CLASSIFICATION OF OCEAN VESSELS FROM LOW RESOLUTION SATELLITE SAR IMAGES

by

Rory George Vincent Meyer

Submitted in partial fulfillment of the requirements for the degree
Master of Engineering (Electronic Engineering)

in the

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

UNIVERSITY OF PRETORIA

April 2017

SUMMARY

CLASSIFICATION OF OCEAN VESSELS FROM LOW RESOLUTION SATELLITE SAR IMAGES

by

Rory George Vincent Meyer

Supervisor(s):	Dr W. Kleynhans
Department:	Electrical, Electronic and Computer Engineering
University:	University of Pretoria
Degree:	Master of Engineering (Electronic Engineering)
Keywords:	Synthetic Aperture Radar, Automatic Identification System, Support Vector Machine, Regression, Classification, Maritime Domain Awareness, Sentinel-1, Low Resolution

In the long term it is beneficial to a country's economy to exploit the maritime environment surrounding it responsibly. It is also beneficial to protect this environment from poaching and pollution. To achieve this the responsible parties of a country must have an awareness of what is transpiring in the maritime domain. Synthetic aperture radar can provide an image, regardless of weather or light conditions, of the ocean showing most vessels therein. To monitor the ocean, using synthetic aperture radar imagery, at the lowest cost would require large swath synthetic aperture radar imagery. There exists a trade-off between large swath imagery and the image's resolution resulting in the largest swath image having the poorest resolution.

Existing research has shown that it is possible to use coarse resolution synthetic aperture radar imagery to detect vessels at sea, but little work has been done on classifying those vessels. This research aims to investigate the coarse resolution classification information gap. This is done by using a dataset of matching synthetic aperture radar and ship transponder data to train a statistical classification algorithm in order to classify or estimate the length of vessels based on features extracted from their synthetic aperture radar image.

The results of this research show that coarse resolution (approximately 40 m per pixel) synthetic aperture radar imagery is able to estimate vessel size for larger classes and provides insight on which vessel classes would require finer resolutions in order to be detected and classified reliably. The range of smaller vessel classes is usually limited to ports and fishing zones. These zones can be mapped using historical vessel transponder data and so a dedicated surveillance campaign can be optimised to use higher resolution products in these areas. The size estimation from the machine learning algorithm performs better than current techniques.

LIST OF ABBREVIATIONS

AIS	Automatic Identification System
CFAR	Constant False Alarm Rate
EEZ	Exclusive Economic Zone
ESA	European Space Agency
ETA	Estimated Time of Arrival
GMSK	Gaussian Minimum Shift Keying
GRDM	Ground Range Detected, Medium Resolution
HDLC	High Level Data Link Control
LRCS	Length Radar Cross Section
LRIT	Long Range Identification and Tracking
MDA	Maritime Domain Awareness
MPA	Marine Protected Area
MMSI	Maritime Mobile Service Identity
NMEA	National Marine Electronics Association
SAR	Synthetic Aperture Radar
SOTDMA	Self Organising Time Domain Multiple Access
SVC	Support Vector Classification
SVM	Support Vector Machine
SVR	Support Vector Regression
TRCS	Total Radar Cross Section
UTC	Coordinated Universal Time
VHF	Very High Frequency
VIIRS	Visible Infrared Imaging Radiometer Suite
VMS	Vessel Monitoring System
WRCS	Width Radar Cross Section

TABLE OF CONTENTS

CHAPTER 1 INTRODUCTION

1.1 PROBLEM STATEMENT

1.1.1 Context of the problem

The Exclusive Economic Zone (EEZ) is defined as the region of the ocean extending 200 nautical miles (nm) from a country's coastline. This region of the ocean is the area where a country has special rights regarding the exploration and use of any resources found in the water column or ocean floor. This includes fishing, oil or gas extraction, energy generation and other economic activities. Combined South Africa's large coastline and the Prince Edward Islands result in an EEZ of approximately 1.5 million square kilometres. This is split between the mainland EEZ (1,068,659 km^2) and the Prince Edward Island group (466,879 km^2) and is shown in Figure 1.1.

Approximately 2% of South Africa's gross domestic product is derived from fishing activities and a further 35% can be linked to global trans-oceanic trade [1]. There are also large reserves of natural gas south of Mossel Bay [2]. Being situated at the lowest point of Africa also results in a significant portion of the world's cargo and petroleum shipping moving through South Africa's EEZ. The EEZ of the Prince Edward Island group also contains South Africa's largest Marine Protected Area (MPA), home to large reserves of Patagonian toothfish [3]. The monitoring and enforcement of South Africa's laws proves to be especially difficult in this remote MPA. There are currently no coastal radars, vessel transponder receivers or patrol boats stationed near the Price Edward Islands.

All these factors indicate that South Africa has a legal and economic interest in activities taking place in the EEZ and that deriving more economic benefit from the ocean would require knowledge of current

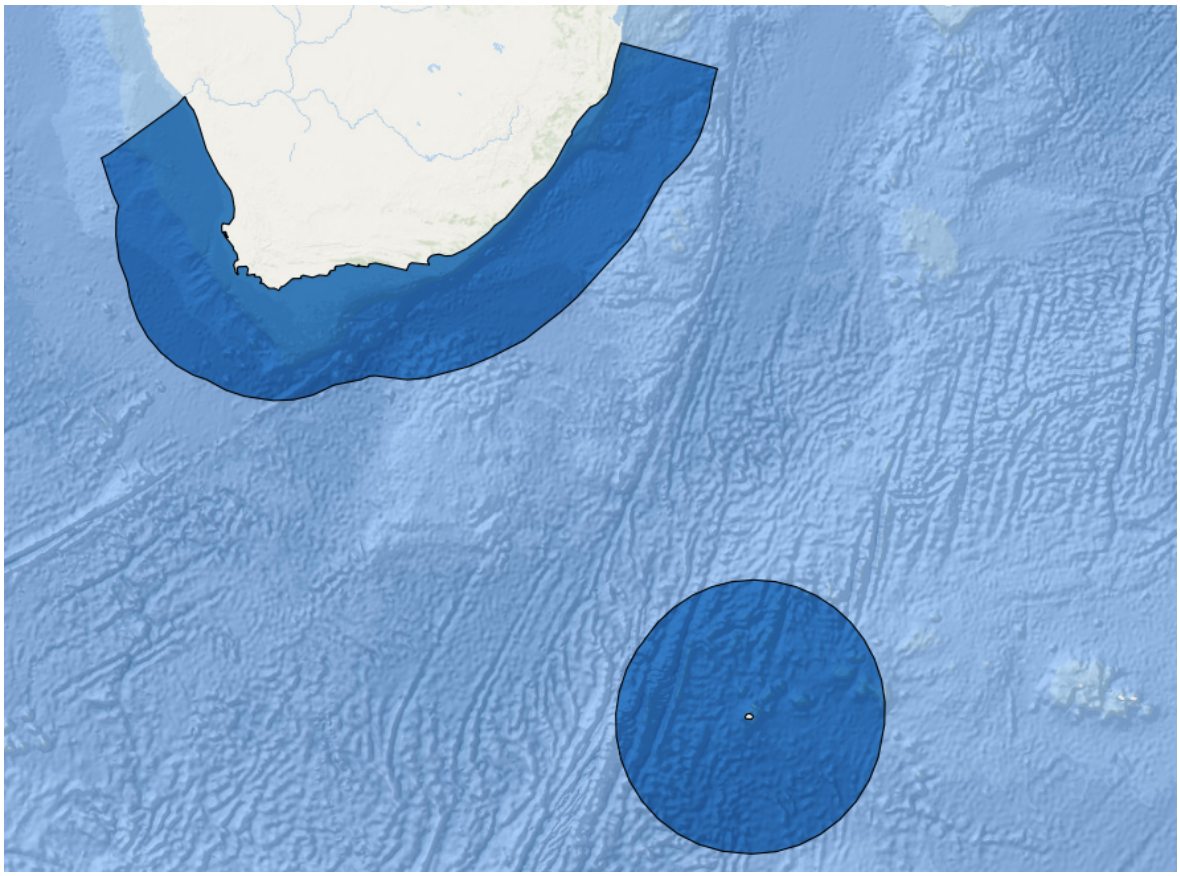

Figure 1.1. South Africa's two EEZ zones; the mainland and Prince Edward Islands EEZ.

activities and trends. Understanding shipping and fishing trends in the ocean would help regulators to better draft and enforce legislation.

Operation Phakisa is a government initiative to improve South Africa's ocean economy. This requires an improvement in stakeholder involvement and the ability to quickly and efficiently use data to make tactical decisions that impact multiple departments and users [4]. Currently there are multiple departments that have interest and responsibilities related to ocean governance. Operation Phakisa aims to integrate the requirements and goals of these users to unlock the potential of the ocean economy. To achieve this, the South African government will provide an integrated framework for governance, protection and marine spacial.

Maritime Domain Awareness (MDA) can be defined as the information required for enforcement, surveillance and pollution monitoring of a country's EEZ, coastal zones and MPAs. In South Africa this is currently performed using patrol vessels, radio transponders and self-reporting by civilian vessels.

This results in limited coverage, increased cost, and increased complexity due to uncoordinated schedules and functions of enforcement agencies [4].

MDA requires the ability to observe large portions of the EEZ and notify stakeholders should an event of interest occur, a task usually performed using vessel transponder data. Transponder data currently form the principal monitoring mechanism for deep ocean areas.

It has been shown that synthetic aperture radar (SAR) can be used to detect vessels at sea [5]. Extracting additional identifying information from the imaged vessels, while still imaging the largest possible area per image, would be beneficial to stakeholders concerned with illegal fishing, pollution monitoring, piracy, vessel traffic or search-and-rescue operations.

1.1.2 Research motivation

While transponder data can be used for MDA, it has limitations. Vessels can change information or disable their transponders relatively easily, but SAR can detect these uncooperative vessels. Other countries use SAR as a method for MDA [6], but currently South Africa only uses SAR for MDA in a limited and unofficial manner. It would be beneficial for South Africa, and other countries, to develop a method of using coarse resolution, large swath SAR imagery to monitor the ocean. Coarse resolution SAR imagery, in this context, refers to images with a pixel spacing of approximately 40×40 m or larger. Other researchers have demonstrated that it is possible to detect vessels with SAR, but it would be of value to also be able to determine a vessel's class or size from the SAR imagery [5].

Current research into this field is focused on high resolution imagery, this is largely due to the researchers having greater access to satellite platforms or smaller areas to monitor. Classification of ocean vessels using coarse resolution SAR imagery would be novel and directly applicable to MDA activities in South Africa.

1.2 RESEARCH QUESTIONS

The research questions are summarised as following:

- Can coarse resolution SAR images be used to classify ships around South Africa?
- Do coarse resolution SAR images contain enough information to estimate a vessel's defining characteristics?
- Which vessel classes are detectable and classifiable with this data?
- How many examples would be required to train a machine learning algorithm to classify or estimate vessel parameters?
- Can this be used in an operational system to provide information to users?

1.3 HYPOTHESIS AND APPROACH

The hypothesis for this research is:

There exists enough information in coarse resolution SAR images that a detected vessel's class can be determined, and physical characteristics such as size and bearing can be estimated.

This can be broken into several smaller parts:

- An algorithm exists to estimate a vessel's class (for example: fishing, cargo or tanker) from coarse resolution SAR imagery that performs better than random guessing.
- An algorithm exists to estimate a vessel's length from SAR imagery that performs better than pixel size estimations.
- An algorithm exists to estimate a vessel's bearing and position from SAR imagery.

To answer these questions, the following objectives must be met:

- A dataset must be created of various vessel types and classes that contains a vessel's SAR image and identifying information. This dataset must be representative of the vessel traffic seen in South Africa's EEZ.

- Features that describe a vessel must be extracted from the SAR images in the matched dataset.

- These features must be used to train an algorithm to predict or classify vessel labels based on features extracted from the associated SAR image.

- The performance of the algorithm will be evaluated using the vessels class and reported length.

1.4 RESEARCH GOALS

The goals of this research are to determine what extra information could be extracted from a vessel detected in a coarse resolution, large swath SAR image of the ocean. This information could be the vessel's class or physical attributes such as size or gross tonnage.

1.5 RESEARCH CONTRIBUTION

The successful completion of this research would add to the body of knowledge of using SAR imagery for MDA. This research is novel due to the focus on classifying vessels from large swath, coarse resolution SAR imagery instead of using high resolution imagery. A functional system based on this research would provide additional information to operators concerned with monitoring South Africa's EEZ.

1.5.1 Conference Papers

The following conference papers have been submitted and accepted and are a result of the research detailed in this dissertation:

- Meyer, Rory G. V., Waldo Kleynhans, and Colin P. Schwegmann. "Small ships don't shine: Classification of ocean vessels from low resolution, large swath area SAR acquisitions." Geoscience and Remote Sensing Symposium (IGARSS), 2016 IEEE International. IEEE, 2016.

- Meyer, Rory G.V., Waldo Kleynhans, and Colin P. Schwegmann. "Vessel Classification Features Using Spatial Bayesian Inference from Historical AIS Data." Geoscience and Remote Sensing Symposium (IGARSS), 2017 IEEE International. IEEE, 2017.

- Meyer, Rory G.V., Waldo Kleynhans, and Colin P. Schwegmann. "The Best of a Bad Situation: Optimising and Algorithm to Match Course Resolution SAR Vessel Detections to Sparse AIS Data." Geoscience and Remote Sensing Symposium (IGARSS), 2017 IEEE International. IEEE, 2017.

1.5.2 Journal Papers

The following journal paper has been submitted and is awaiting review:

- Meyer, Rory G. V., Waldo Kleynhans, and Colin P. Schwegmann. "Ocean Vessel Size Estimation from Coarse Resolution SAR Imagery." IEEE Geoscience and Remote Sensing Letters. Submission number GRSL-01021-2017.

1.6 OVERVIEW OF STUDY

This dissertation is split into several sections. Chapter Two is a literature study detailing the current methods used in MDA including the current state-of-the-art methods of detecting and classifying vessels from SAR imagery. Chapter Three describes the methods used in this research while Chapter Four contains a description of the results of the classification and a regression algorithms used to estimate the class and length of vessels. Chapter Five details the conclusions that can be derived from this research and how these tie into the problem statement described in this introductory chapter.

CHAPTER 2 LITERATURE STUDY

2.1 CHAPTER OBJECTIVES

The objective of this chapter is to provide an introduction to the principles and technologies currently used for MDA. This will include highlighting various shortcomings and trade-offs that must be accounted for in any research that involving these technologies. SAR and the ocean vessel Automated Identification System (AIS) are examined in greater detail as both were used in this research.

2.2 MARITIME DOMAIN AWARENESS

Maritime Domain Awareness can be defined as the knowledge of anything associated with the global or local maritime environment that could affect the security, safety, economy, or environment of the area concerned. The goal of MDA is to identify and intercept or mitigate possible threats before they can significantly impact a country. This would require timely interaction and information-sharing between various stakeholders often situated in different countries or continents. Since most maritime activities are performed by ocean-going vessels, the quick identification of vulnerabilities and threats related to vessels is necessary as well as increased transparency regarding cargo, passengers and crew [7].

2.2.1 Ocean zones

The laws applying to an ocean-going vessel are not as simple as those applying to a person or vehicle on land. A vessel in international waters, namely, the area beyond 12 nautical miles of a country's coast, is subject to the laws of the country with which that vessel is registered. This leads to a practice known as "flags of convenience" which is the practice of registering with countries that historically

do not enforce laws at sea. The only time a foreign flag vessel is subject to another country's laws is when it is within that country's territorial waters (within 12 nautical miles of its coast) or when it is exploiting economic resources within that country's EEZ. The EEZ is a region of 200 nautical miles around a country where that country has exclusive use of any natural resources that can be found above, within or on the floor of the ocean.

There are also specific zones of interest within a country's EEZ that apply to all vessels within that area. MPA's are one of these zones and function as nature reserves where specific rules regarding fishing or travel can be enforced.

2.2.2 Aspects of MDA

The main aspects of MDA can be separated into three categories:

- **Economic concerns**: The economic value of the ocean is composed of the combined value of fishing, trade and natural resource extraction as well as the services that support these industries. Having more information on these activities would allow a country to better exploit and protect them.

- **Maritime safety**: Smuggling and human trafficking can pose a risk to a country's citizens and must be monitored for the good of a country and its neighbours. Search-and-rescue operations rely on many of the same technologies used in MDA and have many of the same stakeholders.

- **Environmental safety**: Harbours, the EEZ, and marine protected areas all need to be continuously monitored for pollution events (such as oil slicks or bilge dumps), environmental events (such as harmful algal blooms) and illegal activities (such as illegal fishing, travelling through a no-go zone or pollution dumps) in order to intervene and minimise any damage caused.

2.3 CURRENT MONITORING TECHNOLOGIES

Current technologies used to monitor ocean vessels can be split into two categories; cooperative systems, where a vessel actively broadcasts its location and status, and non-cooperative systems, where

a sensor detects a ship regardless of any participation from the target vessel. There are cost and accuracy trade-offs involved in using either type of system.

2.3.1 Cooperative monitoring technologies

These methods often provide more information than non-cooperative technologies but require a vessel to accurately and consistently transmit its identity and status. Several technologies currently in use provide vessel monitoring capabilities to interested stakeholders. These technologies are used for specific vessel classes or activities rather than universally. For example, the Vessel Monitoring System (VMS) is used only by fishing vessels while AIS is not required for smaller vessels not involved in international travel.

2.3.1.1 Coastal AIS stations

AIS is used by vessels to automatically provide information to nearby vessels and coastal stations. The system was created to supplement marine radar in preventing collisions [8].

A vessel transmits an AIS message every few minutes, depending on activity, providing good temporal resolution. The message is transmitted through Very High Frequency (VHF) radio which prevents it from propagating over the horizon [9]. Thus, coastal stations usually cannot detect vessels that are more than 40 nautical miles away [10], limiting their ability to monitor the majority of the EEZ and deep ocean.

As with other opt-in systems, AIS is vulnerable to malicious alterations to the reported data by vessels that are engaged in illegal or dubious activities [11]. It is exactly these vessels that need to be monitored to ensure maritime security.

2.3.1.2 Satellite AIS receivers

Satellite AIS receivers overcome the limited range of coastal stations by having a moving reception footprint that can observe portions swaths of the ocean. Multiple satellites with multiple overpasses

each day provide near-continuous, global coverage of vessels that would otherwise be undetectable by coastal stations.

The satellite receiver's footprint often overlaps many vessels that are unaware of each other. The AIS protocol is self-organising to prevent time-domain message collisions, but only if vessels receive notifications of the timeslot that other vessels will transmit in. The large reception footprint often results in multiple vessels transmitting their message in the same time slot. The satellites are often in low-earth orbit and the amount of time that the receiver's footprint spends over a site is limited, so few AIS messages are received from each ship. Message collisions and short time periods spent over a site result in some vessels going undetected during a single satellite overpass. In a scene of 3000 vessels, the likelihood of a particular vessel being detected in a single overpass is under 15% [10].

Satellite receivers allow monitoring over the entire ocean but, as with coastal receivers, are vulnerable to spoofing. Should a vessel not transmit any AIS messages, due to either not having an AIS or purposefully disabling it, then AIS receivers are unable to monitor these vessels.

2.3.1.3 VMS

The Vessel Monitoring System (VMS) is a system used to track the location and activity of commercial fishing vessels. It is used to monitor the health of fish stocks and marine protected areas, and to track the catch reports and quotas used by fishing vessels. The information is collected in a fisheries monitoring centre and distributed to national institutions. The data collected contains confidential information and so is not available to the general public. Vessels have to register with the local monitoring centre before any information is shared.

South Africa's VMS is run by the Department of Agriculture, Forestries and Fisheries (DAFF) and provides information on the location, haul and activity of fishing vessels in the South African EEZ [12]. For foreign flag vessels to obtain a permit to fish within the EEZ, they must register with the Cape Town monitoring centre. The ship's VMS system sends data back to the central server using several different methods depending on the vessels requirements. The system transmits a log of activity so that historical data is accurate even in situations where messages are not received.

This system is used to monitor cooperative fishing vessels and cannot be used to monitor non-cooperative poachers, unregistered foreign flag vessels or non-fishing vessels.

2.3.1.4 LRIT

The Long Range Identification and Tracking (LRIT) system provides global identification and tracking of cooperative vessels. The obligations of vessels and participating governments is laid out in the 1974 of Safety of Life at Sea Convention [13, 14].

Ocean vessels propelled by mechanical means and more than 300 gross tonnes or engaged in international voyages should transmit LRIT information if the government of the flag they fly requires it [15]. This requirement to transmit can be bypassed by registering the vessel with a "flag of convenience" state that does not require vessels to transmit LRIT messages.

The messages are transmitted via on-board equipment through a service providers network. This is usually, but not always, performed using satellite internet.

2.3.2 Non-cooperative monitoring technologies

2.3.2.1 Coastal radar

Coastal radar systems are used, in conjunction with AIS, as a vessel-tracking system used by port authorities [11]. These systems are typically non-coherent maritime radars based on S-, X- and recently K-band technologies . The detection range varies from 5 nautical miles (nm) for small rubber duck vessels, to 20 nm for large tanker and cargo vessels.

The main limitation of these coastal radar systems is the small observed area, variable coverage, and weak returns from targets as they move out of range [16]. It should be noted that the focus of this study is the classification of vessels from large scale images and not from coastal radar sites.

2.3.2.2 Space based radar

Satellite based SAR can detect non-cooperative vessels. It can provide some information on the physical structure and bearing of the ship. It can also detect oil-on-water events, such as spills and bilge dumps [17]. A very large area can be imaged but there exists a trade-off between swath-width and pixel resolution. This results in the images of the largest area having pixel resolutions of several tens of meters.

SAR platforms cannot be used for real time monitoring as they have a revisit time of several days. The Sentinel-1 constellation consists of two platforms in the same orbit but delayed by 180 degrees. This halves the revisit time to several days for a position in the South African EEZ [18].

2.3.2.3 Space based optical

There are several optical and multi-spectral satellite platforms available that could be used to monitor the oceans around South Africa. VIIRS, MODIS and Sentinel-2 are examples of optical and multi-spectral satellite platforms. They are, however, vulnerable to cloud cover and low-light conditions.

Large swath images exist but offer poor resolution, so ships are often not visible. The vessels, when imaged optically, do not have as large a contrast between ocean and vessels as seen in SAR images. Day-night-band images can show brightly illuminated ships at night if moon conditions and cloud cover permit.

High resolution optical images could be used for MDA but require that the position of the vessel be known first so that the small swath area of the acquisition can be positioned over the vessel. This reduces the utility of high resolution optical images being used to monitor large areas for unexpected behaviour.

2.3.3 Summary

Table 2.1 shows a breakdown of the strengths and weaknesses of current MDA technology. Given the ability of SAR to detect non-cooperative vessels at any time or under any weather condition, it

Table 2.1. Technical abilities of various MDA technologies

	Coastal AIS	**Satellite AIS**	**LRIT**	**VMS**	**Coastal Radar**	**Space Based SAR**	**Space Based Optical**
Detect Non-cooperative vessels	No	No	No	No	Yes	Yes	Yes
Weather Independent	Yes	Yes	Yes	Yes	Yes	Yes	No
Day/Night Independent	Yes	Yes	Yes	Yes	Yes	Yes	No
Temporal Resolution	Minutes	Hours	Minutes	Minutes	Minutes	Days	Days
Detection Range	40 nm	Global	Global	Global	5-20 nm	Global	Global

becomes important to develop the ability to use it for MDA. AIS, LRIT and VMS are the current methods of monitoring vessels, but VMS is restricted to fishing vessels and is generally not available to the public. Typically coarse resolution SAR is only used for vessel detection, but this research aims to determine what extra information, useful to MDA stakeholders, can be extracted from these SAR detected vessels.

The remainder of this study will focus on AIS and SAR.

2.4 AUTOMATIC IDENTIFICATION SYSTEM

This section give a more detailed description of the AIS protocol, the information that can be extracted from it and the shortcomings of using it as a ground-truth for maritime activity analysis.

2.4.1 Introduction

AIS was originally designed as a system for ocean vessels to detect and avoid collisions with nearby vessels. The location of a vessel is transmitted every few minutes and this update rate is increased when it is manoeuvring. The protocol is a self-organised, time-division, multiple access system.

AIS is a legal requirement for vessels with a gross tonnage greater than 300 tonnes in international waters or 500 gross tonnes for passenger vessels and vessels not in international waters, but is often used by vessels falling below these thresholds [19]. The gross tonnage of a vessel is a non-linear measure of the total internal volume of the vessel and is not a measure of mass. AIS messages contain identity, location, activity and vessel class information, and are transmitted near real-time. AIS packets are transmitted via VHF radio and so reception is usually limited to line-of-site receivers.

2.4.2 Radio protocol

The AIS radio protocol is summarised in Table A.1 that can be found in Appendix A. The AIS protocol uses two channels to transmit but receivers can operate on either a single channel or both.

2.4.3 Message types

There are 27 different message types that are used in AIS. The most useful, for this application, are types 1, 2, 3 and 18 (Position Reports) and types 5 and 24 (Static and Voyage Related Reports) [20]. These provide information on the identity and position of the transmitting vessel. An exhaustive list of AIS message types, RF protocols used, and voyage and position report descriptions is shown in Appendix A, Tables A.1, A.2, A.3 and A.4.

2.4.4 AIS vessel classes

The AIS static and voyage related reports contain a "ship class" field that is used to identify a vessel's class. This is useful information to group vessels based on their size or expected behaviour. Table 2.2 shows the vessel groups. The vessel categories are also split based on the type of cargo that they are carrying but for brevity this is not shown in the table. When a vessel does not report its class or when that message was not received, it is classified as an "Unspecified Class".

2.4.5 Factors to consider when using AIS

2.4.5.1 Time accuracy

The AIS protocol was designed for real-time, or near real-time, use. The time a message is received is assumed to be within a few seconds, given retransmissions, of it being transmitted. As such, only a Coordinated Universal Time (UTC) second field is included in messages. The date, hour and minute are not transmitted. This is not usually an issue with coastal AIS receivers that can receive, record and transmit a message to a database through the internet within seconds. It can safely be assumed that the difference between the time a message was transmitted and its arrival at the database is insignificant given the speeds of ocean vessels. This allows the messages to be time-stamped on their arrival in a the local database.

This is no longer true when using satellite AIS receivers. Satellite receivers generally move in low Earth orbits and record AIS messages as they pass. They download these messages to ground stations

Table 2.2. Table of AIS class numbers with descriptions.

Code	Reported Ship Class
0	Unspecified (default)
1-19	Reserved for future use
20-29	Wing in ground (WIG), all ships of this type
30	Fishing
31	Towing
32	Towing: length exceeds 200m or breadth exceeds 25m
33	Dredging or underwater ops
34	Diving ops
35	Military ops
36	Sailing
37	Pleasure Craft
38-39	Reserved
40-49	High speed craft (HSC), all ships of this type
50	Pilot Vessel
51	Search and Rescue vessel
52	Tug
53	Port Tender
54	Anti-pollution equipment
55	Law Enforcement
56-57	Spare - Local Vessel
58	Medical Transport
59	Noncombatant ship according to RR Resolution No. 18
60-69	Passenger, all ships of this type
70-79	Cargo, all ships of this type
80-89	Tanker, all ships of this type
90-99	Other Type, all ships of this type

when their orbit takes them over them. The delay between the reception of the message and the recording at the database is approximately 45 minutes [21].

When using a combination of coastal- and satellite-based receivers to build up AIS tracks it becomes difficult to order the messages correctly. This could result in tracks that jump back and forth; giving incorrect position and track length estimates.

2.4.5.2 Reception frequency

Terrestrial AIS stations receive and record position reports constantly. This provides a dense map of vessel tracks and reduces any positional ambiguity that could result in mismatches or errors when combined with other data. Satellite receivers can only record a few messages from each vessel as they pass over. There are multiple satellites in commercial AIS constellations but the point density achieved still does not compare to that achieved by a coastal station.

Figure 2.1 shows AIS tracks created using both coastal and orbiting AIS receivers. There are significant gaps in the tracks of vessels in the deep ocean. This is usually acceptable as these vessels tend to travel in straight lines but this is not true for all vessels, such as working fishing vessels which do not tend to travel in straight lines when operating.

2.4.5.3 Spoofing

The problem with self-reporting systems is that they are vulnerable to various kinds of malicious spoofing [11]. Vessels can change their identity to a less suspicious class to avoid inspections while position information can be changed to make it appear as if the vessel did not move through a marine protected area. Alternatively the vessel's AIS can be switched off completely to hide it. These possibilities require an additional measure to ensure the validity of position and voyage reports. SAR can be used to detect the position of vessels at sea and can possibly, with further research, be used to determine the class, size and bearing of a vessel. This additional information can aid in identifying vessels that are spoofing their AIS class or position.

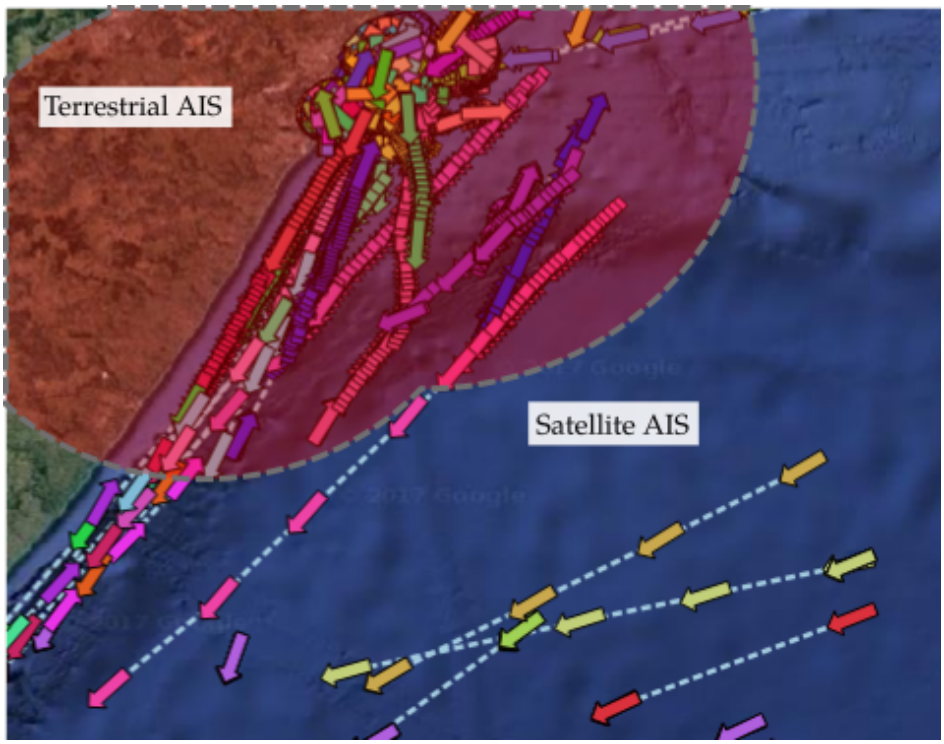

Figure 2.1. Satellite and coastal received AIS messages near Durban harbour received over a 4 hour period. Each message is coloured by its Mobile Maritime Service Identity (MMSI) and rotated by reported course.

2.5 SYNTHETIC APERTURE RADAR

2.5.1 Basic SAR operation

SAR is an active sensing method where a transceiver transmits several radar pulses at a target area and builds an image from the returns. SAR uses the properties of a moving antenna to provide finer resolution than could be expected from a beam steering radar. With SAR an area can be imaged in spite of any cloud cover or lack of light. Optical images measure the reflectivity and colour of objects while SAR platforms are sensitive to an object's physical structure and dielectric properties.

The range, R, from the SAR platform to the imaged object is determined by the normal radar equation using the speed of light (c) and time (t):

$$R = \frac{1}{2}ct \tag{2.1}$$

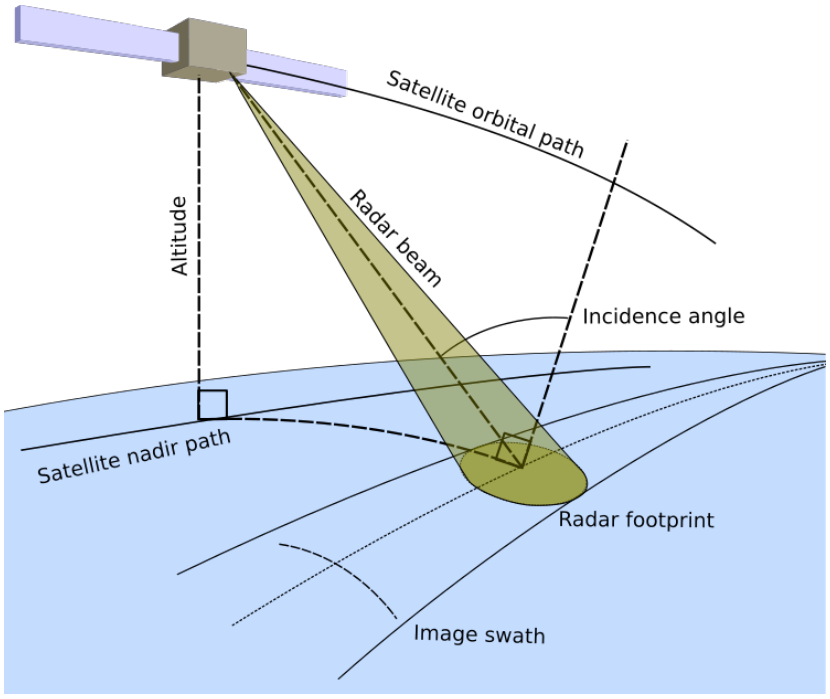

Figure 2.2. Figure with terms used to describe a SAR platform traveling over the curved surface of the planet.

The azimuth of the object, namely it's position ahead or behind the SAR platform, is determined by the Doppler shift of the return pulses. To avoid large ambiguities in azimuth and range, the satellite must look to the side instead of directly down. It is worth noting that this results in the image obtained from SAR platforms being in the slant range, or imaged from the perspective of the satellite. The image must be re-projected into the ground range, as if imaged from directly above, before comparing it to existing topographic or other ground projected data sources.

Figure 2.2 shows the definition of some basic terms used in describing SAR platforms and images.

2.5.2 Wave polarisation

The polarisation, either horizontal or vertical, of the transmitted and received radar waves have different interactions with imaged objects and can be sensed independently. This results in 4 possible polarisation modes; VV, HH, VH and HV. Cross polarisation is defined as VH or HV.

Research shows that cross polarisation has the highest ship-sea contrast for medium- to low-incidence angles [22]. For high-incidence angles (> 60 degrees) HH performs better than cross polarisation for the detection of ocean vessels. Sentinel-1 has an incidence angle limited to 47 degrees for the ground range detected, medium resolution (GRDM) product and therefore using the cross-polarisation product should detect the most vessels in a given image.

2.5.3 SAR resolution

Figure 2.3 shows vessels and harbour scenes at various resolutions. These images were taken from different satellite platforms using different bands and imaging modes. The trade-off between swath size and vessel resolution becomes apparent when comparing these different products. The amount of information that could be extracted from the vessels is also severely reduced at coarse resolutions.

2.5.4 SAR imaging bands

SAR platforms can be broadly separated into groups based on the centre frequency of the radar pulse. The most common, commercially available bands are X-band (8 - 12 GHz), C-band (4 - 8 GHz) and L-band (1 - 2 GHz). The wavelength of the band affects the way the pulse interacts with physical structures in a similar size range. X-band is more affected by atmospheric moisture conditions [23], while L-band has trouble imaging oil-on-water events [17], which typically have a very low return. For general maritime surveillance, C-band is a good choice as it provides high contrast between the ocean surface and vessels [24].

2.5.5 Imaging phenomena with respect to objects in the ocean

The aim of this section is not to provide a comprehensive description of SAR operation but rather to show examples of some artefacts that could cause errors when processing SAR images for MDA. The purpose of this research is aimed solely at maritime monitoring and therefore SAR phenomena that occur when imaging cities, vegetation or mountains are not discussed.

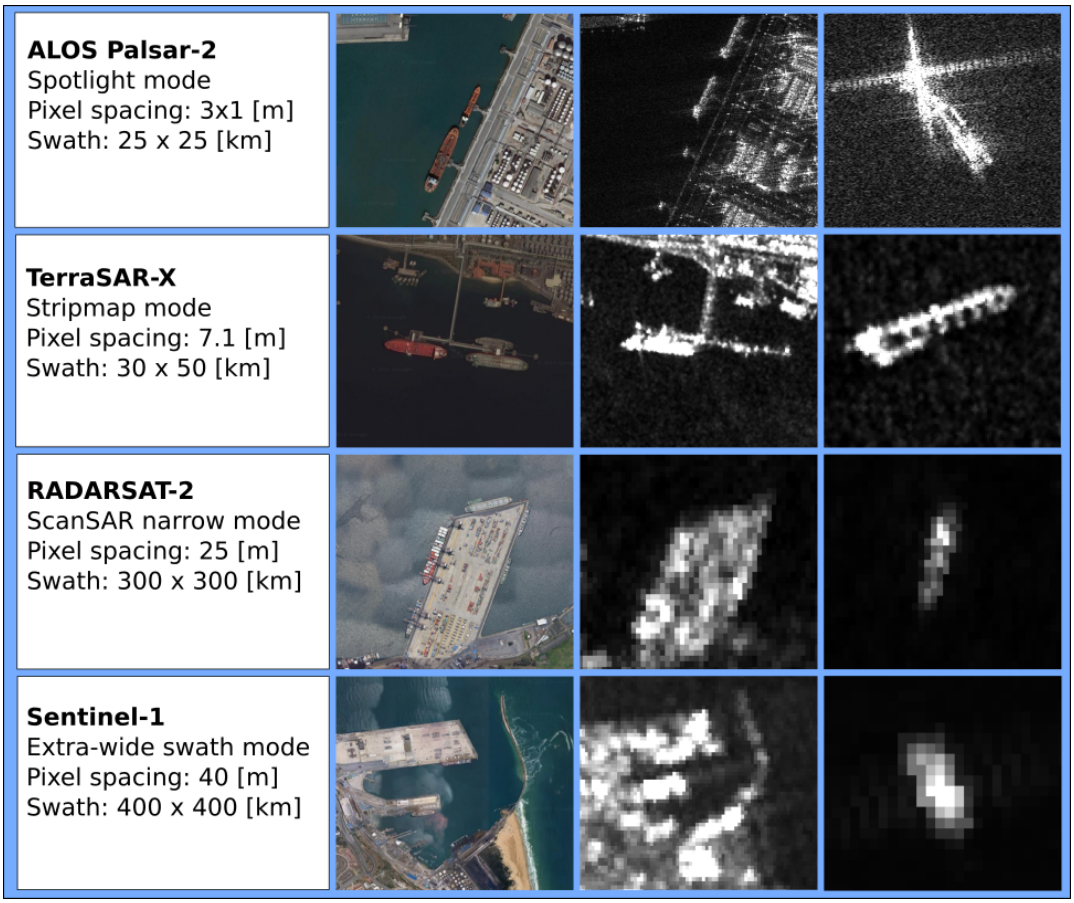

Figure 2.3. Comparison of various SAR resolutions. The images, from left to right, include a Google Earth harbour scene, the associated SAR scene and a vessel at sea. The optical and SAR scenes were not imaged at the same time.

2.5.5.1 Speckle

Speckle is a noise-like variation in brightness seen in all pixels of a SAR image. On a flat surface with uniform backscatter properties it would appear in as a brightness variation with a constant mean over the pixels. Mathematical analysis shows that this noise like signal has well defined statistical properties that can be compensated for by averaging the brightness over several neighbouring pixels. The number of statistically independent pixels averaged this way is called the number of looks, N. It can be shown that the signal standard deviation, S_N, is related to the mean signal power, P:

$$S_N = P/\sqrt{N} \tag{2.2}$$

A greater number of looks improves the radiometric value of a pixel but reduces the spacial resolution [25].

2.5.5.2 Range and azimuth ambiguity

Range ambiguity results in copies of bright objects in the range direction of a SAR image and is caused by overlapping return signals from different pulses that return to the receiver antenna at the same time. This usually occurs when there are high backscatter returns in a side lobe (mountains or buildings) while there are low levels of return in the main lobe (ocean or water scenes). SAR processing cannot correct this effect.

Azimuth ambiguity causes copies of bright objects in the azimuth directions of a SAR image. This occurs when there are multiple returns from a single pulse with multiple Doppler frequencies. This usually occurs in conditions where a side lobe has a high level of returns while the main lobe has a low backscattering level [26].

Figure 2.4 shows examples of range and azimuth ambiguity in a 2016 RadarSat 2 image that was acquired near Durban harbour. Sub-images A, B and C are of the same buildings south of the main harbour but with different contrast levels, while sub-image E is a close-up of the buildings. Sub-image D is a section of open ocean east of the harbour and can be seen to include an echo of E caused by range ambiguity. In A and B, some azimuth ambiguity can be seen as there are echoes of the buildings in the ocean, south of the harbour.

These effects can cause false vessel detections when an area of low-backscatter ocean is imaged where a portion of the swath includes high-backscatter buildings or mountains.

2.5.5.3 Radar incidence angle

The ocean has a large dielectric constant and so is a good reflector of radar pulses. The reason that the ocean usually appears dark in a SAR image is that most of the radar energy is reflected away from the satellite platform. When a vertical structure, such as a ship, is imaged in a reflective horizontal plane, such as the ocean, double bouncing of the radar signal can be seen as a reflection of the vessel

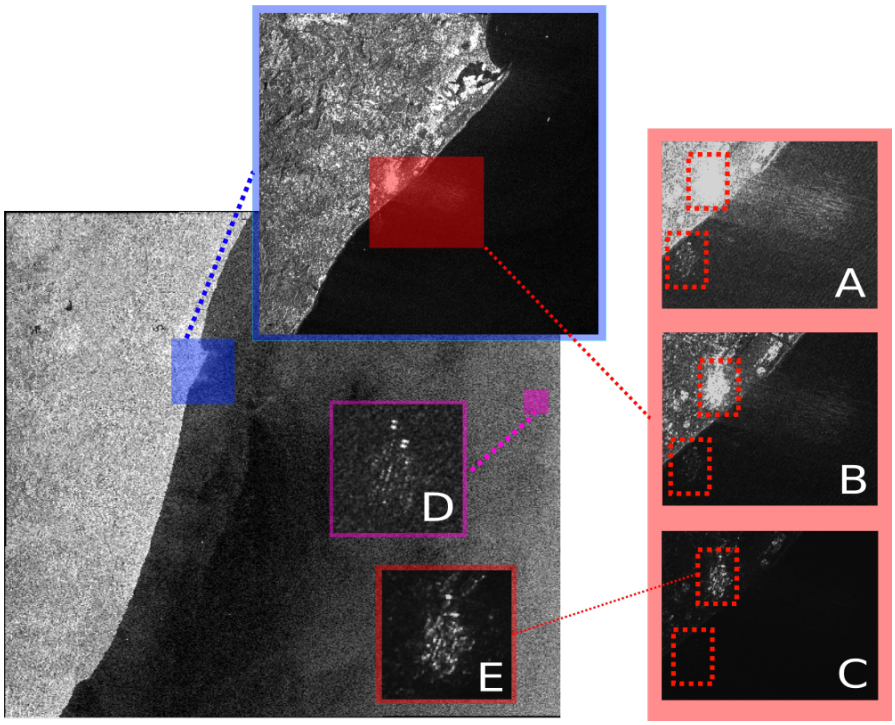

Figure 2.4. Radarsat 2 image of Durban harbour showing some examples of SAR artefacts. Sub-images A, B and C show some buildings near the port with an echo effect in the ocean to the South caused by range ambiguity. Sub-image D shows an echo effect caused by azimuth ambiguity and, by comparing it with subimage E, is due to the strong backscatter return from the building near the port.

in the ocean. This would exaggerate the length or width of the vessel, depending on the bearing of the satellite and vessel. This effect can be seen in Figure 2.5.

2.5.5.4 Motion of the ocean (and vessels too)

An SAR acquisition is performed over several minutes and using many pulses. If a target moves in the slant range during acquisition, this results in azimuth shifts due to the local Doppler estimation being incorrect. For ocean vessels, this becomes more complex because of the six different degrees of motion due to sea surface dynamics. The most significant effect is caused by pitch and can result in azimuth length errors exceeding 100% [27]. Errors in the estimation of the Doppler frequency due to vessel pitching can also result in "ghosts" along the azimuth of the radar image. This can be seen in Figure 2.6.

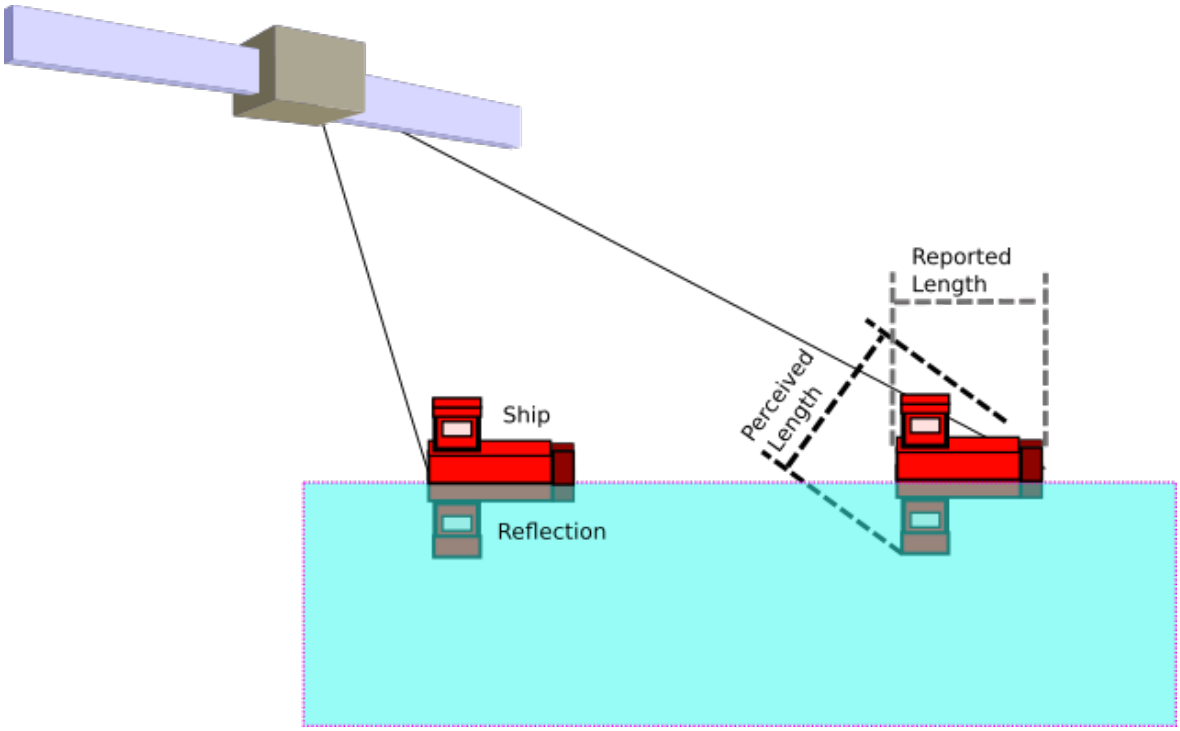

Figure 2.5. Exaggeration of vessel dimensions due to radar incidence angle.

2.5.6 Current SAR platforms

The goal of this research is to extract, in addition to its location , any additional information related to a detected vessel from a SAR image with the largest possible swath. There are several commercial and research SAR platforms that could meet this goal, but the cost of obtaining multiple images that overlap areas that have AIS data can be prohibitively expensive. Table 2.3 shows several examples of SAR platforms.

The ideal platform would:

- have a very large image swath,
- have regular acquisitions around South Africa,
- operate in the C-band,

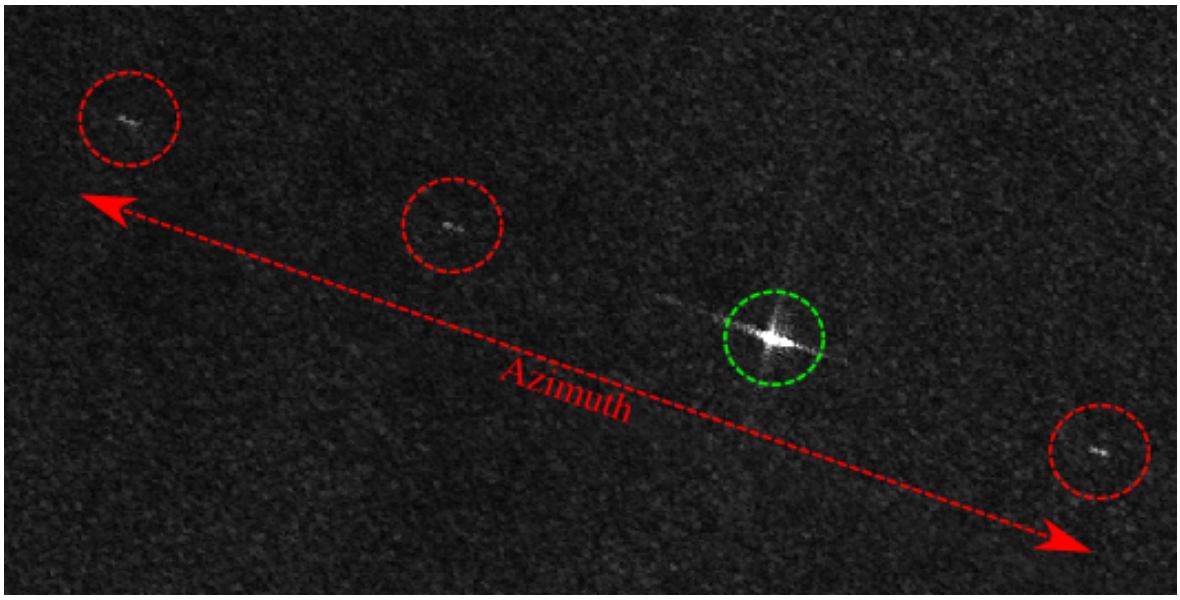

Figure 2.6. SAR detected vessel (green circle) with ghosts along azimuth (red circles).

- and have cost effective pricing on multiple images.

Given that Sentinel-1 uses South Africa as a calibration site, with free data access and has a large image swath, the decision was made to exclusively use Sentinel-1 imagery in this study.

2.6 SHIP DETECTION WITH SAR

SAR is useful for ship detection because the backscatter return from vessels is usually orders of magnitude larger than that from the ocean surface. This makes detecting vessels in the ocean relatively easy to automate. The most common method of detecting vessels in a SAR image is by using a Constant False Alarm Rate (CFAR) [5, 24].

2.6.1 Basic CFAR operation

The detection of vessels using a CFAR method requires selecting a threshold to keep the false alarm rate constant over the whole image. A sliding spatial window is used to test each portion of the image. The window is used to select the background area and cell under test throughout the image. An example

Table 2.3. Comparison of several commercial and scientific SAR satellite platforms

		Sentinel-1	ALOS PalSAR-2	COSMO-SkyMed	RADARSAT-2	TerraSAR-X	EnviSAT
Satellites in Constellation		2	1	4	1	2	1
Launch Date		2014-2016	2014	2007-2010	2007	2007-2010	2002-2012
Public Data Access		Yes	No	No	No	No	Yes
Best Resolution	**Pixel Spacing [m]**	5	7	1	<1	<1	30
	Swath [km]	80	40	10	8	4	58
Best Swath	**Pixel Spacing [m]**	40	100	100	50	40	150
	Swath [km]	400	350	200	500	270	405
Band		C-Band	L-Band	X-Band	C-Band	X-Band	C-Band
Taskable		No	Best Effort	Yes	Yes	Yes	No
Managing Body		ESA	JAXA	ISA	CSA	DLR	ESA

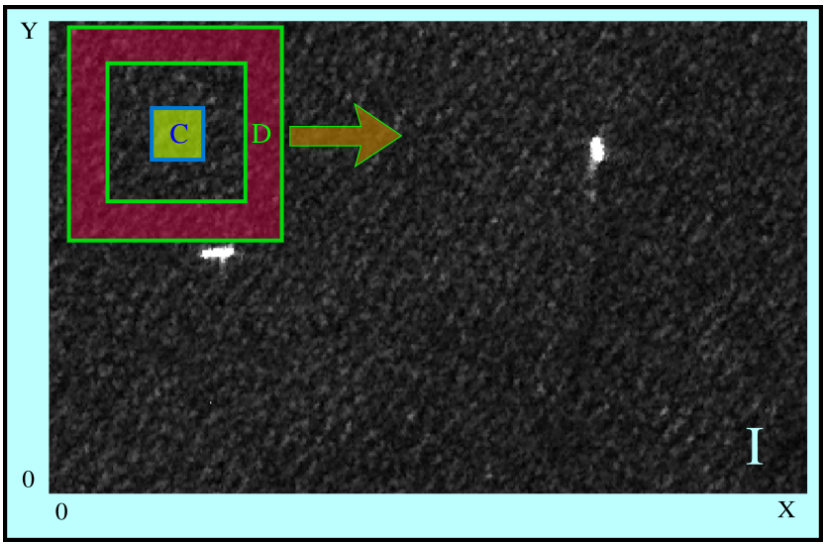

Figure 2.7. CFAR window moving over the SAR image I. Centre window cell (C) is the region of interest while the red box (D) contains the neighbouring pixels providing a background area.

where a vessel is in the region of interest is shown in Figure 2.7. The cell under test and background area can be represented by multiple pixels [5].

Let the SAR input be represented by a backscatter intensity image, I, which is $X \times Y$ pixels large. Also let the cell under test be represented by C and be $X_{ROI} \times Y_{ROI}$ pixels large. Let the background area be D, an area of pixels surrounding C with a guard window between the two regions.

Each pixel of the image represents the total radar backscatter received at the SAR platform's antenna.

$$I = \begin{bmatrix} I_{11} & I_{12} & \dots & I_{1Y} \\ I_{21} & I_{22} & \dots & I_{2Y} \\ \vdots & \vdots & \ddots & \vdots \\ I_{X1} & I_{X2} & \dots & I_{XY} \end{bmatrix} \tag{2.3}$$

The CFAR produced output image, $J(I,x,y,T)$, is a binary image where "true" values represent detected vessel pixels and "false" values represent background ocean or masked land. T is the CFAR threshold and is inversely proportional to the number of permissible false alarms. The binary CFAR image is produced with:

$$J(I,x,y,T) = \begin{cases} True & B(x,y) > T \\ False & \text{otherwise} \end{cases} \tag{2.4}$$

where $B(x,y)$ is the CFAR function defining the interaction between the cell under test and the background area. The method used to describe the interaction between C and D defines the CFAR method. The "Greatest Of", "Least Of" or "Cell Average" CFAR algorithms calculate $B(x_n,y_n)$ from the surrounding area using the following equations:

$$\text{GO-CFAR: } B(x,y) = \frac{Max(C)}{Max(D)} \tag{2.5}$$

$$\text{LO-CFAR: } B(x,y) = \frac{Min(C)}{Min(D)} \tag{2.6}$$

$$\text{CA-CFAR: } B(x,y) = \frac{Mean(C)}{Mean(D)} \tag{2.7}$$

Figure 2.8 shows the binary output of the CFAR algorithm. To avoid single bright pixels being detected as vessels the detections are grouped together and groups with too few connected pixels are discarded.

Using metadata included with the geocoded SAR image the pixel locations of the CFAR detections, (x,y), can be converted to spacial coordinates.

2.6.2 Factors to consider when using CFAR

Using a CFAR method to detect vessels in SAR images has drawbacks. This section describes some of them.

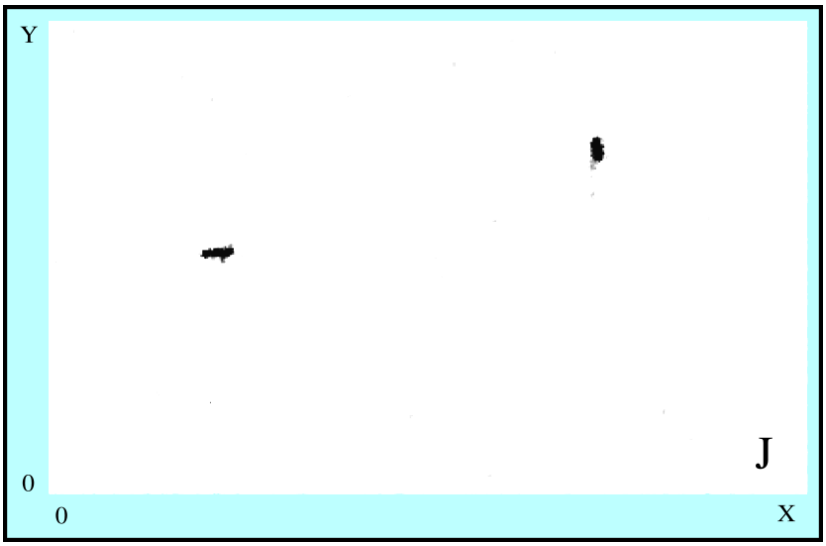

Figure 2.8. Binary output, *J*, of the CFAR algorithm. "True" values are shown as black while "false" are shown as white.

2.6.2.1 Vessel size

SAR speckle might result in some very bright single pixels being detected as vessels when using a low CFAR threshold. To minimise these a minimum detection size must be applied to groups of pixels detected as possible ship pixels. CFAR detections are only accepted if they extend to several pixels. This reduces false detections but also reduces the ability to detect small vessels.

2.6.2.2 Vessel backscatter

Smaller vessels have smaller reflective surfaces and so reduced backscatter. Large vessels that are angled to reflect radar energy away from the sensor also appear dimmer. To detect these vessels, the CFAR threshold must be decreased but this is likely to increase the number of false detections.

2.6.2.3 SAR artefacts

Range, azimuth and Doppler ambiguities introduced in the SAR processing appear as bright spots in a SAR image. The CFAR algorithm cannot distinguish between these artefacts and standard vessel detections. There are methods to distinguish between these positive detections and false positives

using deep neural networks [24]. This method will have to be considered when using a classification algorithm in real-world applications.

2.7 SHIP CLASSIFICATION AND IDENTIFICATION

The automatic detection, classification and identification of targets from SAR images is a popular field of research in military applications. There are many different techniques that are applied to the problem but the essentials can be distilled down to a few steps:

- Detect possible vessels and discriminate between vessels and lookalikes that appear in the SAR image [5, 13, 16, 22, 24].
- Extract features describing the vessel from the vessel's SAR image [28, 29, 30].
- Match detected vessel with known vessel or vessel class, usually using cooperative transponder data [30].
- Use SAR features to train an algorithm to classify or estimate characteristics of the detected vessels automatically. Compare these results with matched vessels to evaluate algorithm performance [22, 28, 29, 30, 31, 32].

The features extracted by the various studies shared some similarities. This is a list of features chosen by various studies:

- Segmented radar cross section [28, 29, 30]
- Angle of incidence of radar beam [22, 28, 29, 31, 32]
- Bearing of vessel [22, 28, 29, 31, 32]
- Width and length of vessel [22, 29, 30, 31]

- Ship main structure position [22, 31]

- Extract features in multiple polarisations [22, 32]

The basis for most of the classifiers was that each class of vessel had large physical structures in similar positions along their length that could be used to identify that class. This required segmenting the radar cross-section to identify the position of large scatterers. By combining this with the width and length of the vessel, the vessels could be grouped by size and then by position of main structure. The bearing, angle of incidence and multiple polarisation features were used to correct for effects caused by the SAR processing.

The features extracted and the machine learning methods may vary from study to study but the common thread is that the aforementioned research was done on high resolution images. Extracting information or classifying vessels with coarse resolution data has not been a priority for the SAR community. A 2006 study [33] stated that reliably extracting the length of a SAR detected vessel would require a pixel resolution better than 10 m.

The effects caused by imaging a non-stationary target situated above a non-ideal reflective surface along with the artefacts introduced in the SAR processing chain, requires a classification algorithm that is able to use linearly inseparable features. Given the scarcity of freely available SAR data and the high cost of tasking satellites to obtain desired data it is desirable that the algorithm not require many training examples to achieve good performance. Having few examples would prevent neural network or deep learning algorithms from being used, while linearly inseparable features require a method that can operate in multi-dimensional feature space. For these reasons, a Support Vector Machine (SVM) was chosen as the algorithm to use in this study. It is probable that other algorithms, such as the Random Forest Classifier or Gradient Boosting, could perform as well or better than an SVM but an SVM was chosen to be the initial testing benchmark. Future research will compare the SVM algorithm with other options, but currently determining to what extent coarse resolution SAR is useful for vessel classification and MDA purposes in general is valuable.

This research focuses on SAR features extraction, comparing those features to relevant AIS labels, and determining the ability of those features to predict a vessels class or size when used in an SVM algorithm.

2.8 SUPPORT VECTOR MACHINE

2.8.1 Linear separable classification

An SVM is a statistical method of separating classes by determining a hyperplane that would best separate the known classes based on their position in feature-space. Figure 2.9 shows a linearly separable 2-class problem. The two classes, class-1 and class-2, have features described by the matrix x. The classes are separable by the hyperplane defined by the equation:

$$(w \cdot f_n) + b = 0 \tag{2.8}$$

The classes are separated by whether they fall above or below the hyperplane. The hyperplane fit is optimised by maximising the margin M between the points p_1 and p_2 from the two class samples [34]. The equations for two margin lines can be rescaled to become:

$$(w^T \cdot x) + b = 1 \tag{2.9}$$

and

$$(w^T \cdot x) + b = -1 \tag{2.10}$$

giving a margin equation of:

$$m = \frac{2}{||w||} \tag{2.11}$$

The margin equation should classify each point correctly into its relevant class; y_i:

$$y_i(w^T + b) \geq 1 \forall i \in (1, 2 \ldots n) \tag{2.12}$$

where n is the number of points in the feature space.

The hyperplane boundary can be solved by the following constrained optimisation problem:

$$\underset{w}{\operatorname{argmin}}\left(\frac{1}{||w||^2}\right) \quad \text{while} \quad y_i(w^T + b) \geq 1 \quad \forall \quad i \in (1, 2 \ldots n) \tag{2.13}$$

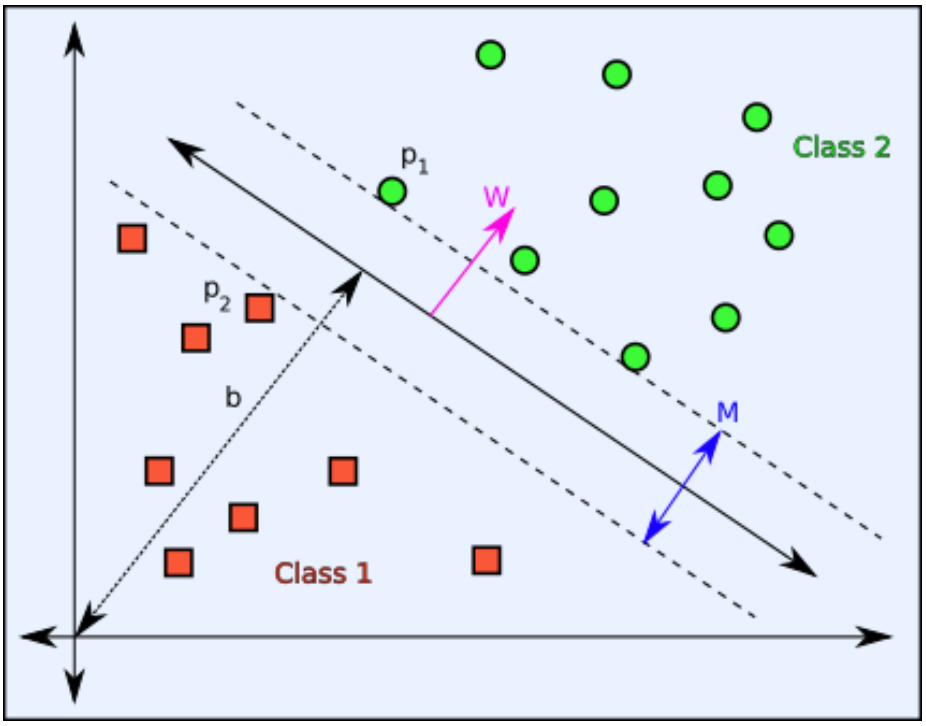

Figure 2.9. Linearly seperable classes shown in feature-space.

2.8.2 Kernel trick

Should the classes not be linearly separable, as in Figure 2.10, the SVM can implement a kernel to explicitly discover decision boundaries with arbitrary shapes by examining the data in a higher dimension [35].

This is done by mapping the n dimensional data to $(n+1)$ dimensions. The two dimensional problem above can be mapped to the third dimension with:

$$[x_1, x_2] \rightarrow [z_1, z_2, z_3] = [x_1^2, \sqrt{2}x_1x_2, x_2^2] \tag{2.14}$$

Equation (2.14), when combined with Equation (2.8), and assuming $b = 0$ for brevity, will result in a hyperplane boundary defined by a paraboloid, Equation (2.15), shown in Figure 2.10.

$$w_1x_1^2 + w_2\sqrt{2}x_1x_2 + w_3x_2^2 = 0 \tag{2.15}$$

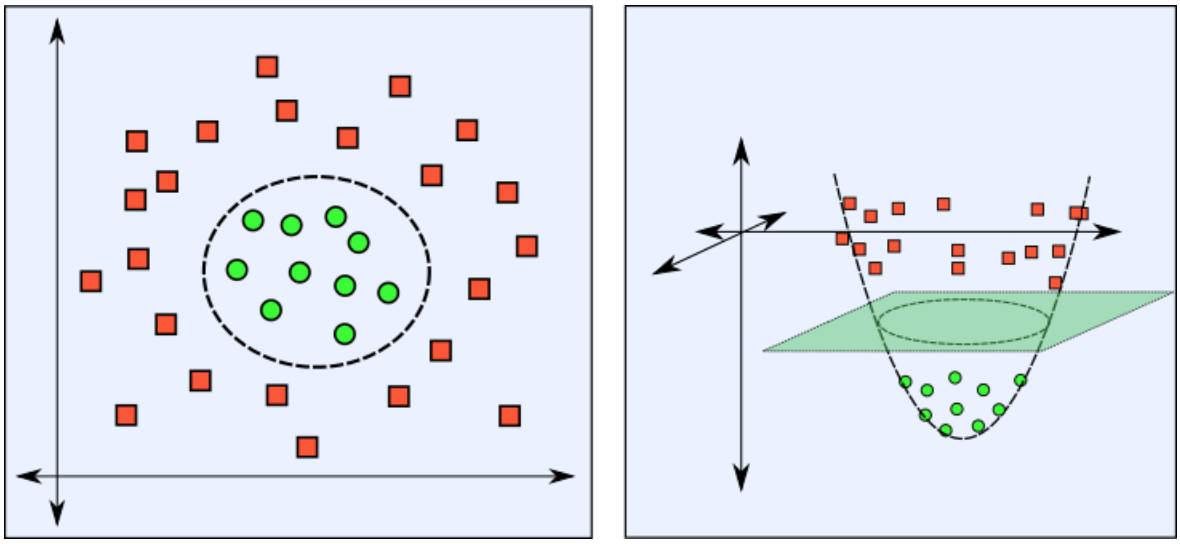

Figure 2.10. Linearly inseparable classes shown in feature-space.

This transformation can be written in matrix notation as:

$$K = \begin{vmatrix} x_1^T x_1 & x_1^T x_2 & \cdots & x_1^T x_p \\ x_2^T x_1 & \ddots & & \\ \vdots & & & \\ x_n^T x_1 & \cdots & & x_n^T x_p \end{vmatrix} = XX^T \tag{2.16}$$

This method of linearly mapping an n -dimensional state vector to $(n+1)$ dimensions is called the linear kernel, $K_{Linear}(X, X^T)$.

The Radial Basis function is a popular kernel used in SVMs:

$$K_{RBF}(X, X^T) = exp\left(-\frac{||X - X^T||^2}{2\sigma^2}\right) = exp(\gamma||XX^T||^2) \text{ for } \gamma > 0 \tag{2.17}$$

There are many other types of kernel that can be used with an SVM:

$$K_{Linear}(X, X^T) = XX^T \tag{2.18}$$

$$K_{Sigmoid}(X, X^T) = tanh\left(\gamma XX^T + r\right) \tag{2.19}$$

$$K_{Polynomial}(X, X^T) = \left(\gamma XX^T + r\right)^d, \gamma > 0 \tag{2.20}$$

2.8.3 Multiple classes

SVM is often used for multi-class problems where a single hyperplane threshold would not work. There are two main methods of overcoming this problem: "one-vs-all" and "one-vs-one" SVM methods [36]. The "one-vs-all" method constructs k SVM models where k is the number of classes. Each model is trained using one class as "positive" and all other classes as "negative" labels. When evaluating a sample using a trained "one-vs-all" model the class of the sample is determined by the SVM model that provides the largest value of its decision function.

The "one-vs-one" method creates a classifier that separates each possible combination of classes:

$$\text{Number of SVM models } = k(k-1)/2 \tag{2.21}$$

When considering a sample using the trained "one-vs-one" SVM model, the sample is given to each of the SVM models and the result of each model is tallied. The class that receives the most votes is used as the classification result.

2.8.4 Regression

Support Vector Regression (SVR) is a method of using an SVM to estimate the value of a sample, based on its support features. This is similar to the classification algorithm but, instead of calculating a hyperplane function that separates the classes, a function is calculated that fits the data as linearly as possible [37]. The function does not have to be linear in the same dimensions as the samples and, by using the kernel trick, the linear function can be folded into another dimension. This is shown in Figure 2.11, where ζ is the Euclidean distance from the fitted function that is above the tolerance variable, ε.

The equations for the SVR are similar to the equations used in a Support Vector Classification (SVC) algorithm. For a problem with n, (x_i, y_i) points, the regression function is obtained by minimising the following function [37]:

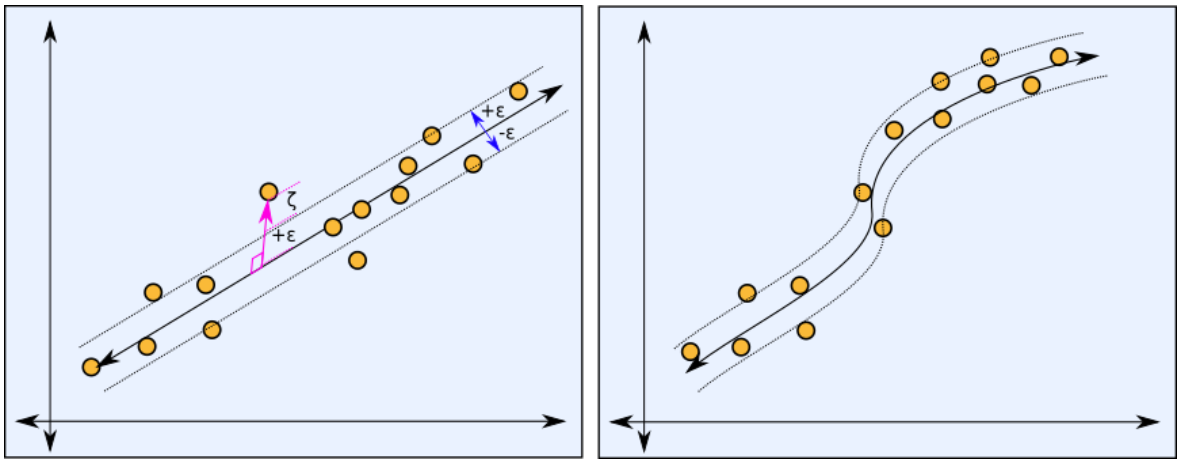

Figure 2.11. Regression examples with SVR

$$\underset{w}{\operatorname{argmin}}\left(\frac{1}{||w||^2}+C\sum_{i=1}^{n}(\zeta_i+\zeta_i^*)\right)$$
$$\text{subject to}\begin{cases} y_i - wx_i - b \leq \varepsilon + \zeta_i \\ -y_i + wx_i + b \leq \varepsilon + \zeta_i^* \\ \zeta, \zeta^* \geq 0 \end{cases} \tag{2.22}$$

The constant C, being greater than 0, determines the trade-off between the flatness of the regression function and the extent that deviations larger than ε are tolerated.

CHAPTER 3 METHODS

3.1 CHAPTER OVERVIEW

This chapter details the method used to create, train and evaluate the classification and regression algorithms. These algorithms were used to obtain more information on ocean vessels that have been imaged by satellite SAR.

Figure 3.1 shows the path taken from the raw data to the machine learning results. Each step will be discussed in this chapter.

3.2 SENSORS

3.2.1 Sentinel-1 satellite platform

Sentinel-1 is a constellation of two C-band, SAR imaging platforms that form part of the European Space Agency's (ESA) Copernicus Earth observation programme. The satellites orbit the earth 180 degrees apart and can image the whole planet every six days. Sentinel-1 data was used for this research because it is freely available, uses the South African coast as a calibration site and has many images of the South African EEZ available.

105 images with both VV and VH polarisations were used in this study. The images were acquired between October 2014 and November 2015. For consistency only the Extra-Wide, Ground-Range Detected, Medium resolution (EW-GRDM) product was used. The EW-GRDM product was chosen as there were many images acquired around southern Africa at the time of this study. This also helps

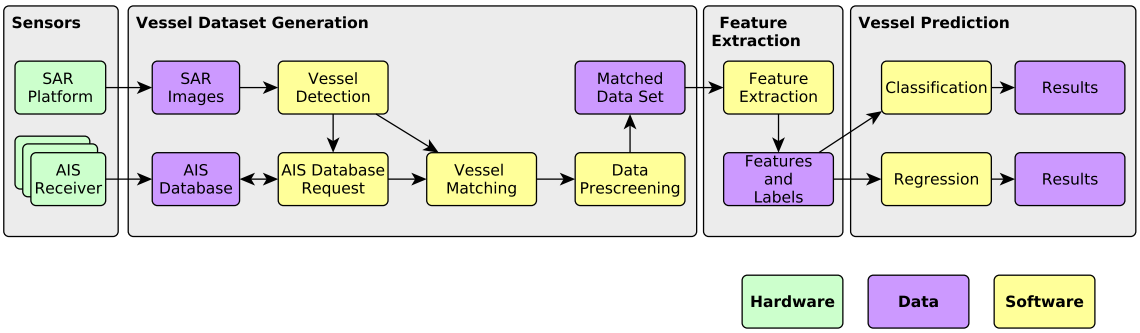

Figure 3.1. Flowchart of the method used to transform sensor data into vessel predictions.

Table 3.1. Sentinel-1 product parameters

Parameter	Value
Product	EW-GRDM
Swath Width [km]	400
Incidence Angle Range [deg]	18.9 - 47.0
Resolution [Rg x Az][m]	93 x 87
Pixel Spacing [Rg x Az][m]	40 x 40
Number of Looks [Rg x Az]	6 x 2
Effective Number of Looks	10.7

by reducing classification inaccuracies caused by features of varying quality or resolution. Further research can be done to determine the viability of classification for each SAR product. Table 3.1 lists the parameters of the Sentinel-1 EW-GRDM product.

3.2.2 Coastal and satellite AIS sensors

This study makes use of coastal and satellite AIS receivers. The coastal receivers are situated in ports and lighthouses around the coast of South Africa and feed AIS messages back to the AIS database within minutes of the message being transmitted by a ship. These provide good coverage along the coast, specifically near ports and harbours, with some gaps south of Durban and north of Langebaan.

The satellite AIS data was obtained from two commercial providers. The data had a variable delay of up to 45 minutes between transmission and arrival at the AIS database. The satellite AIS feed was limited to messages within 400 nautical miles of South Africa.

Together, the AIS data amounted to approximately 18 million messages per day, many of which were repeat messages or contained little new information.

3.3 VESSEL DATASET GENERATION

This section deals with how the training and testing dataset was generated from SAR images and AIS messages. Vessels were detected and extracted from the SAR images and then the AIS dataset was used to identify a portion of these vessels. Features were extracted from the SAR images of the matched vessels and then used to train and evaluate the performance of several learning algorithms. It must be emphasised that the focus of this research was not the detection or matching of the vessels but rather to use previously matched vessels, from SAR and AIS data, to build a classification and regression system.

3.3.1 SAR vessel detection

Figure 3.2 shows the method used to detect and extract the vessels detected in the SAR images along with the image metadata.

Sentinel-1 SAR images were downloaded where the swath fell over the ocean near South Africa. These images were radiometrically calibrated using the GAMMA software package [38] to transform the radar reflectivity into β_0, the radar brightness coefficient aligned to the slant range. Then the image was geocoded using GDAL [39] to warp the image into a ground range projection.

A CFAR algorithm was run with the parameters shown in Table 3.2. See Section 2.6.1 for a description of the CFAR algorithm. These parameters were shown in [5] to produce good results. The location of the detected vessel's centroid and the pixel values for a 1×1 km area around the vessel were stored. The AIS database request was compiled from the SAR images metadata and is shown in Table 3.3.

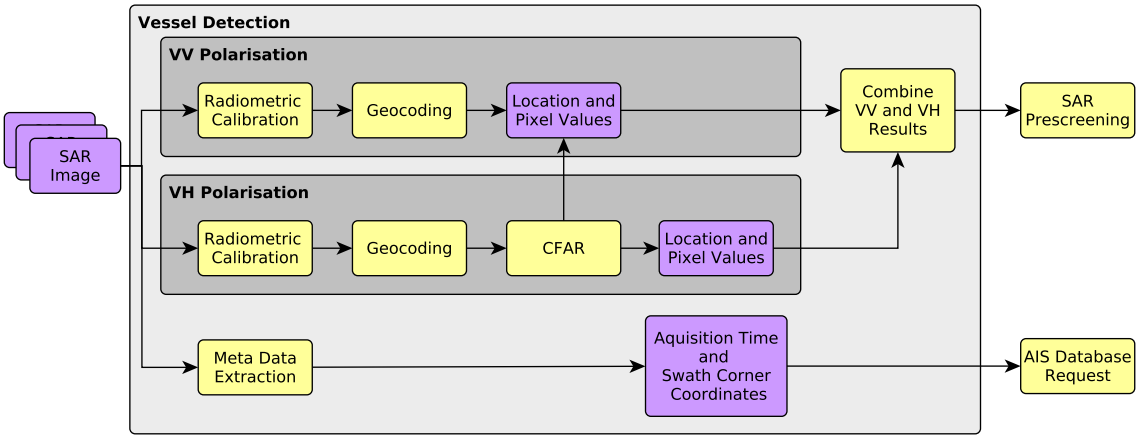

Figure 3.2. Vessel detection flowchart

Table 3.2. CFAR algorithm parameters

Parameter	Value
CFAR Method	GO-CFAR
Threshold	3.5
Background Window Size	29x29 pixels
Guard Pixels	27x27 pixels
Region of Interest	1 pixel

The vessel detection was performed for VH polarisations but the pixel results, for the same location from both polarisations, were stored.

3.3.2 AIS detections

When comparing the density of vessels in SAR and AIS datasets, it is desirable to account for better AIS receiver performance near harbours. If a density plot of all received AIS messages was taken as a true representation of vessel activity then it would seem that the vast majority of vessels were stationary in harbours.

To overcome this, "AIS detections" were created in this study. An AIS detection is defined as a unique vessel, determined by Maritime Mobile Service Identity (MMSI), within a grid location re-evaluated

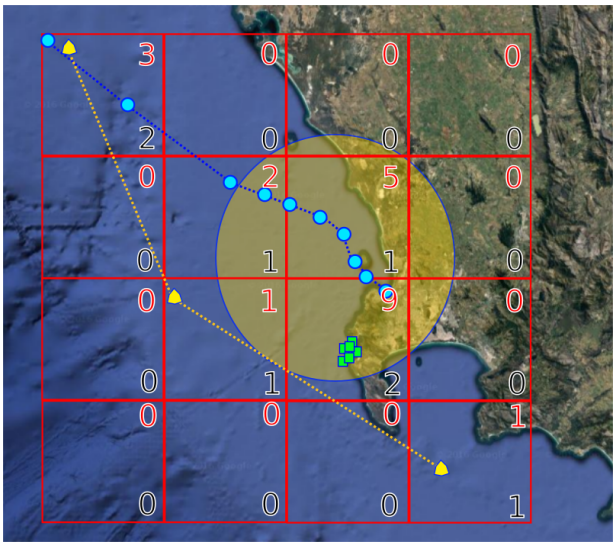

Figure 3.3. Example of generating AIS detections from AIS points

weekly. The use of a grid vastly speeds up analysis and processing of the approximately 18 million messages received per day.

Figure 3.3 shows a simplified plot of AIS messages and AIS detections. Three vessels are shown in this image and coastal AIS reception is indicated by the large yellow circle. The green vessel is anchored in a harbour and transmits messages regularly. The yellow vessel is only detected by satellite AIS and is moving through the EEZ. The blue vessel starts in Cape Town harbour and moves away from the coastal AIS receiver. The AIS message count is indicated by the red numeral in the top right corner of each grid box. The AIS detection count is in the bottom right corner.

It can be seen that the AIS detection number compensates for the repeat messages received by the coastal receiver. This method of analysing ship traffic is not without its own problems. The yellow vessel is not represented in some grids that it has moved through but did not transmit in. A method to overcome this would be to generate tracks from the AIS points and use these to analyse the vessel traffic in a region.

3.3.3 AIS database request

Generally, the timestamp of coastal receivers along South Africa's coast is more reliable than timestamps from satellite AIS receivers. Vessel matching requires a large time window when query-

Table 3.3. AIS request parameters

Parameter	Value
Time Window	SAR Acquisition Time $\pm$ 12 hours
Spacial Window	SAR Swath $\pm$ 100 km

Table 3.4. AIS parameters requested for all vessels within spacial window during time window

Parameter	Value
IMO	International Maritime Organisation identification number
MMSI	Maritime Mobile Service Identity number
Positions	Latitude and Longitude for all AIS messages
Class	AIS code for ship type
Length	Reported length of ship
Width	Reported width of ship

ing the AIS database to ensure all messages that occurred around the SAR acquisition time were retrieved.

Further information on vessels can be requested from the AIS database. Table 3.4 shows all the vessel parameters requested from the database for each vessel.

3.3.4 Vessel matching

SAR detections are matched to AIS messages by finding a single possible match that falls within certain predefined thresholds. These thresholds are described in Table 3.5 and an image showing SAR detections along with AIS messages and SAR-to-AIS matches, is shown in Figure 3.4. These thresholds were chosen to be conservative so as to minimise the number of duplicate matches in high-traffic areas and thus increase the certainty of a successful match.

The first step in the matching algorithm builds tracks from the individual AIS points based on the ship's identity, message time and position. Using linear interpolation, the position of the vessel at the time of SAR acquisition was estimated. SAR detections that lay close enough to the estimated position

Table 3.5. Various thresholds for SAR-to-AIS matching

Parameter	Value	Description
Δ_P	1 km	Maximum distance to nearest AIS point
Δ_T	2 km	Maximum distance to nearest AIS track
Δ_t	1 hour	Maximum time difference between AIS message and SAR acquisition

or track were considered matched. A SAR detection, SAR_i, can be matched to an AIS detection, AIS_j, when:

$$SAR_i \text{ matches } AIS_j = \begin{cases} True : (Pos_i < Pos_j + \Delta_P) \text{ and } (Time_i < Time_j + \Delta_t) \\ True : (Pos_i < Track_j + \Delta_T) \\ False : Otherwise \end{cases} \tag{3.1}$$

Figure 3.4 shows several vessel tracks, AIS messages, estimated positions and SAR detections. SAR detection C is close enough to be within the "Estimated Point" threshold to be accepted but the nearest AIS message was received too far outside of the "AIS-SAR Time Difference" threshold to match C in this manner. Detection A and C are matched using the track instead of the AIS points. If no detections are matched to a vessel, in this way, the detections are matched to tracks using the "Track Distance" threshold. Multiple detections that are matched to the same track, such as E and F, are discarded. Detections that do not fall within any thresholds (such as B, D and G) are also discarded. D is considered a "Dark Target" since there are no AIS messages that could possibly be associated with the SAR detection. The other AIS tracks fall outside the SAR swath and are ignored. The thresholds were chosen to be strict to reduce the occurrences of false matches that would pollute the dataset used in the training and evaluation of the machine learning algorithm.

3.3.5 Dataset prescreening

Not all the AIS data that was used in this study can be assumed to be representative of actual vessel characteristics. The potential is high for poor AIS parameters and false positive SAR detections to creep into the dataset. To reduce the impact bad data would have on the training of the machine

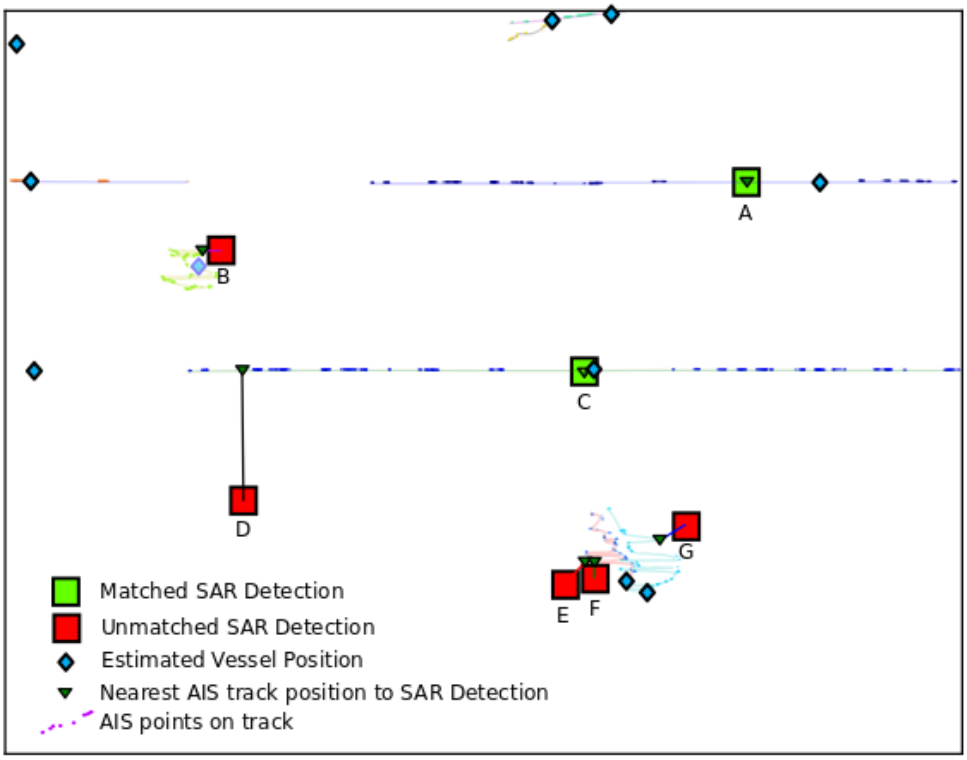

Figure 3.4. SAR detections matched to AIS tracks with estimated vessel positions from SAR image A298.

learning algorithm the data must be filtered, checked and potentially bad data must be discarded. This step was only performed for the matched SAR-to-AIS dataset rather than the whole AIS database and SAR detection sets.

The information in the AIS database was compared with a third party that maintains a more comprehensive, global database of vessel information. The matched detection was discarded where vessel information, such as vessel type, length or gross tonnage, was significantly different from the third party database. In cases where a vessel did not specify the AIS vessel class, the matched detection was discarded.

It must be remembered that all SAR detections that did not match to an AIS detection were not included in the matched dataset. Non-matches were assumed to be due to poor AIS coverage, false SAR detections, or dense traffic with multiple possible matches in the area of interest. It must not be discounted that the behaviour of certain classes, such as fishing vessels, might preclude them from

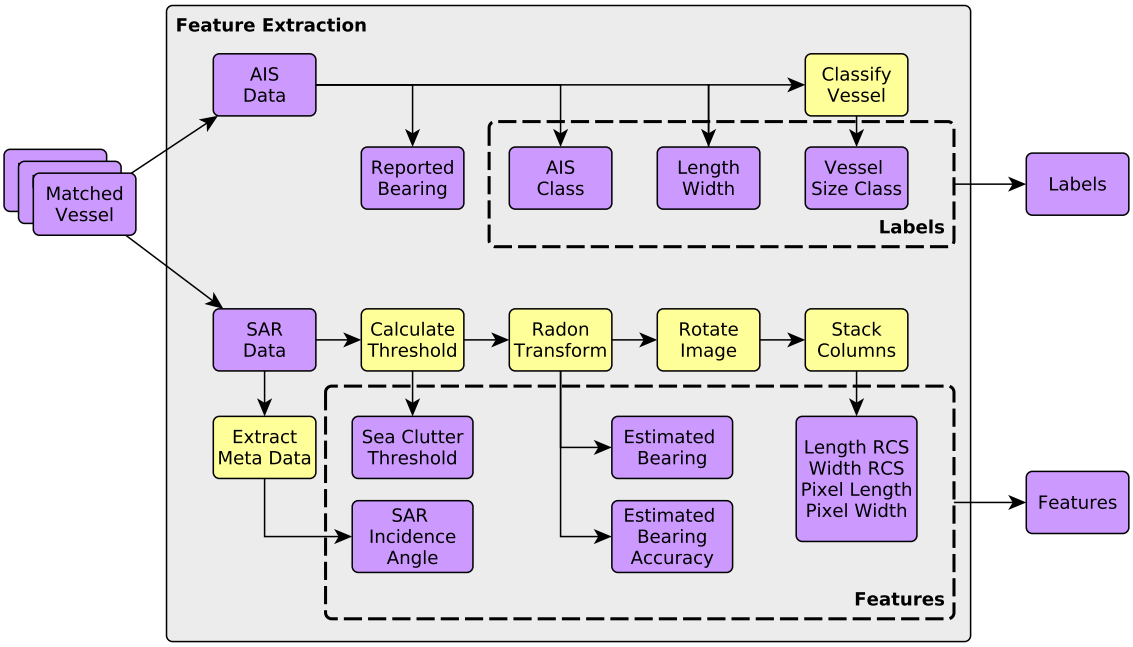

Figure 3.5. Feature extraction and vessel labelling flowchart

being matched using the current matching algorithm.

3.4 FEATURE AND LABEL EXTRACTION

Extracting features and labels from the data is required for supervised training of machine learning algorithms. This section describes how this was achieved using the matched dataset from the previous section. Figure 3.5 shows the method used.

3.4.1 Labels

The various labels are derived from the AIS data associated with the SAR-detected vessel. The supervised learning algorithm can be trained to predict the following from the SAR data:

- AIS Class
- Vessel length or width in meters

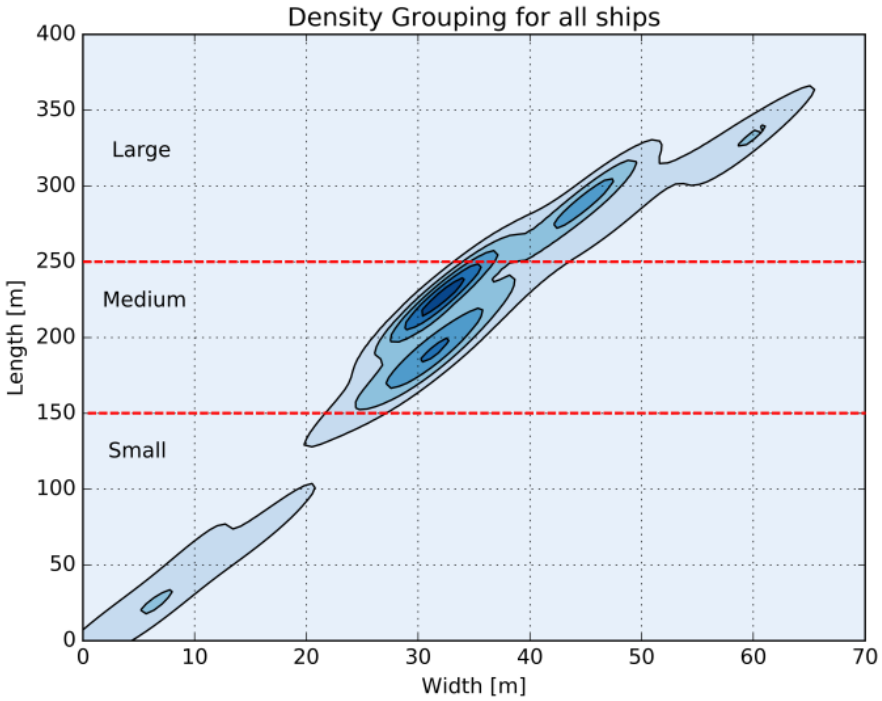

Figure 3.6. AIS width and length distribution for AIS detections for 2015. The vessels were grouped into Small, Medium and Large categories based on their reported lengths.

- Vessel size class (small, medium or large)

The reported bearing, near the SAR-acquisition time, can also be extracted to check the validity of the Radon estimated bearing.

The "vessel size class" groups vessels into three categories based on a vessel's reported length.

- Small Vessels: Length < 150 m
- Medium Vessels: 150 m < Length < 250 m
- Large Vessels: Length > 250 m

These groups were chosen based on the naturally occurring grouping found in the AIS data, as shown in Figure 3.6.

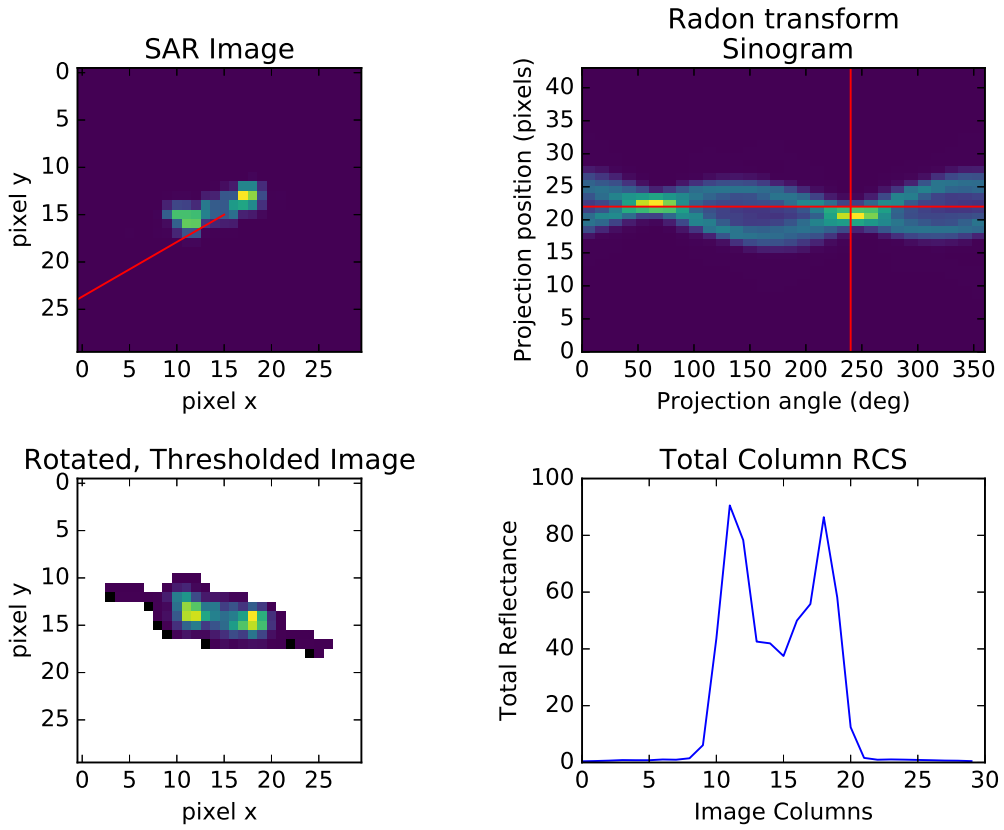

Figure 3.7. Feature extraction process: the bearing of the SAR-detected vessel is extracted with the Radon transform and this is used to extract the width and length RCS.

3.4.2 Features

Figure 3.7 shows a SAR image of a vessel going through the feature extraction process. The vessel is centred, the bearing determined using the Radon transform and the Length Radar Cross Section (LRCS) determined from the sum of the total reflectance for each column along the vessels length.

3.4.2.1 Sea backscatter threshold

The first step in extracting vessel features from the SAR image is to separate the vessel pixels from the sea pixels. The sea backscatter has a different distribution to vessel backscatter and, on the small scale, can be removed by a simple threshold [25].

The threshold, T, is calculated using the 99.9th percentile of the sampled sea pixels in the image;

$$\int_0^T F(x)_{sea} = 99.9\% \tag{3.2}$$

where $F(x)_{sea}$, the sea clutter probability density function, is obtained from sea pixels nearby. All pixels below that threshold are removed. Remaining non-ship pixels are due to side lobe interference from the ship or very high speckle.

The sea clutter can also be used as a feature for the training of predictive algorithms since the sea clutter, as well as other features, is a function of SAR incidence angle . Various vessel sizes are better detected depending on the SAR incidence angle [22].

3.4.2.2 Vessel feature extraction

Each pixel in the SAR image represents the total radar backscatter of that area which is dependent on the size, shape, motion and orientation of any scatterers. Let the backscatter at pixel (x,y) be $f(x,y)$.

A Radon transform is performed on the image;

$$S(\rho,\theta) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\delta(\rho - x\cos\theta - y\cos\theta)\,dx\,dy \tag{3.3}$$

where

$$\rho = x\cos\theta + y\sin\theta \tag{3.4}$$

and δ is the Dirac delta function.

3.4.2.3 Estimated bearing and bearing accuracy

The rotation angle, $\theta_{bearing}$, that results in the maximum value of the sinogram, $S(\rho,\theta)$, regardless of ρ is used as the vessels estimated bearing [30].

The variance of the normalised sinogram, B_v, is also calculated to estimate the confidence of the estimated bearing;

$$B_v = var(S(\rho, \theta_{bearing})) \tag{3.5}$$

Larger vessels have a more rectangular signature, which results in the estimated bearing corresponding better to the reported AIS bearing with the variance in the sinogram higher. Smaller vessels can appear as squares or circles leaving the sinogram with a similar value throughout rotation and thus a lower sinogram variance. It is postulated that this confidence variable will be able to inform on the reliability of the Radon bearing which would also be an indication on the size of the vessel.

3.4.2.4 Radar cross sections

Three radar cross sections are used as features in this study. The Length Radar Cross Section (LRCS), Width Radar Cross Section (WRCS) and Total Radar Cross Section (TRCS) are calculated from the sinogram of the image. After applying the threshold, the remaining pixel values are then rotated by the estimate bearing so that a radar cross section along the vessel's width and length can be calculated.

$$RCS_{Length}(x) = \sum_{y=0}^{n} f(x,y)_{rotated} \tag{3.6}$$

$$RCS_{Width}(y) = \sum_{x=0}^{n} f(x,y)_{rotated} \tag{3.7}$$

The LRCS is shown in the bottom right panel of Figure 3.7. The TRCS is the sum of all the backscatter in the image after the sea pixels are removed:

$$RCS_{Total}(y) = \sum f(x,y) \,|\, (f(x,y) > T) \tag{3.8}$$

3.4.2.5 Pixel width and length

The pixel width and length are calculated as the amount of columns in the LRCS and WRCS that are greater than $2T$.:

$$Length_{px} = \sum_{i=0}^{n} [RCS_{Length}(i) > 2T] \tag{3.9}$$

$$Width_{px} = \sum_{i=0}^{n} [RCS_{Width}(i) > 2T] \tag{3.10}$$

$2T$ was chosen to minimise the size exaggeration to side-lobe effects, as seen in the bottom left image of Figure 3.7. This value was determined to be $2T$ by trial and error to provide the best length accuracy. Vessel width and length can be obtained by multiplying the pixel value by the image's pixel spacing which in Sentinel-1 GRDM images is about 40 meters.

3.4.2.6 SAR incidence angle

The incidence angle at which the SAR beam hits the surface of the ocean can be calculated from the Sentinel-1 metadata. The metadata contains a grid of incidence angles for the image, and the angle for each vessel is calculated by interpolating between the points on the grid. The SAR incidence angle might be an important feature as, along with bearing, it could be used to account for width and length exaggerations, as shown in Figure 2.5.

3.5 VESSEL PARAMETER PREDICTION

This section details how the algorithms that predict information from the SAR images were created, optimised and evaluated. The flow chart for this is shown in Figure 3.8.

The SVM algorithm was used for both prediction and classification algorithms by changing the desired prediction parameter and the machine class used.

The matched vessels were separated into classes based on their reported AIS class or length. An algorithm then predicted the classes based on the extracted SAR features. This section describes the algorithms, parameters, scoring methods and classes selected for the classification of the vessels.

3.5.1 Algorithm scoring

Various scoring methods were provided to the optimisation algorithm to evaluate the differences in parameters that resulted in an optimal algorithm for each scoring method. This section describes each of the scoring methods used for classification results and regression results.

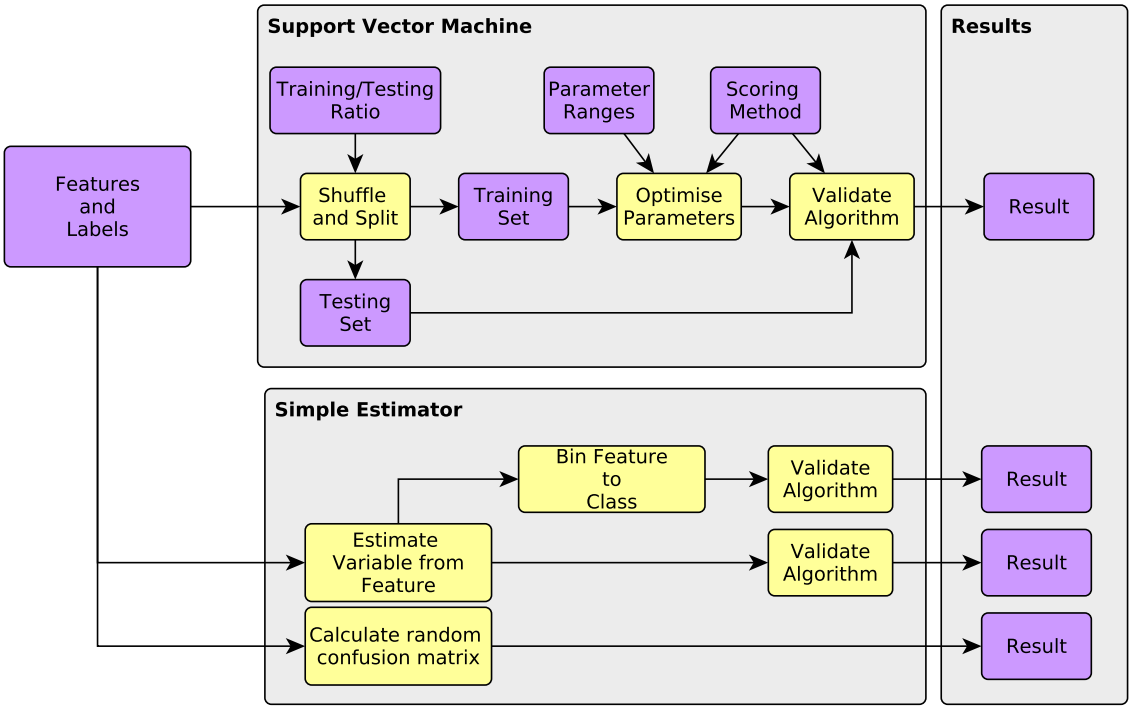

Figure 3.8. Vessel information prediction flowchart

The confusion matrix is often used to describe the performance of a classification model. For the binary case it is represented by the Table 3.6. This can easily be extended to the multi-class case and can provide valuable information on which classes are being mistaken for others.

True Positives (TP), False Positives (FP), True Negatives (TN) and False Negatives (FN) are the possible outcomes. A perfect classifier would be described by an identity matrix.

The following equations describe the Accuracy, Precision, Recall and combined F1 Scores of the confusion matrix where n is the total number of objects to be classified.

$$n = TP + TN + FP + FN \tag{3.11}$$

$$Accuracy = \frac{TP + TN}{n} \tag{3.12}$$

$$Precision = \frac{TP}{TP + FP} \tag{3.13}$$

Table 3.6. Binary confusion matrix

	Predicted False	Predicted True
Actual False	TN	FP
Actual True	FN	TP

$$Recall = \frac{TP}{TP + FN} \tag{3.14}$$

$$F_1 = 2\frac{Precision \cdot Recall}{Precision + Recall} \tag{3.15}$$

Each score is calculated per class and the average is taken as the score of the algorithm.

When performing regression the algorithm scoring was determined by the adjusted R^2 algorithm:

$$R^2_{adjusted} = 1 - \frac{(1 - R^2)(n-1)}{(n-p-1)} \tag{3.16}$$

where p is the number of predictors in the model and R^2:

$$R^2 = 1 - \frac{\sum(y_{actual} - y_{predicted})^2}{\sum(y_{actual} - y_{mean})^2} \tag{3.17}$$

and y_{actual}, $y_{predicted}$ and y_{mean} is the actual, predicted and mean values of the training data.

3.5.2 Support vector machine optimisation

An SVM was used to predict and classify the vessels based on the SAR features and AIS labels. The support vector machine has several parameters that define how it operates. These are described in Table 3.7 and the ranges that were used in the optimisation of the algorithm are given.

These parameters are the same for the SVC, used for predicting a class from features, and the SVR, used for estimating a variable from features.

The SVM algorithm is optimised by performing an exhaustive grid search by calculating the score of the classifier for all parameters given in Table 3.7. The parameters that cause the SVM to score highest

Table 3.7. SVM algorithm parameters

Parameter	**Description**	**Optimisation Range**
C	Penalty parameter of the error term.	1e-3 to 1e5
Kernel	Kernel type to use in algorithm	Linear, Polynomial, RBF or Sigmoid
Degree of Polynomial	Only used for Polynomial Kernel	1 to 5
Coefficient 0	Independent term used in Polynomial and Sigmoid	1e-3 to 1e5
Gamma	Kernel Coefficient for RBF, Polynomial or Sigmoid	1e-3 to 1e5
Class Weight	Weigh C inversely proportionally to class frequency	Boolean

are chosen. This is performed using a k-fold feature selection using 4 folds. This is done to reduce the risk of any overfitting or distortion due to feature outliers or mislabelled vessels.

The algorithms are validated using the testing data. These data were never used to train the algorithm and provide an indication of how the algorithm would behave with real examples.

CHAPTER 4 RESULTS AND DISCUSSION

4.1 CHAPTER OVERVIEW

This section details the various datasets, as they appeared in the previous chapter, and discusses their various shortcomings and imperfections. The results of the machine learning algorithms and a discussion of the impacts and applicability of the dataset, feature or algorithm are also included.

4.2 SAR DETECTIONS

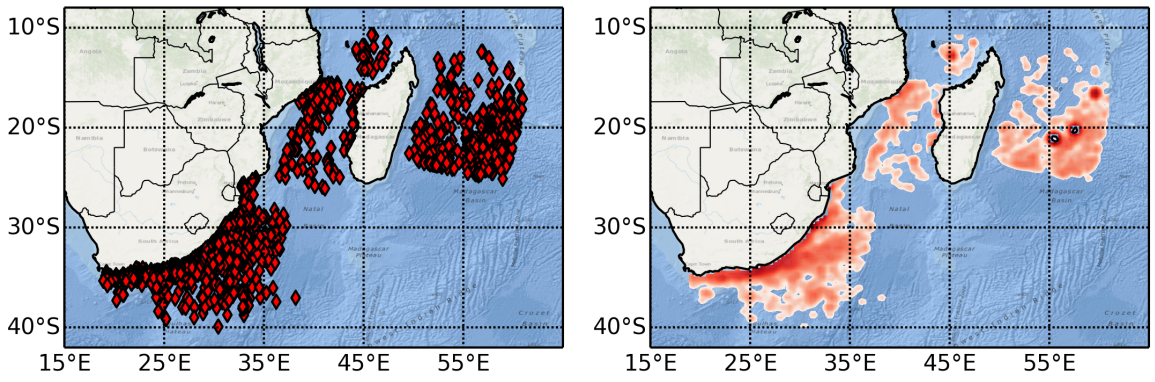

(a) SAR Detections Points

(b) SAR Detections Density

Figure 4.1. The locations of all 2830 Sentinel-1 GRDM, VH, SAR detections for 2015 used in this study.

Figure 4.1 shows that the locations of SAR detections are distributed over a wide area; both near and far from various coastlines. Figure 4.1(b) shows that the vessel detections are more numerous around harbours, and near coastlines and cargo lanes. When comparing this to AIS density plots, as shown in Figure 4.2, it can be seen that the SAR detections follow the same location distribution as the AIS

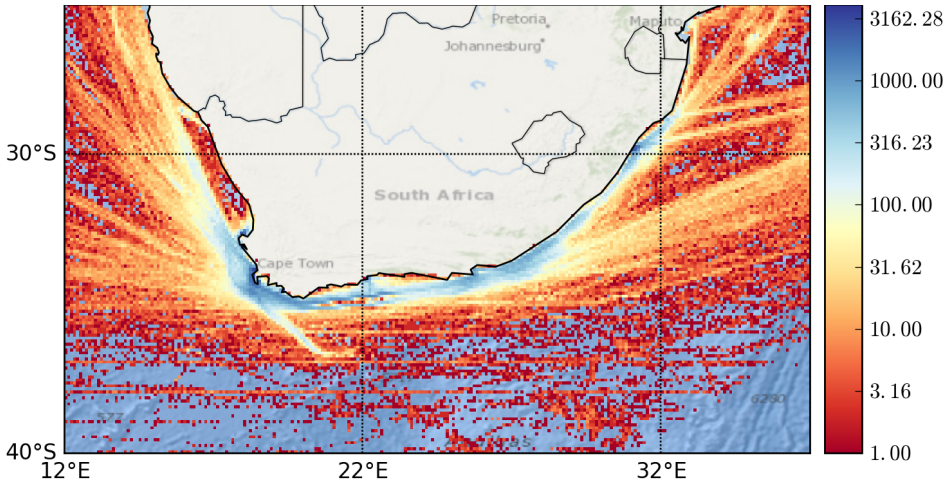

Figure 4.2. AIS detection density for all vessels in 2015 around southern Africa

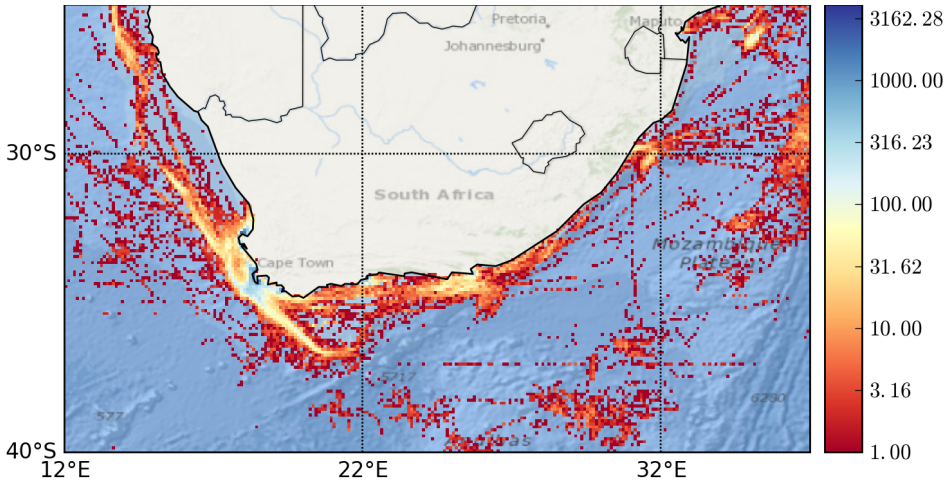

Figure 4.3. AIS detection density for fishing vessels in 2015 around southern Africa

detections around southern Africa. It should be noted that there were no GRDM SAR acquisitions to the west of Cape Town in the 2014-2015 period.

Any algorithm developed using the SAR detections should work for the general case since detections are not limited to specific area types or locations.

4.3 AIS DETECTIONS

Figure 4.2 shows all AIS detections for the ocean around southern Africa for 2015. These detections are used to evaluate whether the SAR-to-AIS matched data set is representative of the vessel traffic monitored by AIS.

Figure 4.3 shows the AIS detections density for fishing vessels in 2015. Fishing vessels are generally smaller than tanker and cargo vessels and so could be under represented in coarse resolution SAR images. It is important to determine whether fishing vessels fall into the areas that have been imaged to determine whether any under-representation is due to their location, radar visibility or AIS behaviour. The location of fishing vessel activity, with the exception of the area to the west and north-west of Cape Town, overlaps with the SAR detection distribution in Figure 4.1.

4.3.1 Vessel classes

Table 4.1 is a breakdown, by AIS class, of all AIS detections in 2015 and also shows a breakdown of SAR-to-AIS matches. It can be seen that the bulk of vessel traffic around southern Africa was cargo vessels, tanker vessels, fishing vessels or vessels of unspecified class.

4.3.2 Vessel sizes

Figure 4.4 shows a histogram of the vessel length for each of the three largest classes. The unspecified class was ignored since it is made up of vessels that have diverse sizes, roles and behaviours. The figure shows that the majority of fishing vessels are below 100 m in length and that cargo and tanker classes have overlapping length distributions.

4.4 AIS-TO-SAR MATCHES

This section deals with the dataset of SAR detections that have been matched to AIS detections and were used in training and testing the prediction algorithms. There were 701 AIS-to-SAR matches extracted from the 105 SAR images used. This number could have been improved by using a more

Table 4.1. Description of AIS class representation in 2015 AIS data and SAR-to-AIS matched data

AIS Class	Description	Number of AIS Detections	Percentage of AIS Detections	Number of SAR Matches	Percentage of SAR Matches
0	Unspecified Ships	78530	11.34	35	4.99
30	Fishing	88586	12.79	17	2.43
31	Towing	2866	0.41	4	0.57
32	Big Tow	89	0.01	1	0.14
33	Dredge	2163	0.31	1	0.14
35	Military	662	0.1	0	0.00
37	Pleasure Craft	2393	0.35	0	0.00
50	Pilot	2550	0.37	0	0.00
51	Search & Rescue	707	0.1	0	0.00
52	Tug	9476	1.37	7	1.00
55	Law Enforcement	17244	2.49	1	0.14
60-69	Passenger	1707	0.25	3	0.43
70-79	Cargo	373115	53.88	480	68.47
80-89	Tanker	112360	16.23	152	21.68
Total		**692975**	**100**	**701**	**100**

reliable and more comprehensive AIS dataset that had accurate timestamps on the AIS messages from multiple satellite AIS receivers, but such a dataset was not available at the time of research. Having a denser AIS dataset would also reduce vessel position uncertainty when matching vessel classes that tend not to move in straight lines.

4.4.1 Location of acquisitions

Figure 4.6 shows the locations and density of the matched detections. The geographic distribution of the matched detections is similar to that of the SAR detections east of Cape Town. There were no SAR acquisitions to west of Cape Town. Figure 4.3 shows that a significant portion of fishing vessels can be found north and west of Cape Town.

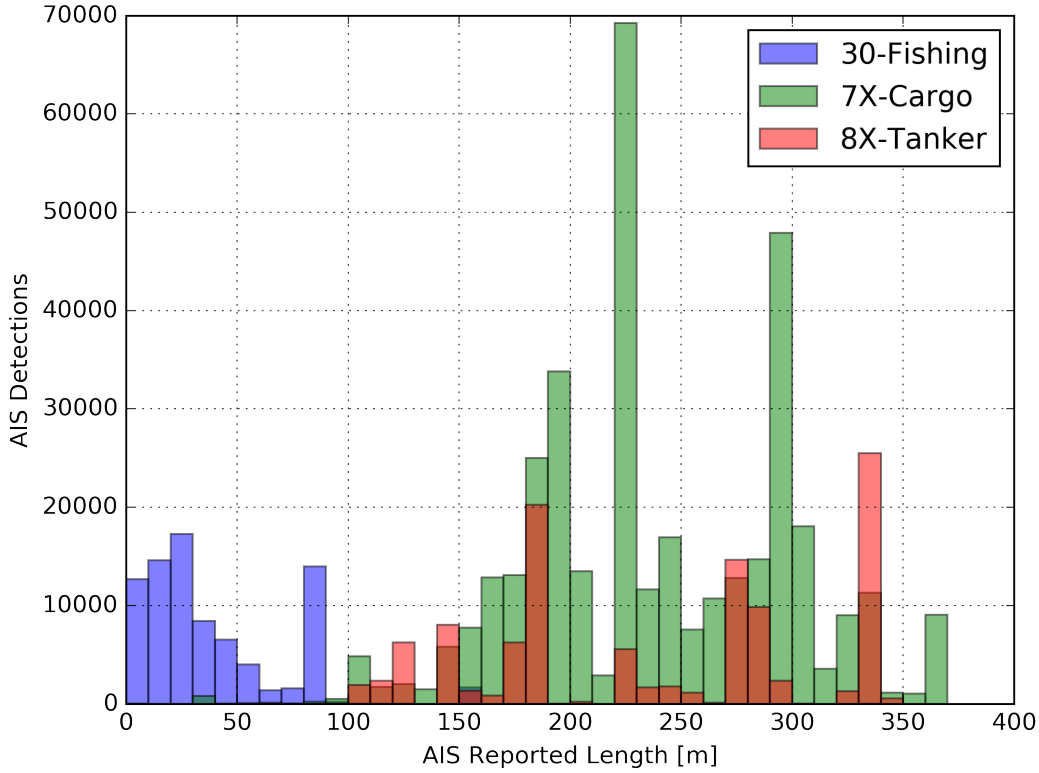

Figure 4.4. Vessel length histogram for various AIS classes derived from 2015 AIS data.

4.4.2 AIS class distribution

Table 4.1 shows a breakdown of the matched detections by AIS class. When compared to the 2015 AIS detections it can be seen that physically smaller classes (such as fishing vessels, tug-boats and dredges) are under-represented in the matched dataset. The proportion of fishing vessels in the matched dataset is 81% lower than in the AIS detections dataset. The proportions of cargo and tanker vessels increases by 27% and 29% respectively.

4.4.3 Length distribution

Figure 4.5 shows a histogram of the length distribution of various AIS classes in the matched dataset. When compared to Figure 4.4 it is clear that fishing vessels are almost completely missing from the matched dataset.

There are a few reasons that could explain why fishing vessels are missing from the matched dataset.

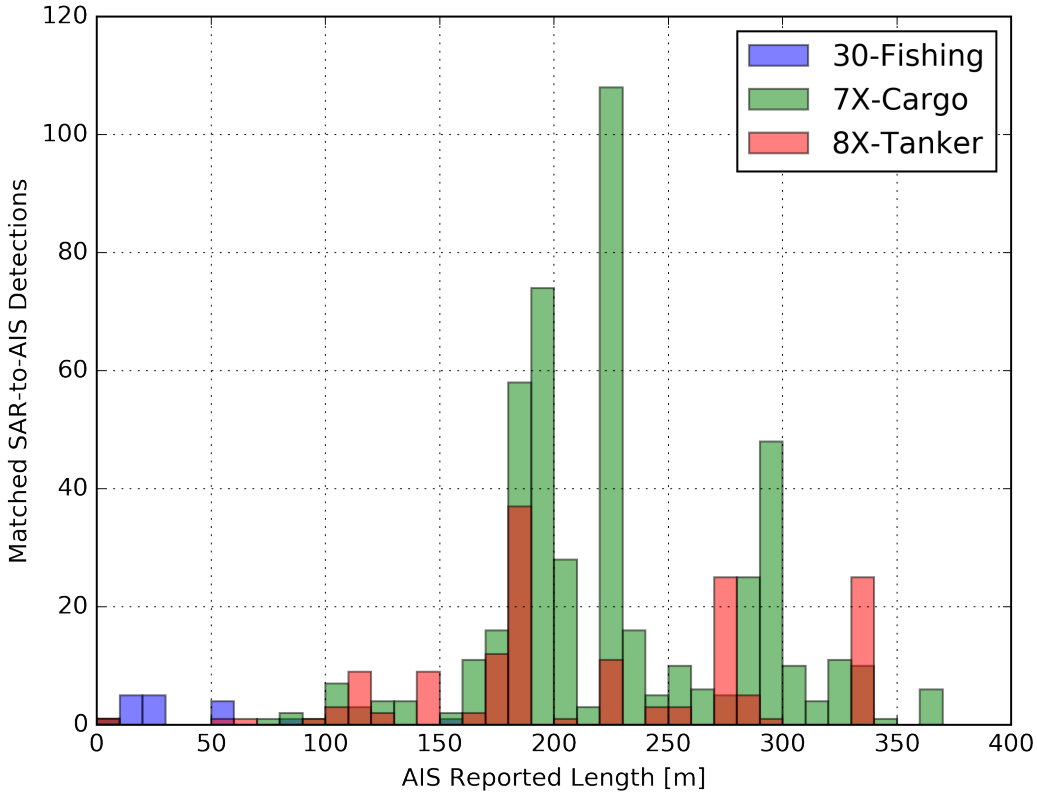

Figure 4.5. Vessel length histogram for various AIS classes derived from SAR-to-AIS matched data.

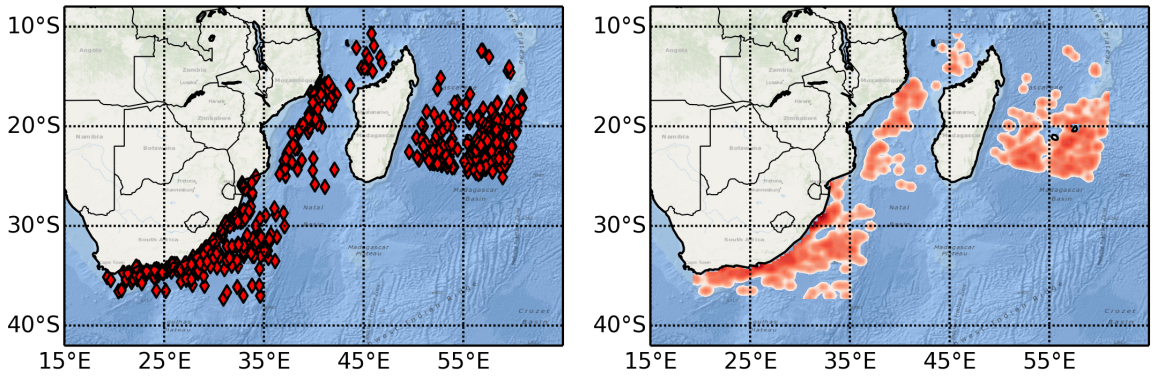

(a) SAR-to-AIS Matched Detections Points

(b) SAR-to-AIS Matched Detections Density

Figure 4.6. The locations of all matched SAR-to-AIS detections used in this study.

The vessels are generally smaller than the other main classes and so might not be detected using the CFAR algorithm with coarse resolution SAR data. This is due to a fishing vessels Total Radar Cross Section (TRCS) not being as high as that in other vessel classes, and the pixel grouping method, used

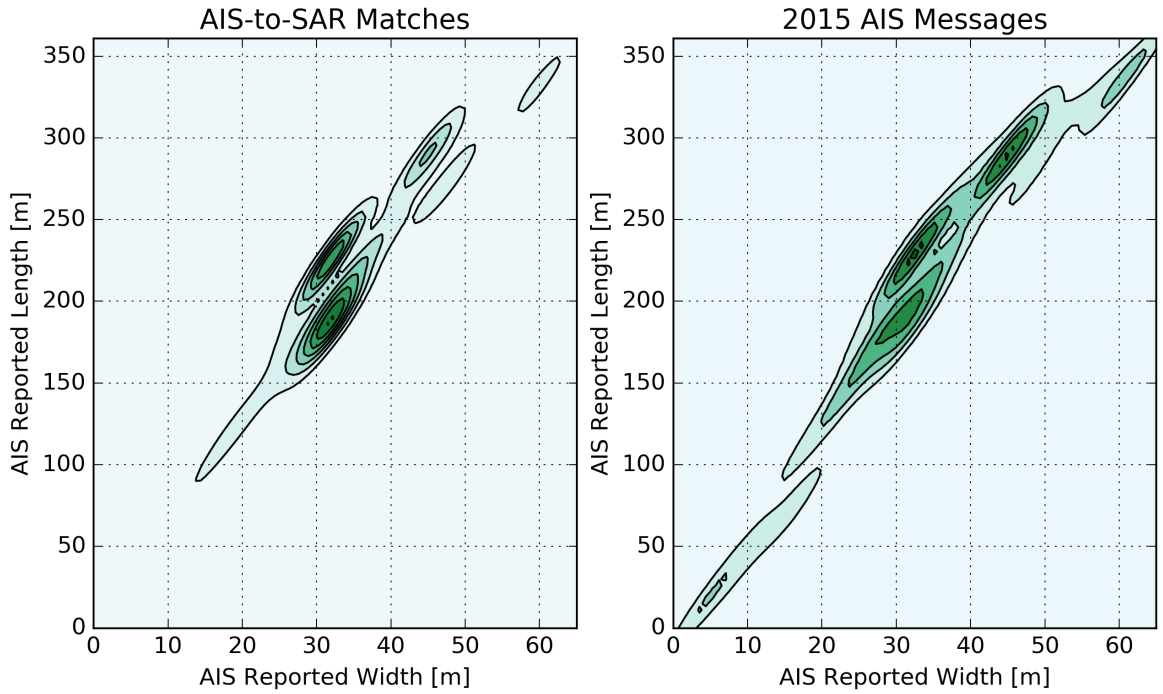

Figure 4.7. Size distribution of SAR detected vessels compared to AIS detected vessels

in the CFAR process to eliminate high value single pixels, would eliminate small vessels that produce small detection groups. Another possibility is that fishing vessels have a low radar cross section due to the materials used in producing smaller vessels. Fibre glass hulls are often used in smaller vessels to reduce costs of production. The non-linear movement of fishing vessels, as seen in Figure 3.4, combined with the infrequent satellite AIS messages could have resulted in fewer SAR detected fishing vessels being matched to AIS data.

Figure 4.7 is a kernel density estimation of AIS reported widths and lengths of vessels for the matched dataset and AIS messages from 2015. Along with Figure 4.4 this image also shows that smaller vessels are not present in the matched dataset.

4.4.4 Shortcomings of SAR-to-AIS matched dataset

The previous sections show that fishing vessels, due to their size or construction, are not being proportionally represented in the SAR-to-AIS matched dataset. Figures 4.1 and 4.3 show that fishing grounds were imaged and, given a high enough SAR resolution with matching AIS data, should contain

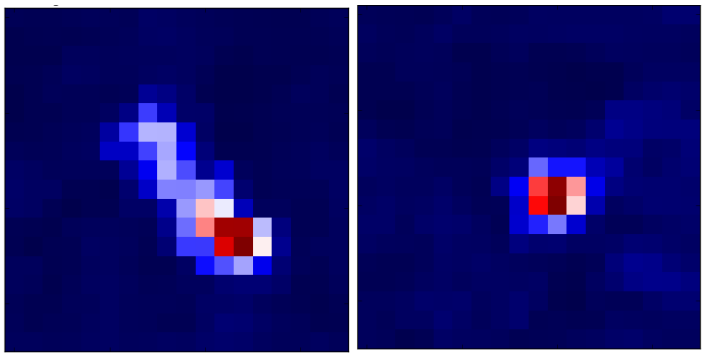

(a) Tanker vessel with reported length of 292 m

(b) Fishing vessel with reported length of 23 m

Figure 4.8. Two SAR imaged vessels matched to AIS.

a representative sample of fishing vessels.

This result shows that the AIS matching combined with the CFAR method currently used on coarse resolution SAR imagery does not, generally, detect vessels under 100 m in length. This would reduce the utility of coarse resolution SAR imagery in monitoring fishing grounds or marine protected areas for poaching by small fishing vessels. Figure 4.3 shows the historical locations of fishing vessels and so when planning a SAR surveillance campaign these areas should be imaged with a higher resolution SAR product.

4.5 SAR FEATURES

This section deals with discussing the SAR extracted features and comparing them to expected values where possible. This is done to ensure the validity of the input features.

4.5.1 Sea clutter and SAR incidence angle

The sea clutter feature is the average backscatter from the sea surface surrounding the detected vessel. It is calculated from the edges of the 1×1 km area around the detected vessel. The average sea pixel backscatter, along with the average vessel backscatter, for each detection is shown in Figure 4.9. It can be seen that the sea clutter differs from the vessel backscatter, although there is some overlap. This

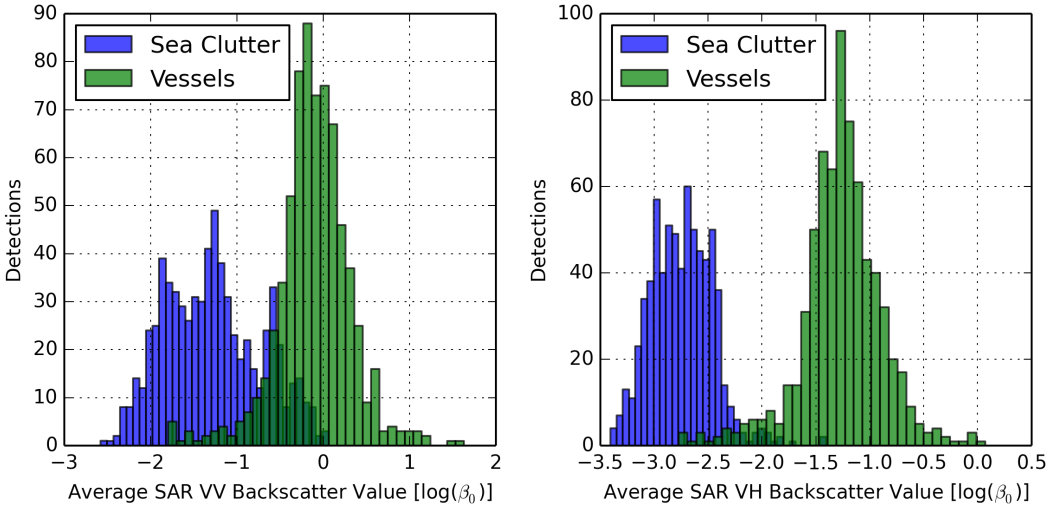

(a) VV polarisation **(b)** VH polarisation

Figure 4.9. Histogram of average SAR backscatter values for sea pixels and vessel pixels for VV and VH polarisations.

overlap is worse in VV images which is why VH is favoured for vessel detection and separating vessel pixels from sea pixels. The sea clutter was chosen as a feature since it has relationships with the SAR incidence angle, sea state and any length or width exaggerations, as seen in Figure 2.5, in the SAR image.

The sea clutter and vessel backscatter is plotted against the SAR incidence angle in Figure 4.10. This shows that the sea clutter decreases as the SAR incidence angle increases and that the worst overlap occurs near the nadir in VV polarisations. These results indicate that for vessel detection images with VH polarisation should be used.

4.5.2 Radar cross sections

Figures 4.11 and 4.12 show the LRCS, in both VV and VH, for two vessels along with a photograph of a vessel with a similar structure. The backscatter from the fishing vessel is much lower than that of the tanker vessel. While the magnitude of the returns is similar for VV and VH polarisations, the structure of the vessels is better represented in VV. For the tanker vessel, the VV LRCS of the vessel shows the

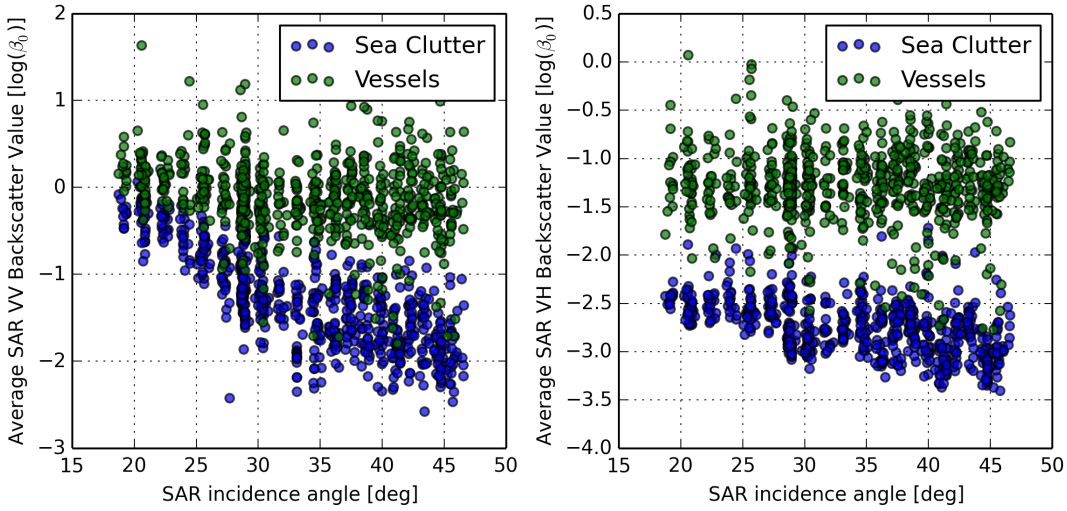

(a) VV polarisation **(b)** VH polarisation

Figure 4.10. Average SAR backscatter values for sea and vessel pixels against SAR angle of incidence.

position of the bridge at the rear of the vessel and the crane towards the front. The physical structure of the fishing vessel is less apparent in the LRCS.

The gross tonnage of a vessel is a non-linear value derived from the total internal volume of that vessel. The relationship between the gross tonnage of the matched vessels and the total vessel backscatter, as measured using the TRCS feature, is shown in Figure 4.13. It can be seen that there is a non-linear relationship between the gross tonnage and vessel backscatter. This figure shows that the TRCS has some value when discriminating between vessels with a high and low gross tonnage although having more examples of vessels with low gross tonnage would be desirable.

The RCS features used in the machine learning algorithms consisted of a vector of 20 values, each describing the total backscatter along the vessels width or length.

4.5.3 Pixel length

The estimated length of a vessel is calculated by multiplying the number of pixels that represent a vessel in a SAR image by their pixel spacing. This method is used in most length estimation methods

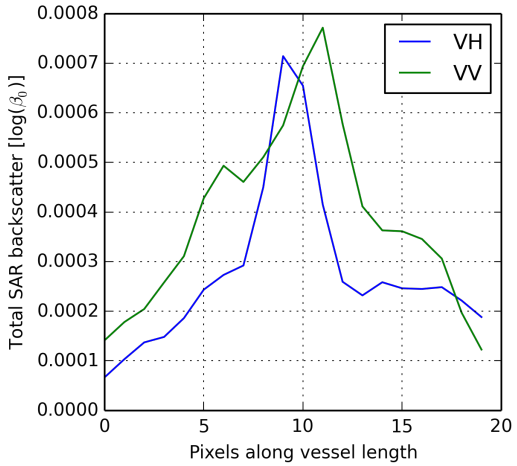

(a) LRCS of fishing vessel

(b) Fishing vessel

Figure 4.11. LRCS of fishing vessel and photograph of similar vessel (Photograph by Dimitris Vetsikas licensed under Creative Commons Zero licence from https://pixabay.com/en/motor-ship-fishing-ship-fishing-sea-1432432/)

[31, 33]. This gives a rough estimate of size and is used as a comparison for more advanced methods. This method provides reasonably accurate results when using high resolution SAR imagery.

Figure 4.14 shows the SAR pixel estimated length against the AIS reported length. The R^2 score for this estimation is -4.614. Some length values are shown as zero due to the vessel not appearing in both image polarisations.

4.5.4 Estimated bearing

The Radon-estimated bearing is compared to the AIS-reported bearing in this section and the error between the two is examined. A correct bearing estimation will ensure that the SAR image is rotated correctly to extract the width RCS and length RCS.

Errors in estimated bearing can occur where the SAR-imaged vessel is small and appears square rather than rectangular. An example of this is shown in Figure 4.8. Errors can also occur where AIS messages are received infrequently or the vessel is making rapid bearing changes. The reported AIS bearing is taken from the AIS message with the nearest reported coordinates.

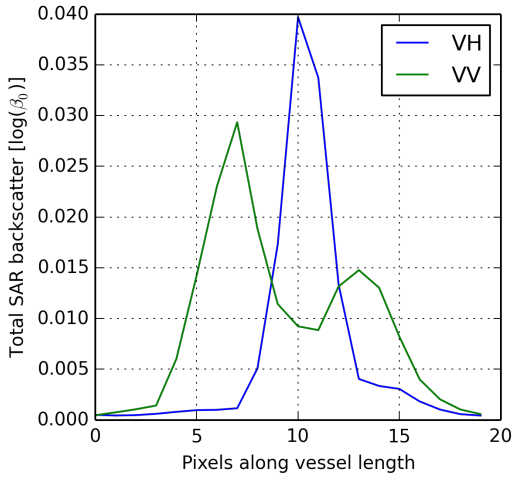

(a) LRCS of tanker vessel

(b) Tanker vessel

Figure 4.12. LRCS of tanker vessel and photograph of similar vessel (Photograph by Alf van Beem licensed under Creative Commons Zero licence from http://maxpixel.freegreatpicture.com/Ship-Transport-Vessel-Logistics-Armia-Ludowa-883895)

Validating the SAR-estimated bearing was done by comparing it to the nearest AIS reported bearing. In cases where the AIS messages are received infrequently or when the vessel is changing direction, the AIS reported bearing is possibly incorrect.

It has been assumed that the rear of the vessel usually has a greater backscatter than the middle or front of vessels due to the physical structure of the bridge which, in cargo and tanker vessels, is usually situated at the rear of the vessel [30].

Figure 4.15(a) shows a histogram of the estimated bearing against the reported bearing. There are peaks at 0 and 180 degrees showing that the Radon estimation function, using coarse SAR imagery, often cannot distinguish between the front and rear of the vessel based on the assumption in [30].

Figure 4.15(b) shows the error when the reported and estimated bearings are limited from 0 to 180 degrees. This will remove any ambiguity coming from front-rear errors in bearing estimation. The histogram shows that the majority of estimated bearings are within 20 degrees of the reported bearings.

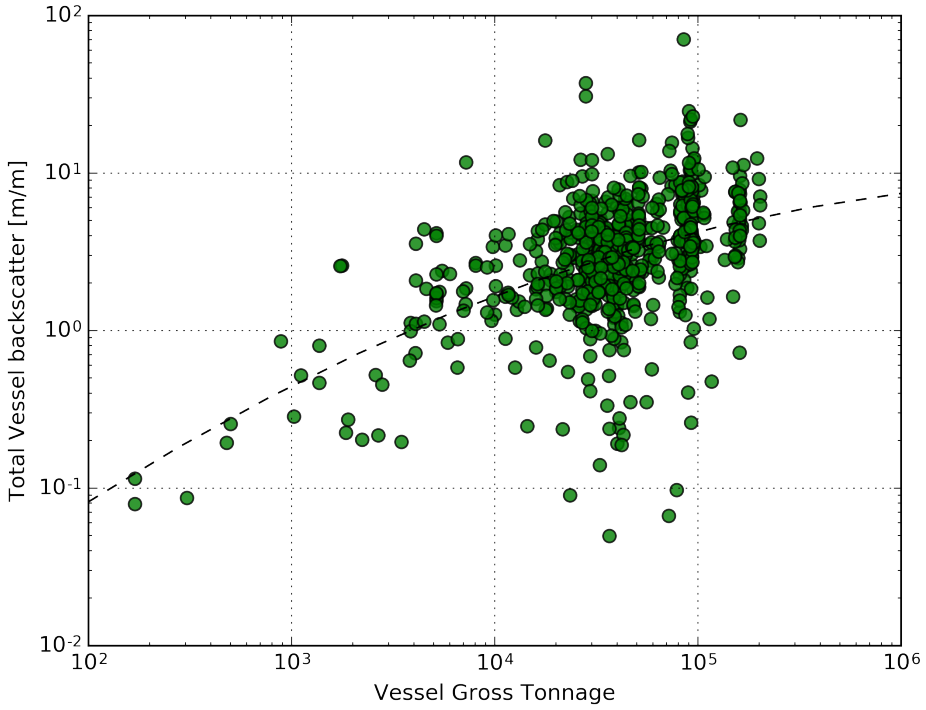

Figure 4.13. Vessel gross-tonnage against total vessel backscatter.

Figure 4.16 shows the estimated bearings against the reported bearings with the size of the marker scaled to the AIS reported length of the matched vessel. A kernel density estimation was also run over the points to give an indication of density. The bearings were again limited to 0 to 180 degrees.

This figure shows that the majority of vessels were travelling between 0 and 90 degrees. Given the locations of the SAR detections in relation to the South African coast, as seen in Figure 4.1, this is to be expected. There is also some "snapping" to 90 degrees and 180 degrees of estimated bearings. This is an artefact of using the Radon estimation process with coarse resolution SAR images. The figure also shows that there are multiple large vessels that have large differences between their reported and estimated bearings. These errors could be due to a failure to report accurate bearings in AIS, either due to user error or a large timespan between the AIS message and SAR acquisition. The errors could also be due to AIS-to-SAR mismatches.

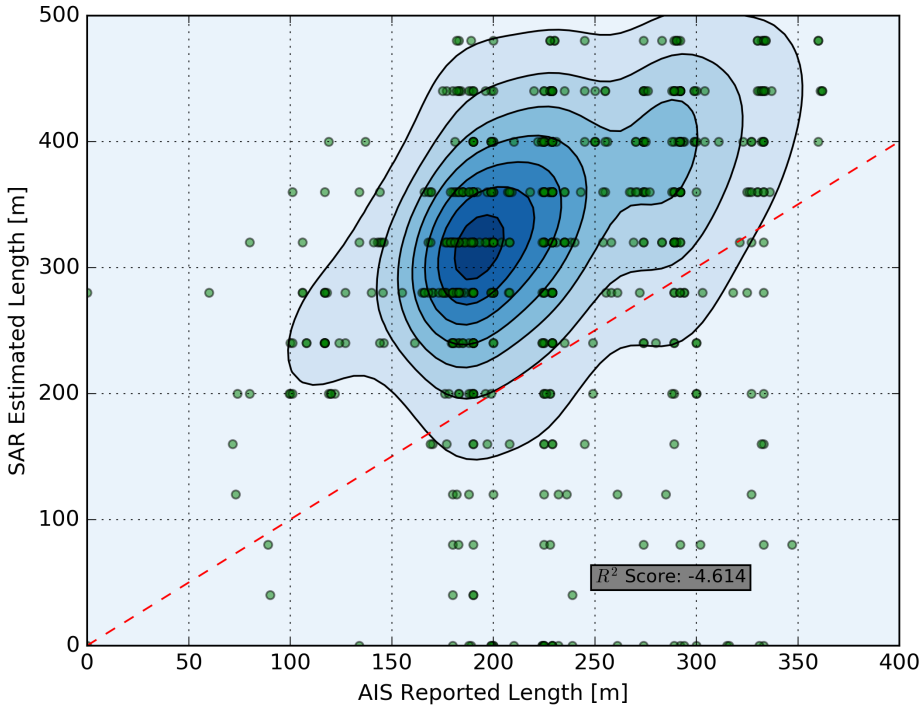

Figure 4.14. Simple estimator vessels length against AIS reported lengths

4.5.5 Estimated bearing accuracy

Figure 4.17 shows the estimated bearing confidence against the estimated bearing error. This feature was designed to give an indication of the difference between the sinogram values for the estimated bearing and the estimated bearing + 90 degrees. This was to give an indication of the confidence that could be expected due to larger vessels producing more accurate estimations. The figure shows that the confidence is high when the error is low but is not low when the error is high. The markers are scaled to the AIS reported vessel length and the inaccurate vessel confidence is not limited to small or large vessels.

This feature gives no indication of when an estimated bearing might be incorrect. It might still prove useful when used with other complementary features.

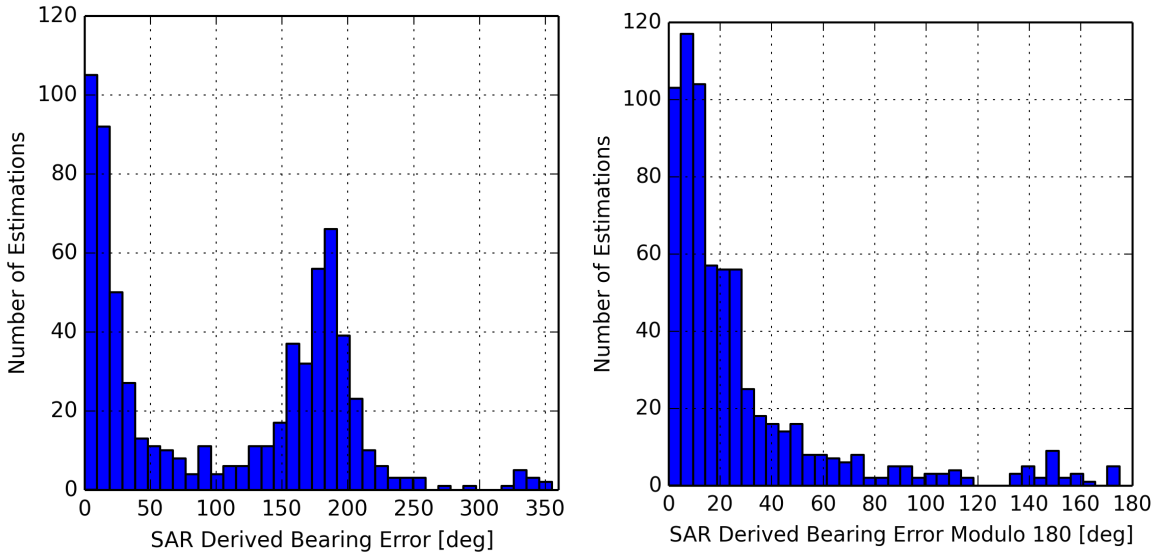

(a) SAR Bearing - AIS Bearing

(b) (SAR Bearing mod 180) - (AIS Bearing mod 180)

Figure 4.15. The error in estimated and detected bearings for detections used in this study.

4.5.6 Feature summary

Table 4.2 shows the selected features for each polarisation. Each value in the TRCS vector can be considered a single feature. A total of 51 features were selected for this study.

4.5.7 Feature selection

The derived SAR features could be added or removed from the training set before training the SVM. An experiment was performed to determine which features, if any, had a negative impact on the results of the SVC and SVR. For the SVR, one feature was randomly selected, the SVR trained, and the results stored. Another feature was randomly selected and the process repeated. This was done until all features were used. This whole process was repeated 100 times and the average (solid blue line) and standard deviation (blue area) of the SVR score was plotted in Figure 4.18. For the matched dataset, this figure shows that adding more features would increase the average score while decreasing the standard deviation of the score. The training and testing set was compiled with k-fold validation to reduce the risk of data over-fitting. The effect of using more features in the training of the SVR had a negligible effect on the processing time.

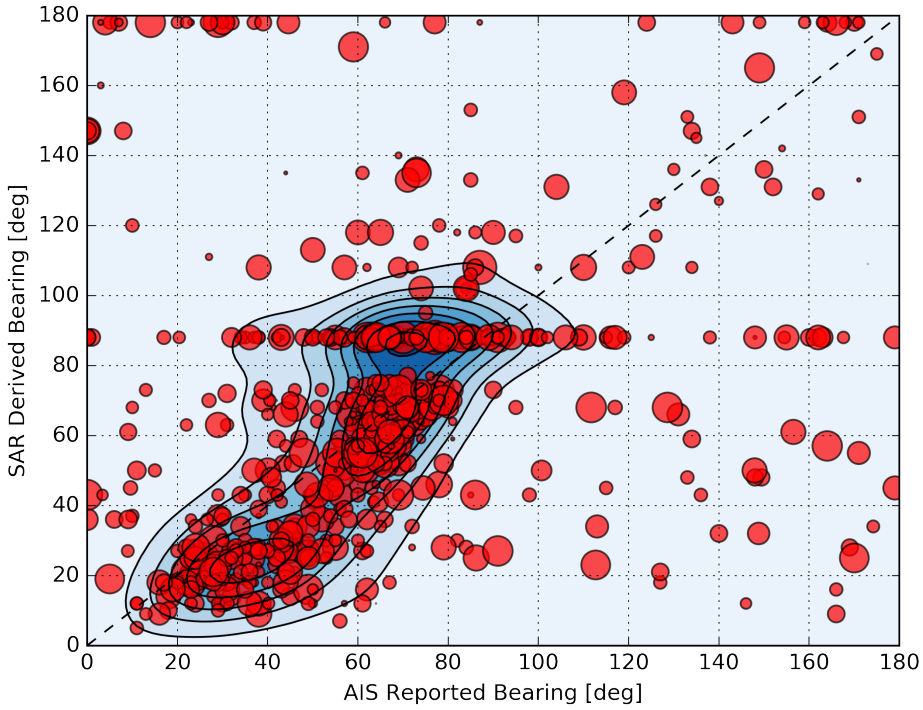

Figure 4.16. Radon estimated bearing against AIS reported bearing for all matched vessels. Markers are scaled to AIS-reported vessel length.

Obtaining the feature weighting of the features using SVM with non-linear kernels is not possible and does not intuitively describe which features are "better". The features are used in combination with each other to separate vessels classes.

The training and testing ratio was then examined for the SVR to determine the required amount of data necessary to train the algorithm while avoiding over-fitting. Figure 4.19 shows how the SVR score changes with regards to the portion of the dataset used for training. This graph was created by training an SVR algorithm on a portion of the dataset and testing with the remainder. This was performed without k-fold validation. This graph shows that relatively few examples are needed to train the SVR to an acceptable level of performance. As the training set becomes larger the performance gradually decreases due to over-fitting of the data. The poor performance where the SVM is trained with almost all the data is because the SVR score becomes dependant on the results of a few samples.

This figure shows that good training is obtained relatively quickly and is a strong indicator that, given the same processing chain to obtain new samples, the algorithm should perform the same with real world samples. This is also a strong indication that increasing the dataset alone won't improve

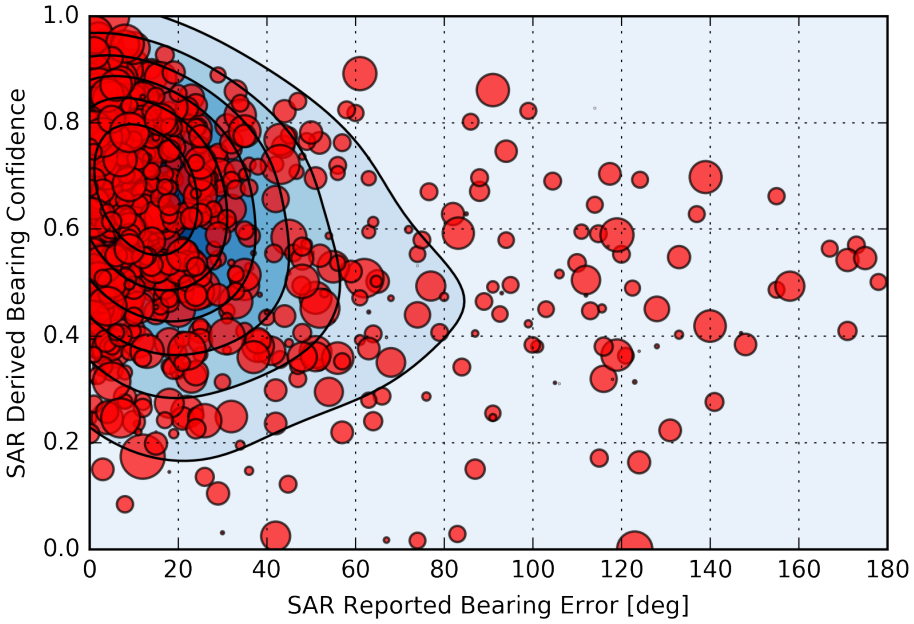

Figure 4.17. SAR-estimated bearing confidence against estimated bearing error. Markers scaled to AIS-reported vessel length.

performance; the feature space would have to be expanded to conceivably improve the score of the SVR.

Figure 4.20 shows how the classification algorithm performs when randomly adding features to the training and testing set. A number of features were randomly selected, the SVC algorithm trained and tested uses K-fold cross validation with 5 folds, and the results evaluated using the F1 score. This was repeated 100 times and the average and standard deviation of the scores for each class were used to generate this image. The figure shows that the F1 score for each classes approaches a score limit that is not improved by adding more features. The variance of the scores is improved by adding more features.

4.6 PREDICTION RESULTS

This section lists the optimal parameters, examines the results and discusses the shortcomings of the prediction algorithms used in the research. All previously mentioned features were used to train all SVC and SVR algorithms.

Table 4.2. SAR feature description

	VV Features	VH Features	Units	Total
Sea Clutter	1	1	m^2/m^2	2
SAR Incidence Angle	0	1	deg	1
Length Radar Cross Section	20	20	m^2/m^2	40
Total Radar Cross Section	1	1	m^2/m^2	2
Vessel Pixel Length	1	1	# pixels	2
Radon Estimated Bearing	1	1	deg	2
Estimated Bearing Accuracy	1	1	-	2
				51

Table 4.3. Classification results when classifying by AIS class.

	Precision	Recall	F1-score	# Vessels
30-Fishing	0.07	0.33	0.12	6
7X-Cargo	0.7	0.52	0.6	185
8X-Tanker	0.31	0.42	0.35	69
Avg / Total	0.58	0.49	0.52	260

4.6.1 Classification of AIS class

This section deals with the classification of vessels by their reported AIS class. This is the most obvious label to use from the matched dataset. Vessels that were not fishing vessels, cargo vessels or tanker vessels were removed. The test was run using k-fold cross validation with the number of folds being 5. The average and standard deviation of the confusion matrix were used to plot the results.

Table 4.3 shows the results obtained from the SVM algorithm. Figure 4.21 shows the normalised confusion matrix of the results. The results indicate that it is difficult to separate cargo vessels from tanker vessels given the extracted SAR features. Fishing vessels appear to overlap with cargo vessels, but by examining Table 4.3 there are not enough examples of fishing vessels to adequately determine this.

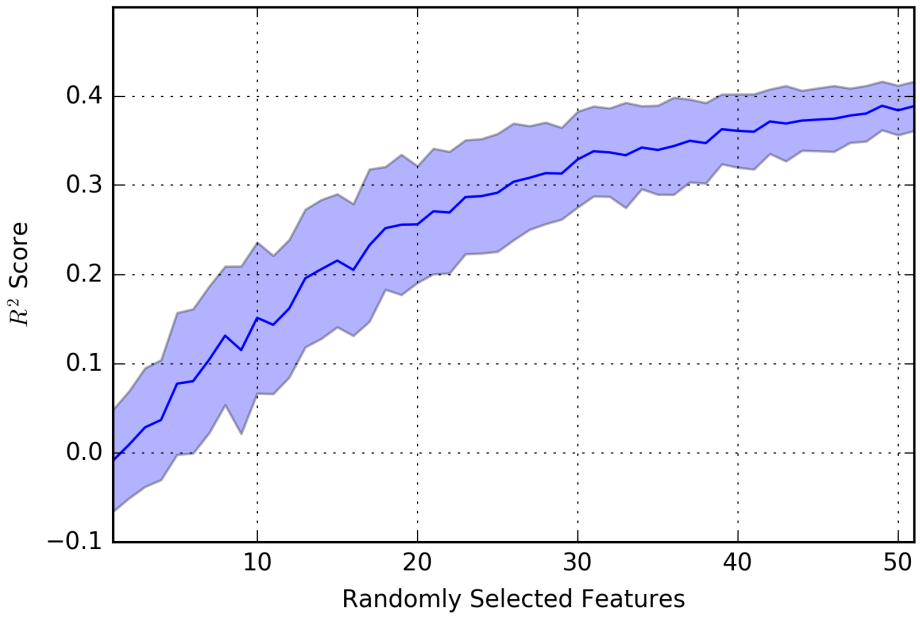

Figure 4.18. Number of randomly selected features against mean (solid blue line) and standard deviation (blue area) of score.

4.6.2 Classification of vessel size

The next label used in the SVM classification algorithm was the size label described in Figure 3.6. Small vessel classes, such as tug boats and law enforcement vessels, were included in the algorithms training set to increase the amount of small vessels available to the learning algorithm. The results are shown in Figure 4.22 and Table 4.4. The test was run using k-fold cross validation with the number of folds being 5. The average and standard deviation of the confusion matrix were used to plot the results.

These results show a stronger relationship between the SAR-extracted features and the reported vessel length. Misclassifications in this algorithm mostly fell into the next size class; large vessels were seldom classified as small vessels and vice-versa.

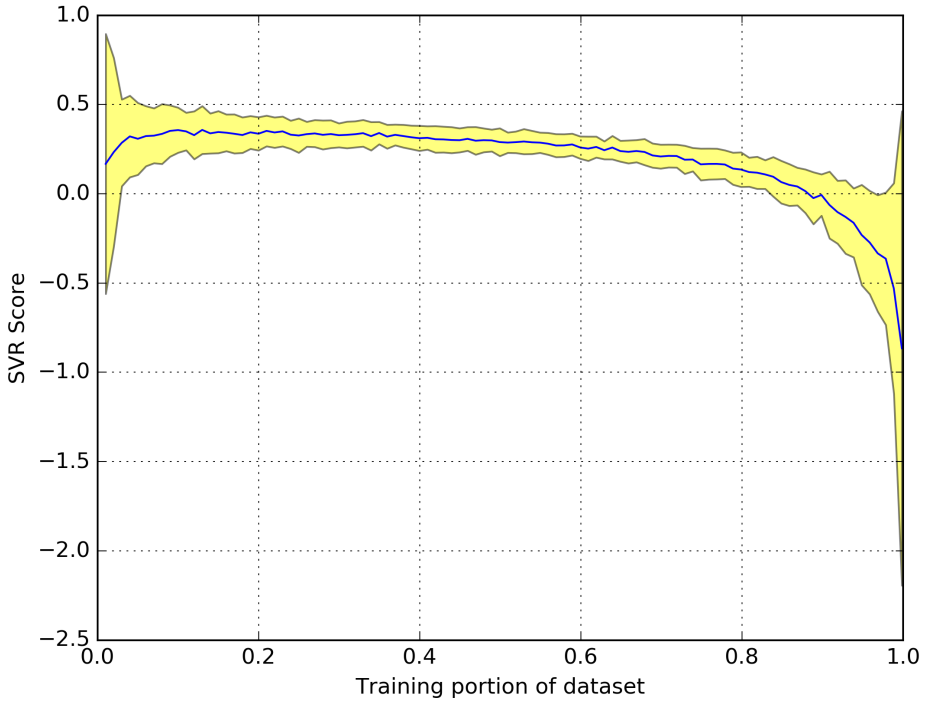

Figure 4.19. SVR score average (blue) and standard deviation (yellow) against the portion of dataset used for SVR training.

Table 4.4. Classification results when classifying by size class.

	Precision	**Recall**	**F1-score**	**# Vessels**
Large	0.61	0.71	0.66	38
Medium	0.79	0.74	0.76	80
Small	0.6	0.57	0.59	21
Avg / Total	0.71	0.71	0.71	139

4.6.3 Classification of combined AIS class and vessel size

In this test vessels were split into size categories and class categories. Cargo and tanker AIS classes, that have a large size variation, were split into small, medium and large categories as described in Figure 3.6. Other AIS classes that traditionally have small sizes were grouped into a single "Small Vessel" class. The test was run using k-fold cross validation with the number of folds being 5. The average and standard deviation of the normalised confusion matrix were used to plot the results. Table 4.5 and Figure 4.23 show the results.

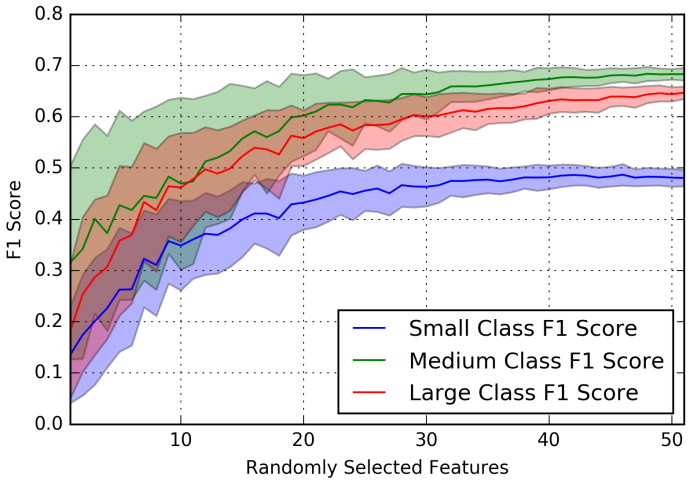

Figure 4.20. F1 score average and standard deviation for each vessel size class against number of randomly selected features.

The non-normalised confusion matrix is dominated by the amount of medium sized cargo vessels. It should be kept in mind that, because each element in the confusion matrix is an average value, the rows might not sum to 1 in the normalised confusion matrix. The standard deviation in the elements of the normalised confusion matrix are, on average, much higher than the standard deviations of the previous two tests. This is due to the small size of some classes in the testing fold used in the k-fold cross validation.

As seen in the Figure 4.21 the algorithm has problems distinguishing between cargo and tanker vessels of the same size. Small vessels are also often misclassified as small cargo, and in some cases, large cargo vessels. Given the small sample set, it is possible that this error is due to one or more large cargo vessels spoofing as a smaller class.

4.6.4 Vessel length regression

The length of vessels is an important parameter that is used to identify their class or identity. The simple estimator shows that simply counting the pixels in a SAR image and multiplying by the pixel spacing is not an accurate way to estimate a vessels length when using coarse resolution SAR imagery.

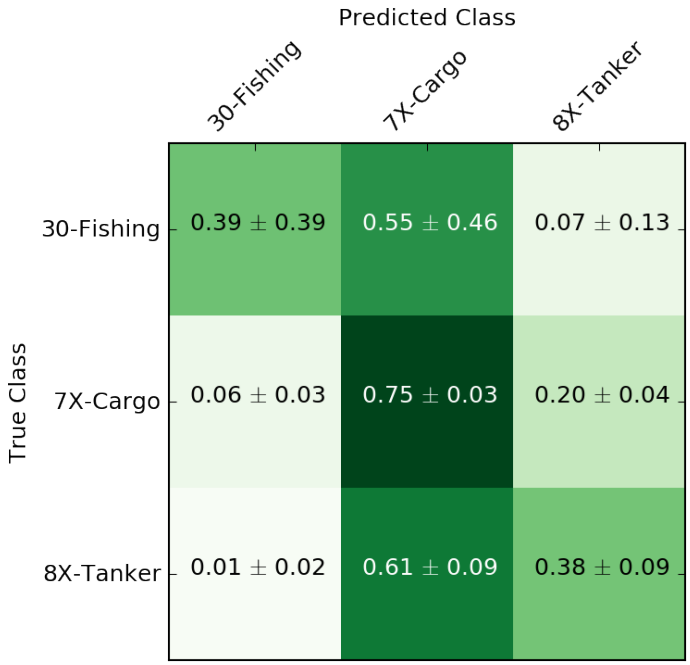

Figure 4.21. Normalised confusion matrix of average detection ± standard deviation.

Figure 4.24 shows the SVR-estimated length of the vessel. The adjusted R^2 score for this method was 0.414 which is a significant improvement on the -4.614 of the simple estimator.

This method overestimated the length of small vessels while underestimating the length of large vessels. This could be due to the under-representation of smaller vessels in the dataset, as shown in Figure 4.5.

4.7 DISCUSSION

Results from the classification and regression tests show that coarse resolution SAR cannot distinguish between AIS classes for vessels with a similar size. The information contained in the coarse SAR imagery does not contain enough information to determine the differences in physical structure of cargo and tanker vessels. Results do show that the current method of estimating a vessels size from coarse resolution SAR can be improved upon by using an SVM algorithm to either classify a vessel into a size class or to estimate its length by using regression.

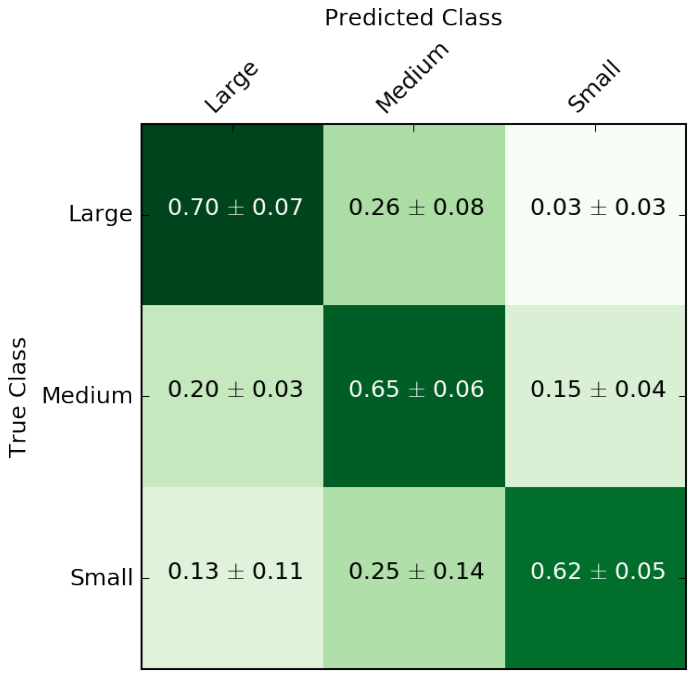

Figure 4.22. Normalised confusion matrix of average detection ± standard deviation for vessels size classification.

Certain features have also been determined to be accurate by comparing them to the AIS reported values; a vessel's bearing and position can be determined from coarse resolution SAR.

The current dataset does not contain enough examples of small ships to adequately determine performance for small vessel classes. This is due to the resolution of the SAR product, the detection method, and matching method not being optimised for small vessels.

Table 4.5. Classification results for the combined size and AIS class algorithm

	Precision	**Recall**	**F1-score**	**# Vessels**
Large Cargo	0.64	0.57	0.6	28
Large Tanker	0.35	0.57	0.43	14
Medium Cargo	0.7	0.37	0.49	70
Medium Tanker	0.04	0.1	0.06	10
Small Cargo	0.1	0.33	0.15	3
Small Tanker	0.14	0.2	0.17	5
Small Vessel	0	0	0	2
Avg / Total	0.56	0.4	0.45	132

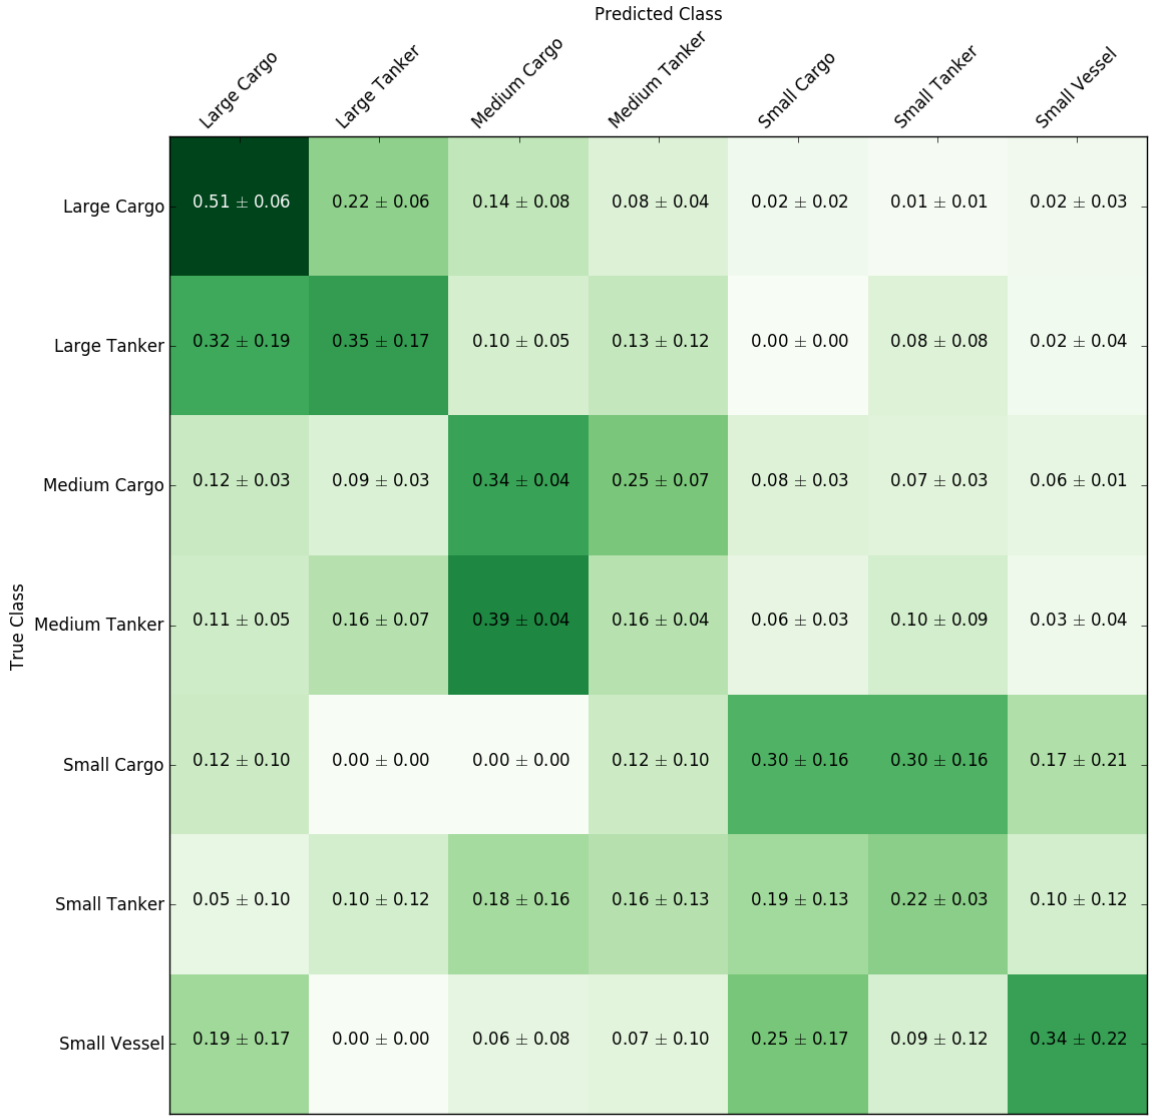

Figure 4.23. Normalised confusion matrix for classes derived from size and AIS class.

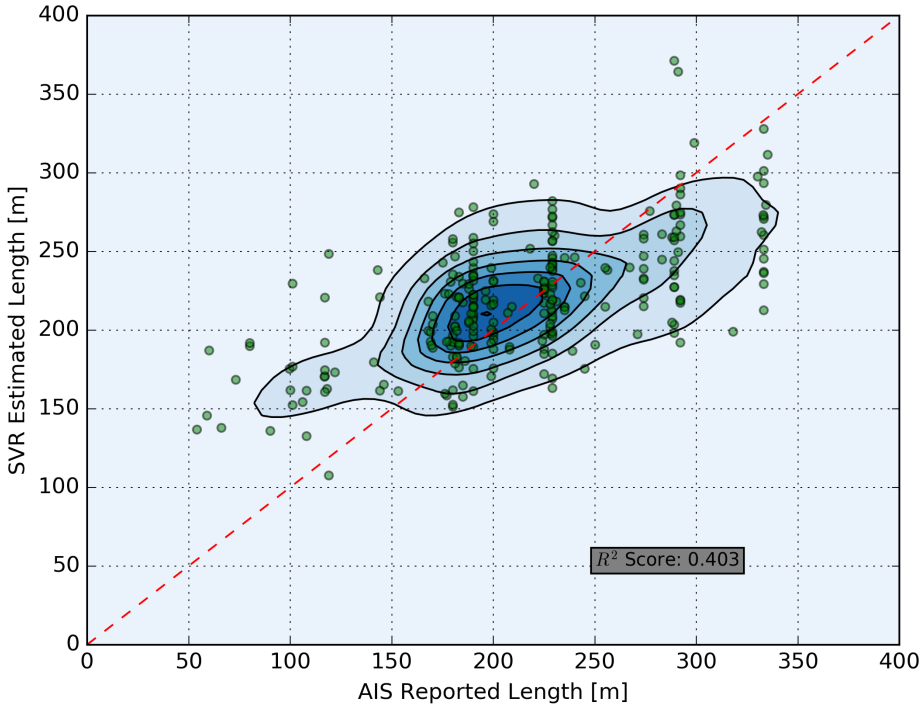

Figure 4.24. SVR estimated length against AIS reported length

CHAPTER 5 CONCLUSION

This section provides a summary of the research performed and the implications that can be derived from this. Several avenues for future work are also suggested.

5.1 SUMMARY OF RESEARCH

Vessels detected in low resolution SAR images were combined with AIS messages to obtain a data set of matched SAR and AIS detections. Several features were extracted from the SAR imaged vessels to be used to train several classification and regression algorithms to determine whether it is possible to extract any extra meaningful information from the SAR imagery and whether it would be possible to classify vessels in the absence of matching AIS messages. The dataset was also examined to determine whether it was representative of actual vessels found around South Africa's EEZ.

The conclusion of the research shows that some classes of similar sizes are indistinguishable with low resolution SAR. There is not enough information contained in the handful of ship pixels to detect differences in vessel construction that are indicative of a vessels class. The coarse resolution SAR, with a trained classification or regression algorithm, can be used to determine a vessel's size. The size of a vessel is indicative of its function and AIS class. While the size estimation of vessels will always be more accurate when using higher resolution SAR products the methods described in this thesis can be used with course resolution products to determine vessel locations, bearings and sizes over a much larger area.

5.2 RESEARCH AIMS AND FINDINGS

The aim of this research was to determine to what extent coarse resolution SAR imagery can be used for MDA purposes. Currently it is only used to detect ocean vessels, but it is not used for classification or identification purposes. Due to the lack of research in coarse resolution SAR classification, it was not the goal of this research to find the optimal solution to classifying or estimating vessel parameters, but instead to see whether this was possible at all.

The findings from this research can be summarised as follows:

- Small vessels (under 100 m long) are under-represented in coarse resolution SAR imagery, due to a combination of the resolution of the SAR image, the multi-looking used to reach the final product and the vessel detection algorithm used.

- The majority of traffic around RSA's EEZ consists three vessel classes: fishing, cargo and tanker vessels.

- An imaged vessel's bearing can be estimated using the Radon transformation and this estimation is generally accurate to within 20 degrees.

- The length of a vessel can be estimated using an SVR algorithm and this method outperforms standard pixel-counting methods.

- Vessels can be split into several size classes using an SVC algorithm and where misclassifications occur, the error usually estimates the vessel into a neighbouring size category.

- Cargo and tanker vessels are generally indistinguishable in coarse resolution SAR imagery.

- For this research a dataset of 700 examples is enough to train and test an SVM algorithm.

- The performance of the algorithms could possibly be improved by adding extra features containing complementary information.

5.3 IMPLICATIONS AND SIGNIFICANCE OF RESEARCH

This research shows that several features of a vessel can be extracted from coarse resolution SAR imagery. The location, bearing and size can be determined to a greater or lesser extent from coarse resolution SAR. small vessels were generally not detected. When assuming that vessels act as rational agents, areas that have economic interests for smaller vessels should be imaged using higher resolution SAR imagery. Limiting the use of higher resolution SAR products to certain areas would reduce the cost of a surveillance campaign and allow more acquisitions for a fixed budget.

The value of SAR imagery is that it can complement real time transponder data, such as AIS, in areas where reception of the transponder data is poor or accurate detection is critical. In areas where there is good terrestrial AIS coverage and responses to events can be quick and cheap, SAR data has less value.

By considering the zones that are generally visited by small vessels and the availability of accurate transponder data near ports some specifications can be drawn up for a future surveillance campaign using SAR imagery.

Using coarse resolution SAR imagery, with the algorithm developed in this research, the bearing and size of a vessel can now be estimated. This would provide more information to operators concerned with MDA events.

5.4 LIMITATIONS AND RECOMMENDATIONS FOR FUTURE WORK

In this research one of the most accurate features of SAR imagery was not considered: the detected vessel's position. Vessels acting as rational agents would follow trends that can be determined by historical data. Figure 5.1 shows AIS data for fishing and tanker vessels near Cape Town. It can be seen that they have different positional distributions and it would be reasonable to assume that they have different spatially distributed size and bearing probability density functions.

Extra features, that can be used in machine learning algorithms, can be derived for a specified location from historical AIS data. Determining that detected vessels have a certain class probability based on

their size, bearing and location would be valuable information.

Large amounts of AIS data could be used to identify shipping lanes, fishing zones, lanes to and from fishing zones, harbour parking areas and other maritime zones defined by ship behaviour. These zones would allow simple heuristics to be used to estimate a vessel's activity based on its location and bearing derived from a SAR image.

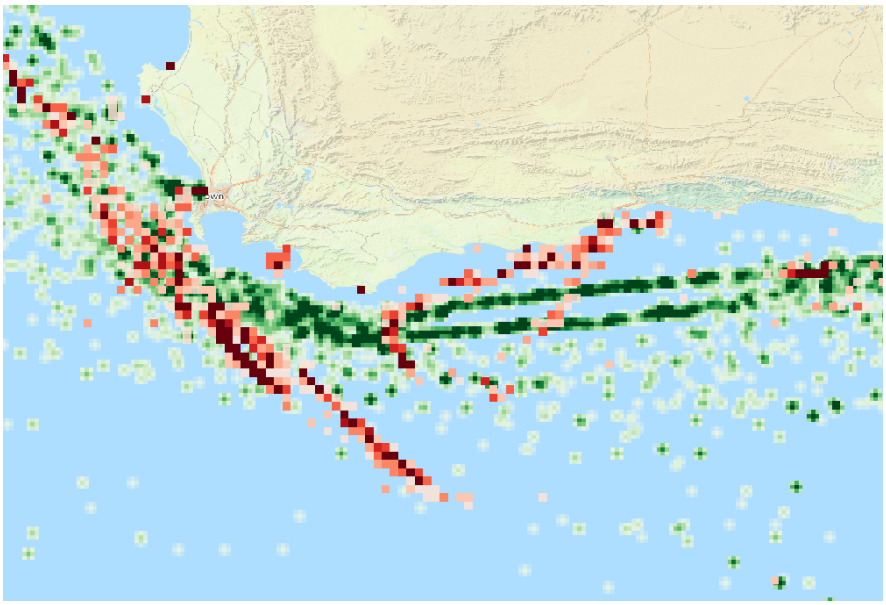

Figure 5.1. Density of fishing vessel AIS messages (red) and cargo vessel AIS messages (green) for a 1 month period.

Future research should include examining alternative machine learning algorithms to find an optimal algorithm to estimate vessel sizes and bearings would also be of value. Estimating vessel sizes using SAR products with better resolutions should be investigated. The swath trade-off with high resolution SAR products would limit their use in large scale MDA surveillance campaigns.

The detection of smaller vessel classes, such as fishing vessels, remains a problem when using coarse resolution SAR images. To mitigate this in a SAR surveillance campaign, historical AIS data should be used to identify locations that would require higher resolution images. This would allow any SAR surveillance campaign to be cost-optimised by using different SAR products based on location requirements.

There will always be some shortcoming when focusing on a single sensor to solve a problem. Using AIS alone would make any monitoring campaign vulnerable to non-cooperative vessels. This research has shown that SAR provides a good method of detecting non-cooperative vessels over large portions of South Africa's EEZ. The classification of these vessels has proven a difficult problem to solve using SAR data alone. SAR images from a single platform would have large revisit times. To monitor a large area would require a platform agnostic system that combined data from multiple sensor types with historical data and highlighted events and irregularities to operators.

This research showed that coarse resolution SAR cannot currently be used to monitor fishing vessels due to their size. Persistent monitoring using SAR is an expensive exercise and could be optimised by careful consideration of historical trends, SAR product resolution and vessel sizes.

REFERENCES

[1] C. N. Snyman, *Maratime Doctrine for the SA Navy*, R. A. J. B. Donkin, Ed. Director Naval Policy and Doctrine, 2006.

[2] A. de Klerk, "Mossgas facility," *Fischer-Tropsch Refining*, pp. 217–230, 2011.

[3] A. Lombard, B. Reyers, L. Schonegevel, J. Cooper, L. Smith-Adao, D. Nel, P. Froneman, I. Ansorge, M. Bester, C. Tosh *et al.*, "Conserving pattern and process in the Southern Ocean: designing a Marine Protected Area for the Prince Edward Islands," *Antarctic Science*, vol. 19, no. 01, pp. 39–54, 2007.

[4] M. P. Services and G. Lab, "Unlocking the economic potential of South Africa's oceans marine protection services and governance executive summary," 2014.

[5] C. P. Schwegmann, W. Kleynhans, and B. P. Salmon, "Manifold adaptation for constant false alarm rate ship detection in South African oceans," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 7, pp. 3329–3337, July 2015.

[6] R. DeAbreu, M.-F. Gauthier, and W. Van Wychen, "Sar-based oil pollution surveillance in canada: operational implementation and research priorities," in *Proceedings OceanSAR 2006–third workshop on coastal and marine applications of SAR, St. JohnâĂŹs*, 2006.

[7] B. J. Tetreault, "Use of the automatic identification system (AIS) for maritime domain awareness (MDA)," in *Proceedings of OCEANS 2005 MTS/IEEE*. IEEE, 2005, pp. 1590–1594.

[8] A. Harati-Mokhtari, A. Wall, P. Brooks, and J. Wang, "Automatic identification system (AIS): data reliability and human error implications," *Journal of navigation*, vol. 60, no. 03, pp. 373–389, 2007.

[9] American Radio Relay League, *The ARRL Handbook for Radio Communications*, ser. Publication no. 6 of the Radio amateur's library. ARRL, 2005. [Online]. Available: https://books.google.co.za/books?id=jSFTAAAAMAAJ

[10] T. Eriksen, G. Høye, B. Narheim, and B. J. Meland, "Maritime traffic monitoring using a space-based AIS receiver," *Acta Astronautica*, vol. 58, no. 10, pp. 537–549, 2006.

[11] M. Balduzzi, A. Pasta, and K. Wilhoit, "A security evaluation of AIS automated identification system," in *Proceedings of the 30th Annual Computer Security Applications Conference*. ACM, 2014, pp. 436–445.

[12] D. Mather, T. Hecht, W. Sauer, and P. Britz, "An economic and sectoral study of the South African fishing industry," 2003.

[13] 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 *2013 IEEE International Geoscience and Remote Sensing Symposium-IGARSS*. IEEE, 2013, pp. 1521–1524.

[14] International Maritime Organisation. (2016) Long-range identification and tracking (LRIT). [Online]. Available: http://www.imo.org/en/OurWork/Safety/Navigation/Pages/LRIT.aspx

[15] International Maritime Organization, *Guidance in relation to certain types of ships which are required to transmit LRIT information on exemptions and equivalents and on certain operational matters*, ser. IMO publication. International Maritime Organization, December 2008.

[16] M. Vespe, M. Sciotti, F. Burro, G. Battistello, and S. Sorge, "Maritime multi-sensor data association based on geographic and navigational knowledge," in *2008 IEEE Radar Conference*. IEEE, 2008, pp. 1–6.

[17] H. Hühnerfuss, A. Gericke, W. Alpers, R. Theis, V. Wismann, and P. A. Lange, "Classification of sea slicks by multifrequency radar techniques: New chemical insights and their geophysical implications," *Journal of Geophysical Research: Oceans*, vol. 99, no. C5, pp. 9835–9845, 1994.

[18] European Space Agency. (2016) SENTINEL-1 SAR user guide. [Online]. Available: https://sentinel.esa.int/web/sentinel/user-guides/sentinel-1-sar

[19] International Maritime Organization, *Solas: Consolidated Text of the International Convention for the Safety of Life at Sea, 1974, As Amended*, ser. IMO publication. International Maritime Organization, 2014. [Online]. Available: https://books.google.co.za/books?id=JsUeogEACAAJ

[20] J. K. Tunaley, "Utility of various AIS messages for maritime awareness," *contract*, vol. 9, p. 7009458, 2013.

[21] M. A. Cervera and A. Ginesi, "On the performance analysis of a satellite-based AIS system," in *2008 10th International Workshop on Signal Processing for Space Communications*. IEEE, 2008, pp. 1–8.

[22] R. Touzi, F. Charbonneau, R. Hawkins, and P. Vachon, "Ship detection and characterization using polarimetric SAR," *Canadian Journal of Remote Sensing*, vol. 30, no. 3, pp. 552–559, 2004.

[23] V. Wismann, "Radar signatures of mineral oil spills measured by an airborne multi-frequency radar and the ERS-1 SAR," in *Geoscience and Remote Sensing Symposium, 1993. IGARSS'93. Better Understanding of Earth Environment., International*. IEEE, 1993, pp. 940–942.

[24] C. Schwegmann, W. Kleynhans, B. Salmon, L. Mdakane, and R. Meyer, "Very deep learning for ship discrimination in synthetic aperture radar imagery," in *Geoscience and Remote Sensing Symposium (IGARSS), 2016 IEEE International*. IEEE, 2016, pp. 104–107.

[25] J. van Zyl, *Synthetic Aperture Radar Polarimetry*, 1st ed. New York: John Wiley and Sons, 2011.

[26] I. Elizavetin, "Radiometric artifacts on SAR images," in *10th International Scientific and Technical Conference-From Imagery to Map: Digital Photogrammetric Technologies, RACURS*, 2010.

[27] G. Margarit, J. Mallorqui, and X. Fabregas, "Study of the influence of vessel motions and sea-ship interaction on classification algorithms based on single-pass polarimetric SAR interferometry," in *2006 IEEE International Symposium on Geoscience and Remote Sensing*. IEEE, 2006, pp. 75–78.

[28] M. Inggs and A. Robinson, "Ship target recognition using low resolution radar and neural networks," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 35, no. 2, pp. 386–393, 1999.

[29] S. Jiang, F. Wu, C. Wang, and B. Zhang, "Civilian vessel classification with COSMO-SkyMed images based on feature analysis," in *Computer Vision in Remote Sensing (CVRS), 2012 International Conference on*. IEEE, 2012, pp. 279–284.

[30] G. Margarit and A. Tabasco, "Ship classification in single-pol SAR images based on fuzzy logic," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 49, no. 8, pp. 3129–3138, 2011.

[31] L. Gagnon and R. Klepko, "Hierarchical classifier design for airborne SAR images of ships," in *Aerospace/Defense Sensing and Controls*. International Society for Optics and Photonics, 1998, pp. 38–49.

[32] G. Margarit and J. J. Mallorqui, "Assessment of polarimetric SAR interferometry for improving ship classification based on simulated data," *Sensors*, vol. 8, no. 12, pp. 7715–7735, 2008.

[33] H. Greidanus and N. Kourti, "Findings of the DECLIMS project: Detection and classification of marine traffic from space," *Proc. Advances in SAR Oceanography from Envisat and ERS Missions*, 2006.

[34] V. N. Vladimir and V. Vapnik, "The nature of statistical learning theory," 1995.

[35] M. I. Jordan and R. Thibaux, "The kernel trick," *Lecture Notes*, 2004.

[36] C.-W. Hsu and C.-J. Lin, "A comparison of methods for multiclass support vector machines," *IEEE transactions on Neural Networks*, vol. 13, no. 2, pp. 415–425, 2002.

[37] A. J. Smola and B. Schölkopf, "A tutorial on support vector regression," *Statistics and computing*, vol. 14, no. 3, pp. 199–222, 2004.

[38] C. Werner, U. Wegmüller, T. Strozzi, and A. Wiesmann, "Gamma sar and interferometric processing software," in *Proceedings of the ERS-Envisat symposium, gothenburg, Sweden*, vol. 1620. Citeseer, 2000, p. 1620.

[39] O. S. G. Fundation, "GDAL-OGR: Geospatial data abstraction library/simple features library software," 2008.

[40] E. S. Raymond. (2015) AIVDM/AIVDO protocol decoding. [Online]. Available: http://catb.org/gpsd/AIVDM.html

APPENDIX A ADDITIONAL INFORMATION

A.1 AIS PROTOCOL

This section contains Tables describing the AIS reporting protocol.

Table A.1. AIS protocol

Parameter	Value
Channel A Frequency	161.975 MHz
Channel B Frequency	162.025 MHz
Bandwidth	25 KHz
Modulation	GMSK
Message Protocol	HDLC
Message Encoding	NMEA 0183
Multiplexing	SOTDMA
Data Rate	38400 Baud

Table A.2. Type 1,2 or 3 position report

Bit Field	Description	Units
0-5	Message Type	Constant: 1-3
6-7	Repeat Indicator	Message repeat count
8-37	MMSI	9 decimal digits
38-41	Navigation Status	Integer value for LUT
42-49	Rate of Turn (ROT)	Integer value for LUT
50-59	Speed Over Ground (SOG)	Tenths of a knot
60-60	Position Accuracy	Boolean: True if accuracy <10 m
61-88	Longitude	Minutes/10000
89-115	Latitude	Minutes/10000
116-127	Course Over Ground (COG)	Relative to true north, to 0.1 degree precision
128-136	True Heading (HDG)	0 to 359 degrees, 511 = not available.
137-142	Time Stamp	Second of UTC timestamp
143-144	Maneuver Indicator	Indicates a "Special Maneuver"
145-147	Spare	Not used
148-148	RAIM flag	Indicates whether Receiver Autonomous Integrity Monitoring is used
149-167	Radio status	Radio diagnostic information

Table A.3. Type 5 Voyage report

Bit Field	Description	Units
0-5	Message Type	Constant: 5
6-7	Repeat Indicator	Message repeat count
8-37	MMSI	9 digits
38-39	AIS Version	0=[ITU1371], 1-3 = future editions
40-69	IMO Number	IMO ship ID number
70-111	Call Sign	7 six-bit characters
112-231	Vessel Name	20 six-bit characters
232-239	Ship Type	See Table 2.2
240-248	Dimension to Bow	Meters
249-257	Dimension to Stern	Meters
258-263	Dimension to Port	Meters
264-269	Dimension to Starboard	Meters
270-273	Position Fix Type	Used for LUT describing navigation method used to determine position
274-277	ETA month (UTC)	1-12, 0=N/A (default)
278-282	ETA day (UTC)	1-31, 0=N/A (default)
283-287	ETA hour (UTC)	0-23, 24=N/A (default)
288-293	ETA minute (UTC)	0-59, 60=N/A (default)
294-301	Draught	Meters/10
302-421	Destination	20 6-bit characters
422-422	DTE	0=Data terminal ready, 1=Not ready (default).
423-423	Spare	Not used

Table A.4. AIS Message Types [40]

Type Number	Type Description
1	Position Report Class A
2	Position Report Class A (Assigned schedule)
3	Position Report Class A (Response to interrogation)
4	Base Station Report
5	Static and Voyage Related Data
6	Binary Addressed Message
7	Binary Acknowledge
8	Binary Broadcast Message
9	Standard SAR Aircraft Position Report
10	UTC and Date Inquiry
11	UTC and Date Response
12	Addressed Safety Related Message
13	Safety Related Acknowledgement
14	Safety Related Broadcast Message
15	Interrogation
16	Assignment Mode Command
17	DGNSS Binary Broadcast Message
18	Standard Class B CS Position Report
19	Extended Class B Equipment Position Report
20	Data Link Management
21	Aid-to-Navigation Report
22	Channel Management
23	Group Assignment Command
24	Static Data Report
25	Single Slot Binary Message,
26	Multiple Slot Binary Message With Communications State
27	Position Report For Long-Range Applications