

A Synthetic Aperture Radar ship dataset for detection, discrimination and analysis

C. P. Schwegmann, W. Kleynhans, B. P. Salmon, R. Meyer,

Abstract—Synthetic Aperture Radar is a monitoring solution which is especially well fitted to maritime surveillance. The large swath widths, time-independent, weather resistant observations can be very useful for the detection of ships typically invisible to other forms of monitoring. A dataset which contains 43 Sentinel-1 Extended Wide Swath images and 3 RADARSAT-2 ScanSAR Narrow images between the 6th October 2014 to the 22 July 2015 is described. These images cover a large percentage of the South African Exclusive Economic Zone using multiple polarisations and three different resolutions. A detailed description of the dataset is given including: the entire generation process from the original compressed data to the geocoded images; the dataset organisational structure tailored to each step of the ship detection process; ship referencing and attribute extraction procedures; supplementary Automatic Identification System transponder information and matching; and ship reference and Automatic Identification System matched ship attribute analysis. A number of novel components are introduced, including a three-class discrimination dataset component which helps provide a more granular approach to ship discrimination performance analysis. The aim of the dataset is to introduce a single Synthetic Aperture Radar dataset which covers a wide range of possible requirements for scientific ship detection method analysis using the newest Synthetic Aperture Radar imagery available.

Keywords—Marine technology, Synthetic aperture radar (SAR), Machine Learning

I. INTRODUCTION

SYNTHETIC Aperture Radar (SAR) satellite imagery has been used for the improvement of a country's Maritime Domain Awareness (MDA) for many years [1]–[3]. The ability to penetrate cloud cover, day or night, ensures that ships at sea can be tracked independent of their cooperation. The cost to observe a large area is reduced by employing SAR satellite imagery, specifically for the case of medium resolution imagery with large swath widths. This type of imagery provides the ability to track ships that occur near or beyond a country's Exclusive Economic Zone (EEZ) border, usually well beyond the reach of conventional ship tracking methods such as terrestrial Automatic Identification System (AIS).

Larger swath width SAR imagery comes at the expense of reduced image resolution which in turn prohibits the minimum

detectable size of ship. This may cause a problem for ships closer to the shore but ships further out to sea typically need to be larger ships to survive further out. A number of recent ship detection and discrimination studies are more focused on high resolution imagery [4]–[7] rather than medium resolution imagery [7], [8]. Furthermore, the manner of reporting ship detection results often combines detection and discrimination which can unintentionally obscure where methods perform well and where they fail.

This paper presents a new, large SAR dataset focused on medium resolution imagery used for ship detection, discrimination and analysis. The paper will detail the entire dataset creation process, what and how certain decisions were made for the dataset creation and will provide an extensive analysis on the dataset to determine operation parameters for new methods design and operational systems. The paper also introduces a new discrimination dataset component that allows for a more granular approach to ship discrimination performance evaluation whilst making use of similar machine learning dataset structures. The intent of this is to produce a dataset that has consistent terminology in the long term for SAR ship detection and allow for more detailed comparisons between new and old ship detection and discrimination methods.

The paper structure is as follows: Section II describes the SAR imagery and some of the details regarding exactly which satellites and resolutions are used, their coverage over the oceans and a general overview of the dataset. Section III provides the preprocessing steps applied to the imagery to create the common dataset format from which the rest of the paper is derived. Given the preprocessed images, section IV describes the referencing procedure used to generate the reference ships against which results from ship detection and discrimination methods can be compared. This section also describes how reference ships are matched to AIS transponder messages. Section V describes the two main organisational structures of the dataset, namely for ship detection and discrimination. Section VI provides a through analysis of the dataset including distribution of ships across the satellites, resolutions and polarisations as well as reference and matched ship attribute analysis. Finally, section VII provides closing remarks and a description of possible future work on the dataset.

II. SAR IMAGERY

To effectively build a operational ship detection and tracking system a scientific exploration of available data needs to be undertaken. The extents of the data need to be tested thoroughly in order to determine the best parameters when designing operational systems. To aid in doing this a dataset

C. P. Schwegmann, W. Kleynhans and R. Meyer are with the Department of Electrical, Electronic and Computer Engineering, University of Pretoria as well as the Remote Sensing Research Unit, Meraka Institute, CSIR, Pretoria, South Africa. Email: cschwegmann@csir.co.za. Telephone: +27 12 841 3207.

B.P Salmon is with the School of Engineering, University of Tasmania, Australia as well as the Remote Sensing Research Unit, Meraka Institute, CSIR, Pretoria, South Africa.

has been created that allows for rigorous scientific evaluation and provides a platform to allow for comparison among methods. The dataset needs to be large enough to encompass most, if not all, possible detection scenarios and strict reference points need to be included to determine when and how methods fail. To this end this paper presents a uniquely large SAR dataset, covering two separate SAR satellites using three separate resolutions. The geographic extent of the SAR imagery used in this dataset is shown in Fig. 1. The area covers approximately 1.17 million km². This means the entire dataset covers approximately 76% of the South African Exclusive Economic Zone (EEZ) and 90% of its territorial waters. The data set was acquired between the 6th October 2014 and the 22nd July 2015. The proposed SAR dataset has two main subsets of data from two different satellites, namely Sentinel-1 and RADARSAT-2. The dataset covers 1596 ships all across the South African EEZ and further out and the distribution of these ships can be seen in Fig. 2.

A. Sentinel-1

The Sentinel-1A satellite was launched on 3 April 2014 and became operational on the 6th October 2014 with a C-band sensor. The satellite is a Sun-synchronous, near-polar circular orbiter with a repeat cycle of 12 days. Sentinel-1B was launched on the 25 April 2016 and is expected to be operational through the course of 2016. The current 12 day repeat cycle of Sentinel-1A will be improved to a 6 day repeat cycle when Sentinel-1B becomes operational. The first imagery freely from the Sentinel-1 constellation was released on the 6 October 2014. Two of these images were located over South Africa and were used as the area of study in [9], [10]. All imagery used in this data set are from the Extra Wide (EW) swath acquisition mode with Incidence Angles between 19.0° and 47.0° and a swath width of 400 km. At the time of writing, no Single Look Complex (SLC) imagery over South Africa in EW mode was available, specifically over the EEZ waters of South Africa. All of the data was Ground Range Detected (GRD) imagery using either the High (GRDH) or Medium (GRDM) resolution class. GRDH images have a resolution of 50 x 50 m and pixel spacing of 25 x 25 m (in range and azimuth respectively). Similarly, GRDM imagery has a resolution of 93 x 97 m with a 40x40 m pixel spacing. The data set covers the three of the possible four EW polarisation acquisition modes, namely: SDH (HH+HV), SDV (VV+VH) and SSV (VV). There are a total of 22 Sentinel-1 acquisitions, 21 containing two polarisation images (either SDH or SDV) and one with a single VV polarisations, for a total a total of 43 Sentinel-1 images. A summary of the imagery within the data set is given in Table I.

B. RADARSAT-2

The Canadian RADARSAT-2 satellite launched on 14 December 2007 with a C-band sensor. It has a repeat cycle of 24 days within a Sun-synchronous, near-polar circular orbit. The RADARSAT-2 imagery was acquired for a number of internal uses and all three were acquired near Cape Town, South Africa.

TABLE I: Dataset information per image sensor and polarisation.

Attributes	Sentinel-1		RADARSAT-2
	GRDH	GRDM	SCNA
Acquisitions	6	16	3
Images Total	12	31	3
Type	EW	EW	SGF
Incidence Angle (°)	19.0 – 47.0	19.0 – 47.0	20.0 – 39.0
Swath Width (km)	400	400	300
Resolution [rg x az] (m)	50 x 50	93 x 97	81 x 30
Pixel Spacing [rg x az] (m)	25 x 25	40 x 40	25 x 25
Number of Looks [rg x az]	3 x 1	6 x 2	2 x 2

All three are ScanSAR Narrow (SCNA) imagery of the non-SLC type (intensity only). The imagery has a resolution of 81 x 30 m and pixel spacing of 25 x 25 m (in range and azimuth respectively). The images have a swath width of 300 km and have an incidence angle range of between 20.0° and 39.0°. All three SAR images were acquired in the HH polarisation. A summary of the RADARSAT-2 imagery is given in Table I.

III. PREPROCESSING

The following section will describe the preprocessing steps applied to the image before it is analysed to determine if ships are present. Following this a detailed description of the geocoding and radiometric calibration steps will be given. Both Sentinel-1 and RADARSAT-2 acquisitions data are received as a compressed zip files. The decompressed zip files are fed into a collection of tools (namely either Geospatial Data Abstraction Library (GDAL) or GAMMA based tool-chains) to generate interoperable Geographical Tagged Image File Format (GeoTIFF) files and associated Extended Markup Language (XML) meta-data files. Each acquisition and polarisation within (for Sentinel-1 images with multiple polarisations per acquisition) are processed independent from one another. The final output of the preprocessing steps is a geocoded GeoTIFF for a consistent data access strategy.

A. Radiometric Calibration

The first step of the dataset preprocessing is the conversion from digital numbers into normalised Radar Cross Section (RCS) values using Radiometric Calibration [11], [12]. To perform geographical parameter extraction radiometric calibration is an absolute necessity, even more so if comparisons are made between objects across acquisitions. To measure and compare ships across the different acquisitions radiometric calibration of SAR imagery is a required preprocessing step. The Sentinel-1 products provides an updated Look Up Table (LUT) for its Level-1 products defined as

$$A_{\sigma} = \sqrt{\frac{A_{DN}^2 K}{\sin(\alpha)}} \quad (1)$$

where A_{DN} is an unsigned 16 bit LUT which defines the scaling from internal SLC to GRD product, K is the

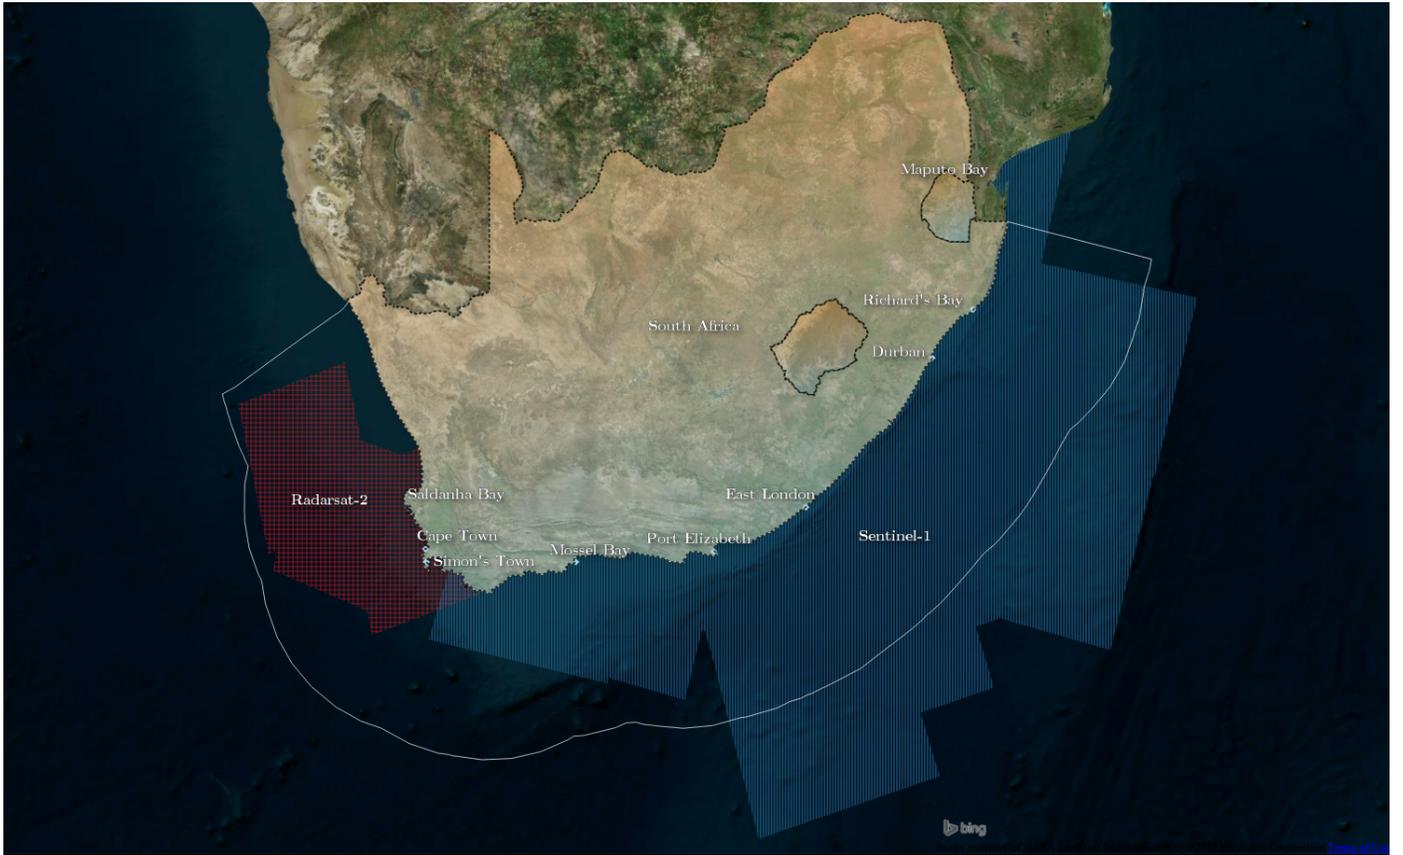


Fig. 1: A map of South Africa and the extent that the Sentinel-1 and RADARSAT-2 images cover. The white border indicates South African EEZ. Approximately 80% of the EEZ is covered by the SAR data set. In addition to this a number of images cover the Maputo Bay area of Southern Mozambique and form part of the data set.

single calibration constant for all final products and α is the local incidence angle [12]. Finally, the average backscatter coefficient $\sigma^0(i)$ [13] at digital number $DN(i)$ is

$$\sigma^0(i) = \frac{E\{DN(i)\}^2}{A_\sigma^2} \quad (2)$$

where $E\{DN(i)\}$ is the mean pixel amplitude digital number values over a small area (typically 3×3), taken directly from the measurement (GeoTIFF) file. A similar procedure is followed for RADARSAT-2 radiometric calibration. For a complete overview of radiometric calibration in SAR see [11].

B. Georeferencing and warping

The next preprocessing step involves georeferencing the input SAR image. This assigns real-world coordinates to each pixel of the image using Ground Control Points (GCPs). These GCPs are inherent within the GeoTIFFs in both Sentinel-1 and RADARSAT-2. More precise estimations of these pixel longitudes and latitudes can be acquired using each satellites orbital state vector meta-data [14] but for the purpose of this study the ones provided with the GeoTIFF files were deemed

sufficient given that the dataset is comprised of medium resolution imagery and the corrections should not make a significant impact to detection locations within reasonable error tolerances.

To provide a uniform access format all of the images are then warped to be north-facing (geocoded) images. For the case of this study all images were warped to the World Geodetic System 1984 or "WGS84" projection. This was selected due to being the same datum used for Global Positioning System (GPS) coordinates which align with transponder-based acquisitions used in the data set later on.

One important aspect to note is that with warping requires interpolation when the image is resampled. The simplest method, nearest neighbour, was used as the resampling method to reduce the amount of alterations done to pixels. Other resampling methods such as bilinear, cubic, cubic spline and Lanczos resampling are available but were not used due to the possibility of altering ship pixels and ocean backscatter pixels significantly (from smoothing, sharpening, changing the distribution of the clutter, etc).

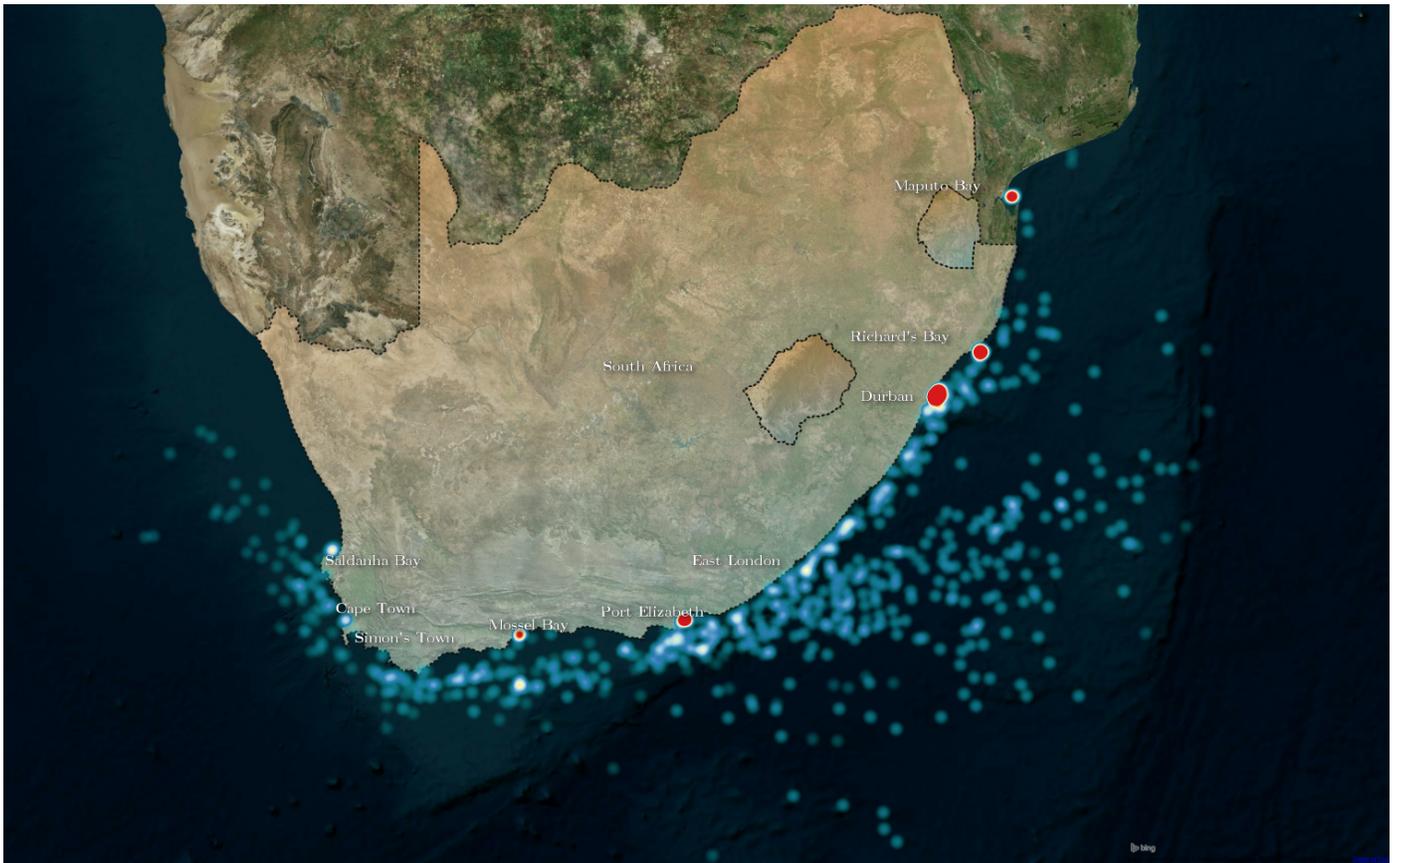


Fig. 2: A distribution map of all the 1596 ships in the data set. The Durban and Richard's Bay ports show significantly higher ship density than most of the South African coast. A number of factors cause this including the fact that most of the Sentinel-1 images occur over these areas and because these ports are the closest to the South African economic hub, Gauteng.

C. Land Masking and Artifact removal

The final step of preprocessing is to remove any land within the image. The complexity of the land removal step is based on the needs of the dataset and can range from simple landmasking using a shapefile to advanced coastline extraction procedures [14]–[17]. Due to the nature of this dataset such high accuracy land removal near to the coast was not necessary and a buffer was used to ignore ships too close to the shore. This aligns with the fact that any ships close enough to the coast can be detected using terrestrial transponder-based systems and the necessity for highly accurate SAR detections that near to the coast is reduced.

IV. SHIP REFERENCING

In this section the definition of a ship is provided as well as the process that was undertaken to identify all ships in this dataset's SAR imagery. By identifying the positions and shapes of the ships within the dataset, automatic methods can be tested against these reference ships to determine the quality of their output.

A. Ship definition

In SAR intensity imagery covering the ocean, a ship is defined as any object that is sufficiently brighter than its surrounding ocean backscatter [1]–[3]. For the sake of simplicity and due to the resolution of the SAR imagery for this dataset it was assumed that a ship is an area of ocean that has significantly higher backscatter than its neighbours and also is at least 2 pixels in length. Ships that are smaller than these dimensions are beyond the scope of this dataset and would be better analysed using higher resolution SAR imagery. This definition, therefore, ignores single pixels with high backscatter due to speckle noise. The intent of this dataset is to cover most areas of the EEZ and thus wider swath widths are necessary which in turn reduces the available resolution of the dataset. To this end, any objects that do not meet these criterion are deemed non-ships or ocean clutter and are defined as false positive areas of interest.

Each ship within the dataset has four associated images, namely a “ship patch”, “reference patch”, “ship subimage” and “reference subimage”. Patches are images which cover a large area of pixel (101×101) centred on the ship. Subimage are smaller images (21×21) with little to no other information

but the ships found at the centre of the image. The large size of the patches ensure ships of all possible sizes are captured within the image and the subimage’s smaller size is based on statistical evaluations of all ships present in the dataset (see section VI). Patch images provide local context to evaluate where the ship is (near a harbour, far out at sea, within a crowd of other ships). This context is important to help identify why detection errors may occur relative to various features around the ship. This may be necessary if an algorithm is made to make use of the ship’s surrounding environment. Finally, by providing these larger images around the ship the need to scan the large SAR images is reduced. Subimages are smaller images around the ship designed to emulate machine learning datasets such as the MNIST [18] or CIFAR-10 and CIFAR-100 datasets [19] and a detailed discussion of subimages is discussed in section V.

B. Reference images

One the intentions of this paper is to study how well the SAR data relates to the physical ship positions and sizes. To this end, “ground truth” or “reference” images need to be created. For the purpose of this study, a reference image is a binary image which indicates “true” for pixels most likely associated with the ship at the centre and “false” for all other pixels. Reference images sizes were set to match the same size as the SAR ship patches and subimages (101×101 and 21×21 respectively).

C. Patch and subimage creation

To effectively identify ships within the SAR imagery and then to reference them correctly the following procedure was taken:

- 1) Each of the 46 SAR images are scanned and targets of high backscatter are identified using both the basic Geo-TIFF and a contrast stretched version (which enhances darker areas and darkens brighter, near-nadir areas of the SAR image). Each potential bright area’s latitude and longitude was noted. This process was repeated twice for each image to ensure no potential bright target areas were missed.
- 2) For each potential target, a patch of the SAR image centred on the target was extracted. For each target the patch was reexamined using the context provided by the patch and marked as either a ship or not. Known geographical features such as islands, sea-platforms, rough sea areas were used to eliminate targets with high backscatter that were in fact false alarms.
- 3) Another identification step occurred by matching ships across polarisations. Ships that appeared within one pixel (25m or 40m depending on sensor resolution) for a single acquisitions were deemed the same ship across different polarisations. Ships that did not match across polarisations was still kept if they matched the criteria as discussed above.
- 4) For each identified ship patch an associated reference patch was created as described above. It should be noted that any other potential ships or objects not at the centre were ignored within each reference patch.

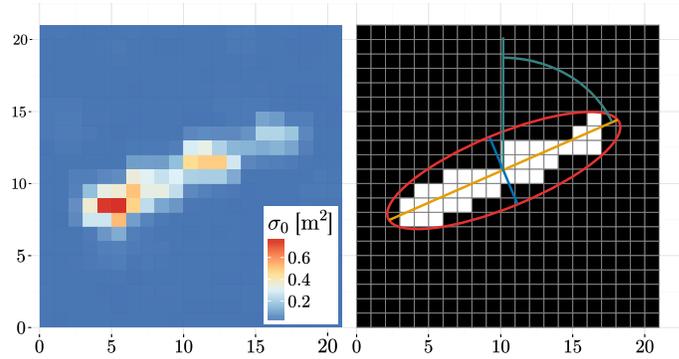


Fig. 3: A radiometrically correct SAR ship subimage and its associated reference subimage (ship number 160 of 1596). The image also shows the four main attributes of the reference image, namely the ship’s centroid \mathbf{C} , Major Axis Length L_{maj} , the Minor Axis Length L_{min} and ship orientation or heading θ extracted using the Minimum-Volume Enclosing Ellipsoid.

- 5) A final stage of ship identification matched the identified centres of the SAR ships against a dataset of known AIS positions. The exact details of the AIS referencing is discussed in section IV-E.
- 6) Each reference patch was analysed to determine the ship’s centroid. For each centroid a 21×21 sized subimage was extracted from both the ship patches and reference patches to create the ship and reference subimages.

The final collection of 1596 ship subimages is shown in Fig. 4a with associated reference subimages shown in Fig. 4b.

D. Ship attributes

For each reference area, a list of attributes are extracted. These attributes provide valuable information to help guide the design of detection methods. For the purpose of this paper, four attributes are calculated namely the ship’s centroid \mathbf{C} , Major Axis Length L_{maj} , the Minor Axis Length L_{min} and ship orientation or heading θ . Within the dataset itself, it is possible to extract more attributes but for the sake of brevity this paper only discusses these four.

The ships centroid is the location of the centre of mass of the reference image. Given a finite set \mathcal{S} of k reference points with a “true” value such that $\mathcal{S} = \{\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_k\}$ in \mathbb{R}^2 where $\mathbf{P}_1 = (x_1, y_1)$ is the first point’s x and y coordinates, the centroid $\mathbf{C} = (x_{\text{centroid}}, y_{\text{centroid}})$ of the reference image can be calculated as [20]

$$\mathbf{C} = \frac{\mathbf{P}_1 + \mathbf{P}_2 + \dots + \mathbf{P}_k}{k}. \quad (3)$$

To calculate the other three ship attributes the Minimum-Volume Enclosing Ellipsoid (MVEE) can be used [21]. For the sake of completeness it should be stated that the method is defined for higher dimensionalities but for the case of the 2-D reference images the method actually calculates a minimum-area enclosing ellipsoid. As defined in [21], a full-dimensional

ellipsoid $\mathcal{E}_{Q,C}$ in \mathbb{R}^2 is specified by a 2×2 symmetric positive-definite matrix Q with centre \mathbf{C} is defined as

$$\mathcal{E}_{Q,C} = \{\mathbf{P} \in \mathbb{R}^2 : (\mathbf{P} - \mathbf{C})^T Q (\mathbf{P} - \mathbf{C}) \leq 1\}. \quad (4)$$

The area of this ellipsoid is therefore calculated as $\text{Area}(\mathcal{E}_{Q,C}) = \eta \det(Q)^{-1/2}$ where η is the area of the unit ball in \mathbb{R}^2 and \det is the determinant applied to Q [21]. The MVEE of a set of points \mathcal{S} is defined as $\text{MVEE}(\mathcal{S})$ and will satisfy the following

$$(1/2)\text{MVEE}(\mathcal{S}) \subseteq \text{conv}(\mathcal{S}) \subseteq \text{MVEE}(\mathcal{S}). \quad (5)$$

The notation $\text{conv}(\mathcal{S})$ is the convex hull of \mathcal{S} and the left-hand side ellipsoid is scaled by $1/2$ around its centre. Assuming \mathcal{S} is the set of vertices for the full-dimensional polytope $\mathcal{P} \subseteq \mathbb{R}^2$ (the reference ship shape), then $\text{MVEE}(\mathcal{S})$ yields a rounded approximation of \mathcal{P} . For exact details on the calculation of the MVEE, see [21].

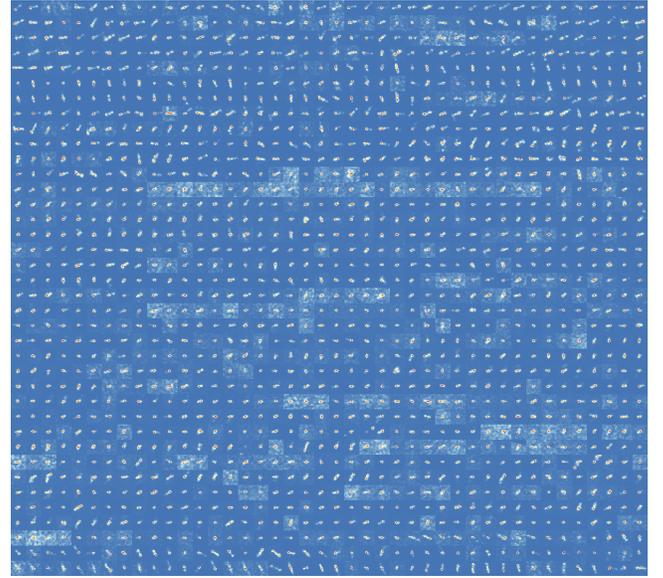
The Major Axis Length L_{maj} is the approximate length of the ship and can be calculated by finding the furthest two points that lie on the MVEE. A ship's approximate width, Min Axis Length L_{min} , is calculated as the line perpendicular to the Major Axis. Finally, the ships heading θ is calculated as the angle between the Major Axis Length and north at 0° . The attributes for a single reference subimage are shown on the right in Fig. 3.

E. AIS Ship Identification

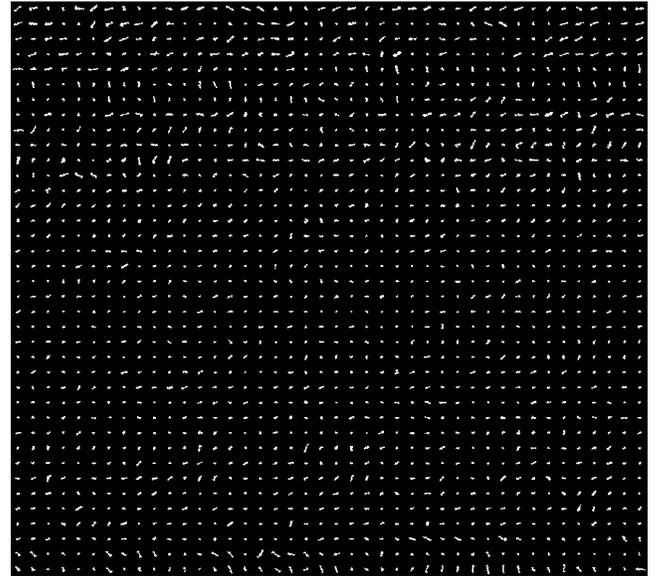
Automatic Identification System (AIS) is a transponder system attached to ships which transmits their positions at regular intervals. These positions can be tracked over a period to form ship tracks. AIS messages are received in one of two ways, either via coastal-based terrestrial AIS receivers and inland satellite based receivers. Coastal-based AIS receivers have an estimated range of approximately 74 km from the coast and are received on a near real-time basis whereas satellite-based AIS messages are received every 6-12 hours but have a near-global range [22]. The dataset of AIS points that covered the SAR dataset includes approximately 220 million AIS messages obtained from the 6th October 2014 to the 22 July 2015 and covers the entirety of South Africa's EEZ.

The last step of the referencing procedure is the matching of transponder-based ship positions to that of the ships within SAR image and this is a vital component in ship detection literature [7], [14], [23]–[25]. The process concerns itself with matching the centroid Latitudes and Longitudes of SAR ships to positions received from the ships themselves. The primary problem with matching AIS to SAR ships is a temporal one. AIS messages may or may not line up with SAR acquisition times and so it is rare that all ships in a SAR image transmit an AIS message at the same time the image was taken by the satellite. To combat this AIS tracks are built using the history of the ship's AIS messages and positions to build a map of the likely area a ship was at any given moment.

For this dataset, all AIS messages 12 hours before and after a SAR acquisition were used to build AIS tracks. An AIS track is a collection of AIS messages built-up over time to indicate the history of a ship and will have either an



(a) All ship subimages



(b) All reference subimages

Fig. 4: All 1596 ship (a) and reference (b) subimages within the data set. These cover the two satellites at three resolutions. Each ship subimage has been normalised so that the brightest point is at unity for display purposes. Each reference subimage is analysed using the MVEE to calculate attributes such as SAR ship length, width and orientation.

interpolated AIS position at the time of a SAR acquisition or a fixed AIS message at that time. An interpolated AIS position is the estimated position of a ship at the SAR time of acquisition. This position is created using the time and

position of the last AIS message before and first AIS message after a SAR acquisition. If an SAR ship had a single AIS message (interpolated or fixed) within 500m of its position it was deemed to have a “hard” AIS match. If more than one AIS position existed for a single SAR ship (therefore multiple tracks) within the 500m window then a ship was said to have a “soft” AIS match. For this dataset, only hard matches were considered to ensure that SAR-AIS matches were of the highest quality with little to no matching ambiguities. Across all 1596 identified SAR ships 494 ships were matched to AIS messages. Of these 494, 437 of them were unique matches whereas the additional 57 were AIS matches for the same ship across different polarisations. Unfortunately, the AIS dataset available did not cover any of the RADASAT-2 images near Cape Town harbour and therefore only Sentinel-1 GRDM and GRDH images had AIS coverage.

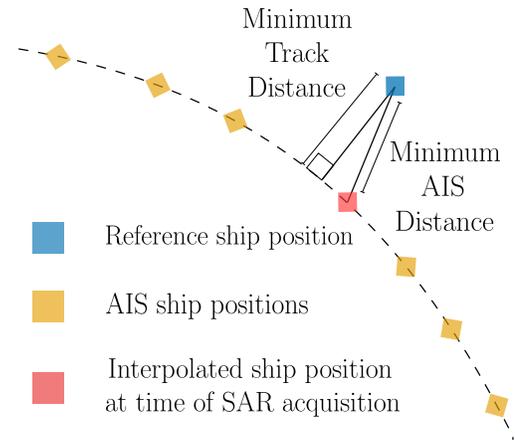
V. DATA ORGANISATION

The purpose of creating such a dataset as described above is to facilitate repeatable, verifiable experiments within a fixed and known extent of data. To this end this dataset has been split into two organisational structures, namely the SAR images themselves with associated ship positions used for the process of ship detection and a collection of subimages and patches with associated metadata for ship discrimination and analysis. In most of the literature, the process of ship detection and discrimination is combined to form a single ship detection stage. For the context of this dataset, ship detection refers to methods that can detect bright areas within SAR images and extract these positions (which may or may not be ships). Ship discrimination refers to methods that directly process these detections to identify which are actually ships and which are not. This distinction is important as methods that perform well in one task may not necessarily perform well for the other task and splitting the ship detection process into separate entities for detection and discrimination allows for greater focus on each task [26]. The following sections describe the two data organisational structures which help facilitate the creation and evaluation of methods for each task.

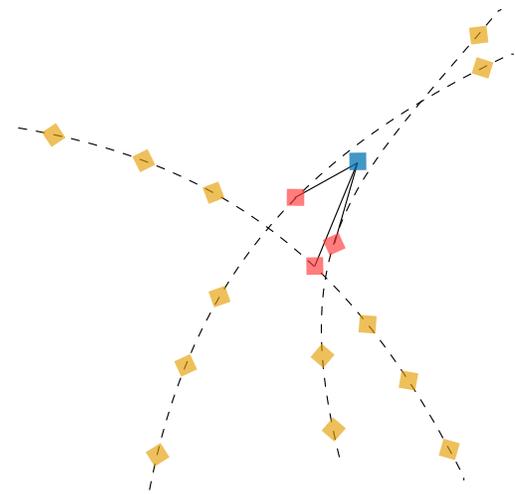
A. SAR imagery for ship detection

Almost all ship detection literature uses a known collection of ship positions to determine performance. These are created using either a referencing procedure similar to section IV [24], [26] or by using the AIS matched ships only and ignoring all other detections [23]. For this dataset, a comma-separated text file is stored with each SAR image and contains the positions of all the ships identified within that image. Each row within this text-file provides the ship’s position in Latitude/Longitude coordinates and geocoded image-specific row/column coordinates. This representation helps when comparing ship detection methods which only provide detected ship positions and helps simplify SAR-AIS ship matching.

Ship detection methods such as constant false alarm rate (CFAR) prescreening [1], [2], [23] and Wavelet processing [25] results are then compared to these known ship positions to determine ship detection accuracy performances. This type of



(a) Hard SAR-AIS match



(b) Soft SAR-AIS match

Fig. 5: Two subfigures showing the different between a Hard (a) and Soft (b) AIS match. When a single AIS tracks aligns within a reference ship patch image matches and the track distance, time difference and AIS distance were below the specified minimums a match was recorded. For multiple AIS tracks within the reference patch image area and within the specified parameters a soft match was assigned to the target. Only hard matches were considered in this dataset.

comparison provides vital information for system designers to determine if the implemented methods provide adequate performance for the system. To get a deeper understanding of the data, the dataset is further split into a discrimination dataset in order to highlight method errors that a simpler detection versus no detection dataset would not.

B. SAR imagery for ship discrimination

Ship discrimination is a broad task which tries to use some additional processing after ship prescreening and de-

tection to eliminate false alarms and highlight only highly probable ships [1], [3], [26]. This step can range from an additional CFAR prescreening stage [27] to machine-learning based elimination techniques [26], [28]. The focus of most of these methods have been to classify detections as either ships or non ships (ocean) [1], [2], [23], [24], [26], [28]. While sufficient to identify targets, this broad classification makes it difficult to identify exactly how methods are failing and so our discrimination dataset has been created using three classes rather than two. Specifically, the subimages containing ships discussed in section IV were combined with subimages containing ship-like areas and ocean areas to create a three class dataset of positives, false positives and negatives respectively (see Fig. 6). By evaluating methods with three classes we can highlight obvious errors such as when a method detects an ocean area as a ship and we can highlight more subtle errors such as ship-like areas being classified as ships. This more granular approach to ship discrimination evaluation allows for the examination of which parameters affect which type of error and also allows for the design of methods that can avoid one type of error more than the other based on user requirements (ideally both would be avoided). In addition to this, the discrimination dataset has been created by mimicking other machine learning datasets such as MNIST [18] and CIFAR-10 and CIFAR-100 datasets [19]. All subimages are 21×21 with the ships and ship-likes centred and hence are very similar in composition to MNIST's digits of width 24×24 . The rationale behind making the dataset similar to those is that methods known to work well on these other machine learning datasets such as deep learning can be applied to this dataset [29] and metrics known to work well with multi-class problems can be used to evaluate and identify errors in ship discrimination.

The 1596 positive examples were identified as described above while the false positive subimages were generated by selecting areas that did not contain ships but were incorrectly detected by a low-threshold cell-averaging constant false alarm rate (CA-CFAR) prescreening method [2], [26]. This low-threshold approach was done to generate as many samples as possible for future experiments. Of the approximate 500 000 false positive sub-images extracted across the 46 images, 3192 false positive subimages were selected at random. This is twice as many as the positive examples and represents the distribution of false positives to positives that would be encountered with careful prescreening threshold design and selection. Finally, 1596 subimages that did not contain either a ship or false positive were selected as negative ocean samples. It should be noted that technically the false-positives and negative samples can be selected randomly from all the possible samples present in the dataset but to better equalise the distribution of samples across the three classes they were selected with the above values. All of these subimages were combined with the patch and reference images, ship attributes and AIS matches into a single ship discrimination dataset containing the bulk of the information required for advanced discrimination method design and evaluation.

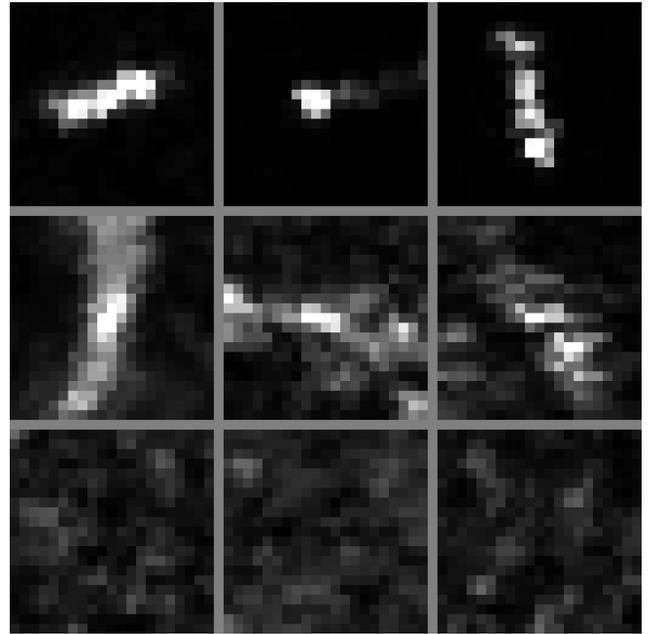


Fig. 6: Nine sub-images from the SAR dataset. The rows consist of positive (ships), false positive (ship-likes) and negative examples (ocean). The first two columns are Sentinel-1 GRDH and GRDM examples and the last column is from RADARSAT-2.

VI. SAR DATA ANALYSIS

This section of the paper details the analysis of the dataset. The analysis is split into two main sections. The first section details some statistics about the ships, their distribution across sensors, resolutions and polarisations. A statistical overview of the ship characteristics are then presented along with how these can be used to estimate common parameters for ship detection methods. The second section provides an analysis of the AIS-SAR matched ships and details how well the SAR reference ships compare to their real-world ships.

A. Overall ship distribution analysis

Table II provides an overview of the dataset in terms of SAR ships and their distribution across sensors, resolutions and polarisations. The dataset presented in this paper was not created to cover every single possible combination of SAR sensor, resolution and polarisation but to provide a more general data set that covers more than just a single instance or type of image. These images were chosen from available data at the time and hence the dataset contains no Sentinel-1 GRDM HH and HV imagery and no RADARSAT-2 HV,VV or VH imagery.

A confirmation of previous work [30] is the discrepancy between the number of ships in cross-polarised images and co-polarised images. Across the Sentinel-1 dataset, the cross-polarised imagery tends to have more detectable ships than those in the co-polarised imagery. The GRDM VH polarisation

TABLE II: Reference ship distribution per image sensor and polarisation.

Sensor	Sentinel-1 (GRDM)				Total	Sentinel-1 (GRDH)				Total	RADARSAT-2 (SCNA)				Total	Combined Total
	HH	HV	VV	VH		HH	HV	VV	VH		HH	HV	VV	VH		
Total SAR images	0	0	15	16	31	4	4	2	2	12	3	0	0	0	3	46
Total Ships	0	0	466	551	1017	131	149	93	106	479	100	0	0	0	100	1596
Total AIS Matches	0	0	185	213	398	44	46	3	3	96	0	0	0	0	0	494

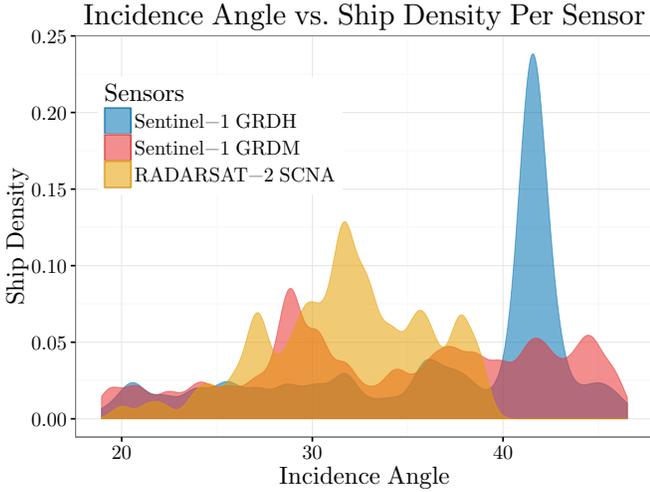


Fig. 7: Graph showing incidence angle per ship count density for each sensor. The sensors cover approximately the same incidence angle range for all 3 types but the peaks for each type seem to coincide at different incidence angles covering the entire spectrum available.

had 85 more ships and 28 more AIS matched ships than VV polarised images. Similarly, GRDH HV imagery produced 18 more ships and 2 more AIS matched ships whilst VH had 13 more ships in it than VV but had the same number of AIS matches.

Looking more broadly at the distribution of ships, the average number of ships per image type was as follows: GRDM had 32.8 ships per image; GRDH had 39.91 ships per image; and SCNA had 33.33 ships per image. A number of factors lead into why GRDH had the highest number of ships. GRDH images had track numbers 16, 43 and 45 which cover two of South Africa’s largest ports (East London and Durban) at almost precisely the same incidence angle ($41.0^\circ - 42.0^\circ$). This causes the large spike in ship distribution per incidence angle as shown in Fig. 7. Furthermore, GRDH’s resolution is almost twice that of GRDM allowing it to detect smaller ships that might not be seen on GRDM. And finally, the number of looks performed is doubled in the case of GRDM images which could contribute to smaller ships being removed and thus not appearing on the GRDM images.

B. Ship attribute analysis per sensor/resolution

The statistics of the four main ship attributes can be seen in Table III, IV and V. These tables show the calculated values of the various ship attributes across the various images. As noted previously, no GRDM HH, HV, SCNA HV, VV or VH images were part of the dataset, and these table values are shown as zero and not used in any calculations.

1) *Incidence Angle*: The incidence angle across all three resolutions remained within their given ranges, with comparable standard deviations per sensor. Ships in GRDM and SCNA were approximately equally distributed over the possible incidence angles. The GRDH median incidence angle was significantly greater than mean incidence angles which indicates a negatively skewed distribution of incidence angles. This result confirms the spike in incidence angle density shown in Fig. 7 and implies fewer low incidence angle ships compared to the high end of the scale. The implications of this may be two-fold: this skewed distribution may have unintended effects on size estimations (section VI-C); and knowledge of this bias for higher incidence angles for these orbits can be used as additional parameters for adaptive threshold methods such as the method introduced in [26]. The method allows for arbitrarily adapting thresholds (or alternatively false alarm rates) across the image which, in this case, could be varied according to the known incidence angle bias for these orbits to possibly improve CA-CFAR method results.

2) *Orientation*: The orientation or headings of the ships all varied between $0.0^\circ - 178.8^\circ$. At the time of writing, there was no reliable way of detecting which part of SAR ships were the front and back and hence a 180° uncertainty is expected. The mean and median orientations for the GRDM had similar values indicating an approximately balanced set of orientations with an average value of 102.2 with a standard deviation of 31.17. This was the smallest deviation across all three image resolutions and is expected due to the lower resolution which decreases the discrimination in all facets (size and rotation deviation). The GRDH and SCNA imagery had opposing distributions - GRDH orientations were skewed negatively when comparing mean and median whereas SCNA was skewed positively. This means GRDH had more southern facing ships than northern facing ships whereas the opposite was true for SCNA. *All of the SCNA images were taken over the Cape Town harbour. Distribution of AIS activities around the harbour indicate ships often headed north for the areas where the imagery was taken. Similarly, at the East London and Durban harbours (where most of the GRDH ships occur) ships most often head south towards the southern tip of South Africa and this may be the cause of the southerly orientated*

TABLE III: Sentinel-1 GRDM reference ship attribute statistics (mean values with standard deviation in brackets). The right hand side of the table indicates the mean, median, minimum and maximum for each ship attribute across all polarisations.

Sensor Polarisation	Sentinel-1 (GRDM)				Mean	Median	Minimum	Maximum
	HH	HV	VV	VH				
Incidence Angle (degrees)	0.000 (0.000)	0.000 (0.000)	34.97 (7.013)	33.59 (7.720)	34.22 (7.433)	34.55	18.95	46.22
Orientation (degrees)	0.000 (0.000)	0.000 (0.000)	100.4 (29.67)	103.8 (32.33)	102.2 (31.17)	101.1	0.000	178.6
Major Axis Length (pixels)	0.000 (0.000)	0.000 (0.000)	5.997 (1.799)	6.277 (1.831)	6.149 (1.821)	5.808	3.237	15.61
Minor Axis Length (pixels)	0.000 (0.000)	0.000 (0.000)	3.182 (0.443)	3.286 (0.452)	3.238 (0.451)	3.196	2.000	5.690

TABLE IV: Sentinel-1 GRDH reference ship attribute statistics (mean values with standard deviation in brackets). The right hand side of the table indicates the mean, median, minimum and maximum for each ship attribute across all polarisations.

Sensor Polarisation	Sentinel-1 (GRDH)				Mean	Median	Minimum	Maximum
	HH	HV	VV	VH				
Incidence Angle (degrees)	36.75 (7.086)	35.02 (8.088)	38.60 (5.300)	37.00 (7.149)	36.63 (7.221)	41.07	18.96	46.52
Orientation (degrees)	82.76 (42.75)	82.97 (46.37)	103.9 (38.80)	100.9 (38.69)	90.93 (43.29)	96.02	0.000	178.8
Major Axis Length (pixels)	8.867 (3.722)	8.375 (3.776)	7.759 (3.046)	9.326 (3.332)	8.600 (3.563)	8.021	3.266	16.90
Minor Axis Length (pixels)	3.560 (0.506)	3.441 (0.520)	3.368 (0.415)	3.692 (0.644)	3.515 (0.539)	3.463	2.211	6.101

TABLE V: RADARSAT-2 SCNA reference ship attribute statistics (mean values with standard deviation in brackets). The right hand side of the table indicates the mean, median, minimum and maximum for each ship attribute across all polarisations.

Sensor Polarisation	RADARSAT-2 (SCNA)				Mean	Median	Minimum	Maximum
	HH	HV	VV	VH				
Incidence Angle (degrees)	31.78 (4.211)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	31.78 (4.211)	31.81	20.08	38.99
Orientation (degrees)	71.50 (48.68)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	71.50 (48.68)	60.59	0.124	177.0
Major Axis Length (pixels)	8.072 (3.529)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	8.072 (3.529)	7.175	3.333	16.09
Minor Axis Length (pixels)	3.395 (0.463)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	3.395 (0.463)	3.341	2.211	4.575

ships.

3) *Major Axis Length*: Major axis length indicates the longest section of the reference ships and thus the maximum possible single dimension of a ship. This, in turn, gives us an indication of the pixel sizes of ships to appear even in as-of-yet unseen images. Analysing Tables III, IV and V we note that GRDM imagery had the smallest mean and standard deviation of major axis length across all images with 6.149 and 1.821, followed by SCNA with 8.072 and 3.529 and then GRDH with 8.600 and 3.563. The standard deviation of the GRDM ships is almost half that of the other two methods and indicates that while GRDM produces smaller, less detailed ships their sizes are more consistent and thus easier to predict adjust parameters for. The largest ships are found in GRDH VH images with a mean major axis length of 9.326 pixels. This implies that not only are the more ships visible in VH imagery, their average sizes are also larger. Both GRDM and GRDH images had overall mean major axis lengths that matched well with their median lengths. The SCNA imagery's overall mean major axis length 8.072 was greater than the median 7.175 implying positively skewed lengths with most of the major axis lengths below the mean and therefore more smaller ships than larger ones.

This attribute can help determine what the maximum and expected ship sizes are across all three platforms. This in

turn can help guide the selection of necessary parameters. The maximum ship size of 16.9 (which is an overestimation) indicated that the maximum expected ship size should always be smaller than 17 pixels. Accordingly, subimage sizes were chosen as 21×21 so that they could cover the largest possible ships across all images and would be centred with a minimum of 2 pixels on all sides as buffers. Similarly, the CFAR method requires selecting a guard and background window size [1]–[3]. If we assume that the guard window should always cover the ship, even at either end of the ship, then twice the size of the maximum median ($8.021 \times 2 = 16.041$) should provide adequate coverage. Hence, a guard window of at least 16×16 should be chosen (and a background window greater than that).

4) *Minor Axis Length*: Looking towards the minor axis lengths, all three of the overall mean lengths were approximately the same, with the GRDM having the lowest mean minor axis lengths across all its imagery with 3.238(0.451) pixels. As expected the minimum minor axis length was 2.0 pixels (also GRDM), as expected by how ships were defined in section IV. Interestingly, the mean minor axis lengths were all approximately the same across the three resolutions around 3.2 – 3.5 pixels in length. This indicates an overall agreement for the width across these images. However, as discussed in the next section, despite these similarities the ships' widths are often overestimated.

TABLE VI: Comparison between SAR reference ship sizes and those of AIS matched targets.

Sensor Polarisation	SAR				AIS			
	Mean	Median	Minimum	Maximum	Mean	Median	Minimum	Maximum
Incidence Angle (degrees)	34.10 (7.510)	34.36	19.01	46.27	0.000 (0.000)	0.000	0.000	0.000
Orientation (degrees)	92.93 (36.01)	100.6	0.458	168.9	79.47 (35.11)	74.00	1.000	178.0
Major Axis Length (m)	240.2 (84.43)	229.9	86.60	487.9	231.9 (64.12)	229.0	60.00	334.1
Minor Axis Length (m)	113.2 (28.03)	113.3	60.03	222.6	36.88 (10.91)	32.30	13.00	60.05

C. Referenced ship vs AIS analysis

The final analysis is the comparison of how well the reference ship sizes compare to their matched AIS ships. As noted before, of the 1596 ships only 494 of them were hard AIS matches. This collection of 494 ships (across GRDM and GRDH images only) were grouped together to form the ‘‘SAR’’ ship group and their associated AIS reported details. The statistics for these measurements can be seen in Table. VI and are measured in meters. The conversion of pixels to meters was handled by multiplying the ships size in pixels by the pixel spacing in that ship’s imagery (GRDM was 40 m and GRDH was 25 m).

We begin our analysis of the matched ships by looking at the orientations of table VI which show that the orientations which do not align, on average, very well up. The difference of approximately 14.4° might indicate a few things. The first is the quality of the MVEE for each ship. The method can provide orientations that differ by $\pm 10^\circ$ with the addition or subtraction of certain pixels for smaller ships. This can in turn cause mismatches in orientation. In general the differences in orientation are important for estimating dead reckoning of ships and hence its an important to note the difference in mean orientation between SAR ships and their matched AIS counterparts.

Moving on to the major axis length we get what appears to be a much better representation here of the ship length. The means of the two groups differ by only 9 m which is well within the pixel spacing size (25 to 40 m) used to calculate the size of the ships. More to the point, the medians match up within 1 meter and it is a good indication that, on average, the ships size can be estimated quite well. However, this does not paint the entire picture correctly. Firstly, the minimum and maximum ship sizes are off quite significantly compared to the AIS results. In a similar vein, the minor axis lengths are significantly overestimated. The maximum ship width is 60 m from the AIS data, this would imply either 2.4 pixels in a GRDH image or 1.5 pixels in a GRDM image. Despite this, these sizes match up relatively well with the minimum size defined for ships.

To better understand the differences for the sizes we turn to Fig. 8. This figure shows all matched ships which had lengths (a) and widths (b) that were 2 pixels (50 m or 80 m) greater or less than their matched AIS. The figure is arranged so that the first 97 ships are GRDH with the rest being GRDM. In Fig. 8a we notice their are more differences in length in GRDH than in GRDM. Surprisingly, the vast majority of mismatches in GRDH is underestimates whereas GRDM

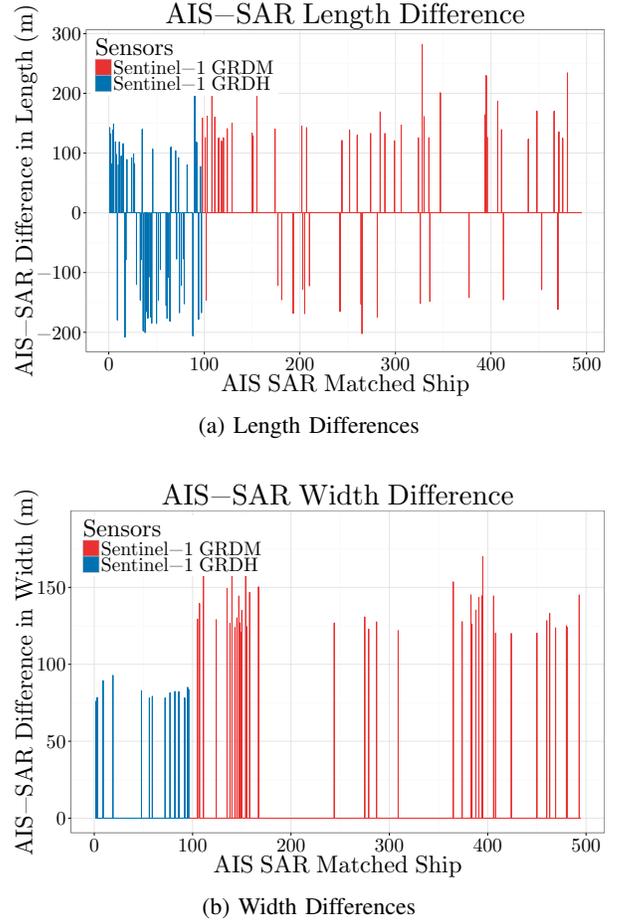


Fig. 8: Two subfigures showing the differences between sizes estimated using SAR reference ships and AIS reported sizes where (a) shows the differences in lengths and (b) the differences in widths. All values above the line indicate ships that were overestimated (too large) and all below indicate ships that were underestimated (too small).

mostly has overestimates. This is contrary to what is expected due to the higher resolution present in GRDH imagery. The underestimates provided by the GRDH are of note because they all seem to be similar in magnitude (around 200 m too small). Furthermore, these underestimates are approximately

equal to the underestimates of the GRDM. One likely reason for the concentration of underestimated sizes in this dataset is the incidence angle with which most of these occur which is, again, $41.0^\circ - 42.0^\circ$ as shown in Fig. 7. This indicates a possible significant trade-off. While more ships seem to be detectable given the cross-polarisation and higher incidence angle, their estimated sizes seem to be underrepresented. By taking into account that the incidence angle could have an effect on both the visibility of the ships as well as their output representation we can design ship detectors and discriminators accordingly. Looking towards Fig. 8b we notice more predictable behaviour. Both GRDM and GRDH overestimate ship widths but GRDM overestimates more often and with greater magnitude than that of GRDH. This is a direct consequence of selecting ships that are at least 2 pixels in width/length.

The matching of ship sizes from SAR medium resolution imagery to AIS is a complicated topic. The above results indicate that orientation estimates often differ between reference SAR ships and their AIS matches. Similarly, for large ships and small ships the SAR reference ships either overestimate or underestimate the ships more often than not. Table VI showed that for average sized ships (ships around 220 m) we could get a reliable estimate of its length if it is not a ship within GRDH imagery around $41.0^\circ - 42.0^\circ$. A number of other factors could lead to the mismatches seen above. The subswaths in Sentinel-1 imagery have varying incidence angles and variations in resolutions that cannot easily be summarised with one pixel value for every ship across the image. The shapes of the ship could be distorted by the preprocessing applied to transform the image into the ground-range. The referencing procedure which highlights specific pixels might require different representations of the image to highlight the correct pixels (such as log-scale adjustments or compensating for the linear variation of reflectivity in slant-range [25]). All of these may cause small errors on each but the culmination of them leads towards a conclusion that without further study, the estimation of ship parameters from SAR medium resolution imagery will be imprecise.

VII. CONCLUSIONS AND FUTURE WORK

The creation of a scientific dataset is a process of iterative adjustment. Datasets undergo constant change as they are analysed and a dataset that has the same composition as when it was created is often a poorly understood dataset. This paper introduced a complete SAR dataset for ship detection, discrimination and analysis across Sentinel-1 GRDM and GRDH imagery as well as RADATSAT-2 SNCA imagery. The dataset proposed here is presented after a number of iterations, whereby some initial assumptions were replaced in favour of others after the analysis of the dataset. Various parameters were also chosen so that the dataset can be reprocessed with new methods in the future. This paper also presents a thorough description of the creation of the dataset, how and why certain dataset characteristics were chosen as well as an analysis of the dataset given these assumptions.

One of the driving factors for the creation of this dataset was the lack of coherence among medium resolution SAR ship

detection studies. Ship detection studies sometimes combine components which can cause difficulties when comparing different methods. By splitting the detection process up into steps, and designing components of the dataset for each step, a better understanding of how and why methods fail can be garnered. To this end, the dataset is split into detection-only components (ship positions in the SAR imagery), discrimination components (three-class MNSIT-like subimages), shape analysis (reference ships and associated attributes) and matched components (SAR-AIS matched ships).

Conventional ship detection studies often process SAR imagery to detect ships and classify pixels as either ships or sea-water. This study introduced a more granular approach to discrimination by providing a three class discrimination dataset component namely: positive (ships), false-positive (ship-like areas) and negatives (ocean backscatter). By being more descriptive of the possible discrimination classes, methods that fail outright can be separated from those that only have trouble with false alarms. Furthermore by designing the discrimination component to be similar to wide-spread machine learning datasets such as MNIST and CIFAR-10 and CIFAR-100, similar methods can be used to perform discrimination (classification) on the dataset. This also introduces the possibility of using metrics known to work well for method evaluation in machine learning to be applied to ship discrimination methods.

The paper also presented a description of the creation of reference ships and their associated characteristics. In addition to this, reference ships were matched to their closest AIS counterparts. A thorough analysis of the reference ships, their distribution across the SAR imagery, their attributes and AIS matched attributes was undertaken. This analysis revealed a number of operational parameters such as which polarisation is the best for ship detection, maximum and expected ship sizes and how the ship attributes are affected by resolution and incidence angles.

Future work for the dataset includes an additional round of referencing and analysis with the new referenced SAR ships under varying circumstances including image log scaling and compensating for the linear variation of reflectivity. Furthermore, many matched SAR-AIS ships have associated classes (cargo, tanker, fishing, etc.) which was not explored in this paper. Evaluation of the possibility of ship type classification within the dataset using machine learning is another possible area of focus. In closing, by providing a complete dataset including a description of the creation process the hope is that a better understanding for ship detection in SAR imagery is gained and applied to images of the same types in the future.

REFERENCES

- [1] D. J. Crisp, "The State-of-the-Art in Ship Detection in Synthetic Aperture Radar Imagery," Australian Department of Defence, Edinburgh, Australia, Tech. Rep. DSTO-RR-0272, 2004.
- [2] K. El-Darymli, P. McGuire, D. Power, and C. Moloney, "Target detection in synthetic aperture radar imagery: a state-of-the-art survey," *Journal of Applied Remote Sensing*, vol. 7, no. 1, pp. 071 598–071 598, Mar. 2013.

- [3] C. P. Schwegmann, W. Kleynhans, and B. P. Salmon, "Ship Detection in South African oceans using SAR, CFAR and a Haar-like Feature Classifier," in *Geoscience and Remote Sensing Symposium, 2014. IGARSS 2014. IEEE International*, 2014, pp. 557–560.
- [4] J. M. de Nicolas, P. Jarabo-Amores, N. del Rey-Maestre, P. G. del Hoyo, and J. L. Barcena-Humanes, "Robustness of a Generalized Gamma CFAR ship detector applied to TerraSAR-X and Sentinel-1 images," in *EUROCON 2015 - International Conference on Computer as a Tool (EUROCON)*, IEEE, Sept 2015, pp. 1–6.
- [5] P. Iervolino, R. Guida, and P. Whittaker, "A novel ship-detection technique for Sentinel-1 SAR data," in *Synthetic Aperture Radar (AP SAR), 2015 IEEE 5th Asia-Pacific Conference on*, Sept 2015, pp. 797–801.
- [6] D. Velotto, C. Bentes, B. Tings, and S. Lehner, "Comparison of sentinel-1 and terrasar-x for ship detection," in *Geoscience and Remote Sensing Symposium (IGARSS), 2015 IEEE International*, July 2015, pp. 3282–3285.
- [7] R. Pelich, N. Longepe, G. Mercier, G. Hajduch, and R. Garello, "Performance evaluation of Sentinel-1 data in SAR ship detection," in *Geoscience and Remote Sensing Symposium (IGARSS), 2015 IEEE International*, July 2015, pp. 2103–2106.
- [8] P. Iervolino, R. Guida, and P. Whittaker, "A novel ship-detection technique for Sentinel-1 sar data," in *Synthetic Aperture Radar (AP SAR), 2015 IEEE 5th Asia-Pacific Conference on*, 2015, pp. 797–801.
- [9] C. P. Schwegmann, W. Kleynhans, B. P. Salmon, and L. Mdakane, "A CA-CFAR and localized wavelet ship detector for Sentinel-1 imagery," in *Geoscience and Remote Sensing Symposium, 2015. IGARSS 2015. IEEE International*, 2015, pp. 3707–3710.
- [10] —, "Ship detection in Sentinel-1 imagery using the H-dome transformation," in *Geoscience and Remote Sensing Symposium, 2015. IGARSS 2015. IEEE International*, 2015, pp. 3710–3714.
- [11] D. Small, "Flattening Gamma: Radiometric Terrain Correction for SAR Imagery," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 8, pp. 3081–3093, Aug. 2011.
- [12] E. S. A. (ESA), "Sentinel-1 Radiometric Calibration of Products," PDF, European Space Agency (ESA), Tech. Rep., 2015, ESA-EOPG-CSCOP-TN-0002. [Online]. Available: <https://earth.esa.int/documents/247904/685163/S1-Radiometric-Calibration-V1.0.pdf>
- [13] K. El-Darymli, P. McGuire, E. Gill, D. Power, and C. Moloney, "Understanding the significance of radiometric calibration for synthetic aperture radar imagery," in *Electrical and Computer Engineering (CCECE), 2014 IEEE 27th Canadian Conference on*, 2014, pp. 1–6.
- [14] H. Greidanus and C. Santamaria, "First Analyses of Sentinel-1 Images for Maritime Surveillance," PDF, European Commission Joint Research Centre (JRC), Tech. Rep., 2014, 978-92-79-44715-0. [Online]. Available: <http://publications.jrc.ec.europa.eu/repository/handle/JRC92666>
- [15] M. T. Alonso, C. Lopez-Martinez, J. J. Mallorqui, and P. Salembier, "Edge Enhancement Algorithm Based on the Wavelet Transform for Automatic Edge Detection in SAR Images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 1, pp. 225–235, Jan. 2011.
- [16] F. Baselice and G. Ferraioli, "Unsupervised Coastal Line Extraction from SAR Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 10, no. 6, pp. 1350–1354, Nov. 2013.
- [17] Z. Liu, F. Li, N. Li, R. Wang, and H. Zhang, "A Novel Region-Merging Approach for Coastline Extraction from Sentinel-1A IW Mode SAR Imagery," *IEEE Geoscience and Remote Sensing Letters*, vol. 13, no. 3, pp. 324–328, Mar. 2016.
- [18] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [19] A. Krizhevsky, "Learning multiple layers of features from tiny images," University of Toronto, Toronto, Canada, Tech. Rep. Tech. Rep. 001, 2006.
- [20] M. H. Protter and C. B. Morrey, *College calculus with analytic geometry*, 3rd ed. Boston, MA: Addison-Wesley, 1970.
- [21] P. Kumar and A. Yildirim, "Minimum-Volume Enclosing Ellipsoids and Core Sets," *Journal of Optimization Theory and Applications*, vol. 126, no. 1, pp. 1–21, Jul. 2005.
- [22] T. Eriksen, G. Høyve, B. Narheim, and B. J. Meland, "Maritime traffic monitoring using a space-based AIS receiver," *Acta Astronautica*, vol. 58, no. 10, pp. 537–549, May 2006.
- [23] R. Pelich, N. Longepe, G. Mercier, G. Hajduch, and R. Garello, "Performance Evaluation of Sentinel-1 Data in SAR Ship Detection," in *Geoscience and Remote Sensing Symposium, 2015. IGARSS 2015. IEEE International*, 2015, pp. 2103–2106.
- [24] D. Velotto, M. Soccorsi, and S. Lehner, "Azimuth Ambiguities Removal for Ship Detection Using Full Polarimetric X-Band SAR Data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 52, no. 1, pp. 76–88, Jan. 2014.
- [25] G. Margarit, J. A. B. Milans, and A. Tabasco, "Operational Ship Monitoring System Based on Synthetic Aperture Radar Processing," *Remote Sensing*, vol. 1, no. 3, pp. 375–392, Sep. 2009.
- [26] C. Schwegmann, W. Kleynhans, and B. 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. 1939–1404, Jul. 2015.
- [27] F. Bi, F. Pang, B. Zhu, and L. Chen, "A cascaded false-alarm elimination method for accurate ship detection in SAR images," in *Radar Conference 2013, IET International*, 2013, pp. 1–4.
- [28] J. R. A. Morillas, I. C. Garca, and U. Zolzer, "Ship detection based on SVM using color and texture features," in *Intelligent Computer Communication and Processing, 2015. ICCP 2015. IEEE International Conference on*, Sep. 2015, pp. 343–350.
- [29] C. P. Schwegmann, W. Kleynhans, and B. P. Salmon, "Very Deep Learning for Ship Discrimination in Synthetic Aperture Radar Imagery," in *Geoscience and Remote Sensing Symposium, 2016. IGARSS 2016. IEEE International*, 2016, pp. 1–4, pre-published.
- [30] W. Kleynhans, B. P. Salmon, C. P. Schwegmann, and L. Mdakane, "Use of Sentinel-1 data for vessel detection in South African oceans: Early results," in *IEEE Radar Conference, 2015. Johannesburg., 2016*, pp. 431–434.