Wavelet prescreening methods respectively and two ship discrimination methods. The format of the results in all tables are DA and FAR in parenthesis. Note that for the results in Table 4.9 were obtained using the recommended guard and background sizes of $S_g = 5$ and $S_b = 7$ respectively.

Real SAR image processing using the Global prescreening method system configuration (Table 4.8) presents a system with a high DA and high FAR. As with the synthetic image, when the threshold is increased DA is decreased. The higher the threshold is the higher the likelihood of the threshold being above certain ship pixels thus misclassifying them as background pixels. This system configuration, irrespective of ship discrimination used, had the lowest FAR of the three methods tested. For instance, with a mean FAR of 0.1%, an average SAR image of size of 10000×10000 will have approximately 100000 falsely detected pixels. The Mean-Shift ship discrimination method provides slightly higher FAR than the Connect Component ship discrimination method.

Real SAR image processing using the CFAR prescreening method system configuration (Table 4.9) presents a system with a medium DA and low FAR. The first point to note is that at a low threshold this system configuration correctly identifies a very high percentage of all the ships in each image with a mean DA of 100% when the prescreening method threshold is low ($T = 1.80$). It has the highest DA among all of the prescreening methods across all of the SAR images tested. It should be noted that the mean DA decreases slightly when the threshold value is increased to $T = 2.90$ and, as with the Global prescreening method, when the threshold is increased past the medium the detection accuracy dramatically decreases. The overall mean DA of the Mean Shift is lower than that of the Connect Component method.

The CFAR-based system configuration also has the lowest FAR values as well as the lowest standard deviations for the performance metrics among the methods tested. The FAR is 100-times better when the threshold is increased from $T = 1.80$ to $T = 2.90$ for a very reasonable loss in DA. This loss of DA can be mitigated by selecting a more appropriate threshold between $T = 1.80$ and $T = 2.90$ on a per-image basis. When the threshold is set high then very few to no pixels are false detected. Take note that once again the Mean Shift ship discrimination method has slightly lower mean FAR compared to the Connect Component method.

Finally, the Wavelet prescreening system configuration results are shown in Table. 4.10. The Wavelet DA results are the lowest out of all three methods tested with a mean DA of 77.2% even at the low threshold. The DA across all images is relatively stable for the low threshold but had a far greater standard deviation as the threshold was increased. The FAR of this system configuration was second to the CFAR-based system configuration. Varying the threshold value has a comparatively small effect on the FAR in comparison to how the other methods react to varying the threshold. The differences in both DA and FAR between the two ship discrimination methods is almost negligible with the Connect Component method having a slightly higher mean DA and the Mean-Shift having slightly lower mean FAR.

4.7.5 Results discussion

The entire ship detection system was tested with two sets of SAR imagery: synthetically generated SAR images and actual SAR images. The following two sections discuss the results of each of these and give a more in-depth look at the results described previously.

4.7.5.1 Synthetically generated SAR system testing discussion

The proposed ship detection system was tested using a variety of synthetic SAR images. These images were K-distributed synthetic SAR images with a range of noise levels. The noise levels ranged from a level of noise in which the ships were indistinguishable to a level of noise where ships were perfectly distinguishable. Testing the system with a range of images allowed for a number of benefits: the system performance at the extremes could be determined as well as the expected performance of the system at the mean SNR dispersion measured previously. This imposes a lower bound on the performance of the system with

better results being more likely.

The ship detection system performance increased rapidly for all configurations when the noise levels in the synthetic SAR images were above SNR $= -6$ dB. In essence this means that any images (synthetic or real) where the SNR < -6 dB could not attain a 100% DA and a 0% FAR. It is possible to achieve either of these (and possibly both) but the methods tested here cannot be expected to *reliably* do so for every image input into the system, over a number of images, when the noise level is at or below SNR $= -6$ dB. This test also validated that with an input image with almost no noise all of the components together for each configuration performed as expected.

Two further points need to be discussed about the first part of the system test: the comparatively lagging DA performance of the Wavelet prescreening method and how the ship discrimination method chosen affects the results. It is apparent from the results that the Wavelet DA results lag behind the other system configurations. The most likely reason for this is that the spatial correlation used in the Wavelet method tends to highlights too many pixels nearby to the pixels that are actually ships. This causes too many pixels to be highlighted as ships and in combination with the ship discrimination method causes the Wavelet based system configuration to fail when the noise in the image is high. A possible remedy for this would be to apply a different way to correlated the four images generated by the system as well as by applying a local threshold algorithm to the correlated image $\mathbf{R}$ rather than a global threshold as is in the system tested.

The ship discrimination methods, when appropriate grouping sizes g_p or bandwidths h are chosen, provide very similar results to one another. Therefore performances cannot be selected on these results alone as these vary and do not provide sufficient evidence to select one above the other. The Connect Component ship discrimination method has a lower computational complexity but can only process ship pixels that are integer distances apart. While being more computationally complex, the Mean-Shift algorithm affords the ability to process ship pixels and the associated groups at the sub-pixel level. This ability allows for a more flexible ship discrimination stage that can adapt to different range and azimuthal resolutions and it for this reason that the Mean-shift algorithm is selected as the ship discrimination method to use in an operational system.

4.7.5.2 Actual SAR system testing discussion

The final experiment used the six ENVISAT ASAR WSM images as input and tested the system with the same configurations as the previous experiment. The first point to note is that all of the system configurations performed ship detection to a reasonable level of DA and FAR. This indicates a functional system irrespective of the configuration of the system. Furthermore, the results obtained from the real SAR image compares well with the synthetic images.

Each of the three main system configurations (different prescreening methods) can be analysed separately. Surprisingly, the Global prescreening method did not have the best DA across all configurations. This indicates that the lower threshold, $P_t = 33$, was still too high to highlight the darker ship pixels within the SAR images consistently. However, as before, the FAR of the Global-based system configuration was consistently high. The expected level of noise around a ship pixel was determined to be approximately approximately SNR $= -6.54\,\text{dB}$. Checking the synthetic experimental results near this SNR and with a low threshold $P_t = 33$, the synthetic results show a DA of 100% and FAR of 21% (Table 4.5). This result, if both DA and FAR are weighted equally, is actually worse than that of the real SAR image results with mean DA of 90% and FAR of 1.41×10^{-1}% (Table 4.8). This indicates that the calculated SNR dispersion around the ships is most likely higher (less noisy) than anticipated. This is validated by the fact that increasing the SNR to SNR $= 2\,\text{dB}$ improved the performance results dramatically and allows the results to fall more in line with the results found for the real SAR images. Finally, the two ship discrimination methods show very little difference in both DA and FAR. The Connect Component ship discrimination based systems show slightly better results with a higher standard deviation than those of the Mean-Shift based systems.

The CFAR-based system configuration (irrespective of ship discrimination method used) performed the best of all of the system tested in terms of all metrics: mean DA, mean FAR and standard deviation. As expected, the low threshold was low enough to produce the highest DA and had a reasonable FAR and increasing the threshold decreased both the DA and FAR (as is the case for all of the systems). When compared to the Global-based system configurations, the CFAR's FAR was 100 and 10 000 times better at the low and high thresholds respectively. Unlike the synthetic results, however, the DA result is better than the

Global-based systems. Two reasons for this are possible with the first being discussed before that the low threshold on the Global prescreening method is most likely still to high. The other possibility is that a single threshold value is simply not enough to correctly discriminate ship pixels from sea pixels. The CFAR method uses the neighbourhood of a pixel to determine if that pixel is bright rather than using a single value for all pixels to determine if the pixel is bright. This seems to work better for large actual SAR images as the variations within the image are likely drastic on a local scale and thus local processing is required rather than a single, global processing method. Similar to the Global system configurations, the Mean-Shift method had slightly lower DA and FARs than the Connected Component but with better standard deviation values.

The Wavelet-based system configurations had the lowest DA of all the systems tested but had the second best FAR across all the images tested for the real SAR images. The results seem to indicate that the prescreening method often confused incorrectly correlated noise with ship pixels across orientations thus decreasing the number of correctly detected ships. This is reinforced by the fact that the FAR is relatively low and hence the low correlation between noisy pixels causes them to be discarded thus improving FAR. Again this can most likely be attributed to the inherent "global-ness" of the implemented Wavelet method. It uses a single value to determine if the correlation between pixels among the orientations is high and does not use the local neighbourhood to determine this.

Among the two stages tested, the prescreening stage accounts for the most significant variance in performance among system configurations tested. It, therefore, makes sense that a prescreening method specifically designed to account for a well-known sea noise distribution (K-distribution) in SAR images would perform the best. By assuming that the sea noise can be modelled by the K-distribution and using a simplified version of a density estimation of each pixel's neighbours the CFAR-based ship detection system performed well in terms of both DA and FAR. With regard to the ship discrimination stage the two methods performed similarly and which method to use should be selected on factors other than DA and FAR. For this reason the Mean-Shift method was chosen for its ability to more readily adapt to various SAR geometry resolution changes due to the possibility of sub-pixel accuracy for this ship discrimination method.

Table 4.5: System test results using the six K-distribution images with the Global prescreening method and the two ship discrimination methods.

SNR (dB)	**Global** $P_t = 33$	**Global** $P_t = 88$	**Global** $P_t = 142$	**Global** $P_t = 33$	**Global** $P_t = 88$	**Global** $P_t = 142$
	Connected Component $g_p = 1$			**Mean Shift** $h = 4.0$		
-30	100 (100)	100 (93)	100 (79)	100 (100)	100 (93)	100 (79)
-22	100 (93)	100 (51)	100 (16)	100 (93)	100 (51)	100 (16)
-14	100 (73)	100 (10)	100 (1.92×10^{-1})	100 (73)	100 (10)	100 (1.91×10^{-1})
-6	100 (21)	100 (2.05×10^{-3})	93 (0)	100 (21)	100 (2.05×10^{-3})	93 (0)
2	100 (4.92×10^{-2})	93 (0)	87 (0)	100 (4.92×10^{-2})	93 (0)	87 (0)
10	100 (0)	100 (0)	93 (0)	100 (0)	100 (0)	93 (0)
μ	100 (48)	99 (26)	96 (16)	100 (48)	99 (26)	96 (16)
σ	0 (42)	3 (35)	5 (29)	0 (42)	3 (35)	5 (29)

Table 4.6: System test results using the six K-distribution images with the CFAR prescreening method and the two ship discrimination methods.

SNR (dB)	**CFAR** $T = 1.80$	**CFAR** $T = 2.90$	**CFAR** $T = 7.73$	**CFAR** $T = 1.80$	**CFAR** $T = 2.90$	**CFAR** $T = 7.73$
	Connected Component $g_p = 1$			**Mean Shift** $h = 4.0$		
-30	0 (3.69×10^{-3})	0 (0)	0 (0)	0 (3.69×10^{-3})	0 (0)	0 (0)
-22	100 (16)	20 (2.31×10^{-1})	0 (0)	100 (16)	20 (2.30×10^{-1})	0 (0)
-14	100 (1.93×10^{-1})	100 (2.64×10^{-1})	0 (0)	100 (1.93×10^{-1})	100 (2.62×10^{-1})	0 (0)
-6	93 (0)	100 (2.69×10^{-1})	53 (0)	93 (0)	100 (2.66×10^{-1})	53 (0)
2	87 (0)	100 (2.79×10^{-1})	87 (0)	87 (0)	100 (2.77×10^{-1})	87 (0)
10	93 (0)	100 (2.71×10^{-1})	100 (0)	93 (0)	100 (2.66×10^{-1})	100 (0)
μ	79 (3)	70 (2.19×10^{-1})	40 (0)	79 (3)	70 (2.17×10^{-1})	40 (0)
σ	36 (6)	43 (9.91×10^{-2})	42 (0)	36 (6)	43 (9.81×10^{-2})	42 (0)

Table 4.7: System test results using the six K-distribution images with the Wavelet prescreening method and the two ship discrimination methods.

SNR (dB)	**Wavelet** $c = 1.80$	**Wavelet** $c = 2.90$	**Wavelet** $c = 7.73$	**Wavelet** $c = 1.80$	**Wavelet** $c = 2.90$	**Wavelet** $c = 7.73$
	Connected Component $g_p = 1$			**Mean Shift** $h = 4.0$		
-30	27 (2)	13 (5.72×10^{-1})	0 (9.81×10^{-2})	27 (2)	13 (5.66×10^{-1})	0 (9.81×10^{-2})
-22	60 (2)	33 (6.88×10^{-1})	0 (7.76×10^{-2})	60 (2)	33 (6.82×10^{-1})	0 (7.72×10^{-2})
-14	67 (2)	60 (5.55×10^{-1})	47 (6.24×10^{-2})	67 (2)	60 (5.47×10^{-1})	47 (6.20×10^{-2})
-6	93 (1.53×10^{-1})	73 (4.19×10^{-2})	67 (2.05×10^{-3})	93 (1.53×10^{-1})	73 (4.19×10^{-2})	67 (2.05×10^{-3})
2	87 (4.10×10^{-4})	87 (4.10×10^{-4})	80 (0)	87 (1.23×10^{-3})	87 (8.21×10^{-4})	80 (0)
10	87 (0)	73 (0)	60 (0)	80 (4.10×10^{-4})	67 (4.10×10^{-4})	60 (0)
μ	70 (1)	57 (3.10×10^{-1})	42 (4.00×10^{-2})	69 (1)	56 (3.06×10^{-1})	42 (3.99×10^{-2})
σ	23 (1)	26 (2.99×10^{-1})	31 (4.07×10^{-2})	22 (1)	25 (2.96×10^{-1})	31 (4.06×10^{-2})

Table 4.8: Test Results using the six ASAR images with the Global prescreening method and the two ship discrimination methods.

Image Number	**Global** $P_t = 33$	**Global** $P_t = 88$	**Global** $P_t = 142$	**Global** $P_t = 33$	**Global** $P_t = 88$	**Global** $P_t = 142$
	Connected Component $g_p = 1$			**Mean Shift** $h = 4.0$		
Image 1	88 (8.09×10^{-2})	82 (1.44×10^{-1})	24 (5.85×10^{-2})	82 (1.44×10^{-1})	82 (8.09×10^{-2})	18 (5.86×10^{-2})
Image 2	84 (1.42×10^{-1})	91 (3.37×10^{-2})	7 (1.20×10^{-1})	86 (3.37×10^{-2})	79 (1.42×10^{-1})	21 (1.20×10^{-1})
Image 3	85 (1.60×10^{-1})	75 (8.31×10^{-2})	0 (1.27×10^{-1})	86 (8.31×10^{-2})	81 (1.60×10^{-1})	19 (1.28×10^{-1})
Image 4	94 (1.16×10^{-1})	100 (1.93×10^{-1})	33 (7.23×10^{-2})	83 (1.93×10^{-1})	78 (1.16×10^{-1})	22 (7.23×10^{-2})
Image 5	100 (1.83×10^{-1})	100 (1.51×10^{-1})	0 (1.15×10^{-1})	100 (1.51×10^{-1})	100 (1.84×10^{-1})	0 (1.15×10^{-1})
Image 6	88 (1.49×10^{-3})	78 (2.58×10^{-3})	20 (1.56×10^{-3})	85 (2.59×10^{-3})	81 (1.49×10^{-3})	20 (1.56×10^{-3})
μ	90 (1.14×10^{-1})	88 (1.01×10^{-1})	14 (8.23×10^{-2})	87 (1.01×10^{-1})	83 (1.14×10^{-1})	17 (8.24×10^{-2})
σ	6 (5.97×10^{-2})	10 (6.74×10^{-2})	13 (4.41×10^{-2})	6 (6.75×10^{-2})	8 (5.98×10^{-2})	8 (4.41×10^{-2})

Table 4.9: System test results using the six ASAR images with the CFAR prescreening method and the two ship discrimination methods.

Image Number	**CFAR** $T = 1.80$	**CFAR** $T = 2.90$	**CFAR** $T = 7.73$	**CFAR** $T = 1.80$	**CFAR** $T = 2.90$	**CFAR** $T = 7.73$
	Connected Component $g_p = 1$			**Mean Shift** $h = 4.0$		
1	94 (8.54×10^{-3})	88 (1.76×10^{-4})	18 (3.03×10^{-6})	94 (8.54×10^{-3})	88 (1.76×10^{-4})	24 (3.02×10^{-6})
2	98 (4.35×10^{-3})	91 (8.43×10^{-4})	9 (1.01×10^{-5})	95 (4.35×10^{-3})	91 (8.41×10^{-4})	23 (1.01×10^{-5})
3	95 (1.02×10^{-2})	80 (3.02×10^{-4})	5 (6.22×10^{-6})	95 (1.02×10^{-2})	91 (3.02×10^{-4})	24 (6.21×10^{-6})
4	100 (3.76×10^{-3})	100 (1.61×10^{-4})	17 (1.32×10^{-6})	94 (3.75×10^{-3})	89 (1.61×10^{-4})	22 (1.32×10^{-6})
5	100 (6.30×10^{-3})	100 (1.58×10^{-4})	0 (0)	100 (6.29×10^{-3})	100 (1.58×10^{-4})	0 (0)
6	100 (1.88×10^{-3})	95 (2.12×10^{-4})	42 (7.46×10^{-7})	95 (1.88×10^{-3})	90 (2.12×10^{-4})	22 (7.46×10^{-7})
μ	98 (5.85×10^{-3})	92 (3.09×10^{-4})	15 (3.58×10^{-6})	96 (5.84×10^{-3})	91 (3.08×10^{-4})	19 (3.57×10^{-6})
σ	2 (2.86×10^{-3})	7 (2.44×10^{-4})	13 (3.56×10^{-6})	2 (2.86×10^{-3})	4 (2.43×10^{-4})	9 (3.56×10^{-6})

Table 4.10: System test results using the six ASAR images with the Wavelet prescreening method and the two ship discrimination methods.

Image Number	**Wavelet** $c = 1.80$	**Wavelet** $c = 2.90$	**Wavelet** $c = 7.73$	**Wavelet** $c = 1.80$	**Wavelet** $c = 2.90$	**Wavelet** $c = 7.73$
	Connected Component $g_p = 1$			**Mean Shift** $h = 4.0$		
1	47 (3.48×10^{-2})	24 (4.24×10^{-2})	18 (7.88×10^{-3})	77 (3.48×10^{-2})	59 (4.24×10^{-2})	12 (7.87×10^{-3})
2	79 (5.08×10^{-2})	70 (6.17×10^{-2})	0 (7.24×10^{-4})	79 (5.08×10^{-2})	61 (6.16×10^{-2})	12 (7.23×10^{-4})
3	75 (5.82×10^{-2})	50 (9.00×10^{-2})	10 (8.26×10^{-4})	76 (5.82×10^{-2})	62 (8.99×10^{-2})	10 (8.25×10^{-4})
4	89 (4.22×10^{-2})	50 (5.31×10^{-2})	33 (2.38×10^{-3})	78 (4.21×10^{-2})	61 (5.30×10^{-2})	11 (2.38×10^{-3})
5	100 (5.86×10^{-2})	100 (6.98×10^{-2})	0 (3.76×10^{-4})	100 (5.86×10^{-2})	50 (6.96×10^{-2})	0 (3.75×10^{-4})
6	73 (1.63×10^{-3})	37 (2.40×10^{-3})	17 (1.14×10^{-3})	78 (1.63×10^{-3})	61 (2.40×10^{-3})	12 (1.14×10^{-3})
0	77 (4.11×10^{-2})	55 (5.32×10^{-2})	13 (2.22×10^{-3})	81 (4.10×10^{-2})	59 (5.31×10^{-2})	9 (2.22×10^{-3})
0	16 (1.96×10^{-2})	25 (2.71×10^{-2})	12 (2.61×10^{-3})	8 (1.95×10^{-2})	4 (2.70×10^{-2})	4 (2.61×10^{-3})

CHAPTER 5

CONCLUSION

5.1 SUMMARY OF RESEARCH

The research presented here began by placing South Africa and its MDA within the global MDA context. By highlighting various aspects of MDA such as global MDA participants, concepts such as a country's EEZ and MDA's international importance, the study showed how improving MDA is important to South Africa. The methods most commonly used to monitor ships (and hence the maritime environment) were discussed in detail. It was also discussed how these methods can fail and how a satellite-based observational tool could be used to supplement maritime ship monitoring systems.

The role of SAR satellites and its associated imagery as an emerging and important resource for countries to improve their MDA was discussed. A number of benefits of SAR satellite observation were discussed with the most relevant to monitoring maritime regions being the fact that it can be used to monitor vast areas of the Earth, irrespective of weather condition. The massive area covered by a single SAR image helps to reduce the monitoring cost per square kilometre and from this the cost of ship detection can also be reduced if SAR imagery is used as the observational tool.

A large section of the dissertation was dedicated to detailing the most prevalent literature regarding SAR-based ship detection. It highlighted how ship detection systems are composed of a number of steps seen throughout the literature. Each step constitutes a different aspect of the ship detection system from geo-locating the images to comparing the detected ships to known positions.

Following the general overview of the field, the research went into detail regarding the selected methods for each step of the ship detection system. A detailed description of how SAR images are acquired and formatted for processing was given along with mathematical details pertaining to each method selected.

Chapter 4 presented a number of experiments that investigated several aspects of the ship detection system. It began by investigating which of the most prominent noise distribution models best suits the sea noise found in ENVISAT ASAR imagery. Following this, two experiments were performed in order to refine some of the assumed parameters and synthetic image variables for the testing of the methods in the ship detection system. This included a general, single stage experiment to benchmark the performance of the system without any parameter adjustments. The two prominent stages in the ship detection system, the prescreening and ship discrimination steps, were each tested with their own experiment. These experiments aimed to highlight the fact that the steps were working as intended and their robustness was investigated using synthetic SAR imagery. The final experiment involved testing the ship detection system as a whole with various configurations and parameter values on both real and synthetic SAR imagery to determine the general performance of the system. A detailed discussion of the results obtained is provided for all of our findings.

5.2 RESEARCH AIMS AND FINDINGS

This study investigated the plausibility of a South African focused ship detection system using ENVISAT ASAR data. It also investigated whether assumptions made for other parts of the world about certain sea-noise distributions would hold true for the sea waters found in ASAR imagery off of the coast of South Africa.

A number of published works ([4, 49, 37, 35]) state that the K-distribution best models the ocean noise distribution found in SAR images. These findings were all based on SAR imagery for specific regions in the world. This assumption has major implications on the methods used and thus it was prudent to first test this assumption on the ENVISAT ASAR imagery located near the South African coast. The findings in this study suggest that, in general for the ENVISAT ASAR WSM VV-polarised images, the most appropriate noise distribution model for sea-noise in coastal waters around South Africa is the K-distribution. This confirmed the first part of the hypothesis presented in the introduction.

The ship detection system with the best performance found in our experiments used a CA-CFAR-based ship prescreening method with any of the suitable ship discrimination methods. The CFAR-based system configurations not only kept a consistently low FAR across synthetic and actual SAR image results but also had a correspondingly high DA. This confirmed the second and third hypothesis that CFAR-based ship detection system will outperform all other configurations.

The selection of which of the two tested ship discrimination methods to use was was difficult as both methods performed equally well in terms of both DA and FAR if the correct parameter values were selected. The selection of which method therefore need to be based on metrics other than these two performance metrics. The Connect Component ship discrimination method had a computational complexity of approximately $O\left(n^2\right)$ whereas the Mean-Shift method had a complexity of approximately $O\left(n^3\right)$. This is a significant difference, when computing millions (and even billions) of pixels. Despite this, the flexibility afforded to the Mean-Shift method due to sub-pixel accuracy (which theoretically allows the method to scale to different SAR image resolutions) and the possibility of speed improvements means it is the more future-proof of the two methods and provides a ship discrimination method that can better adapt to various image resolutions.

In closing, the CA-CFAR based prescreening method with the Mean-Shift ship discrimination method provides the best DA and FAR results irrespective of if the SAR imagery is real or synthetic and is the system recommended to be implemented in an operational SAR-based ship detection system for South African coastal waters.

5.3 IMPLICATIONS AND SIGNIFICANCE OF RESEARCH

The evidence from this study suggests that it is possible to create an ASAR processing ship detection system with a high level of DA and a low level of FAR that focuses on ship detection in South African waters. With near real-time processing of SAR images and better accessibility of SAR imagery with shorter revisit times (Sentinel), South Africa can have an effective and fast ship detection system that is able to improve its MDA. This has many far-reaching consequences such as the identification of IUU fishing activities or illegal oil dumping in South African waters [2, 50]. By combining an effective SAR-based ship detection system with auxiliary ship tracking methods such as AIS, it should be possible to provide information

to the authorities to assist in greatly reducing the types of problems related to maritime monitoring for South Africa's oceans.

The findings in the dissertation enhance our understanding of South African waters and the means to effectively monitor it. It was shown that not only can noise in South African waters be effectively modelled using the K-distribution but methods based on the K-distribution (CFAR) are the most effective for detection ships amongst all of the methods tested on ASAR images located off the South African coast.

5.4 LIMITATIONS AND RECOMMENDATIONS FOR FURTHER WORK

The current research was not specifically designed to evaluate factors related to the preprocessing step in the ship detection system. This, in itself, is a complex topic and for this study it was assumed that the images received from the satellite (ENVISAT ASAR WSM images) were in a suitable condition to be processed by the ship detection system. Furthermore, the results obtained in this dissertation only apply to ENVISAT ASAR WSM images. While it is possible to extrapolate the results to other SAR satellite information, further extensive testing needs to be done in order to make this assumption.

The experiments in this dissertation regarding the Wavelet prescreening method mostly tested the method using only K-distributed synthetic SAR images. In the original Wavelet prescreening paper the authors do not assume a K-distribution for the SAR input images but a Rayleigh distributed SAR input image [24]. In this dissertation the method was applied to imagery with a K-distribution sea-noise distribution which accounts for the decrease in DA.

Further experimental investigations are needed to determine an automatic method for selecting thresholds for all prescreening methods. One of the more complex tasks when working with prescreening algorithms is the selection of the threshold parameter used to determine when a pixel is selected as a bright (ship) pixel or not. Usually, thresholds are chosen empirically by selecting a desired FAR and then working out the threshold from that. An area of research is to use auxiliary data such as AIS or LRIT positional data to build a distribution of ship movement in a geographical region. With this ship distribution map it can

be determined what a given threshold's probability metric is and the change in this metric across thresholds can highlight when a threshold is too high. This topic is an advanced topic and requires its own, specialised study to determine how effective this method can be at automatically selecting thresholds [51].

5.5 CLOSING REMARKS

The ability of SAR to observe large areas with a minimal cost per square kilometer coupled with maritime ship monitoring systems can effectively be used as an operational maritime monitoring system to improve South Africa's MDA.

REFERENCES

[1] D. J. Agnew, J. Pearce, G. Pramod, T. Peatman, R. Watson, J. R. Beddington, and T. J. Pitcher, "Estimating the Worldwide Extent of Illegal Fishing," *Public Library of Science ONE*, vol. 4, no. 2, pp. 1–8, Feb. 2009.

[2] S. Moolla, "Contextualising illegal, unregulated and unreported fishing of marine resources in south african waters," Institute for Security Studies South Africa, New Muckleneuk, Pretoria, Tech. Rep., Apr. 2009.

[3] C. Oliver and S. Quegan, *Understanding Synthetic Aperture Radar Images*. Rayleigh, NC: SciTech Publishing, 2004.

[4] D. J. Crisp, "The State-of-the-Art in Ship Detection in Synthetic Aperture Radar Imagery," Australian Department of Defence, Edinburgh, Australia, Tech. Rep. DSTO-RR-0272, 2004.

[5] "National plan to achieve maritime domain awareness for the national strategy for maritime security," White Paper, United States Department of Defense, May 2005.

[6] "Securing an open society: Canada's national security policy," White Paper, Canadian Government, Apr. 2004.

[7] N. Wegge, "Small state, maritime great power? Norway's strategies for influencing the maritime policy of the european union," *Marine Policy*, vol. 35, no. 3, pp. 335–342, May 2011.

[8] "Maritime Doctrine For The SA Navy," White Paper, South African Navy, Department: Defence, Oct. 2006.

[9] International Maritime Organization, *Electronic SOLAS Consolidated Edition*. Reading, MA: International Maritime Organization, 2011.

[10] "Performance standards and functional requirements for the long-range identification and tracking of ships," Piscataway, NJ, Tech. Rep. MSC.210, 2006.

[11] M. I. Skolnik, *Radar Handbook*, 3rd ed. New York, NJ: McGraw-Hill, 2008.

[12] P. M. Mather, *Computer Processing of Remotely-Sensed Images: An Introduction*, 3rd ed. Hoboken, NJ: John Wiley & Sons, 2004.

[13] K. Eldhurst, "Automatic Ship and Wake Detection In Spaceborne SAR Image From coastal Regions," in *Proceedings of the IEEE 1988 International Geoscience and Remote Sensing Symposium (IGARSS)*, vol. 1, Edinburgh, U.K, Sep. 1988, pp. 1529–1533.

[14] W. Lee and A. Kim, "An efficient automatic geo-registration technique for high resolution spaceborne SAR image fusion," in *Proceedings of the IEEE 2011 International Geoscience and Remote Sensing Symposium (IGARSS)*, vol. 1, Vancouver, Canada, Jul. 2011, pp. 3566–3569.

[15] H. Liu and K. C. Jezek, "Automated extraction of coastline from satellite imagery by integrating canny edge detection and locally adaptive thresholding methods," *Internationl Journal of Remote Sensing*, vol. 25, no. 5, pp. 937–958, Mar. 2004.

[16] J. Lee and I. Jurkevich, "Coastline Detection And Tracing In SAR Images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 28, no. 4, pp. 662–668, Jul. 1990.

[17] V. S. Frost, J. A. Stiles, K. S. Shanmugan, and J. Holtzman, "A Model for Radar Images and Its Application to Adaptive Digital Filtering of Multiplicative Noise," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 4, no. 2, pp. 157–166, Mar. 1982.

[18] C. W. G. Li, P. Huang, and W. Yu, "SAR Image Despeckling Using a Space-Domain Filter With Alterable Window," *IEEE Geoscience and Remote Sensing Letters*, vol. 10, no. 2, pp. 663–667, Mar. 2013.

[19] J. Y. Ji, J. M. Zhang, and X. Zhang, "A new CFAR ship target detection method in SAR imagery," *Acta Oceanologica Sinica*, vol. 29, no. 1, pp. 12–16, Jan. 2010.

[20] I. I. Lin, L. K. Kwoh, Y. Lin, and V. Khoo, "Ship and ship wake detection in the ERS SAR imagery using computer-based algorithm," in *Proceedings of the IEEE 1997 International Geoscience and Remote Sensing Symposium (IGARSS)*, vol. 1, Singapore, Singapore, Aug. 1997, pp. 151–153.

[21] M. Ferrara, A. Gallon, and A. Torre, "Improvement in automatic detection and recognition of moving targets in Alenia Aerospazio activity," in *Proceedings of SPIE 3500, Image and Signal Processing for Remote Sensing IV*, vol. 3500, Barcelona, Spain, Sep. 1998, pp. 96–103.

[22] C. Wang, M. Liao, and X. Li, "Ship Detection in SAR Image Based on the Alpha-stable Distribution," *Sensors*, vol. 8, no. 8, pp. 4948–4960, Aug. 2008.

[23] G. Gao, "A Parzen-Window-Kernel-Based CFAR Algorithm for Ship Detection in SAR Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 8, no. 3, pp. 557–561, May 2011.

[24] M. Tello, C. Lopez-Martinez, and J. J. Mallorqui, "A novel algorithm for Ship Detection in SAR imagery based on the Wavelet Transform," *IEEE Geoscience and Remote Sensing Letters*, vol. 2, no. 2, pp. 201–205, Apr. 2005.

[25] A. Gambardella, F. Nunziata, and M. Migliaccio, "A Physical Full-Resolution SAR Ship Detection Filter," *IEEE Geoscience and Remote Sensing Letters*, vol. 5, no. 4, pp. 760–763, Oct. 2008.

[26] H. Leung, N. Dubash, and N. Xie, "Detection of small objects in clutter using a GA-RBF neural network," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 1, pp. 98–118, Jan. 2002.

[27] G. Schwartz, M. Alvarez, A. Varfis, and N. Kourti, "Elimination of false positives in vessels detection and identification by remote sensing," in *Proceedings of the IEEE 2002 International Geoscience and Remote Sensing Symposium (IGARSS)*, vol. 1, Toronto, Canada, Jun. 2002, pp. 116–118.

[28] C. K. I. Williams, "Predicition with Gaussian Processes: From Linear Regression to Linear Predition and Beyond," Aston University, Birmingham, U. K., Tech. Rep. NCRG/97/012, 1997.

[29] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2008.

[30] J. V. Candy, *Bayesian Signal Processing: Classical, Modern and Particle Filtering Methods.* Hoboken, NJ: Wiley-Interscience, 2009.

[31] J. W. Goodman, "Some fundemental properties of speckle," *Journal of the Optical Society of America*, vol. 66, no. 11, pp. 1145–1150, Nov. 1976.

[32] J. B. Billingsley, A. Farina, F. Gini, M. Greco, and L. Verrazzani, "Statistical analyses of measure radar ground clutter data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 35, no. 2, pp. 579–593, Apr. 1999.

[33] S. Demirci, C. Ozdemir, A. Akdagli, and E. Yigit, "Clutter Reduction in Synthetic Aperture Radar Images with Statistical Modeling: An Application to MSTAR Data," *Microwave and Optical Technology Letters*, vol. 50, no. 6, pp. 1514–1520, Jun. 2008.

[34] S. Bocquet, "Calculation of Radar Probability of Detection in K-Distributed Sea Clutter and Noise," Australian Department of Defence, Edinburgh, Australia, Tech. Rep. DSTO-TN-1000, 2011.

[35] E. E. Kuruoglu and J. Zerubia, "Modeling SAR Images With a Generalization of the Rayleigh Distribution," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 527–533, Apr. 2004.

[36] N. J. Redding, "Estimating the Parameters of the K Distribution in the Intensity Domain," Australian Department of Defence, Edinburgh, Australia, Tech. Rep. DSTO-TN-0839, 1999.

[37] D. R. Iskander and A. M. Zoubir, "Estimation of the parameters of the K-distribution using higher order and fractional moments [radar clutter]," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 35, no. 4, pp. 1453–1457, Oct. 1999.

[38] D. A. Abraham and A. P. Lyons, "Reliable Methods for Estimating the K-Distribution Shape Parameter," *IEEE Journal of Oceanic Engineering*, vol. 35, no. 2, pp. 288–302, Apr. 2010.

[39] P. P. Gandhi and S. A. Kassam, "SAR Prescreening using Both Target and Shadow Information," *IEEE Transactions on Intelligent Transportation Systems*, vol. 40, no. 4, pp. 1226–1228, Jul. 1994.

[40] S. Kuttikkad and R. Chellappa, "Non-Gaussian CFAR techniques for target detection in high resolution SAR images," in *Proceedings of IEEE 1994 International Conference on Image Processing (ICIP)*, vol. 1, Texas, USA, Nov. 1994, pp. 910–914.

[41] D. T. L. Lee and A. Yamamoto, "Wavelet Analysis: Theory and Applications," *Hewlett-Packard Journal*, vol. 1, no. 1, pp. 44–54, Dec. 1994.

[42] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. New York, NY: Wiley-Interscience, 2001.

[43] R. M. Haralick and L. G. Shapiro, *Computer and Robot Vision, Volume I.* Boston, MA: Addison-Wesley, 1992.

[44] D. Comanicius and P. Meer, "Mean Shift: A Robust Approach Toward Feature Space Analysis," *IEEE Transactions On Pattern Analysis and Machine Intelligence*, vol. 24, no. 5, pp. 1–18, May 2002.

[45] L. G. Shapiro and G. C. Stockman, *Computer Vision.* Upper Saddle River, NJ: Prentice Hall, 2001.

[46] J. I. Marcum, "A Statistical Theory of Target Detection by Pulsed Radar," The RAND Corporation, Santa Monica, United States of America, Tech. Rep. RM-753, 1948.

[47] M. Rowbotham, *Introduction to Marine Cargo Management*, 2nd ed. New York, NY: Informa Law from Routledge, 2014.

[48] L. Ma and J. Yu, "Texture segmentation based on local feature histogram," in *Proceedings of IEEE 2011 International Conference on Image Processing (ICIP)*, vol. 1,

Brussels, Belgium, Nov. 2011, pp. 3349–3352.

[49] E. Jakeman and P. N. Pusey, "Mean Shift: A Robust Approach Toward Feature Space Analysis," *IEEE Transactions on Antennas and Propagation*, vol. 24, no. 6, pp. 806–814, Nov. 1976.

[50] M. Rouault, J. Johannessen, F. Collard, and S. Bernard, "Synthetic aperture radar products for the African marine environment," *South African Journal of Science*, vol. 105, no. 1, pp. 85–86, Mar. 2009.

[51] W. Kleynhans, B. P. Salmon, C. P. Schwegmann, and M. V. Seotlo, "Ship detection in South African oceans using a combination of SAR and historic LRIT data," in *Proceedings of the IEEE 2013 International Geoscience and Remote Sensing Symposium (IGARSS)*, vol. 1, Melbourne, Australia, Jul. 2013, pp. 1521–1524.

APPENDIX A

EVALUATION TOOLS FOR SYSTEM TESTING

A.1 RECEIVER OPERATING CHARACTERISTIC CURVES

A Receiver Operating Characteristic (ROC) curve is a pattern recognition tool used to determine the performance of a system that uses parameter values that can be varied. The ROC curve plots a method's Detection Accuracy (DA) on the vertical axis and the False Alarm Rate (FAR) on the horizontal axis. A corresponding 2D point is placed on the ROC curve for each threshold for a specific method. This is repeated for all thresholds for the method and this produces a curve within the ROC space.

Curves that lie directly above the diagonal curve have above average performance statistics. ROC curves that touch the top left corner of the graph and do not change as the thresholds are varied have perfect performance. Those that touch the diagonal line for all threshold values are considered to produce random results. The closer the graph is to the top left corner of ROC space the better its performance is across all threshold values. In other words, the Area-Under-Curve (AUC) indicates how close a specific method is to perfect DA and FAR. A method's performance across all thresholds improve as a method's curve AUC approaches unity.

A.2 HISTOGRAM COMPARISONS

Two histograms can be compared to one another using a histogram distance metric. Two common histogram comparison functions are the Bhattacharyya and Quadratic-Chi histogram

distance [48]. The Bhattacharyya distance is a bin-to-bin distance metric that is relatively simple to calculate. Assuming L bins in two normalised histograms p and q, we can be define the Bhattacharyya distance between the two histogram as

$$D_{bhat}(p,q) = \sqrt{1 - \sum_{i=1}^{L} \sqrt{p(b_i)\, q(c_i)}}, \tag{A.1}$$

where b_i and c_i are the ith bins in histograms p and q respectively.

Another popular histogram distance metric is the Quadratic-Chi histogram distance metric. The Quadratic-Chi is a cross-bin histogram distance metric and it is better at highlighting perceptual similarity compared to bin-to-bin metrics such as the Bhattacharyya histogram distance metric [48]. Assuming, once again, two normalised histograms p and q with L and O bins respectively and a bin-similarity matrix A, the Quadratic-Chi distance between two histograms can be defined as

$$D_{QC}(p,q,m,A) = \sqrt{\sum_{i}^{L}\sum_{j}^{O} \left(\frac{p(b_i) - p(c_i)}{\left(\sum_k (p(b_k) - p(c_k))\, A_{ki}\right)^m} \right) \left(\frac{p(b_j) - p(c_j)}{\left(\sum_k (p(b_k) - p(c_k))\, A_{kj}\right)^m} \right)}, \tag{A.2}$$

where m is the scaling factor chosen such that $0 \leq m < 1$ and $A_{ij} = 1 - \frac{min(|i-j|,T)}{T}$ where T is some threshold. For this dissertation, a similarity matrix by choosing $T = 10$ and $m = 0.9$ as selected in [48].