

The SARFish Dataset and Challenge

Connor Luckett, Benjamin McCarthy, Tri-Tan Cao, Antonio Robles-Kelly
Defence Science and Technology Group, Edinburgh, SA 5111, Australia
{*connor.luckett, benjamin.mccarthy5, tan.cao, antonio.robleskelly*}@defence.gov.au

Abstract

In this paper, we present the SARFish challenge and dataset. The challenge focuses on the use of Synthetic Aperture Radar (SAR) imagery for the identification of vessels involved in illegal, unreported and unregulated (IUU) fishing which damages ecological systems and causes losses for fishing industries and governments worldwide. The SARFish dataset is a free and open large-scale complex-valued SAR dataset which is based upon Sentinel-1 imagery and built upon the xView3 labels. We expect this dataset to help in advancing the state of the art in automated ship detection from SAR imagery, contextual representation learning, and the application of deep complex-valued neural networks. We also hope the availability of the SARFish dataset will stimulate developments on other topics of interest that can naturally tackle complex-valued data such as quantum-inspired approaches.

1. Introduction

1.1. Illegal, Unreported, and Unregulated Fishing

IUU fishing is a global problem with severe ecological, economic, and political impacts. It primarily affects developing nations, many of which rely on regional fisheries as a major source of food and income [33, 41]. At a local scale, it depletes the marine environment and directly impacts the sustainability of the ecosystem [41], at a global scale, it is a major contributor to piracy [9, 34] and causes USD 25 billion in economic loss per year [45].

Monitoring, Control and Surveillance (MCS) of fishing practices has been undertaken globally to enforce regulation and combat IUU fishing [12]. Effective enforcement of fishing regulation has been shown to reduce the proportion of IUU fishing catch [1, 16].

Maritime surveillance of fishing activities is a challenging task. First, the scale of the problem is staggering; industrial fishing occurs in greater than 55% of the world's ocean area [25], and in 2009, the extent of IUU fishing was

estimated to constitute 20% of the world's catch or between 11 and 26 million metric tons [1]. Recent studies have also pointed to the difficulty in accurately assessing the extent of IUU activities due to the patchy geographic coverage of current IUU fishing estimates [16].

A second challenge of maritime surveillance is the presence of “dark” vessels. Automatic Identification Systems (AIS) and Vessel Monitoring Systems (VMS) [10, 21] are commonly used by countries to monitor fishing activities [12]. Dark vessels engaging in IUU fishing may disable AIS broadcasts in what is known as “disabling events” to avoid reporting obligations and obscure their illicit activities [4, 26, 36]. A 2022 study estimated that 6% of the total global fishing activity between 2017 and 2019 was obscured by disabling events [45]. AIS and VMS data sources are insufficient to evaluate the full extent of IUU fishing in these cases. As a result, alternate sources of data that are capable of capturing dark vessels must be utilised.

1.2. Synthetic Aperture Radar

One solution to the problem of tracking dark vessels is to detect ships automatically from satellite imagery. The use of Electro-Optical (EO) imagery for ship detection and maritime surveillance has a long history in the literature [24]. Despite this, it does have characteristics that make it unsuitable for use in dark vessel monitoring tasks. These include being limited to daytime operation and susceptibility to being obscured by clouds [22]. In contrast, SAR is an active imaging system that utilises physics and signal processing concepts to form imagery of the Earth's surface and offers an alternative to EO imagery providing day-night and all-weather coverage [43].

The maritime SAR surveillance problem can be decomposed into two stages: 1) detection of maritime objects in the imagery; and 2) subsequent classification of these detections. The output of the detection algorithm can be associated with fishing vessels tracked using a combination of AIS and VMS, with the remainder revealing the activities of dark fishing vessels [12, 26, 36, 47].

There has been an explosion in the study of SAR for surveillance in recent years, primarily driven by two de-

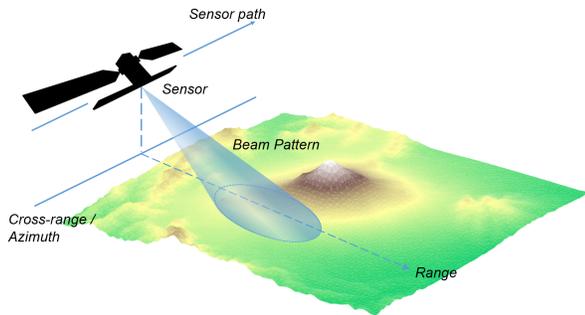


Figure 1. The collection geometry for SAR imaging. The sensor traverses the ground scene and transmits a beam of microwave pulses. The range direction contains information based on the pulse time-delay, and corresponds to the vector from the sensor to the ground as projected onto the ground plane. The azimuth or cross-range direction is orthogonal to the range direction. Information about the scene in cross-range is obtained by coherently processing across multiple pulses.

velopments. The first development was the launch in 2014 of the Sentinel-1 satellite constellation [6], which was the first constellation of the Copernicus Programme; the European Union’s remote sensing program delivered by the European Space Agency [23,26]. The Copernicus Programme provides open access to data collected by a suite of Earth sensing constellations through the Copernicus Open Access Hub¹.

The second driver for the development of SAR surveillance systems began in 2016 with the application of deep learning approaches to the ship detection and classification tasks [39,49]. These developments have addressed two key challenges with MCS. Firstly, the use of remote-sensing imagery provides an additional source of data independent of AIS that is unaffected by disabling events. Secondly, automated processing provides vessel detection and classification data allowing Regional Fisheries Management Organisations (RFMOs) to monitor huge fishery areas in a cost effective manner [10].

The acquisition of SAR imagery differs significantly from that of EO imagery and requires extensive signal processing before an image is formed [7]. As illustrated in figure 1, a SAR sensor passes over a scene illuminating a patch of the Earth’s surface with multiple pulses of transmitted microwave energy. The reflectivity of scatterers in a scene are mathematically modelled as complex-valued wave functions, with the magnitude of the complex quantity representing the energy of the reflected pulse, and the phase representing the change in phase between the transmitted and reflected pulse. The electromagnetic reflectivity of the scene can be reconstructed from the record of the transmitted and reflected pulses called “phase history” and interpreted as a

¹<https://scihub.copernicus.eu>

two-dimensional image [22].

Most modern SAR systems contain multiple antennas, which can be used to transmit radiation in different orientations or “polarisations”. This allows separate images to be produced depending on the various combinations of transmission and reception. The Sentinel-1 sensor can collect dual-polarisation data in either VH+VV or HH+HV configurations. For example VH+VV indicates collection of both VH, which is the result of sending a vertically polarised pulse and receiving a horizontally polarised pulse; and VV, which is the result of sending and receiving a vertically polarised pulse. It has been shown that multi-polar SAR data contains features crucial to improving the performance of ship detection methods [46]. In addition, the appearance of ships is more prominent in cross-polarisation (VH) images than in co-polarisation (VV) images [38,43].

SAR imagery products, such as Sentinel-1 products are provided in a number of different processing levels. Single Look Complex (SLC) products are complex-valued images whose pixels represent the amplitude and phase information of radar returns from a given surface area. Ground Range Detected (GRD) products are real-valued images whose pixel values represent the intensity of the reflected radar return. GRD products are distinguished from SLC products by a number of processing steps including “multi-looking”, which is a type of weighted filtering method that increases the contrast of objects from background clutter, and “detection” that projects the amplitude and phase values of its ancestor products into intensity values [37]. The Sentinel-1 processing pipeline generates both GRD and SLC products from a common ancestor, making it possible to provide coincident complex and real-valued SAR imagery. Figure 2 summarises the relationship between the SLC and GRD products in the Sentinel-1 processing pipeline.

Despite its positive traits, SAR imagery also poses challenges for ship detection. While SAR can be high resolu-

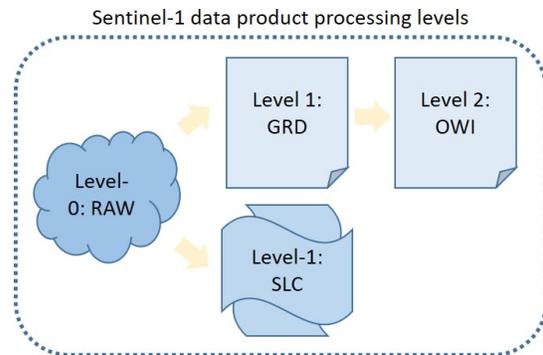


Figure 2. Brief summary of the Sentinel-1 processing pipeline illustrating that the GRD and SLC products share a common ancestor [37].

tion, the application to surveillance tasks requires a compromise between resolution and ground coverage which results in relatively low-resolution imagery and small ship extent [49]. Another challenge is SAR specific imaging artefacts such as noise, speckle, and side-lobing, as demonstrated in figures 5 and 6.

1.3. SARFish Challenge and Dataset

The SARFish dataset, first introduced in 2022 [5], is a SAR imagery dataset for the purpose of training and validating deep learning models on the tasks of detection, classification and length regression of maritime objects. SARFish builds upon the xView3-SAR dataset [35] by providing coincident GRD and SLC products along with labelled ship positions and Bounding Boxes (BBox), three hierarchical imbalanced ship class labels, and vessel length measurements. To the authors' knowledge there is no publicly available dataset which provides a sufficient quantity of dual-polarisation, full-sized, complex-valued SAR imagery with sufficient ground truth ship detection, classification and length regression labels suitable for the purpose of testing and evaluating deep learning based models. In this paper, we announce the public release of the data, along with the associated SARFish challenge which seeks to promote the use of deep learning with Complex-Valued SAR (CV-SAR) imagery for application in maritime surveillance in the context of combating IUU fishing.

The rest of this paper will be organised as follows. In Section 2 we note that there is valuable information in the complex SAR imagery that might be exploited to develop deep learning models that outperform methods using real-valued SAR imagery. We also assert that the running of a machine learning challenge, similar to the xView3-SAR challenge, may be an effective way of bootstrapping research in this field by incentivising competition and providing a set of benchmark models. Finally, we survey the literature and argue that the relative dearth of deep learning methods that utilise CV-SAR data for ship detection and classification is due partly to a lack of appropriate datasets.

Section 3 will briefly describe the SARFish dataset, give examples of the imagery and labelled maritime objects, and describe the attributes that make the SARFish dataset the only openly available dataset suitable for developing and benchmarking deep learning models on full-size CV-SAR imagery. Section 4 will describe the SARFish machine learning challenge and detail the metrics which will be used to measure model performance. Section 5 will summarise the key contributions of this paper.

2. Related Work

2.1. Deep learning for CV-SAR Target Detection

Complex-valued imagery is the natural result of the SAR image formation process. The physics model that underpins SAR utilises complex-valued functions as a mathematically convenient way to represent coherent radar pulses [22]. The relationship between the amplitude and phase components of pulses reflected from a given surface area captures valuable information about the electrical and physical properties of scatterers in a scene that are discarded when projected into the real domain [14, 17]. There are relatively few studies investigating ship detection and classification methods exploiting CV-SAR data. Despite this, deep learning methods for exploiting the information contained in CV-SAR data such as those in [50] and [3] have demonstrated improvements in performance on classification tasks using fully Complex-Valued Convolutional Neural Networks (CV-CNNs) over real-valued CNNs, indicating the potential for study in this field.

In contrast, the application of deep learning models to the task of ship detection and classification in real-valued SAR imagery has been studied extensively [31]. Significant improvements in performance has been achieved [49] since the first application of deep learning to the task in 2016 [39]. This improvement is partly due to the availability of ready-to-use datasets. SAR datasets abstract away the complicated pre-processing steps [43] to acquire analysis-ready pixels so that ML researchers can focus on model development. The lack of CV-SAR datasets suitable for both training deep learning models and for directly comparing methods exploiting GRD and SLC imagery has been identified as a barrier to research in this field [2, 15].

2.2. Success of Deep Learning on Real-Valued SAR Datasets

Li et al. [31] in 2022 surveyed 177 papers on ship detection in real-valued SAR and showed a relationship between the release of publicly available datasets and the number of publications. This demonstrates the impact of ready-to-use data on the pace of research. GRD datasets have evolved from ship-chip datasets [18, 29, 30], to larger image crops showing ships in context [44], and more recently towards abundant full-size imagery [35, 49] suitable for developing deep learning models intended for use on real world data. One reason for this shift towards full-size data is that it has been shown that the smaller the domain gap between the training and application data, the better the model performance [20].

The introduction of the xView3 challenge [35] in 2021 alongside the release of the xView3-SAR dataset had a significant impact on its field. Over two thousand entrants from around the globe competed to develop the best performing

model on the tasks of maritime object detection, classification and length regression, with 50 results being published on the xView3 leaderboard. The xView3-SAR challenge was an effective method for encouraging the study in the field and resulted in an extensive set of publicly released models and benchmark results which future users of the dataset can use to demonstrate an improvement over the state-of-the-art.

2.3. Current State of Publicly Available CV-SAR Datasets

The number of available CV-SAR datasets is substantially lower than that for real-valued SAR datasets. The following is a review of the publicly available CV-SAR datasets. These are summarised in table 1 with a brief explanation of their short-comings which limit their applicability for training deep learning models.

In 2014, Lang et al. [27] released a dataset comprised of 22 dual-polarisation SLC products from the RadarSat-2 and TerraSAR X sensors [6] in Ultrafine, Standard and Wide imaging modes. Labels were generated using an automated AIS method and consist of position and classification labels totalling 712 ground truth. Maritime objects were categorised into four imbalanced ship classes and a sea-clutter class. This is an early example of a CV-SAR dataset and provides high-resolution full-size imagery. However it consists of less than one thousand total labelled ships making it less suitable for training deep learning methods and does not provide coincident GRD products for comparison.

In 2017, OpenSARShip [18] and its updated version: OpenSARShip-2.0 [29] were released. The latter consisted of $34\,528\,256 \times 256$ pixel ship-chips taken from 87 Sentinel-1 images in Interferometric Wide (IW) Swath Mode in a range of locations and sea states. Ships were categorised into four imbalanced ship classes. OpenSARShip has been used by Huang et al., [19] in 2020 to develop a joint spatial and time frequency analysis network which utilised both intensity and phase information to gain an increase in performance over GRD-only methods in classification. It was also used by Zhang et al. [48] in 2020 to develop a squeeze-and-excitation network which utilised the complex-value representation to extract coherent ship features between dual-polarisation channels aiding classi-

fication performance. And Shao et al. [40] developed an information-guided network, by utilising the complex-value representation as one of the inputs to a polarisation channel cross-attention framework achieving state-of-the-art classification performance. The OpenSARShip datasets remain prominent in the field as they were the only public SAR datasets providing coincident GRD and SLC data, and the first CV-SAR datasets suitable for training deep learning classification models. The lack of full-scene imagery however, means it cannot be used to evaluate real-world detection tasks as it removes challenging aspects of detection such as close-to-shore scenarios.

In 2021, Lei et al. [28] released CSRSDD, a detection and classification dataset comprising $514\,1024 \times 1024$ pixel crops of single-polarisation Gaofen-3 [6] SLC imagery in one metre pixel resolution taken in Spotlight mode. The dataset labels 1 958 ships with rotated bounding box (RB-Box) positional labels and six imbalanced ship classes. CSRSDD was used in [51] to develop a fully CV-CNN model which utilises a single complex channel input and complex area max pooling for target detection and classification. This dataset is the highest resolution of the CV-SAR datasets, but is not ideal for open-ocean surveillance applications where the imagery is typically of lower resolution but wider ground coverage. The small number of ships makes it less suitable for training deep learning models. The data is also not dual-polarisation, and does not provide coincident GRD products.

We argue that dearth of research in this field is partly due to a lack of CV-SAR datasets suitable for training deep learning models. We believe that the introduction of the SARFish dataset addresses the shortcomings of previous datasets and is a valuable contribution to the community. We also argue that a machine learning challenge specifically designed around the SARFish dataset may foster the same contribution to the field of CV-SAR exploitation as the xView3 challenge did for real-valued SAR.

3. Dataset Description

The SARFish dataset was constructed specifically to address issues with previous public CV-SAR datasets. SARFish provides a number of key attributes that, taken as a whole, makes it uniquely suitable for developing deep learning ship detection and classification models in CV-SAR imagery. SARFish provides full-size dual-polarisation imagery with large-scale backgrounds, minimally pre-processed coincident GRD and SLC products, a large number of maritime object ground truth, and evaluation code based on the xView3 challenge metrics.

3.1. Data

The SARFish data consists of 753 pairs of full-size coincident Sentinel-1 GRD and SLC products sourced from the

Table 1
RELATED CV-SAR DATASET SUMMARY

Name	Full-scene	Dual-pol	SLC and GRD	Labels	Classes
Hierarchical	True	Partial	False	712	5
OpenSARShip 2.0	False	True	True	34 528	4
CSRSDD	False	False	False	1 958	6
SARFish	True	True	True	143 284 GRD 140 721 SLC	3

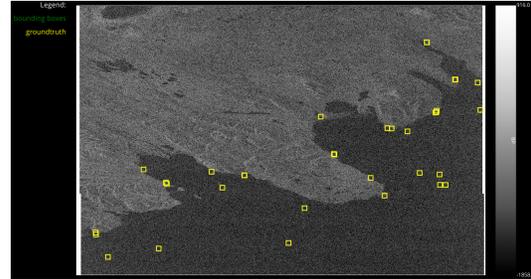
Table 2
SUMMARY OF THE SENTINEL-1 DATA [6, 13].

Platform	Sentinel-1 (A, B)	
Operator	European Space Agency Sentinel-1 Mission Performance Center	
Sensor	C-Band SAR	
Mode	Interferometric Wide Swath (IW)	
Polarisations	VV, VH	
Ground range coverage (km)	~ 250	
Product type	SLC	GRD
Image size (pixels)	~ 23000 × 12000	~ 27000 × 16000
Data type	Complex Int16	Unsigned Int16
Azimuth pixel spacing (m)	14.1	10
Range pixel spacing (m)	2.3	10

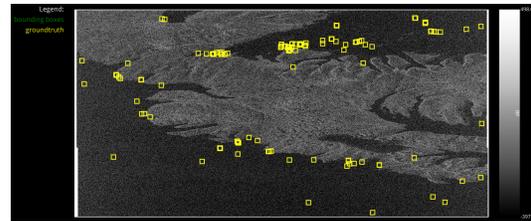
Alaska Satellite Facility (ASF) Distributed Active Archive Center (DAAC) [32]. A summary of the Sentinel-1 data is shown in Table 2. While the xView3-SAR dataset consisted of 754 GRD products, we could only determine the corresponding SLC product for 753. Figures 3a, 3b, and 3c show the three “swaths” of a SARFish SLC product with maritime object labels. The SARFish SLC data is divided into three swaths of approximately 23000 × 12000 pixels in VV+VH polarisation channels. Figure 3d shows the corresponding GRD product. GRD products consist of one image of approximately 27000 × 16000 pixels in VV+VH polarisation channels. The SAR imagery displayed throughout the paper have been post-processed for the purpose of visualisation. For SLC imagery, the complex amplitude and phase data have been mapped into real intensity values. For both GRD and SLC images, a decibel scaling is applied to the imagery.

The xView3-SAR dataset provided GRD products with a significant amount of pre-processing including georectification, filtering and decibel scaling, which may have removed useful information. In contrast, the SARFish dataset provides Sentinel-1 SLC and GRD products with minimal pre-processing. Figure 4 illustrates the relationship between the xView3-SAR and SARFish datasets. Described in detail in [5], the SLC products were deburstured [43] to generate contiguous images. Both SLC and GRD products had no-data values set to distinguish what parts of the image contains valid radar data, and were then flipped to counter the mirroring of the Earth’s surface seen in the

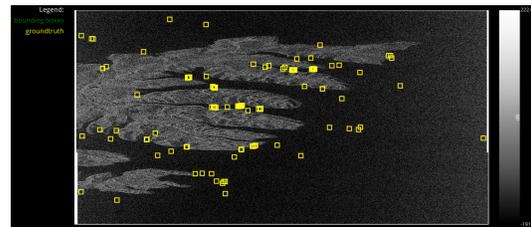
source imagery.



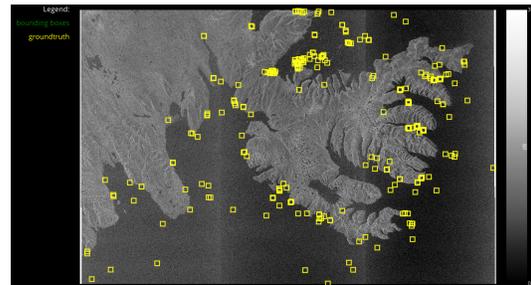
(a) First swath comprising a SARFish SLC product in VH polarisation.



(b) Second swath comprising a SARFish SLC product in VH polarisation.



(c) Third swath comprising a SARFish SLC product in VH polarisation.



(d) Full sized SARFish GRD product in VH polarisation

Figure 3. Example coincident SARFish SLC and GRD product in VH polarisation showing the Westfjords of Iceland with ground truth maritime objects. Both GRD and SLC images have been displayed on decibel scale where the dynamic range is displayed between 15 and 60 dB for visualisation. Note that the SLC products are comprised of 3 swaths shown in figures 3a, 3b, 3c with 2.3×14.1 m pixel spacing. The corresponding GRD product in figure 3d has 10×10 m square pixel spacing.

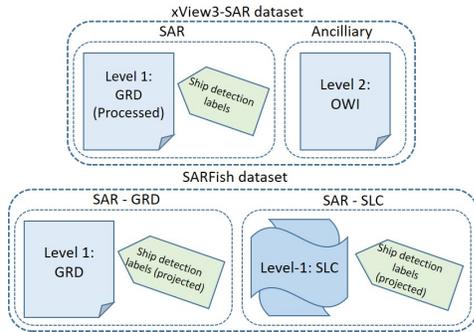


Figure 4. Relationship between the xView3-SAR dataset [35] and the SARFish dataset.

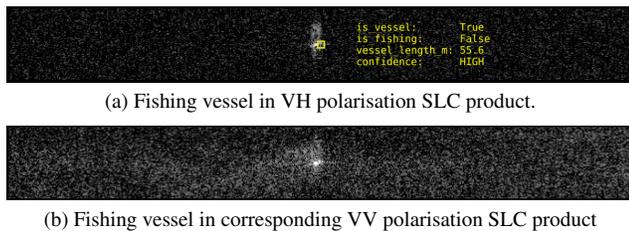


Figure 5. Example labelled fishing vessel in detected and decibel scaled SARFish SLC product in both VH (5a) and VV (5b) polarisations. Note the differences in the speckle and side-lobing artefacts on the vessel between polarisations, the non-square pixel spacing (see table 2), and that the spatial extent of these figures and Figure 6 are the same.

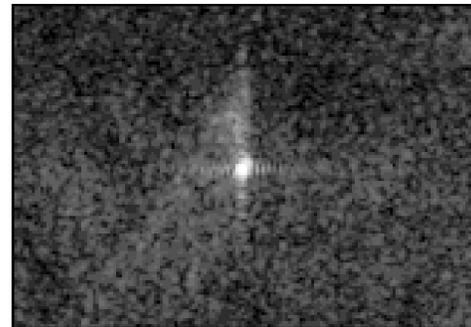
3.2. Labels

The SARFish labels consist of three main types. Firstly, positional BBox labels denote geographic and image positions of maritime objects. Secondly, hierarchical imbalanced classes denote maritime objects as to whether or not they are vessels and vessels as to whether or not they are fishing vessels. Finally, vessel length information is provided. Table 3 shows the total number of unique maritime objects labelled in the dataset. Other ground truth attributes include distance from shore values for evaluating close-to-shore detection tasks (see section 4.2) and a labelling confidence evaluation. Figures 5 and 6 show a close-up of a labelled fishing vessel in a coincident SLC and GRD product.

The labels were generated by projecting the xView3 labels into the pixel space of the SARFish products using range doppler geocoding [8]. Due to the overlap between the swaths of the SLC products, some ground truth appear in multiple products. We also found that in some cases GRD products include ground areas not found in the corresponding SLC products leading to a $\sim 1\%$ discrepancy between the total number of unique labelled maritime objects (see in table 3). Like in the xView3-SAR dataset, ground truth are uniquely identified by a detection ID, allowing for a 1:1



(a) Fishing vessel in VH polarisation GRD product.



(b) Fishing vessel in corresponding VV polarisation GRD product

Figure 6. Example labelled fishing vessel in decibel scaled SARFish GRD product in both VH 6a and VV 6b polarisations. Note the square pixel spacing (see table 2). The spatial extent of this figure and 5 are the same.

Table 3
PARTITIONS OF THE SARFISH CHALLENGE DATASET

Partition	Coincident products	Labels provided	Unique maritime object labels	
			SLC	GRD
train	553	True	63 071	64 054
validation	50	True	18 906	19 222
public	150	False	58 744	60 008
Total			140 721	143 284

comparison between models trained on both imagery types.

The SARFish dataset has been partitioned into training, validation and public partitions as shown in Table 3. Labels are provided only for the train and validation partitions during the SARFish challenge. In total, the three partitions contain 140 721 and 143 284 unique maritime object labels for the SLC and GRD product types respectively.

4. Challenge Description

The SARFish challenge is straightforward. Researchers will compete to develop the best ship detection, classifi-

cation and length regression models on SLC data in three challenge tracks: maritime object detection, maritime object classification, and vessel length regression.

The challenge will be hosted on the ‘‘SARFish Computer Vision Challenge’’ competition on Kaggle ². Participants will train and evaluate models on the train and validation partitions of the dataset defined in table 3. They will then run inference on the public partition of the dataset, and submit their prediction CSV file to the Kaggle challenge for ranking. Participants will be able to submit multiple predictions to the Kaggle competition page until the challenge expires.

4.1. Track 1: Maritime object detection

The maritime object detection track is comprised of the following tasks, which are directly comparable with those from the xView3 challenge.

4.1.1 Maritime Object Detection Task

The objective of this task is to find maritime objects in the SARFish imagery. The detection task is evaluated by assigning model predictions to the ground truth maritime object locations contained in the SARFish labels. Model predictions within 200 meters of a ‘‘HIGH’’ or ‘‘MEDIUM’’ confidence ground truth are assigned as true positive detections. Successful models will have to localise maritime objects in full-size SARFish products and distinguish them from sea-clutter, small-islands and SAR ambiguities. Performance on this task is quantified over all the relevant maritime objects in the public partition using an F1-score [11] denoted $F1_D$.

4.1.2 Close-to-Shore Object Detection Task

The Close-to-Shore detection task is a subset of the maritime detection task. This task is evaluated by assigning model predictions to ‘‘HIGH’’ or ‘‘MEDIUM’’ confidence ground truth within 2000 meters of the shoreline. The SARFish dataset provides two shorelines for evaluation of the Close-to-Shore task. The challenge will utilise the xView3 shoreline in order for the results to be directly comparable with the xView3 challenge. Close-to-Shore detection is a particularly challenging task whereby successful models will have to distinguish maritime objects from land objects with intense radar returns in addition to the challenges posed by standard maritime object detection. Performance on this task is quantified by the F1-score, $F1_S$.

4.2. Track 2: Maritime Object Classification

The maritime object classification track consists of two hierarchical classification tasks. These have also been cho-

²<https://kaggle.com/competitions/sarfish>

sen to be directly comparable with those in the xView3 challenge.

4.2.1 Vessel Classification Task

The objective of this task is: given a maritime object has been detected, distinguish whether it is a vessel or not. This task is evaluated only on the true positive predictions from the overall maritime detection task, and on ground truth detections for which a True or False ‘‘is_vessel’’ label is available. Successful models will have to distinguish between non-vessel maritime objects such as oil-rigs and offshore wind farm turbines and vessels such as oil-tankers and container ships. Performance on this task is quantified by an F1-score denoted $F1_V$.

4.2.2 Fishing Classification

This task is dependent on the results of the vessel classification task. The objective of this task is: given a maritime object has been detected, and the maritime object has been correctly classified as a vessel, distinguish whether or not it is a fishing or non-fishing vessel. This task is evaluated on the true positive classifications from the vessel classification task, and only on the ground truth for which the ‘‘is_vessel’’ label is True or False. Successful models will have to distinguish subtle differences such as those between large fishing vessels and small cargo vessels. Performance on this task is quantified by the F1-score, $F1_F$.

4.3. Track 3: Vessel Length regression

The objective of this task is to accurately predict the length of vessels in the SARFish imagery. The task is evaluated on the true positive detections from the maritime object detection task and only on ground truth for which a vessel length label exists. We note that the correct estimation of the length of maritime objects is a crucial factor that may increase the performance of classification methods, as the length of a vessel alone is a powerful feature for the classifying different types of maritime objects. Successful models will have to overcome challenging SAR artefacts associated with high intensity radar returns from metal-hulled ships such as side-lobing, smearing and multi-path effects. Performance on this task is quantified by the aggregate percentage error which penalises a prediction equally for over or under-estimating vessel length. The percentage error used in this task is defined as:

$$PE_L = 1 - \min\left(\frac{1}{N} \frac{|\min(\hat{l}_n, l_{max}) - \min(l_n, l_{max})|}{\min(\hat{l}_n, l_{max})}, 1\right)$$

Where \hat{l} is the predicted vessel length, l is the actual vessel length, and l_{max} is an upper bound for vessel length prediction corresponding approximately to the size of the largest vessel in history.

4.4. Reference Model and SLC Data Preparation

The SARFish challenge provides a baseline reference implementation of a real-valued deep learning model for the purpose of introducing new users to training and evaluating models on the SARFish SLC data. The intent of the reference model is to provide a starting point for the participants which they may use to build their own solutions. In particular, the sample code demonstrates how to extract the SLC data for use in CNNs, and how to use the SARFish metric framework for evaluation.

In the reference model, the SLC data is split into two channels: a magnitude channel and a phase channel. A total of four channels are generated between the dual VV+VH polarisations: VH_{mag} , VH_{phase} , VV_{mag} , VV_{phase} . Participants will want to explore complex-valued detection and classification models to better exploit the information contained in the complex data.

The reference model uses the predefined PyTorch implementation of FCOS [42]. FCOS was chosen because it uses the concept of “centre-ness”, which we believe is applicable to the maritime objects in this dataset. The reference model is limited in a number of regards:

1. The reference model used version of FCOS the implementation that is limited to 3-channel data. The reference uses three of the four channels: VH_{mag} , VH_{phase} and VV_{mag} ignoring the VV_{phase} channel. Participants interested in fully exploiting the SLC data in an FCOS-based model may need to write an N-channel implementation.
2. The reference model was not trained on background examples not containing maritime objects. To properly train the model, participants may want to include the background examples.
3. The reference model not attempt to evaluate vessel length, it does however, output the required *vessel_length_m* column in its prediction csv file for the metric script to evaluate.

The reference model demonstrates how to use the SARFish metrics during training and evaluation to help inform the development of better performing models. The SARFish metrics are the same as those used to evaluate the performance of models submitted to the challenge.

4.5. How to Get Started

The SARFish dataset is hosted in two locations. The data is hosted in the SARFish repository on huggingface.co³ and the labels are hosted on the xView3 website⁴.

³<https://huggingface.co/datasets/ConnorLuckettDSTG/SARFish>

⁴<https://iuu.xview.us/download-links>

The SARFish challenge provides a GitHub repo⁵ which contains tools to help get a new user working with the SARFish dataset. The repository includes:

1. Scripts for unzipping and checking the md5sum of downloaded SARFish data
2. A Jupyter Notebook containing a demonstration of how to use the SARFish dataset
3. The reference model
4. The metric script used to evaluate models submitted to the challenge
5. A visualisation script for plotting SARFish imagery, ground truth labels and model predictions

5. Conclusions

IUU fishing is a problem of global scale and significance. It causes immense environmental, economic and political damage and affects some of the poorest and most vulnerable societies on the planet. Maritime surveillance utilising automated computer vision methods on real-valued SAR has been shown to be extremely effective at detecting and classifying dark vessels engaging in IUU fishing in cases where other vessel tracking systems such as AIS and VMS fail. Further, deep learning methods for target detection in CV-SAR imagery has demonstrated higher performance on target detection tasks than real-valued SAR, indicating a potential for the development of improved tools for combating IUU fishing.

Previous work such as the xView3 dataset and challenge have made a significant contribution to advancing the state-of-the-art with real-valued SAR ship detection and classification tasks. To foster research into CV-SAR, we have released the SARFish which provides coincident SLC and GRD products. To our knowledge, the SARFish dataset is the only one of its kind suitable for training deep learning models on the tasks of ship detection and classification in full-size dual-polarisation SLC and GRD SAR imagery. We also announce the SARFish challenge which intends to advance CV-SAR research state-of-the-art by encouraging competition and fostering interest in the field within the wider computer vision community.

Acknowledgment

We thank Dr Ben Yip for his help in developing the SARFish reference model, and Dr Tristrom Cooke for reviewing the paper and making valuable suggestions for its improvement.

⁵<https://github.com/RitwikGupta/SARFish>

References

- [1] David J Agnew, John Pearce, Ganapathiraju Pramod, Tom Peatman, Reg Watson, John R Beddington, and Tony J Pitcher. Estimating the worldwide extent of illegal fishing. *PLoS one*, 4(2):e4570, 2009. [1](#)
- [2] Reza Mohammadi Asiyabi, Mihai Datcu, Andrei Anghel, and Holger Nies. Complex-valued end-to-end deep network with coherency preservation for complex-valued sar data reconstruction and classification. *IEEE Transactions on Geoscience and Remote Sensing*, 2023. [3](#)
- [3] Reza Mohammadi Asiyabi, Mihai Datcu, Holger Nies, and Andrei Anghel. Complex-valued vs. real-valued convolutional neural network for polsar data classification. In *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*, pages 421–424. IEEE, 2022. [3](#)
- [4] Priyal Bunwaree. The illegality of fishing vessels ‘going dark’ and methods of deterrence. *International & Comparative Law Quarterly*, 72(1):179–211, 2023. [1](#)
- [5] Tri-Tan Cao, Connor Luckett, Jerome Williams, Tristram Cooke, Ben Yip, Arvind Rajagopalan, and Sebastien Wong. Sarfish: Space-based maritime surveillance using complex synthetic aperture radar imagery. In *2022 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–8. IEEE, 2022. [3](#), [5](#)
- [6] Shankar N Ramasari Chandra, Jon Christopherson, and Kimberly A Casey. *2020 Joint Agency Commercial Imagery Evaluation—Remote sensing satellite compendium*. Number 1468. US Geological Survey, 2020. [2](#), [4](#), [5](#)
- [7] Ian Cumming and John Bennett. Digital processing of seasat sar data. In *ICASSP’79. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 4, pages 710–718. IEEE, 1979. [2](#)
- [8] John C Curlander. Location of spaceborne sar imagery. *IEEE Transactions on Geoscience and Remote Sensing*, (3):359–364, 1982. [6](#)
- [9] Raj M Desai and George E Shambaugh. Measuring the global impact of destructive and illegal fishing on maritime piracy: A spatial analysis. *Plos one*, 16(2):e0246835, 2021. [1](#)
- [10] Emmanouil Detsis, Yuval Brodsky, Peter Knudtson, Manuel Cuba, Heidi Fuqua, and Bianca Szalai. Project catch: A space based solution to combat illegal, unreported and unregulated fishing: Part i: Vessel monitoring system. *Acta Astronautica*, 80:114–123, 2012. [1](#), [2](#)
- [11] Tom Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006. [7](#)
- [12] Iwao Fujii, Yumi Okochi, and Hajime Kawamura. Promoting cooperation of monitoring, control, and surveillance of iuu fishing in the asia-pacific. *Sustainability*, 13(18):10231, 2021. [1](#)
- [13] Guillaume Hajduch, M Bourbigot, Harald Johnsen, and Riccardo Piantanida. *Sentinel-1 Product Specification*. Feb 2022. [5](#)
- [14] Ronny Hänsch. Complex-valued multi-layer perceptrons—an application to polarimetric sar data. *Photogrammetric Engineering & Remote Sensing*, 76(9):1081–1088, 2010. [3](#)
- [15] Jinglu He, Yinghua Wang, and Hongwei Liu. Ship classification in medium-resolution sar images via densely connected triplet cnns integrating fisher discrimination regularized metric learning. *IEEE Transactions on Geoscience and Remote Sensing*, 59(4):3022–3039, 2020. [3](#)
- [16] G Hosch and G Macfadyen. Killing nemo: Three world regions fail to mainstream combatting of iuu fishing. *Marine Policy*, 140:105073, 2022. [1](#)
- [17] Qinglong Hua, Yun Zhang, Yicheng Jiang, and Huilin Mu. Gaussian-type activation function for complex-valued cnn and its application in polar-sar image classification. *Journal of Applied Remote Sensing*, 15(2):026510–026510, 2021. [3](#)
- [18] Lanqing Huang, Bin Liu, Boying Li, Weiwei Guo, Wenhao Yu, Zenghui Zhang, and Wenxian Yu. Opensarship: A dataset dedicated to sentinel-1 ship interpretation. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 11(1):195–208, 2017. [3](#), [4](#)
- [19] Zhongling Huang, Mihai Datcu, Zongxu Pan, and Bin Lei. Deep sar-net: Learning objects from signals. *ISPRS Journal of Photogrammetry and Remote Sensing*, 161:179–193, 2020. [4](#)
- [20] Zhongling Huang, Zongxu Pan, and Bin Lei. What, where, and how to transfer in sar target recognition based on deep cnns. *IEEE Transactions on Geoscience and Remote Sensing*, 58(4):2324–2336, 2019. [3](#)
- [21] International Maritime Organization (IMO). Revised guidelines for the onboard operational use of shipborne automatic identification systems (ais), 2015. [1](#)
- [22] Charles VJ Jakowatz, Daniel E Wahl, Paul H Eichel, Dennis C Ghiglia, and Paul A Thompson. *Spotlight-mode synthetic aperture radar: a signal processing approach*. Springer Science & Business Media, 2012. [1](#), [2](#), [3](#)
- [23] S Jutz and MP Milagro-Pérez. Copernicus: the european earth observation programme. *Revista de Teledetección*, (56):V–XI, 2020. [2](#)
- [24] Urška Kanjir, Harm Greidanus, and Krištof Oštir. Vessel detection and classification from spaceborne optical images: A literature survey. *Remote sensing of environment*, 207:1–26, 2018. [1](#)
- [25] David A Kroodsmma, Juan Mayorga, Timothy Hochberg, Nathan A Miller, Kristina Boerder, Francesco Ferretti, Alex Wilson, Bjorn Bergman, Timothy D White, Barbara A Block, et al. Tracking the global footprint of fisheries. *Science*, 359(6378):904–908, 2018. [1](#)
- [26] Andrey A Kurekin, Benjamin R Loveday, Oliver Clements, Graham D Quartly, Peter I Miller, George Wiafe, and Kwame Adu Agyekum. Operational monitoring of illegal fishing in ghana through exploitation of satellite earth observation and ais data. *Remote Sensing*, 11(3):293, 2019. [1](#), [2](#)
- [27] Haitao Lang, Jie Zhang, Ting Zhang, Di Zhao, and Junmin Meng. Hierarchical ship detection and recognition with high-resolution polarimetric synthetic aperture radar imagery. *Journal of Applied Remote Sensing*, 8(1):083623–083623, 2014. [4](#)
- [28] Songlin Lei, Xiaolan Qiu, Chibiao Ding, and Shujie Lei. A feature enhancement method based on the sub-aperture decomposition for rotating frame ship detection in sar images.

- In 2021 *IEEE International Geoscience and Remote Sensing Symposium IGARSS*, pages 3573–3576. IEEE, 2021. 4
- [29] B. Li, B. Liu, L. Huang, W. Guo, Z. Zhang, and W. Yu. Opensarship 2.0: A large-volume dataset for deeper interpretation of ship targets in sentinel-1 imagery. *2017 SAR in Big Data Era: Models, Methods and Applications (BIGSAR-DATA)*, 2017. 3, 4
- [30] J. Li, C. Qu, and J. Shao. Ship detection in sar images based on an improved faster r-cnn. *2017 SAR in Big Data Era: Models, Methods and Applications (BIGSAR-DATA)*, 2017. 3
- [31] Jianwei Li, Congan Xu, Hang Su, Long Gao, and Taoyang Wang. Deep learning for sar ship detection: Past, present and future. *Remote Sensing*, 14(11):2712, 2022. 3
- [32] University of Alaska. Alaska satellite facility (asf) distributed active archive center. <https://asf.alaska.edu/asfsardaac/>. 5
- [33] Ifesinachi Okafor-Yarwood. Illegal, unreported and unregulated fishing, and the complexities of the sustainable development goals (sdgs) for countries in the gulf of guinea. *Marine Policy*, 99:414–422, 2019. 1
- [34] Ifesinachi Okafor-Yarwood. The cyclical nature of maritime security threats: Illegal, unreported, and unregulated fishing as a threat to human and national security in the gulf of guinea. *African Security*, 13(2):116–146, 2020. 1
- [35] F.S. Paolo, T-t.T. Lin, R. Gupta, B. Goodman, N. Patel, D. Kuster, D. Kroodsmas, and J. Dunmon. xview3-sar: Detecting dark fishing activity using synthetic aperture radar imagery. *arXiv:2206.00897v4 [cs.CV]*, Nov 2022. 3, 6
- [36] Jaeyoon Park, Jungsam Lee, Katherine Seto, Timothy Hochberg, Brian A Wong, Nathan A Miller, Kenji Takasaki, Hiroshi Kubota, Yoshioki Oozeki, Sejal Doshi, et al. Illuminating dark fishing fleets in north korea. *Science advances*, 6(30):eabb1197, 2020. 1
- [37] Riccardo Piantanida, G Hajduch, and J Poullaouec. Sentinel-1 level 1 detailed algorithm definition. *ESA, techreport SEN-TN-52-7445*, 2016. 2
- [38] Guo Rui, Zhang Zhao, Gao Rui, Jing Guobin, Xing Mengdao, and Zhi Yongfeng. Ocean target investigation using spaceborne sar under dual-polarization strip-map mode. In *2021 2nd China International SAR Symposium (CISS)*, pages 1–5. IEEE, 2021. 2
- [39] Colin P Schwegmann, Waldo Kleynhans, Brian P Salmon, Lizwe W Mdakane, and Rory GV Meyer. Very deep learning for ship discrimination in synthetic aperture radar imagery. In *2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 104–107. IEEE, 2016. 2, 3
- [40] Zikang Shao, Tianwen Zhang, and Xiao Ke. A dual-polarization information-guided network for sar ship classification. *Remote Sensing*, 15(8):2138, 2023. 4
- [41] Andrew J Temple, Daniel J Skerritt, Philippa EC Howarth, John Pearce, and Stephen C Mangi. Illegal, unregulated and unreported fishing impacts: A systematic review of evidence and proposed future agenda. *Marine Policy*, 139:105033, 2022. 1
- [42] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. 8
- [43] Ramon Torres, Paul Snoeij, Dirk Geudtner, David Bibby, Malcolm Davidson, Evert Attema, Pierre Potin, Björn Rommen, Nicolas Floury, Mike Brown, Ignacio Navas Traver, Patrick Deghaye, Berthyl Duesmann, Betlem Rosich, Nuno Miranda, Claudio Bruno, Michelangelo L’Abbate, Renato Croci, Andrea Pietropaolo, Markus Huchler, and Friedhelm Rostan. Gmes sentinel-1 mission. *Remote Sensing of Environment*, 120:9–24, 2012. The Sentinel Missions - New Opportunities for Science. 1, 2, 3, 5
- [44] S. Wei, X. Zeng, Q. Qu, M. Wang, H. Su, and J. Shi. Hsrid: A high-resolution sar images dataset for ship detecton and instance segmentation. *IEEE Access*, 8:120234–120254, 2020. 3
- [45] Heather Welch, Tyler Clavelle, Timothy D White, Megan A Cimino, Jennifer Van Osdel, Timothy Hochberg, David Kroodsmas, and Elliott L Hazen. Hot spots of unseen fishing vessels. *Science Advances*, 8(44):eabq2109, 2022. 1
- [46] Xiaowo Xu, Xiaoling Zhang, Zikang Shao, Jun Shi, Shunjun Wei, Tianwen Zhang, and Tianjiao Zeng. A group-wise feature enhancement-and-fusion network with dual-polarization feature enrichment for sar ship detection. *Remote Sensing*, 14(20):5276, 2022. 2
- [47] Darrell L Young. Deep nets spotlight illegal, unreported, unregulated (iuu) fishing. In *2019 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, pages 1–7. IEEE, 2019. 1
- [48] Tianwen Zhang and Xiaoling Zhang. A polarization fusion network with geometric feature embedding for sar ship classification. *Pattern Recognition*, 123:108365, 2022. 4
- [49] Tianwen Zhang, Xiaoling Zhang, Xiao Ke, Xu Zhan, Jun Shi, Shunjun Wei, Dece Pan, Jianwei Li, Hao Su, Yue Zhou, et al. Ls-ssdd-v1. 0: A deep learning dataset dedicated to small ship detection from large-scale sentinel-1 sar images. *Remote Sensing*, 12(18):2997, 2020. 2, 3
- [50] Zhimian Zhang, Haipeng Wang, Feng Xu, and Ya-Qiu Jin. Complex-valued convolutional neural network and its application in polarimetric sar image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 55(12):7177–7188, 2017. 3
- [51] Dandan Zhao, Zhe Zhang, Dongdong Lu, Jian Kang, Xiaolan Qiu, and Yirong Wu. Cvgg-net: Ship recognition for sar images based on complex-valued convolutional neural network. *arXiv preprint arXiv:2305.07918*, 2023. 4