Vessel detection in Synthetic Aperture Radar images using Faster R-CNN models:

Advanced monitoring techniques to improve fisheries management

Universitat Oberta de Catalunya

UNIVERSITAT DE BARCELONA

Pol Carbó Mestre

MU Bioinf. i Bioest.
Statistical Bioinformatics and Machine Learning

Nombre Tutor/a de TF
Romina Astrid Rebrij
Profesor/a responsable de la asignatura
Carles Ventura Royo

20 de Junio de 2023

FICHA DEL TRABAJO FINAL

Título del trabajo:	*Vessel detection in Synthetic Aperture Radar images using Faster R-CNN models: Advanced monitoring techniques to improve fisheries management*
Nombre del autor:	*Pol Carbó Mestre*
Nombre del consultor/a:	*Romina Astrid Rebrij*
Nombre del PRA:	*Carles Ventura Royo*
Fecha de entrega:	*06/2023*
Titulación o programa:	Máster Universitario Bioinformática y Bioestadística UOC-UB
Área del Trabajo Final:	*Statistical Bioinformatics and Machine Learning*
Idioma del trabajo:	*Inglés*
Palabras clave	*Machine Learning, Synthetic Aperture Radar, fisheries management*

Resumen del Trabajo

Tradicionalmente, la gestión pesquera se ha basado en el establecimiento de cuotas y la recolección manual de datos. Sin embargo, el constante deterioro de las poblaciones pesqueras ha requerido la adopción de sistemas de control electrónicos como el Sistema de Monitoreo de Embarcaciones (VMS) y el Sistema de Identificación Automática (AIS), cuyas limitaciones, sin embargo, siguen sin prevenir la sobreexplotación pesquera. Debido a esto, en los últimos años, varios avances en el análisis de imágenes satelitales de Radar de Apertura Sintética (SAR) con técnicas de Aprendizaje Automático (ML) han destacado como prometedoras herramientas para gestionar las actividades pesqueras.

Estos avances están revolucionando el control de la industria marina, reduciendo las limitaciones de las técnicas tradicionales. Este trabajo se centra en la detección de embarcaciones, utilizando métodos de "computer vision", específicamente Redes Neuronales Convolucionales (CNN) del tipo Faster Region-Based (Faster-RCNN). En este proyecto, hemos evaluado el rendimiento de estos modelos, incluyendo la implementación de diversas técnicas de preprocesamiento de imágenes. Además, para demostrar el potencial de este enfoque en la gestión pesquera, hemos aplicado el modelo en imágenes del satélite Sentinel-1 a través de un caso de estudio sobre las pesquerías chilenas y desarrollado un informe interactivo para presentar los resultados.

La integración del análisis de imágenes SAR y las técnicas de ML tiene un gran potencial para mejorar la gestión de pesquerías. La evaluación de Faster-RCNN para la detección de embarcaciones, junto con el análisis comparativo de las técnicas de preprocesamiento, ha proporcionado información valiosa sobre la efectividad de estos métodos. Además, también ha revelado algunas de las limitaciones que estas técnicas presentan, subrayando la necesidad de avances adicionales y enfatizando en la integración de los distintos enfoques, tanto tradicionales como modernos, para una gestión efectiva de las pesquerías.

Abstract

Fisheries management has traditionally relied on catch quotas and manual reporting methods. However, the depletion of fish stocks has required the adoption of electronic reporting systems such as the Vessel Monitoring System (VMS) and Automatic Identification System (AIS), with its own limitations in preventing overexploitation. To address this, advancements in Synthetic Aperture Radar (SAR) satellite imagery analysis, coupled with Machine Learning (ML) techniques, have emerged as promising tools for monitoring fishing activities.

These advancements are revolutionizing marine industry monitoring, closing the data gaps from traditional techniques, and enhancing transparency. This project focuses on vessel detection, employing computer vision methods. Convolutional Neural Networks (CNN), specifically Faster Region-Based CNN (Faster-RCNN), exhibit promising results with reduced detection time and computational costs. We evaluated the model's performance of these models and implemented various image pre-processing techniques to improve them. Furthermore, to demonstrate the potential of this approach in fisheries management, we tested the model using real-world Sentinel-1 images in a case study on Chilean fisheries and developed an interactive report presenting the results.

The integration of SAR-based satellite imagery analysis and ML techniques holds significant promise for enhancing fisheries management. The evaluation of Faster-RCNN for vessel detection, along with the comparative analysis of pre-processing techniques, provides valuable insights into the effectiveness of this method. Furthermore, it also revealed some limitations of these techniques, underscoring the need for further advancements and emphasizing the reliance on combined approaches for effective fisheries management.

Table of Contents

List of Figures

List of Tables

1. Introduction

1.1. Context and justification

Traditionally, fisheries management has focused on setting catch quotas for specific species and areas. These quotas are monitored using landing reports and skipper's logbooks. However, new electronic reporting requirements have been introduced due to the long-term depletion of fishing stocks. These requirements enable authorities to track the movements of individual fishing vessels and their fishing activities (Lemoine et al. 2004). The most crucial of these technologies are the Vessel Monitoring System (VMS) and the Automatic Identification System (AIS), which analysis allows identifying fishing vessels and fishing hours to extrapolate the fishing effort (David A. Kroodsma et al. 2018). Despite their utility, these systems have significant limitations. For instance, some fishing vessels are not required to have them, and transceivers can be disconnected anytime. Additionally, unregistered fishing vessels can operate undetected worldwide. These factors make it difficult to monitor illegal and unreported fishing practices, which contribute to the overexploitation of fish stocks and hinder the recovery of fish populations and ecosystems (Shepperson et al. 2018a). To address these limitations, recent advancements in satellite imagery analysis have allowed the monitoring of fishing activities under those circumstances where VMS and AIS data fail. Synthetic Aperture Radar (SAR) has outstood other types of spaceborne imagery because it operates effectively under all weather conditions. The potential of SAR images for fishing monitoring was already highlighted by Lemoine et al. in 2004. However, only a decade later, with the development and improvement of Machine Learning (ML) techniques, the use of SAR images for vessel detection has reached its higher potential.

The availability of satellite imagery and advancements in ML have paved the way for a new era in monitoring marine industries. This is closing the VMS and AIS data gaps, enabling increased transparency in ocean activities. The use of these technologies in fishing monitoring has progressed through various stages, starting with the development of accurate models to detect marine structures and vessels, followed by the discrimination of fishing vessels, and ultimately, with the integration of SAR-derived information with VMS and AIS data outputs to provide a more comprehensive and accurate picture of fishing activities (Galdelli et al. 2021).

This project will focus on the initial stage, vessel detection. Various methods have been proposed for detecting vessels, primarily relying on computer vision techniques. The state-of-the-art and existing modeling approaches will be discussed, including the pros and cons of the existing modeling frameworks.

One of the more popular methods is the Convolutional Neural Networks (CNN), showing lower detection time and computational costs with minimal performance loss. As such, this project will implement a model based on the object detection architecture Faster Region-Based Convolutional Neural Networks (Faster-RCNN), which has already been used for vessel detection, enabling us to explore this field further. Finally, integrating the model into an interactive report and testing it with real-world Sentinel-1 data will allow us to evaluate the potential of these tools and the data outputs derived from SAR images in a practical setting.

While these techniques have already been tested over the past few years, there is a lack of available projects with sufficient documentation to learn from scratch about these methods applied to fisheries. This project stands out from others as it aims to create accessible methods with well-documented scripts, notebooks, and reports that can serve as resources for anyone interested in learning about this field and these techniques. The project's ultimate goal is not only to develop and test the model on real data but also to create comprehensive documentation that facilitates the learning curve for others, including myself, who found it challenging to get started with these techniques.

1.2. General description

We developed an accurate and reliable model using the Faster-RCNN object detection architecture for detecting vessels from SAR images. Furthermore, the project compared the model's performance by implementing various image pre-processing techniques. The project's final phase involved integrating the model into an interactive report and testing it with real-world Sentinel-1 data from the Copernicus Program. By achieving these goals, we aim to gain hands-on experience with modeling techniques used in the existing literature for vessel detection while reviewing the current state-of-the-art in this field.

1.3. Project goal

1.3.1. Overall objective

Develop a model to detect vessels from SAR images and evaluate them using Copernicus Program data.

1.3.2. Detailed objectives

2. Conduct a comprehensive review of the current state-of-the-art methods for vessel detection using SAR images in fisheries monitoring.
3. Develop a Faster-RCNN model to detect ships in SAR images.
4. Evaluate the performance of the Faster-RCNN model by implementing different image pre-processing techniques.

5. Create an interactive report showcasing the model potential on real-world Sentinel-1 for a fisheries-related case study.

1.4. Impact on sustainability, ethical-social, and diversity

This project focuses on enhancing fisheries management, which plays an important role in environmental and social aspects. From an environmental perspective, effective fisheries management is crucial for ensuring the sustainability and preservation of marine ecosystems by preventing overfishing, protecting endangered species, and maintaining the overall health of our oceans. On a social level, fisheries provide livelihoods for millions worldwide, especially in developing nations where communities often depend on fishing for subsistence and economic stability. Implementing sound management strategies can lead to sustainable fisheries, ensuring food security and stable incomes for these communities. Furthermore, effective management strategies incorporate ethical factors such as fair labor practices and involving communities in decision-making processes. Lastly, fisheries management also plays a crucial role in promoting socio-cultural diversity by recognizing the value of traditional knowledge and supporting artisanal and indigenous fishing practices, which are essential for the sustainable stewardship of our marine resources.

In this regard, this project is clearly framed under the Sustainable Development Goal (SDG) target 14—Life Below Water—as the main objective of the methods applied is to “conserve and sustainably use the oceans, seas, and marine resources for sustainable development.” Furthermore, considering the socio-economic implications of fisheries and any technological advancements aimed at improving them, the project is tangentially related to SDG 8 in "Decent Work and Economic Growth."

1.5. Approach and methods

Although SAR images, in combination with ML techniques, have proven their great potential to assess fishing activities, the progress of these techniques has been hindered by the scarcity of labeled images. Despite the widespread availability of satellite images from open sources like the Sentinel 1 imagery from the Copernicus Project, the development of vessel detectors using SAR has been slow due to the lack of a large volume of annotated datasets (Wang et al. 2019). Labeling images is a significant undertaking involving manually annotating detections or using spatially extrapolated information from VMS and AIS data. While the former method is very time-consuming, the latter is not feasible, given our project constraints. Therefore, we have decided to use pre-annotated datasets available online. Among the existing options, we have chosen the dataset created by Zhang et al. (2020) as it has a sufficient number of images to train our model and achieve appropriate performance while still being small enough to be processed within our computational resource constraints. This dataset consists of 15 large-scale VV polarization SAR images obtained from the Sentinel-1 satellite, with sizes of 24000 × 16000 pixels. The large-scale images were cut into 9000 sub-images for training and evaluation purposes with dimensions of 800 × 800 pixels.

With the availability of an annotated dataset, we can apply the supervised learning algorithm Faster-RCNN, a two-stage object detection algorithm introduced by Ren et al. (2017). The steps to implement Faster R-CNN are as follows:

- Prepare the training dataset: Select a suitable dataset and annotate the images with bounding boxes around the objects of interest. Since annotated datasets are already available online, our task is to identify an appropriate dataset for our specific needs and consider alternative datasets as backups. Potential alternative datasets were the SAR-Ship-Dataset from Wang et al. (2019) and the HRSID by Wei et al. (2020).
- Image preprocessing: To enhance the model's performance, appropriate image preprocessing techniques such as data augmentation and image denoising can be applied. Several studies have shown that image preprocessing techniques can significantly improve the performance of Faster R-CNN on object detection tasks (Yang et al. 2021; Zhao et al. 2022).
- Train the model: Train the Faster R-CNN model, which generates candidate regions for objects in an image using a Region Proposal Network (RPN). The model then performs posterior classification and refinement of those candidate regions to identify the areas where ships are located using a Fast R-CNN.
- Improve the model: improve the model's performance by implementing various image pre-processing techniques, such as data augmentation.
- Test the model: evaluating the performance of the trained model on unseen data.

This workflow is based on the approaches presented in previous works (Wei et al. 2020; Wang et al. 2019a; J. Li, Qu, and Shao 2017; Zhang et al. 2020). Although there are other methods available for ship detection, such as those based on the multilayer Constant False Alarm Rate (CFAR) (Hou, Chen, and Jiao 2015), YOLO (Ting et al. 2021), or RetinaNet (Wang et al. 2019b), among others, our focus will be on the Faster R-CNN approach due to its widespread use and high performance in object detection tasks.

Finally, to illustrate the implementation of the Faster R-CNN approach in vessel detection using SAR images, we will create an interactive report where the model will be tested against real-world Sentinel-1 data from the Copernicus Open Access Hub. This report, produced with Quarto, will clearly and concisely demonstrate the capabilities of SAR imagery and the benefits of using deep learning models for vessel detection tasks. By presenting our findings in an accessible and interactive format, we hope to create a science outreach document to increase awareness and understanding of the potential applications of SAR imagery in maritime surveillance and other related fields.

1.6. Project planning

1.6.1. Tasks

Table 1. Project task details.

Description	Start	End
Work Plan. PEC1.	**1/3/23**	**20/3/23**
Search a collection of labeled vessel SAR images.	1/3/23	7/3/23
Evaluation and selection of the models to be applied.	7/3/23	14/3/23
Designing and writing the work plan.	14/3/23	20/3/23
Submission of the work plan and incorporation of feedback.	20/3/23	27/3/23
Development of the project. Phase 1. PEC2.	**21/3/23**	**24/4/23**
Evaluation of the required methods, framework, and tools.	21/3/23	26/3/23
Assessing model input format and pre-process data.	26/3/23	3/4/23
Designing and testing the workflow to implement the model.	3/4/23	10/4/23
Model training.	10/4/23	17/4/23
Evaluation of the model's performance.	16/4/23	20/4/23
Documenting and submitting the report.	17/4/23	24/4/23
Development of the project. Phase 2. PEC3.	**25/4/23**	**29/5/23**
Model improvement.	25/4/23	15/5/23
Design and creation of the interactive report	11/5/23	10/6/23
Documenting and submitting the report.	22/5/23	29/5/23
Finalize report. PEC4.	**30/5/23**	**20/6/23**
Finalize report.	22/5/23	20/6/23
Preparing the presentation.	12/6/23	20/6/23
Thesis defense. PEC5.	**3/7/23**	**14/7/23**

1.6.2. Calendar

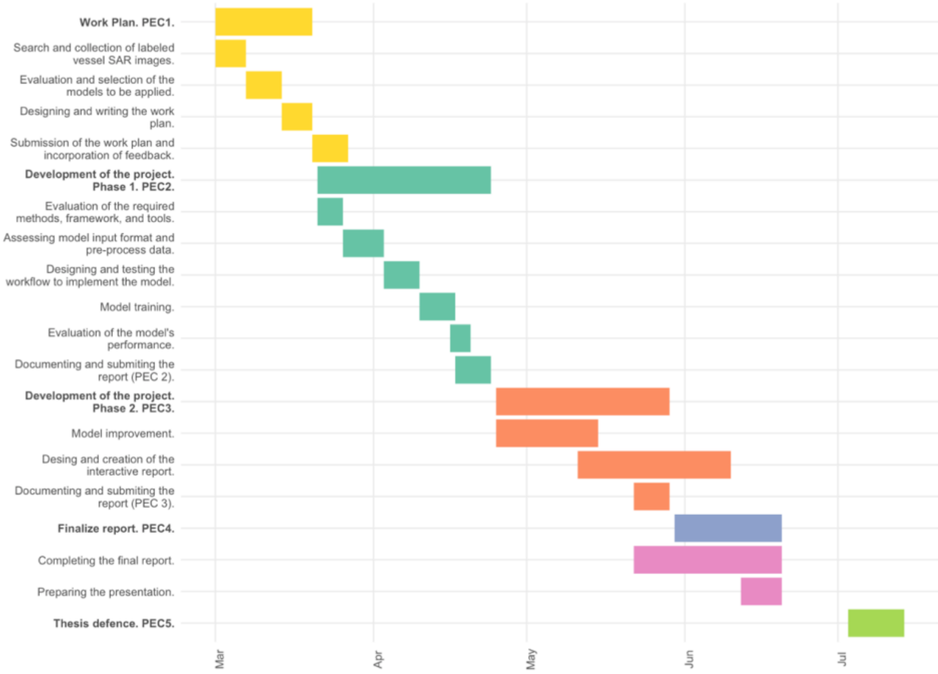

Figure 1. Gantt diagram with the project tasks.

1.6.3. Milestones

Table 2. Milestones details.

Description	Deadline
Work plan submission.	20/3/23
Project development submission. Phase 1.	24/4/23
Project development submission. Phase 2.	29/5/23
Final report submission.	20/6/23
Presentation submission.	20/6/23
Thesis defense.	3/7/23

1.6.4. Risk assessment

Table 3. Risk assessment details.

Risk description	Severity	Probability	Mitigation
Issues with the functioning of Python packages.	High	Low	Researching potential alternative packages.
Not being able to apply the model to the data.	High	Low	Verify the necessary data input characteristics for the chosen framework. Source alternative datasets.

Not achieving optimal model performance.	High	Moderate	Evaluate all factors to optimize the models.
Unable to develop the interactive report	Moderate	Low	Research extensively about Jupiter and Quarto notebooks.
Unable to meet the proposed deadlines for developing the tasks of each phase.	Moderate	Moderate	Tasks could be completed in later stages. Additionally, more time could be allocated to each task.

1.7. Deliverables summary

- **Work plan:** Initial document in pdf format containing the guidelines and the timeline of all the tasks necessary for the project's development and achievement of goals.
- **Final report:** PDF document detailing all the research, methods, results, and conclusions obtained throughout the master's thesis work.
- **Deliverables:** Scripts, datasets, and the final trained model archived on an accessible public repository. Jupyter notebooks describing each step of the process and providing code examples. An interactive report displaying real-world examples of vessel identification using SAR images from the Copernicus Program.
- **Presentation:** Slide presentation to summarize the project and present the outputs.

Such deliverables are available on the following GitHub repository, along with a description of their contents and usage. Additionally, they will be attached to the final report submission.

https://github.com/pcarbomestre/SAR-VesselDetection-FisheriesMonitoring

1.8. Brief description of the other chapters of the thesis

- **State of the art**. A brief exposition is presented on the theoretical concepts and key elements for understanding the project. The chapter contents effectively justify the project's context and goals while highlighting its relevance, emphasizing the biological applications of the technology in fishery science.
- **Methodology**. The development of this project is explained in detail in this chapter.
- **Results**. The obtained results are presented in terms of the model’s performance.
- **Discussion**. The combined results from the previous two sections are discussed and related to the theoretical framework.
- **Case Study**. This chapter demonstrates this approach's potential in fisheries management by testing the model in a case study on Chilean fisheries.
- **Conclusions**. The project's conclusions are presented, and future lines of work are proposed.

2. State of the art

2.1. Introducing the importance of fisheries

The fishing and aquaculture industries play a crucial role in providing food security and nutrition globally and are an essential economic driver in many countries and communities. In 2020, the global capture fisheries production was 90.3 million tonnes—60% of it corresponding to large-scale fisheries—and has remained relatively stable over the past few decades, fluctuating between 86 million tonnes and 93 million tonnes per year since the late 1980s. Half a billion people depend on fisheries for their livelihood, including 38 million direct workers, their household members, and others who engage in subsistence fishing. Additionally, fish is a significant part of the diet of over 3 billion people worldwide, providing at least 20% of the average per capita intake of animal protein (FAO 2022).

Given the economic importance of fisheries and the changing landscape of population growth, market access, technological progress, and economic development of several nations, there has been a dramatic increase in the environmental pressures stemming from fisheries, particularly from industrial fishing, over the last decades.

2.2. Environmental impacts of fisheries

Regardless of a decrease in fishing vessels in the past two decades due to fleet reduction programs in Europe and China, there are still an estimated 4.1 million operative fishing vessels worldwide. The continued decrease in effective catch per unit of effort (CPUE) indicates that the world's fisheries are placing immense pressure on ocean resources (Rousseau et al., 2019). CPUE is an indirect measure of the abundance of a target species, and it responds to the logic of how much effort is invested in catching a fish and how much fish you get in return. Therefore, a decrease in CPUE means that, since 1950—when technological advancements allowed for the creation of modern industrial fishing—the number of fishing hours, distance to fishing grounds, and resources spent have increased for the same amount of fish caught, which translates into more pressure on fishing resources.

According to FAO's assessment, this has resulted in a decrease in biologically sustainable fishery stocks. While in 1974, 90% of fishery stocks were biologically sustainable (i.e., yield does not impair the stock reproductivity, reaching a good balance between human use and ecological conservation), the percentage of stocks fished at biologically unsustainable levels has increased reaching 35.4% in 2019. In the meantime, the underfished stocks have followed a decreasing trend, and today only a tiny percentage of global fish stocks (7.2%) are not fished at their maximum sustainable capacity or are overfished.

Overall, this trend poses a significant threat to marine fish populations worldwide, as approximately 90% of them are fully exploited (i.e., maximally sustainably fished) or overfished. As a result, over the past 50 years, more than 366 fisheries, equivalent to a quarter of all fisheries, have experienced a collapse (Mullon, Fréon, and Cury 2005). Notable examples include the Peruvian anchoveta and the cod fisheries in the North Atlantic waters of the USA and Canada in the 1990s, which have yet to recover.

Despite the abundance of available data and methods used to assess fishing stocks and sustainable catch rates, reliance is placed on the information provided by the fishing industry and the national authorities. Unfortunately, the existing assessments account for great uncertainty due to illegal, unregulated, and unreported (IUU) fishing activities, which hide the actual impacts of global fishing.

2.3. IUU fishing challenges

Fisheries management faces a significant challenge from IUU fishing. This involves a wide range of activities, such as fishing without proper licenses or quotas for specific species, transferring catch to cargo vessels without authorization, falsifying catch reports, keeping undersized or protected fish, exceeding permit limits for catch quantities, fishing in restricted areas or during prohibited seasons, and using prohibited fishing gear. These practices contribute to the widespread issue of IUU fishing, which threatens world fish populations and marine ecosystems. Despite its difficulty to estimate, experts suggest that IUU fishing accounts for 14-33% of the global catch and up to 50% in some areas, with a total value of approximately $10-23 billion annually (Agnew et al. 2009).

Therefore, IUU Fishing represents a significant global problem with economic, social, and environmental impacts. In socio-economic terms, IUU fishing can directly affect the income of fishermen and others in the seafood supply chain and harm nations' economic, income, and tax revenues (Sumaila et al. 2020). Besides the destructive fishing practices involved in IUU, the unreported catch can compromise the status of fish stocks by exceeding capture limits and undermining the accuracy of fisheries data and assessment models, ultimately leading to overfishing. Inaccurate models derived from poor-quality or uncompleted data make it difficult to set effective policies and hinder efforts to achieve proper fisheries management (Pitcher et al. 2002). Ultimately, the decline in fish stocks can have far-reaching effects on habitats and ecosystems, which in turn would negatively impact the long-term sustainability of economically important fisheries and communities that rely on fishing for food and economic security (Temple et al. 2022).

To combat IUU fishing, traceability and catch documentation are essential. The UN Food and Agriculture Organization (FAO) has identified a lack of transparency in fishing vessel activities and registries as a critical factor contributing to ongoing challenges in the industry. In consonance, the EU and the USA have established fisheries policies prioritizing traceability, which involves monitoring the entire seafood supply chain. However, supply chain traceability is limited to shore-side landing sites or storage and processing facilities, which limits the effectiveness of this monitoring approach. To address this, new initiatives are being developed to enhance fisheries transparency by tracking vessel movements (Chuaysi and Kiattisin 2020).

2.4. Existing monitoring tools

Traditionally, fisheries management has centered on setting catch quotas and other restrictions for specific species and areas. These quotas are monitored using landing reports and skipper's logbooks. However, new electronic reporting requirements have been introduced due to the long-term depletion of fishing stocks and the uncertainty associated with IUU fishing. These requirements enable authorities to track the movements of individual fishing vessels and their fishing activities (Lemoine et al. 2004). The most important of these technologies are the Vessel Monitoring System (VMS) and the Automatic Identification System (AIS), which data analysis allows identifying fishing vessels and calculate fishing hours to extrapolate the fishing effort (David A. Kroodsma et al. 2018). Despite their utility, these systems have significant limitations. For instance, some fishing vessels are not required to have them, and transceivers can be disconnected anytime. Additionally, unregistered fishing vessels can operate undetected worldwide. These factors make it difficult to monitor IUU, which contributes to the overexploitation of fish stocks and hinders the recovery of fish populations and ecosystems (Shepperson et al. 2018). To address these limitations, recent advancements in satellite imagery analysis have allowed the monitoring of fishing activities under those circumstances where VMS and AIS data fail.

2.5. Vessel detection using SAR images

Synthetic Aperture Radar (SAR) has outstood other types of spaceborne imagery because it operates effectively under all weather conditions. The potential of SAR images for fishing monitoring was already highlighted by Lemoine et al. in 2004. However, only a decade later, with the development and improvement of Machine Learning (ML) techniques, the use of SAR images for vessel detection has reached its higher potential.

2.5.1. Synthetic Aperture Radar (SAR)

SAR is a remote sensing technology that uses RADAR systems to capture high-resolution images of the Earth's surface. Unlike traditional passive satellite imaging techniques, which measure the energy naturally available on Earth's Surface (i.e., visible light, IR, and other electromagnetic waves), SAR operates by emitting microwaves and measuring the return signals that bounce back after interacting with the surface. This active radar imaging technique enables SAR to acquire data regardless of daylight, cloud cover, or atmospheric conditions (Moreira et al. 2013).

The synthetic aperture concept arises from its operational procedure to obtain images. RADAR imagery largely depends on the antenna size. For high-resolution images such as the ones related to landscape imagery, the necessary antenna dimensions would not be feasible to install on any aircraft or spacecraft. Therefore, unlike conventional radars, SAR systems generate images by combining the echoes received from multiple positions along the flight path, simulating a large antenna and allowing for high-resolution imagery (Jansing 2021).

One of the primary advantages of SAR is its ability to penetrate through clouds, providing higher spatial and temporal Earth image accessibility (i.e., all-weather and day-night imaging), allowing the monitoring of remote and inaccessible regions while offering data consistency among the generated products. Moreover, SAR systems provide a variety of imaging modes/bands—which collect signals in different polarizations and scattering mechanisms—with different resolutions and extensions covered. Fine-resolution SAR modes deliver highly detailed images over narrow areas, while wide-swath modes cover larger regions at slightly reduced resolution. Further, some bands can penetrate certain media, such as vegetation, or interact with the surface in different ways providing information about the topography, land cover, surface deformation, and texture. This flexibility makes SAR well-suited for applications like land use and land cover mapping, environmental monitoring, and maritime surveillance (Moreira et al. 2013).

2.5.2. Vessel detection

The aforementioned unique capabilities make SAR a highly effective sensor for ship detection. Several key factors explain the widespread use of these images in vessel detection. First, SAR can achieve resolutions that match the size of ships—except for small vessels. Second, SAR images cover relatively wide areas while maintaining a constant resolution, enabling efficient coverage of large maritime regions. Third, SAR is not reliant on daylight or cloud cover, making it operational regardless of environmental conditions (Kanjir, Greidanus, and Oštir 2018). Fourth, ships, especially larger ones, are predominantly constructed with metallic materials that exhibit strong radar signal reflections. As a result, ships appear as bright objects in SAR images, facilitating their detection on open waters.

Moreover, the accessibility of SAR images has significantly contributed to its widespread utilization. Since the early 1990s, many SAR systems have been deployed in orbit around the Earth. Notable examples include ESA's Sentinel 1 and COSMO-SkyMed, JAXA's PALSAR, and the latest commercial SAR sensors developed by Capella Space.

However, SAR ship detection has certain limitations. Radar images inherently suffer from noise which can impact the accuracy of ship detection. High wind and sea state conditions also compromise ship detection, limiting its effectiveness. Additionally, detecting small targets, false positives identification, and ship classification remain difficult using SAR imagery.

Despite the mentioned drawbacks, the main reason ship detection from satellite SAR is inadequate is its limited spatiotemporal coverage. Even though the amount of available SAR images has risen in recent years, the current number still needs to meet the requirements for sufficient coverage, as a single medium-resolution SAR scene covers less than 0.1% of the ocean. To achieve comprehensive coverage—at a similar level to the one provided by AIS and VMS—a significantly larger fleet of SAR satellites, potentially numbering in the hundreds, would be necessary, but this is unlikely in the short term due to the high cost associated with SAR products. Another major challenge lies in linking SAR detections to specific broadcasting vessels, making it difficult to distinguish between detections corresponding to broadcasting and non-broadcasting vessels. Consequently, this poses challenges in creating labeled datasets for training machine learning algorithms and validating models. (David Allen Kroodsma et al. 2022).

2.5.3. Machine Learning approaches

While vessels can be discerned on SAR images with the naked eye, it is not feasible to assess vast global areas continuously over time solely relying on human observation. Therefore, several automated approaches have been developed to detect ships, encompassing both traditional-based and ML methods, accounting for 59 different techniques (Yasir et al., 2023). While the Constant False Alarm Rate (CFAR) method stands out among the traditional-based approaches (Hou, Chen, and Jiao 2015; Marzuki et al. 2021), ML has emerged as the predominant choice in ship detection.

Among ML approaches, shallow architectures such as Support Vector Machines (SVM) have been proposed for vessel detection using satellite images (Hwang and Jung 2018; H. Li and Wang 2008). However, while these techniques offer low computational costs and timely results, they are often associated with lower accuracy rates and require larger amounts of data for effective training and testing. On the other hand, despite the higher computational times required, Deep Learning (DL) architectures have demonstrated superior performance in achieving high recognition rates even with relatively small amounts of data.

DL techniques, specifically Convolutional Neural Networks (CNN), have been extensively used in computer vision tasks. Many articles have been published on SAR images DL-based object detection, categorizing the techniques into one-stage and two-stage methods. One-stage methods treat object detection as a regression problem, directly generating bounding box coordinates and class probabilities from image pixels. Prominent examples of one-stage methods for vessel detection include YOLO (Ting et al. 2021), SSD (Liu et al. 2016), and RetinaNet (Wang et al. 2019a). In contrast, two-stage methods generate region proposals as potential bounding boxes and then employ a classifier to categorize them. Popular two-stage methods include Fast R-CNN, Faster R-CNN, Mask R-CNN, (J. Li, Qu, and Shao 2017; Lin et al. 2019; Ren et al. 2017; Zhang et al. 2020). While single-stage algorithms are simpler and faster to train, they may exhibit lower accuracy compared to two-stage techniques (Yasir et al. 2023).

2.5.4. Fishing monitoring

Although AIS and VMS have significantly enhanced the monitoring of industrial fishing vessels, we only have reliable estimates of stock status for fish populations accounting for approximately half of the global catch, and our knowledge of the state of the majority of the world's “unassessed” fish stocks remains highly uncertain (Ovando et al. 2021). The inconsistent implementation of AIS and VMS across regions and fleets, with larger and wealthier nations more likely to adopt these technologies, contributes to this uncertainty. Additionally, the ability of captains to deactivate these systems and operate undetected poses a significant challenge in accurately estimating fishing activity in specific areas, as it fails to consider non-broadcasting vessels. This is where SAR images can play a crucial role in addressing this issue.

The availability of satellite imagery and the advancements in ML have paved the way for a new era in monitoring marine industries, closing the VMS and AIS data gaps and enabling increased transparency in ocean activities. The use of these technologies in fishing monitoring has progressed through various stages, starting with the development of accurate models to detect marine structures and vessels, followed by the discrimination of fishing vessels, and ultimately, with the integration of SAR-derived information with VMS and AIS data outputs to provide a more comprehensive and accurate picture of fishing activities (Galdelli et al. 2021).

As discussed in the preceding section, extensive research has been conducted on vessel detection. However, when explicitly assessing fisheries, the detection results are further analyzed to classify the identified vessels into different fishing categories. Several methods have been proposed for vessel classification to discriminate fishing vessels, such as K-Nearest Neighbour (KNN) (Young 2019; Sasamal and Mallenahalli 2019), Random Forest (RF) (Snapir, Waine, and Biermann 2019), Support Vector Machines (SVM), and C4.5 classifiers (Sasamal and Mallenahalli 2019).

Once fishing vessels have been distinguished from other types of vessels, it is crucial to assess whether certain ships are involved in illegal fishing activities. This evaluation involves comparing AIS data with SAR detections. By establishing associations between both data sources, we can identify ships that are not transmitting their positions, whether intentionally or unintentionally, and detect potential suspicious behaviors (Galdelli et al. 2021; Park et al. 2020; David Allen Kroodsma et al. 2022; Young 2019; Paolo et al. 2022). This aspect is essential for evaluating the impact of unreported fishing on fisheries assessments.

These approaches are revolutionizing fisheries monitoring and have provided significant advancements to accurately estimate fishing stocks, species populations, and the status of fisheries based on gear types. Existing research has demonstrated its effectiveness in various aspects, such as estimating the actual size of fleets that cannot be identified exclusively through AIS, monitoring fleets at a global resolution (David Allen Kroodsma et al. 2022), estimating captures of target species, tracking changes in fleet size over time, and addressing data gaps associated with nearby fisheries targeting the same species (Park et al. 2020). These advancements provide evidence of the transformative impact of these approaches, enabling more accurate assessments and insights into fishery science.

Nonetheless, there are limitations associated with SAR technology. For instance, the low revisit rate hampers the reconstruction of continuous vessel tracks necessary for detecting fishing behavior over time. Additionally, SAR detections of dark ships—non-broadcasting vessels—can only be classified as potential fishing vessels, lacking the named identity required to compare them with known illegal, unreported, and unregulated (IUU) fishing lists (Young 2019). To tackle these challenges, existing research combines AIS, SAR, and other complementary methodologies, to overcome the limitations of individual approaches and enhance the accuracy and effectiveness of detecting and monitoring fishing vessels.

2.5.5. Relevant research examples

Kroodsma et al. (2022): made several significant findings regarding fisheries monitoring using AIS and medium-resolution SAR. Firstly, they discovered that by combining these two data sources, they could estimate the total number of vessels operating, including those not detected by AIS or SAR alone (i.e., real fleet size). Secondly, the study enabled them to determine the actual footprint of longline activity in the Pacific and Indian Oceans, where non-broadcasting and broadcasting vessels showed similar spatial distributions. Specifically, in the Pacific, most vessel activity occurred in the high seas, with minimal activity in the French Polynesian waters or Kiribati. This suggests that no significant "dark fleets" operates inside Economic Exclusive Zones (EEZ). Similar findings were observed in the Indian Ocean. Additionally, the authors' methodology allowed to determine the lengths of non-broadcasting vessels. This information is valuable as vessel size correlates with fishing effort, with larger vessels typically catching more fish. Establishing proxies for fishing efforts using SAR images and correlating them with actual measurements of fishing effort, such as those derived from AIS or landings, poses an ongoing research challenge. Finally, the study identified non-compliant vessels regarding AIS regulations with important policy enforcement implications. Surprisingly, the study revealed that the lengths of non-broadcasting ships were higher than anticipated. Notably, a significant number of vessels over 40 meters in length were not broadcasting AIS signals. Previous research had suggested that nearly all fishing vessels of this size globally were equipped with AIS devices. These findings suggest that either these larger vessels are intentionally disabling their AIS or that there are more non-compliant vessels without AIS than previously acknowledged.

Park et al. (2020): used four satellite technologies—AIS, SAR, the Visible Infrared Imaging Radiometer Suite (VIIRS) sensor, and high-resolution optical imagery—to monitor fishing activity and quantify changes over time. The researchers leveraged the strengths of each technology to mitigate their respective limitations. AIS provides detailed movement and identity information but is utilized by only a fraction of vessels. SAR can identify large metal vessels and penetrate clouds but lacks global coverage. VIIRS sensor has a higher daily global revisit time and can detect ships with bright lights but can only operate at night. High-resolution optical imagery provides the best visual confirmation of vessel activity but is limited by clouds. First, the authors established the actual fleet size and observed its evolution over time, identifying seasonality in fishing activities. Additionally, they estimated an increase in fishing days over the past four years. Moreover, they could differentiate flag vessels based on their distinctive signals. They found that the significant growth in the fleet was primarily attributed to Chinese vessels, which used brighter lighting power, and engaged in pair trawling. Secondly, they estimated the captures for Chinese dark vessels by assuming a catch per unit effort (CPUE) similar to those in nearby waters. The analysis revealed that the squid fisheries under surveillance amounted to approximately 101,300 metric tons of squid worth $275 million in 2017 and 62,800 metric tons of squid worth $171 million in 2018. Furthermore, the researchers were able to establish a correlation between the previously unseen fishing efforts by Chinese vessels and the declining CPUE of South Korean and Japanese fisheries that also target the same stock in nearby waters—for which there is enough data for traditional fish stock and ecology assessments. Lastly, the study was able to map a shift in fishing efforts towards neighboring Russian waters due to the Chinese fleet operating in North Korean waters, where larger trawlers displaced local small-scale fishing boats. Overall, the study underscored the significant challenge presented by the existence of unmonitored vessels, which not only affects stock management but also contributes to escalating tensions surrounding resource sovereignty conflicts.

3. Methods

3.1. Data

The data set selected for training and testing the model is the Large-Scale SAR Ship Detection Dataset v1.0 (LS-SSDD-v1.0) created by Zhang et al. (2020). We selected this dataset due to its exclusive reliance on Sentinel-1 images, encompassing entire regional images featuring multiple annotated vessels, rather than individual offshore vessel chips as in other available datasets (Wang et al. 2019; Wei et al. 2020). This property likely enhanced the model's suitability for detecting vessels on actual Sentinel-1 data for the case study.

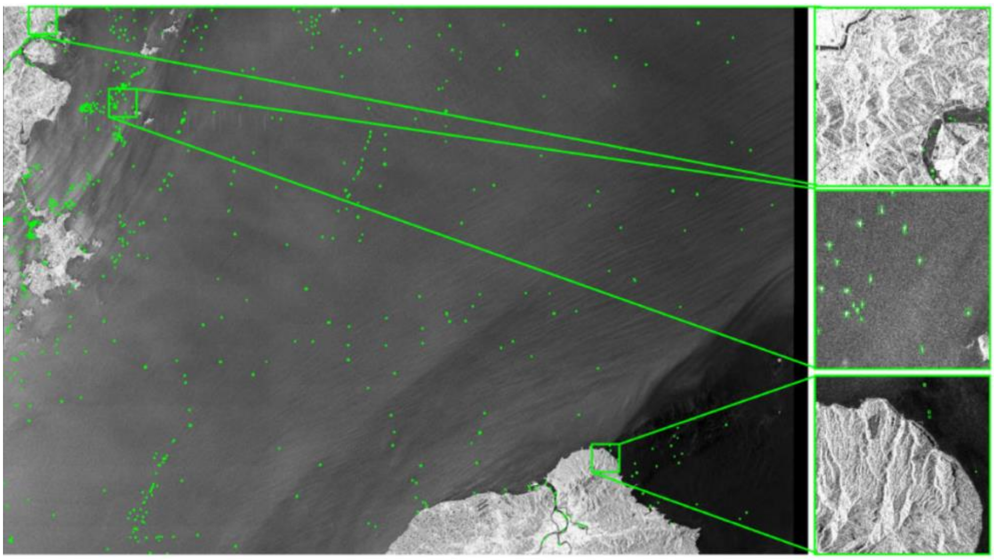

Figure 2. Example of a large-scale image in LS-SSDD-v1.0, extracted from Zhang et al. (2020). Large-scale image (left) and 800x800 sub-images (right). Bounding boxes from annotations are in green.

The LS-SSDD-v1.0 dataset consists of 15 large-scale SAR images from the Sentinel-1. These images have dimensions of 24000 × 16000 pixels. The authors divided the large-scale images into smaller sub-images for training and evaluation. VV polarization band sub-images were downloaded, each measuring 800 × 800 pixels. The vessel annotations in this dataset were available in .xml format, containing the bounding boxes of the vessels. We excluded sub-images without annotated vessels to optimize the training process, resulting in a remaining set of 1859 images for training and testing. The total size of the dataset was 536.4 MB.

3.2. Data environment

Python was the coding language chosen for this project. We decided to run our training and analysis in Google Colab to simplify the setup process and avoid potential environment issues. This workspace offers a pre-configured Python environment, seamlessly integrating with Google Drive. Additionally, it provides access to free computational resources, including GPUs. These factors made Colab a convenient choice for our implementation, ensuring a smooth workflow throughout the project.

We conducted a comprehensive review of different techniques used in constructing Faster RCNN models to determine the most suitable platform for object detection and segmentation. After careful consideration, we chose Detectron2, developed by Facebook AI Research. This library, implemented in PyTorch, stands out for its extensive tutorials, and other resources demonstrating its successful application to satellite imagery. Furthermore, Detectron2 offers state-of-the-art algorithms for object detection and provides a diverse collection of pre-trained baseline models. These pre-trained models served as a baseline for training our own model, enabling us to save computational time while potentially enhancing overall performance.

Additional libraries were used for the training and model evaluation and for the data transformation and representation in the case study. The details of these libraries are included in the Jupyter notebooks associated with this project. Some notable libraries used include torch, torchvision, rasterio, pyproj, geopandas, and matplotlib.

3.3. Data pre-processing

Data preprocessing involved several tasks. As mentioned earlier, we initially removed images without annotated vessels from the original dataset to optimize computational resources. Following the criteria outlined by Zhang et al. (2020), the dataset was split into training and testing sets, resulting in 1395 images for training and 464 for model evaluation. Lastly, the datasets were registered within Detectron2, specifying the metadata such as annotation locations, class names, and annotation formats for the images. These steps ensured that the dataset was appropriately prepared and organized for subsequent model training and evaluation.

3.4. Model training

The Faster R-CNN object detection model was built using the Detectron2 library. For our model, we adopted the architecture and pre-trained baseline model, namely "faster_rcnn_R_101_FPN_3x". Based on the literature review, we selected this configuration from the Detectron2 Model Zoo, which indicates that Faster R-CNN with Feature Pyramid Network (FPN) architecture achieves superior performance for object detection tasks involving SAR images. Multiple training attempts were made, each testing different numbers of iterations. Ultimately, the base model was constructed using 26000 iterations, approaching the limit of the GPU free computational resources available in Colab.

Due to computational constraints, we could not allocate a separate validation dataset for assessing our model's performance during training. Consequently, the learning curves generated from our model object were derived from metrics obtained from the training dataset (Figure 3). While these learning curves provide useful insights into the model's learning and adaptation over iterations, it is important to acknowledge that they do not fully reflect the model's ability to generalize to unseen data. Hence, the initial testing phase was carried out through a trial and error approach to determine the optimal number of iterations and model configuration. The resulting testing models were evaluated on the test dataset described in the next section.

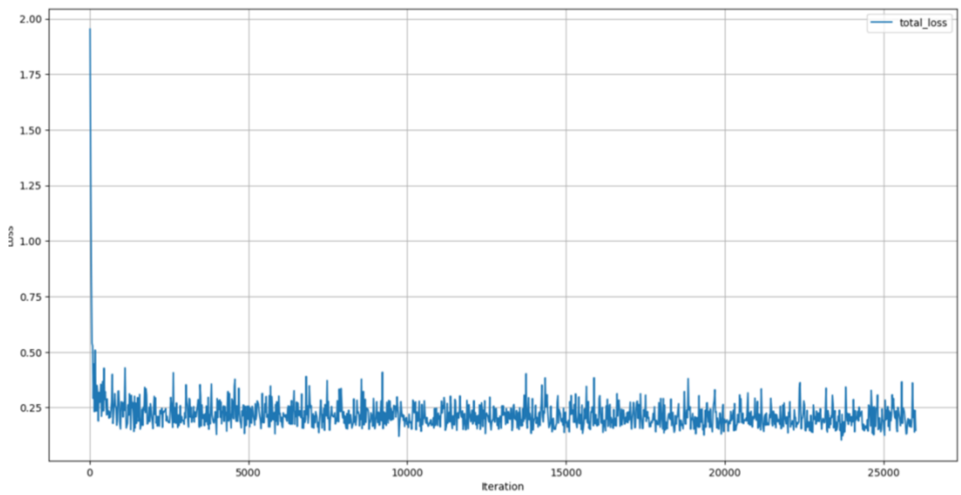

Figure 3. Learning curve depicting the loss over time for the base model, as measured from the training dataset. Additional learning curves can be found in the respective Jupyter notebooks.

3.5. Model evaluation

In this analysis, we evaluated the performance on the entire testing dataset, as well as on two subsets. One subset contained only offshore images—with no coastline, showcasing areas located in the open ocean— and the other consisted exclusively of inshore images, which included land areas.

To assess the performance of the object detection model, we used the Common Objects in Context (COCO) Evaluator, a tool available in Detectron2 that calculates standard metrics. The default metrics provided by this evaluation function include Average Precision (AP), representing the mean precision of the detector at various recall levels, and object size-based metrics like APs, APm, and APl. These measure the Average Precision for small, medium, and large objects, respectively. However, since vessels are exclusively represented as small objects in SAR images, we have omitted them in our results and discussion section. Additionally, the evaluator provides average precision metrics based on different Intersection over Union (IoU) values, such as AP_{50} and AP_{75}, for 50% and 75% IoU respectively. These metrics were valuable to our analysis, as the performance evaluation criteria proposed by Zhang et al. (2020) chose a score threshold of 0.5 and an IoU threshold of 0.5. This means that for a detection to be considered successful, the overlapping area of a predicted box and a ground truth box must be greater than or equal to 50%. Furthermore, we created a custom model evaluator to extract the number of true and false positives to calculate recall and precision metrics.

3.6. Model improvement

Various image preprocessing techniques were explored to improve the model's performance. Overall, a series of transformations were designed to augment the input data with multiple changes to the image, including resizing, flipping, and adjusting brightness, contrast, and saturation. Augmentation techniques help improve the performance of computer vision models by exposing them to a wider variety of training data so they can generalize better on new images and prevent overfitting.

The following transformations were applied:

- Image flipping: Images were subject to random horizontal and vertical flipping during augmentation. The probability of flipping occurrence was set to 0.5—not all images were transformed due to computational constraints. It is important to note that these transformations can potentially result in the loss of certain image portions. Resizing the image to a larger size helped to recover those portions.
- Image resizing: images were resized from their original size of 800x800 pixels to 1400x1400 pixels. By resizing the image, the model can see more details and finer-grained features that might not have been seen before. This can help the model recognize objects and patterns more accurately and robustly, improving its performance on new, unseen images.
- Image tone: Image tone adjustments, including brightness, contrast, and saturation properties, were randomly applied. Each adjustment was performed with a probability of 0.3 by a factor ranging between 0.7 and 1.2.

The specific values used for the sequence of transformations were incorporated into a mapper function, which was subsequently integrated into the Detectron2 training process. The other parameters and training procedures for the improved model were the same as those used for our base model. We chose not to adjust the hyperparameters due to the limited computational resources and time constraints imposed by the available free GPU resources in Colab. Given these constraints, we prioritized data augmentation over hyperparameter tuning, as the former was more likely to yield a comparable or greater improvement in performance, considering our limited computational resources. Finally, the performance assessment followed the same procedure as the one for the baseline model.

3.7. Case study and interactive report

The project's final step involved the creation of an interactive report to showcase the potential use of the model with real-world Sentinel-1 data. To do so, a study location was defined, and yearly imagery was processed.

In order to assess vessel location and generate density maps of vessel presence, the bounding boxes resulting from the model's predictions were transformed into point coordinates. A code pipeline was created to facilitate this process and transform the model's outputs into meaningful spatial data. This encompassed the extraction of Sentinel-1 data from the Copernicus ESA Program, the reprojection of the data, the application of the detection model, and the subsequent transformation of the model's outputs. We used the Google Earth Engine (GEE) Python API to access the data from the Copernicus ESA Program. GEE offered a convenient platform for accessing geospatial data, providing increased access and automation capabilities. Its integration with existing Google tools used in this project, such as Drive and Colab, further streamlined the workflow and enhanced efficiency.

The following summarizes the steps taken:

- <u>Extracting real-world Sentinel-1 data</u>: Access Sentinel-1 band VV data and automate the download process using the GEE API. The automation of data downloads was done using Drive and Colab.
- <u>Validating image shape and metadata</u>: Evaluate the extracted Sentinel-1 images to ensure that their shape and contents were valid. This step allowed us to verify the integrity of the data before further processing.
- <u>Reprojecting the data</u>: Some images had different Coordinate Reference Systems (CRS). Therefore, it was necessary to reproject the extracted data to a desired CRS. This ensured consistency and compatibility among spatial layers.
- <u>Applying the detection model</u>: Once the images were ready, we run the detection model to generate the predictions.
- <u>Transforming model outputs into geographic data</u>: A code was developed to convert the model outputs, specifically the bounding boxes, into geographic data. This transformation process entailed the transformation of the bounding boxes into point coordinates by calculating the centroids.
- <u>Data visualization</u>: We used the generated data to create different visual outputs, including visualizing the bounding boxes on original images, vessel location coordinates, density maps of vessel presence, and plots of vessel detections over time. These outputs provided valuable insights into vessel activity in the study area, supporting the interactive report on the fisheries case study.

The resulting spatial data was used to describe the fishing status of the study area and demonstrate the contributions of technologies like SAR to its assessment. Finally, the information was gathered and included in an HTML interactive report—separate from the final report—using Quarto. Quarto functionalities allow the reader to interact with some of the outputs, enhancing the understanding of the study case and the data displayed.

4. Results

4.1. Base model

The results from our object detection model, using the Detectron2 framework, reveal mixed performance across different categories (Table 4). Overall, for the entire testing set, the model achieved a recall of 79.77% and a precision of 57.69%, indicating that it correctly identified approximately 80% of the annotated vessels, but only approximately 58% of the vessels identified were actually so. This suggests a relatively high rate of false positives in the model's predictions, with 1391 false positive instances against 1897 true positives. The AP_{50} and AP_{75} scores, measuring the average precision at different IoU thresholds, were 73.67% and 10.72% respectively. These metrics show that while the model has good precision at a 50% IoU, its precision drops significantly when the threshold is increased to 75%.

Table 4. Base model performance metrics for the entire test dataset and the subsets.

	TP	FP	Recall	Precision	AP_{50}	AP_{75}
all	1897	1391	79.77%	57.69%	73.67%	10.72%
offshore	1371	445	91.77%	75.49%	87.34%	14.22%
inshore	525	946	59.46%	35.69%	46.39%	4.08%

When evaluating performance on the offshore subset, the model displayed substantially improved results, with a recall of 91.77% and a precision of 75.49%. This means it was able to correctly identify about 92% of vessels in this subset, with about 75% correct identifications. The model exhibited fewer false positives (445) in the offshore environment compared to the inshore, while detecting more true positives (1371). The AP_{50} and AP_{75} values for this subset are also significantly better than the overall scores, standing at 87.34% and 14.22%, respectively.

For the inshore subset, however, the model showed lower performance. The recall fell to 59.46%, meaning the model correctly identified just under 60% existing vessels. Precision dropped even more drastically to 35.69%, indicating that most of the vessels detected where false positives—946 false positives and only 525 true positives. The AP_{50} (46.39%) and AP_{75} (4.08%) values for this subset were considerably lower than the overall and offshore subsets.

Overall, the base model showed a remarkable performance in the offshore category but needs to be improved in precision and inshore category detections.

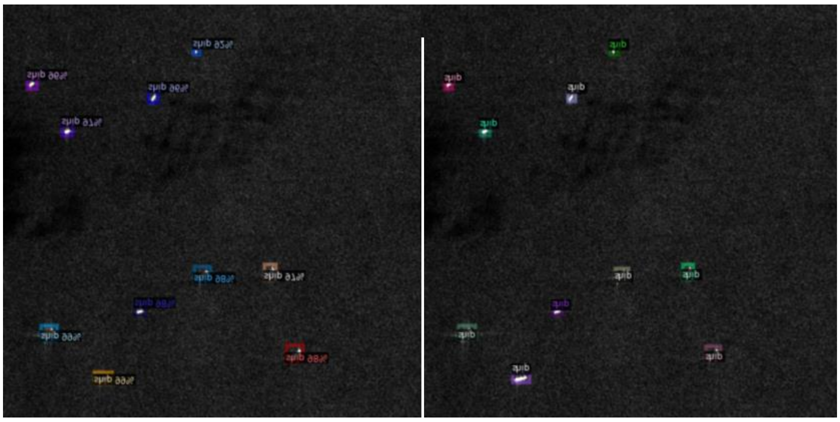

Figure 4. Detections made by the base model (left) and the corresponding original annotations (right) for an offshore image.

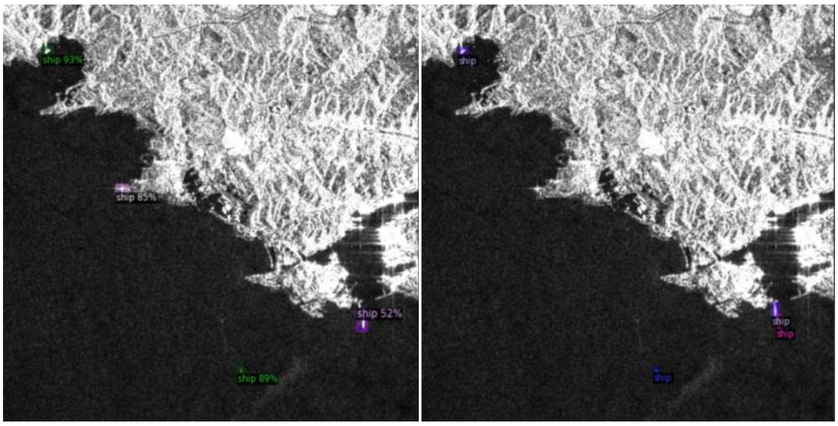

Figure 5. Detections made by the base model (left) and the corresponding original annotations (right) for an inshore image.

4.2. Improved model

After applying image augmentation to train our model, we observed some noticeable shifts in the model's performance. For the entire testing set, the model's recall increased to 84.99% from 79.77%, suggesting an improvement in its ability to identify vessels correctly—the number of true positives increased to 2021 from 1897. However, the model's precision decreased from 57.69% to 46.69%, showing a higher proportion of false positives (2308) among the objects identified as vessels. This suggests that while the model has become more capable of identifying vessels, it has also misclassified a larger number of objects as ships. Relative to the average precisions, the AP_{50} increased slightly to 76.92%, improving at a 50% IoU threshold. The AP_{75}, however, decreased to 8.32%.

Table 5. Improved model performance metrics for the entire test data and the subsets.

	TP	FP	recall	precision	AP_{50}	AP_{75}
all	2021	2308	84.99%	46.69%	76.92%	8.32%
offshore	1392	633	93.17%	68.74%	88.61%	9.79%
inshore	628	1675	71.12%	27.27%	53.81%	5.72%

For the offshore subset, both recall and the number of true positives slightly increased, while precision slightly dropped. This suggests a similar trend to the overall set, where the model has improved in identifying vessels (with 1392 true positives up from 1371) but also falsely identified more irrelevant objects (633 false positives up from 445). The AP_{50} and AP_{75} scores also mirrored the overall set with an increase in AP_{50} (88.61% up from 87.37%) and a decrease in AP_{75} (9.79%, down from 14.22%).

In the inshore subset, the model's recall improved significantly to 71.12% from 59.46%, with an increase in true positives from 525 to 628. However, precision decreased drastically to 27.27%. As with the overall set and offshore subset, the AP_{50} for the inshore subset increased from 46.39% to 53.81%, while the AP_{75} dropped to 5.72% from 4.08%.

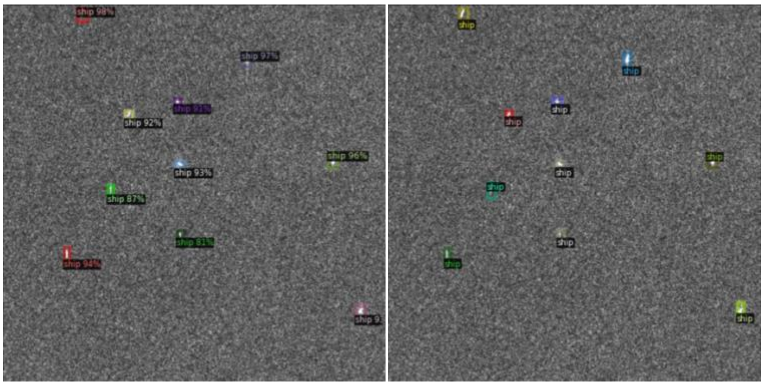

Figure 6. Detections made by the improved model (left) and the corresponding original annotations (right) for an offshore image.

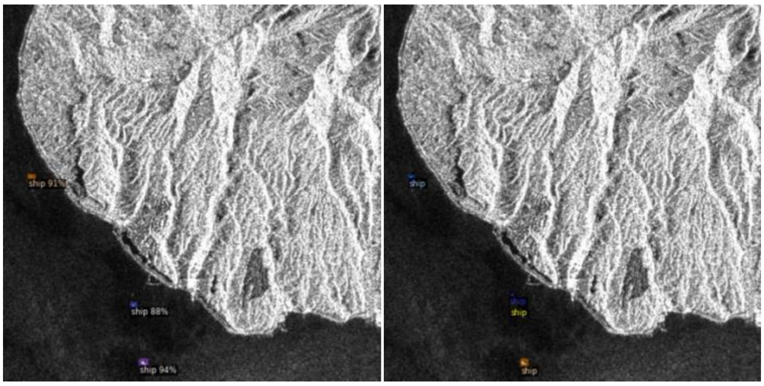

Figure 7. Detections made by the improved model (left) and the corresponding original annotations (right) for an inshore image.

Overall, applying image augmentation techniques has led to trade-offs between different performance metrics. While recall and AP_{50} improved, precision and AP_{75} decreased. Therefore, further refinement of the augmentation strategy might be required. This might include fine-tuning the augmentation parameters or employing selective augmentation to balance the improvement in recall with the preservation of precision and accurate bounding box delineation.

4.3. Interactive report

To illustrate the application of our model on real-world Sentinel-1 data, we have developed an HTML interactive report that showcases the use of these methods in a fisheries-related case study. While section 6 of the report provides comprehensive details on the case study, this section will highlight key elements incorporated in the interactive report.

The HTML allows readers to explore the case study while engaging with the displayed data. Several maps and plots featured in the case study section are accessible in an interactive way within the report. Furthermore, readers can unfold the code that transforms the model outputs into visual representations.

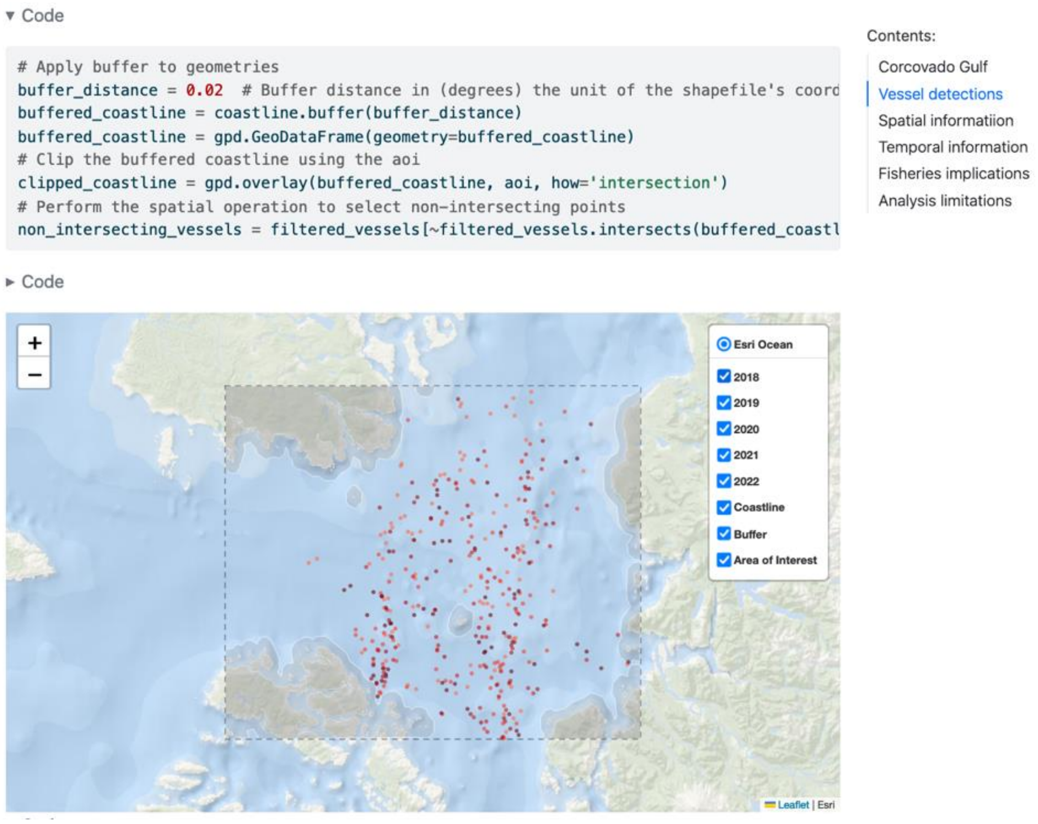

Figure 8. Screenshot of a section of the interactive report, displaying an unfolded code chunk and a map showcasing the vessels detected in the case study.

Among other features, the HTML report allows the reader to zoom in and out of maps and select specific spatial layers of interest (Figure 8). Furthermore, it enables the evaluation of spatial data over time, providing the ability to analyze temporal patterns and changes (Figure 9 and Figure 10). Additionally, the report allows for selecting specific data from the plots, facilitating a detailed representation of the desired information (Figure 11).

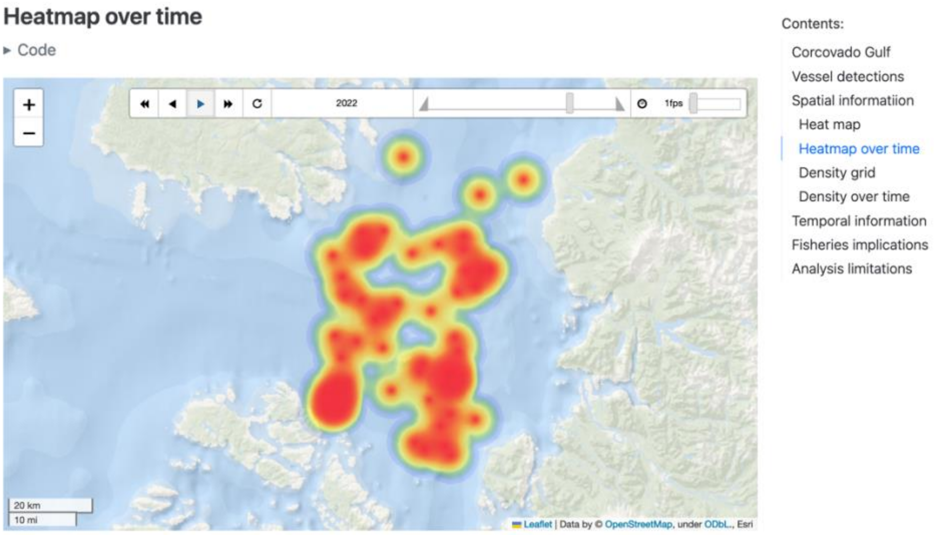

Figure 9. Screenshot of a section of the interactive report, displaying a time-series heatmap of vessel presence.

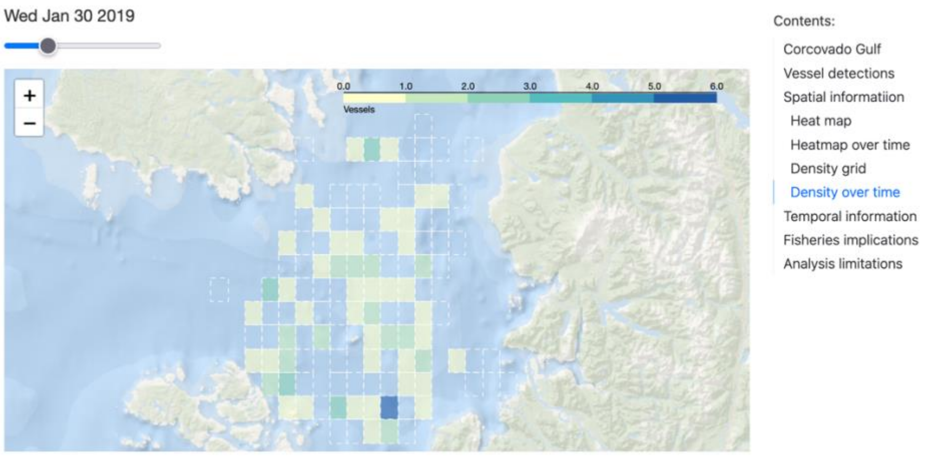

Figure 10. Screenshot of a section of the interactive report, showcasing vessel presence per cell grid, with the option to select the date of interest.

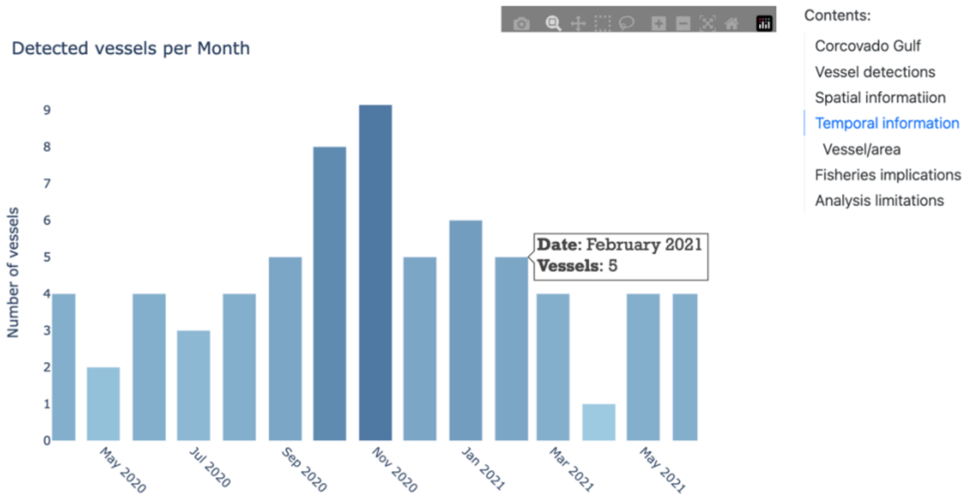

Figure 11. Screenshot of a section of the interactive report, presenting a plot depicting the number of vessels detected over time. The plot lets users zoom in and out on specific date ranges and select the desired data.

Overall, the presentation of our findings in an accessible and interactive format allowed us to create a science outreach document that effectively enhances awareness and understanding of the application of SAR imagery in maritime surveillance and related fields.

5. Discussion

The application of image augmentation techniques to our object detection model, based on the Detectron2 framework, resulted in significant trade-offs across various performance metrics. The most notable improvement was the increased recall across all categories, particularly inshore. This suggests that the augmentation has helped the model generalize better and capture more true positives. However, the precision dropped in all categories, indicating increased false positives. This might be due to the model overgeneralizing from the diversified training data after augmentation. Further, while the model demonstrated better object presence detection, as indicated by the increased AP_{50}, it showed reduced accuracy in bounding box delineation, reflected in the decreased AP_{75}. However, given the goal of the model, the low AP_{75} values are not a concern. Reduced accuracy in the bounding box delineation would only translate into a small change in the vessel coordinate position once the centroids are calculated to establish their location. Hence, low values of AP_{75} do not necessarily mean that the model is performing poorly for our purposes.

Despite some of these trade-offs, our performance values are similar to those of Zhang et al. (2020) for the same dataset and training-testing split (Table 6). In terms of recall, both our base and improved models demonstrate superior performance compared to Zhang et al. Our improved model, in particular, significantly outperforms on offshore and inshore images, suggesting that it identifies vessels in the dataset better. However, regarding precision, our models show a different trend. While Zhang et al.'s model demonstrated a balance between recall and precision, our models struggled with precision. Both the base and improved models showed lower precision across all scenarios compared to Zhang et al., with the drop in precision being especially significant in the improved model. This implies that while our models are better at identifying vessels, they also tend to incorrectly classify non-vessel objects as vessels, which could be challenging in practical applications. In fact, this motivated the use of the base model for the case study instead of the improved one.

As for the mAP/AP_{50} values, the models show a mixed pattern. For the entire dataset and offshore images, our improved model slightly outperforms Zhang et al. However, in the more challenging inshore images where the smoothness of the ocean signal is mixed with the changes in SAR reflectance of the land, both our base and improved models fall behind Zhang et al.'s model. It is important to note that the mAP values from Zhang et al. should not be confused with the Coco Evaluator's APm, which is not included in this analysis. As mentioned in the methods section, the mAP values represent the mean Average Precision, as Zhang et al (2020) defined for a 50% IoU—equivalent to our AP_{50}. On the other hand, the APm metric represents the Average Precision for medium-sized objects.

Table 6. Comparison of our results and those from Zhang et al. (2020).

		Recall	Precision	mAP/AP50
All	Zhang et al. (2020)	77.71%	73.74%	74.80%
	Base model	79.77%	57.69%	73.67%
	Improved model	84.99%	46.69%	76.92%
offshore	Zhang et al. (2020)	91.91%	82.82%	89.99%
	Base model	91.77%	75.49%	87.34%
	Improved model	93.17%	68.74%	88.61%
inshore	Zhang et al. (2020)	53.68%	55.96%	46.76%
	Base model	59.46%	35.69%	46.39%
	Improved model	71.12%	27.27%	53.81%

In summary, while our models demonstrate an improved ability to detect vessels (as evidenced by higher recall values), they fail in precision, leading to more false positives. Our improved model does show an increase in mAP/AP_{50} in the overall and offshore categories and improved recall in the inshore environment compared to our base model, which suggests that image augmentation has had a positive effect. However, these improvements are at the expense of precision, especially in the challenging inshore images, which is a clear area for further improvement. Comparison with Zhang et al.'s results highlights the importance of achieving a balance between recall and precision.

Moving forward, our focus should be on refining our augmentation strategy. However, it is important to acknowledge that data augmentation alone cannot fully address the model's ability to learn ship features based on limited datasets. Hence, expanding the number of labeled datasets should be a priority, as Zhang et al. (2020) indicated. Additionally, exploring other approaches, such as hyperparameter tuning—excluded in this study due to limited computational resources—is crucial to improve model performance.

6. Case Study

This section focuses on implementing our trained model on Sentinel 1 images from the ESA Copernicus program to test its usage on real-world SAR images. The specific marine region under examination is the Corcovado Gulf, located in southern Chile. We chose this area as a case study based on three criteria. Firstly, the Corcovado Gulf is known for active fishing activities, including artisanal and industrial fisheries. The commercial fishing of hake and spider crab is carried out using large fishing vessels with metallic superstructures and hulls, making their signature on SAR images easily distinguishable compared to smaller or wooden boats. Secondly, considering the limited maritime traffic in the area, except for the passenger lanes connecting Quellón to Guaitecas and Port Raul Marin Balmaceda, we can reasonably assume that the model's detections are primarily associated with fishing activity, eliminating the need for an additional classification model to differentiate between fishing and non-fishing vessels—which development was beyond the scope of this project. This allows us to extrapolate the detections to fishing vessels and qualitatively showcase the application of the model through a case study focused on fisheries. Lastly, the Corcovado region benefits from wide Sentinel 1 coverage, providing multiple complete and partial images of the study area on a monthly basis.

6.1. Corcovado Gulf

The area of interest (AOI) encompasses 6615.83 km^2 of sea and coastline at the Corcovado Gulf's entrance. The area is situated between two administrative fishing regions, the X Region of Los Lagos and the XI Region of Aysén, as designated by the Chilean National Service for Fishing and Aquaculture. It also includes a 1019 km^2 protected area, the Tictoc-Golfo Corcovado Marine Park, established in 2022 and located northeast of our AOI.

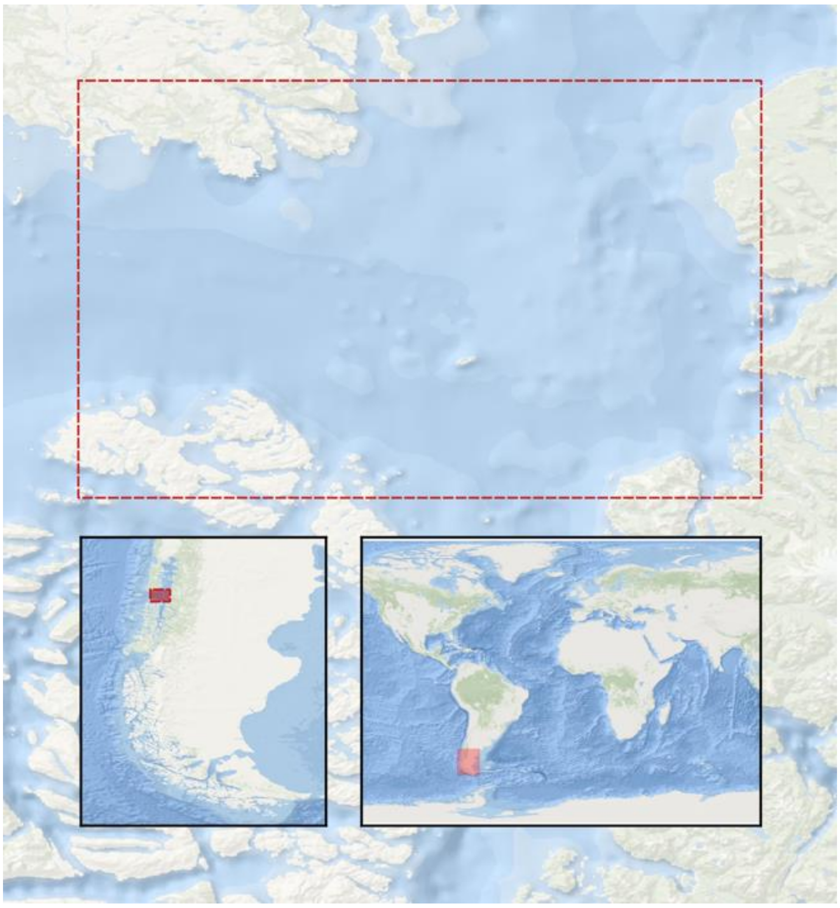

Figure 12. Study area.

Most of the fishing efforts in the AOI are directed towards nine benthic species, including clams, mussels, sea urchins, seaweeds, and crabs from the genus Cancer (Molinet et al. 2011). This activity is primarily carried out by boats ranging from 7 to 15 meters in length, with most of them exceeding the resolution requirements for their detection. However, the area also involves the fisheries of the southern hake (*Merluccius australis*), and southern king crab (*Lithodes santolla*), both from artisanal fishing boats with lengths up to 18 m (Molinet et al. 2020)— above SAR's resolution—, and industrial fishing vessels (Kitts et al. 2020; Molinet et al. 2019). These vessels, and the fishing carriers assisting artisanal vessels in transhipping operations, are the potential targets of the detection model.

Figure 13. Fishing vessel "Seines" (left) and fish carrier "Patagon 7" (right). Picture by Diego Muñoz and Max Peter Schlicher from MarineTraffic.com.

6.2. Vessel detections

A total of 348 Sentinel 1 images, amounting to 93GB of data, were extracted from the Google Earth Engine (GEE) and processed to evaluate a period from January 2018 to December 2022, covering 1024011 km^2. As mentioned in the methods section, the model outputs' bounding boxes were transformed into coordinate points representing the locations of the detected vessels. Each detection was assigned its detection score from the model and recorded with the corresponding date and time of the associated SAR image. Detections with scores below 0.6 were removed from the dataset to exclude potential false positives. Additionally, objects erroneously detected near the shoreline were excluded by clipping them out of a coast shapefile layer obtained from earthworks.stanford.edu. Furthermore, an exclusion buffer of 2km from the coastline was applied to include only vessels operating in open waters, where the fishing activities of interest concentrate. After that preprocessing, a total of 365 vessels were identified and selected for the analysis, with most detections occurring during the initial years of the specified timeframe (Table 7).

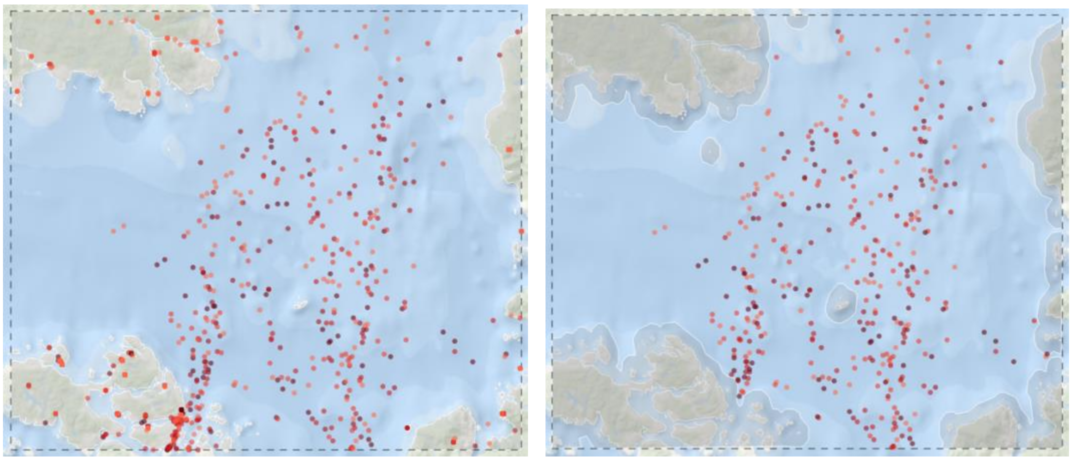

Figure 14. Total vessels detected, including false positives (left). Filtered detections (right). Colors representing different years from 2018 to 2022 (from lighter to darker).

Table 7. Number of vessels detected and selected for the analysis.

Year	2018	2019	2020	2021	2022	Total
Count	103	82	66	59	55	**365**

To assess the spatial distribution of the vessels, we generated a heatmap and a density map using a 5x5km grid. Attending to Figure 15, we can see that all fishing-related activities, including navigation, fishing, and transhipping, are predominantly concentrated in the southern region of our AOI, particularly in the adjacent waters of Guaitecas. Additionally, there is a noticeable absence of any activity on the east-northeast side, precisely where the Tictoc-Golfo Corcovado Marine Park is located.

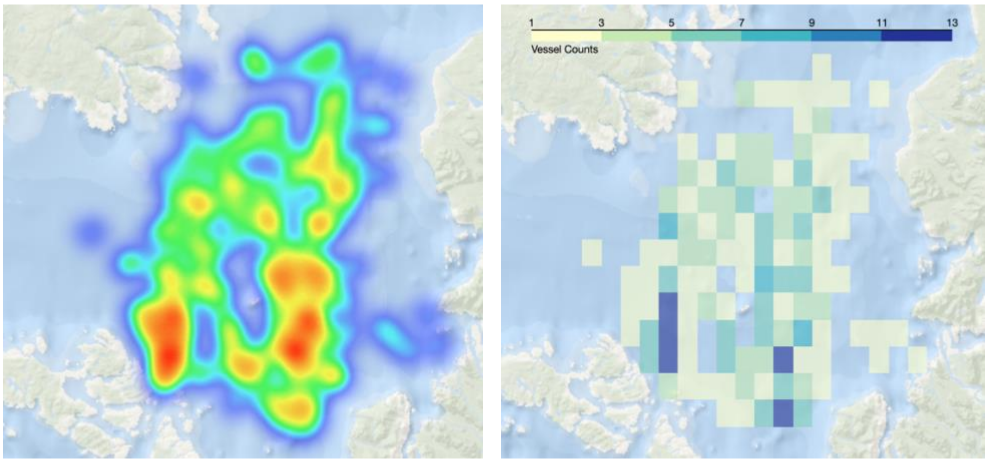

Figure 15. Heatmap and density grid of vessels detected.

Regarding the temporal distribution of the detections, Figure 16 illustrates the number of selected vessels per month. Consistent with the values in Table 7, the graph shows an overall downward trend in the number of detected vessels over the years. However, the variations in vessel counts can be explained by the changes in SAR image availability and area covered rather than actual changes in vessel presence. To address this and to derive meaningful conclusions from the temporal series, we can normalize the number of vessels based on the corresponding area covered in the images processed for each month (Figure 17).

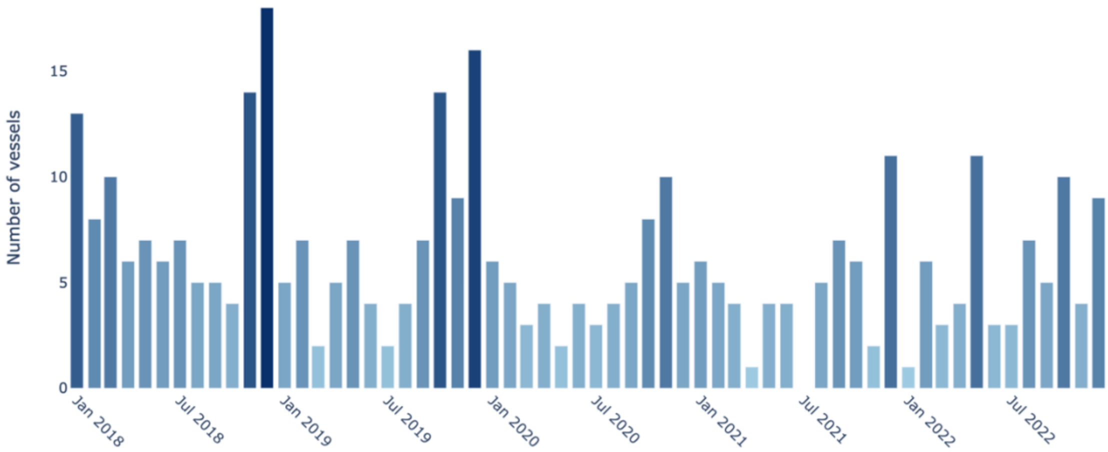

Figure 16. Time series of the number of vessels detected.

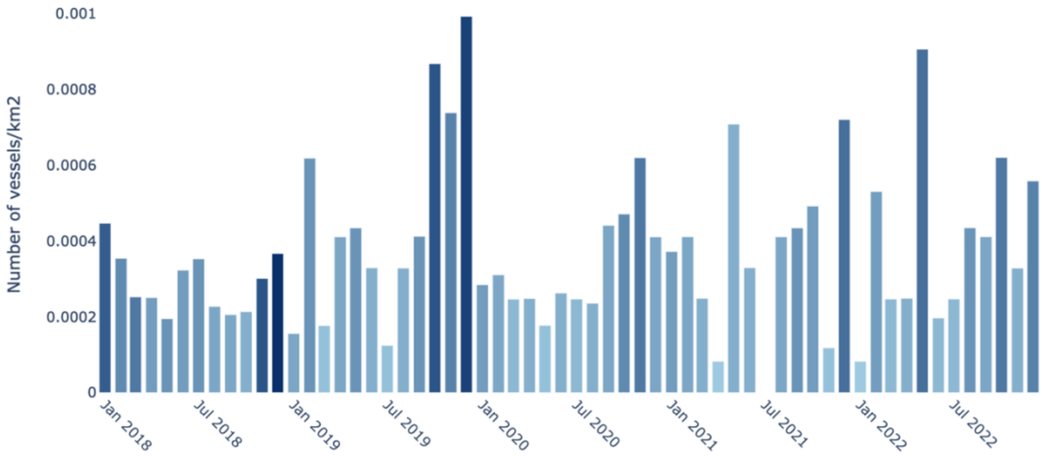

Figure 17. Normalized time series representing the number of vessels by total km2 covered by Sentinel 1.

6.3. Fisheries implications

Despite the limited scope of this analysis, some of its findings could potentially be used for describing the local fisheries. For instance, the spatial pattern detected could be attributed to the location of the primary fishing grounds, which would be situated north of Guaitecas and east of the Queitao islands in the central region of the AOI. This distribution at least highlights areas where fishing activity is not predominant and from which fishery enforcement efforts could be redirected towards higher vessel density areas.

Additionally, the absence of vessel presence within the waters inside the Marine Park is noteworthy. This observation sheds light on the rationale behind designating that specific area as protected in 2022. It is possible that the area's lower interest and reduced conflict among stakeholders facilitated or influenced its selection for protection.

Regarding the temporal distribution, a closer look at Figure 17, which represents the normalized time series, reveals that the previous downward trend is no longer evident. In this case, there isn't a noticeable pattern across the years. However, there are some anomalies in the number of detections, such as a significant increase at the end of 2019. It would be worthwhile to investigate whether this could be linked to a specific event related to fishing activities in the region, such as changes in fishing regulations like quota increases or an abnormality in the stock's growth due to ecological processes.

To further explore the temporal dimension of the data, we can examine if there is a seasonal factor. Referring to Figure 18, it is evident that most detections occurred at the years' ends, particularly in October, November, and December. Additionally, when evaluated by season (Figure 19), fall and winter accounted for the highest number of detections. This pattern could respond to the start of the fishing season.

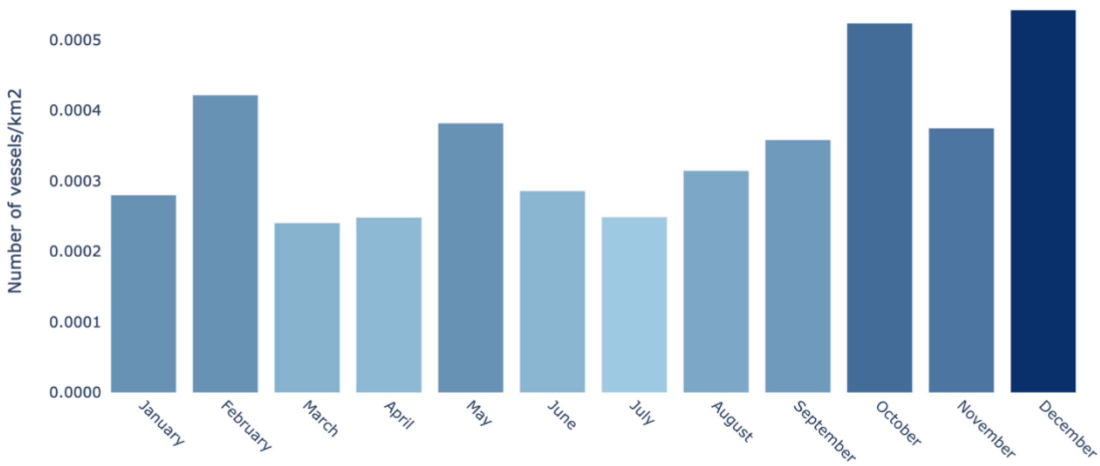

Figure 18. Data grouped on a monthly basis across all years.

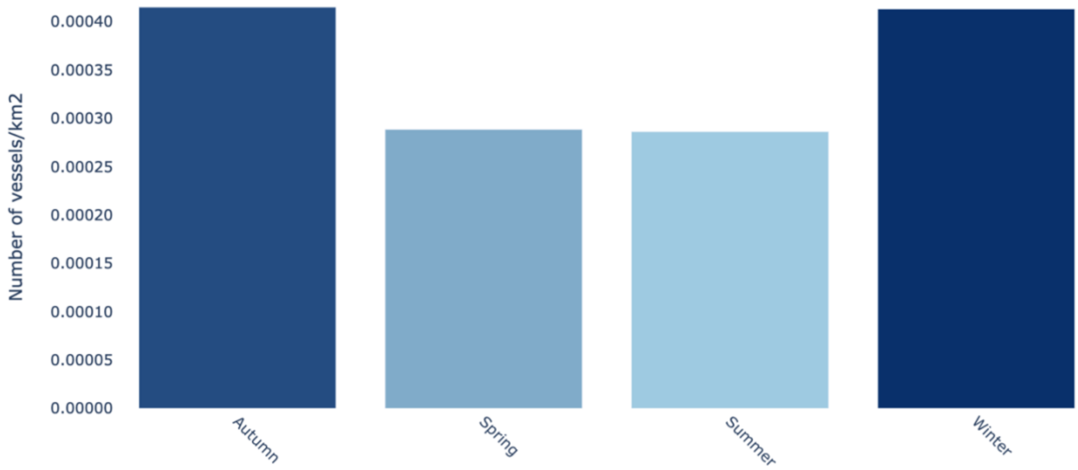

Figure 19. Data grouped on a seasonal basis across all years. Winter: December, January, February; Spring: March, April, May; Summer: June, July, August; Autumn: September, October, November.

6.4. Analysis limitations

It is important to emphasize that this case study does not attempt to analyze a fishery science case thoroughly. Instead, it aims to showcase the practical application of vessel detection models in fisheries management. Its main objective is to comment on the derived data outputs from these detections and establish correlations with existing scientific work that has successfully implemented this technology for fisheries management.

In this particular case, it is important to point out that the data cannot be directly extrapolated for fishing purposes. To achieve this, a classification model is necessary to accurately distinguish between fishing vessels and other types of boats. This differentiation is crucial as it enables the correlation of detections with actual fishing activity. Furthermore, a more comprehensive evaluation of the specific AOI and its fisheries should be conducted. This would allow for the formulation of hypotheses regarding the detections and their potential linkage to regional fishing trends. Additionally, it is essential to incorporate data from other sources to support the observations derived from the detection model. A more robust and comprehensive analysis could be achieved by combining traditional observations that define the regional fisheries with our detections.

While the analysis in this case study had a limited extent, focusing on data extraction, transformation, and representation, it has also served to reveal certain limitations of this technology, including the temporal resolution of SAR images, its constraint for detecting small vessels, and the need of a reliable classification model.

7. Conclusion

This study underscores the potential of machine learning models in detecting vessels within SAR images, offering a promising tool for enhanced fisheries management. We have successfully developed a robust DL model specifically designed for vessel detection, exhibiting outstanding performance on offshore images. Nevertheless, we have also identified several challenges and limitations associated with model training, particularly when it comes to recognizing vessels within inshore images, a common issue encountered in previous studies. While our case study demonstrates the effectiveness of techniques like land masking in improving the application of models to real SAR images, continuous efforts are crucial to refine these models further and bridge the performance gap between inshore and offshore images.

Furthermore, our accomplishments go beyond the development of the model itself, as we have successfully converted the model outputs into meaningful data with potential applications in fisheries management. In the case study, we have identified spatial and temporal patterns that may exhibit correlations with fisheries activity. Importantly, we have also highlighted the constraints of SAR-derived information for that purpose, particularly its limited temporal and spatial resolution. However, the capability of these models to detect vessels in situations where other methods fail is significant and emphasizes the importance of integrating this technology with existing fishing control and enforcement mechanisms.

Considering these findings, we argue that while machine learning models for SAR image analysis cannot replace traditional maritime surveillance techniques, they can enhance and complement existing methods for monitoring fishing activities. This integration has the potential to provide a more comprehensive, precise, and timely understanding of marine activities, thereby contributing to the achievement of SDGs 8 and 14.

7.1. Future research

The current project marks the initial stage of a broader approach to using SAR images in conjunction with machine learning methods to improve fisheries management. Based on the findings presented in this report, several promising steps must be pursued in this line of research.

- Improvement of the detection model, particularly focusing on detecting vessels in inshore images.
- Deriving vessel characteristics from the vessel image chips produced from the model's bounding box outputs, including vessel dimensions and course.
- Development of a classification model to distinguish between fishing and non-fishing vessels.
- Creating a workflow to integrate the SAR-derived information with VMS and AIS data outputs, thereby providing a more comprehensive and accurate picture of fishing activities.
- Exploring potential applications of the outputs for global and regional fishing assessments.

7.2. Project planning assessment

The goals set in the work plan were successfully achieved, with only minor deviations in terms of timing and content adjustments. Completing certain tasks, such as developing and running the code to execute the model on new images, transforming vessel detections into meaningful data, and showcasing its application in fisheries, took longer than expected and extended into project phase 3.

Regarding content, the LS-SSDD-v1.0 dataset (Zhang et al., 2020) was used instead of the initially proposed Wei et al. (2020) dataset. The former was selected since it relies exclusively on Sentinel-1 images and includes entire regional images with multiple annotated vessels, enhancing the model's suitability for detecting vessels on actual Sentinel-1 data. In Phase 2, GEE was used to access data instead of the Copernicus Open Access Hub as initially planned. GEE provided a more convenient platform with increased access and automation capabilities, making integrating with other tools such as Drive and Colab easier.

8. Glossary

Automatic Identification System (AIS): A tracking system that uses transponders on ships to track and monitor vessel movements. The AIS collects data, including the vessel's identity, type, position, course, speed, navigational status, and other safety-related information.

Catch per unit effort (CPUE): A measure of the success rate of catching fish per unit of fishing effort.

Catch quotas: Restrictions set by fisheries management on the amount of a specific species that can be caught within a certain timeframe.

CFAR (Constant False Alarm Rate): A traditional-based technique used in signal processing to detect target signals within background noise, commonly used in radar systems for ship detection.

CNNs (Convolutional Neural Networks): A class of deep learning models most commonly applied to analyzing visual imagery.

Deep Learning (DL): A subset of machine learning techniques that utilizes artificial neural networks with multiple layers between the input and output layer to model and understand complex patterns in datasets.

Economic Exclusive Zones (EEZ): Sea zones over which a state has special rights regarding the exploration and use of marine resources. They extend 200 nautical miles from the coast of the country.

Fast R-CNN: A two-stage deep learning model used for object detection, which improves on R-CNN by only requiring a single pass through the network, making it more efficient.

Faster R-CNN: A model that introduces a Region Proposal Network (RPN) that shares full-image convolutional features with the detection network, thus enabling nearly cost-free region proposals.

Fish stock overexploitation: The unsustainable fishing practice where the fish population is reduced to a level where it is not productive anymore. This is a significant contributor to the degradation of marine ecosystems.

Fishing effort: An estimation of the total amount of fishing activities, often determined by the number of hours a vessel spends fishing, the type of gear used, and the amount of gear deployed.

Flag vessels: Ships that are registered under a particular country's flag. The country under which a ship is registered is responsible for enforcing regulations on the ship, no matter where it operates.

Illegal, Unreported, and Unregulated (IUU) fishing: Unlawful practices in the fishing industry that include fishing without permission, not reporting catches, using illegal methods, and fishing in protected areas.

Mask R-CNN: An extension of Faster R-CNN that adds a branch for predicting an object mask in parallel with the existing branch for bounding box recognition, which allows for pixel-level segmentation in addition to object detection.

Machine Learning (ML): A field of artificial intelligence that uses statistical techniques to give computer systems the ability to learn from data, identifying patterns and making decisions with minimal human intervention.

RetinaNet: A one-stage object detection model that addresses the problem of class imbalance by introducing a loss function that focuses on training samples that are hard to classify.

Synthetic Aperture Radar (SAR): A type of imaging radar used to create two-dimensional images of landscapes that can be used to detect and monitor vessels.

Single Shot MultiBox Detector (SSD): A one-stage object detection algorithm that encodes object location and class probabilities directly in a single shot, making it faster than two-stage detectors.

Support Vector Machines (SVM): A shallow machine learning model that uses supervised learning algorithms for classification and regression analysis.

Vessel Monitoring System (VMS): A type of surveillance system that uses satellite-based tracking devices installed on fishing vessels. The system provides data on a vessel's position, course, and speed, helping authorities to monitor fishing activities.

9. References

Agnew, David J., John Pearce, Ganapathiraju Pramod, Tom Peatman, Reg Watson, John R. Beddington, and Tony J. Pitcher. 2009. "Estimating the Worldwide Extent of Illegal Fishing." *PLOS ONE* 4 (2): e4570. https://doi.org/10.1371/journal.pone.0004570.

Andrew Kitts, Raymon van Anrooy, Sjef van Eijs, and Jesica Pino Shibata. 2020. *Techno-Economic Performance Review of Selected Fishing Fleets in North and South America*. FAO. https://doi.org/10.4060/ca9543en.

Chuaysi, Buncha, and Supaporn Kiattisin. 2020. "Fishing Vessels Behavior Identification for Combating IUU Fishing: Enable Traceability at Sea." *Wireless Personal Communications* 115 (4): 2971–93. https://doi.org/10.1007/s11277-020-07200-w.

FAO. 2022. *The State of World Fisheries and Aquaculture 2022*. FAO. https://doi.org/10.4060/cc0461en.

Galdelli, Alessandro, Adriano Mancini, Carmen Ferrà, and Anna Nora Tassetti. 2021. "A Synergic Integration of AIS Data and SAR Imagery to Monitor Fisheries and Detect Suspicious Activities." *Sensors* 21 (8): 2756. https://doi.org/10.3390/s21082756.

Hou, Biao, Xingzhong Chen, and Licheng Jiao. 2015. "Multilayer CFAR Detection of Ship Targets in Very High Resolution SAR Images." *IEEE Geoscience and Remote Sensing Letters* 12 (4): 811–15. https://doi.org/10.1109/LGRS.2014.2362955.

Hwang, Jeong-In, and Hyung-Sup Jung. 2018. "Automatic Ship Detection Using the Artificial Neural Network and Support Vector Machine from X-Band Sar Satellite Images." *Remote Sensing* 10 (11): 1799. https://doi.org/10.3390/rs10111799.

Jansing, David. 2021. "Introduction to Synthetic Aperture Radar." In *Introduction to Synthetic Aperture Radar: Concepts and Practice*, 1st Edition. McGraw-Hill Education. https://www.accessengineeringlibrary.com/content/book/9781260458961/chapter/chapter1.

Kanjir, Urška, Harm Greidanus, and Krištof Oštir. 2018. "Vessel Detection and Classification from Spaceborne Optical Images: A Literature Survey." *Remote Sensing of Environment* 207 (March): 1–26. https://doi.org/10.1016/j.rse.2017.12.033.

Kroodsma, David A., Juan Mayorga, Timothy Hochberg, Nathan A. Miller, Kristina Boerder, Francesco Ferretti, Alex Wilson, et al. 2018. "Tracking the Global Footprint of Fisheries." *Science* 359 (6378): 904–8. https://doi.org/10.1126/science.aao5646.

Kroodsma, David Allen, Tim Hochberg, Pete Davis, Fernando Paolo, Rocio Joo, and Brian Adrian Wong. 2022. "Revealing the Global Longline Fleet with Satellite Radar," April. https://eartharxiv.org/repository/view/3239/.

Lemoine, G., J. Chesworth, G. Schwartz-Juste, N. Kourti, and I. Shepherd. 2004. "Near Real Time Vessel Detection Using Spaceborne SAR Imagery in Support of Fisheries Monitoring and Control Operations." In *IGARSS 2004. 2004 IEEE International Geoscience and Remote Sensing Symposium*, 7:4825–28 vol.7. https://doi.org/10.1109/IGARSS.2004.1370242.

Li, Heng, and Xinyu Wang. 2008. "Automatic Recognition of Ship Types from Infrared Images Using Support Vector Machines." In *2008 International Conference on Computer Science and Software Engineering*, 6:483–86. https://doi.org/10.1109/CSSE.2008.1647.

Li, Jianwei, Changwen Qu, and Jiaqi Shao. 2017. "Ship Detection in SAR Images Based on an Improved Faster R-CNN." In *2017 SAR in Big Data Era: Models, Methods and Applications (BIGSARDATA)*, 1–6. https://doi.org/10.1109/BIGSARDATA.2017.8124934.

Lin, Zhao, Kefeng Ji, Xiangguang Leng, and Gangyao Kuang. 2019. "Squeeze and Excitation Rank Faster R-CNN for Ship Detection in SAR Images." *IEEE Geoscience and Remote Sensing Letters* 16 (5): 751–55. https://doi.org/10.1109/LGRS.2018.2882551.

Liu, Wei, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. 2016. "SSD: Single Shot MultiBox Detector." In *Computer Vision – ECCV 2016*, edited by Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, 21–37. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-46448-0_2.

Marzuki, Marza Ihsan, Rinny Rahmania, Penny Dyah Kusumaningrum, Rudhy Akhwady, Daud Saputra Amare Sianturi, Yustisia Firdaus, Agus Sufyan, Cecep Ahmad Hatori, and Handy Chandra. 2021. "Fishing Boat Detection Using Sentinel-1 Validated with VIIRS Data." *IOP Conference Series: Earth and Environmental Science* 925 (1): 012058. https://doi.org/10.1088/1755-1315/925/1/012058.

Molinet, Carlos, Nancy Barahona, Beatriz Yannicelli, Jorge González, Alejandra Arevalo, and Sergio Rosales. 2011. "Statistical and Empirical Identification of Multispecies Harvesting Zones to Improve Monitoring, Assessment, and Management of Benthic Fisheries in Southern Chile." *Bulletin of Marine Science* 87 (July). https://doi.org/10.5343/bms.2010.1067.

Molinet, Carlos, Andrés Olguín, Paulina Gebauer, Patricio A. Díaz, Manuel Díaz, Thamara Matamala, Paulo Mora, and Kurt Paschke. 2020. "Upswing and Expansion of the Southern King Crab (Lithodes Santolla) Fishery in Northwest Patagonia: Drivers, Trends and Opportunities for Management." *Regional Studies in Marine Science* 34 (February): 101073. https://doi.org/10.1016/j.rsma.2020.101073.

Molinet, Carlos, María Eugenia Solari, Manuel Díaz, Francisca Marticorena, Patricio A. Díaz, Magdalena Navarro, and Edwin Niklitschek. 2019. "Fragments of the Environmental Hystory of Fjord and Channels North Patagonic System, South Of Chile: Two Centuries of Exploitation Fragmentos de La Historia Ambiental Del Sistema de Fiordos y Canales Nor-Patagónicos, Sur De Chile: Dos Siglos de Explotació." *Magallania*. https://doi.org/10.4067/s0718-22442018000200107.

Moreira, Alberto, Pau Prats-Iraola, Marwan Younis, Gerhard Krieger, Irena Hajnsek, and Konstantinos P. Papathanassiou. 2013. "A Tutorial on Synthetic Aperture Radar." *IEEE Geoscience and Remote Sensing Magazine* 1 (1): 6–43. https://doi.org/10.1109/MGRS.2013.2248301.

Mullon, Christian, Pierre Fréon, and Philippe Cury. 2005. "The Dynamics of Collapse in World Fisheries." *Fish and Fisheries* 6 (2): 111–20. https://doi.org/10.1111/j.1467-2979.2005.00181.x.

Ovando, Daniel, Ray Hilborn, Cole Monnahan, Merrill Rudd, Rishi Sharma, James T. Thorson, Yannick Rousseau, and Yimin Ye. 2021. "Improving Estimates of the State of Global Fisheries Depends on Better Data." *Fish and Fisheries* 22 (6): 1377–91. https://doi.org/10.1111/faf.12593.

Paolo, Fernando, Tsu-ting Tim Lin, Ritwik Gupta, Bryce Goodman, Nirav Patel, Daniel Kuster, David Kroodsma, and Jared Dunnmon. 2022. "XView3-SAR: Detecting Dark Fishing Activity Using Synthetic Aperture Radar Imagery." arXiv. https://doi.org/10.48550/arXiv.2206.00897.

Park, Jaeyoon, Jungsam Lee, Katherine Seto, Timothy Hochberg, Brian A. Wong, Nathan A. Miller, Kenji Takasaki, et al. 2020. "Illuminating Dark Fishing Fleets in North Korea." *Science Advances* 6 (30): eabb1197. https://doi.org/10.1126/sciadv.abb1197.

Pitcher, Tony J, Reg Watson, Robyn Forrest, Hreiðar Þór Valtýsson, and Sylvie Guénette. 2002. "Estimating Illegal and Unreported Catches from Marine Ecosystems: A Basis for Change." *Fish and Fisheries* 3 (4): 317–39. https://doi.org/10.1046/j.1467-2979.2002.00093.x.

Ren, Shaoqing, Kaiming He, Ross Girshick, and Jian Sun. 2017. "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39 (6): 1137–49. https://doi.org/10.1109/TPAMI.2016.2577031.

Rousseau, Yannick, Reg A. Watson, Julia L. Blanchard, and Elizabeth A. Fulton. 2019. "Evolution of Global Marine Fishing Fleets and the Response of Fished Resources."

Proceedings of the National Academy of Sciences 116 (25): 12238–43. https://doi.org/10.1073/pnas.1820344116.

Sasamal, Sasanka Kumar, and Naresh K. Mallenahalli. 2019. "Detection of Fishing Boats in SAR Image Using Linear Moments and Machine Learning." In *2019 IEEE Recent Advances in Geoscience and Remote Sensing : Technologies, Standards and Applications (TENGARSS)*, 21–25. https://doi.org/10.1109/TENGARSS48957.2019.8976043.

Shepperson, Jennifer L, Niels T Hintzen, Claire L Szostek, Ewen Bell, Lee G Murray, and Michel J Kaiser. 2018a. "A Comparison of VMS and AIS Data: The Effect of Data Coverage and Vessel Position Recording Frequency on Estimates of Fishing Footprints." *ICES Journal of Marine Science* 75 (3): 988–98. https://doi.org/10.1093/icesjms/fsx230.

Snapir, Boris, Toby W. Waine, and Lauren Biermann. 2019. "Maritime Vessel Classification to Monitor Fisheries with SAR: Demonstration in the North Sea." *Remote Sensing* 11 (3): 353. https://doi.org/10.3390/rs11030353.

Sumaila, U. R., D. Zeller, L. Hood, M. L. D. Palomares, Y. Li, and D. Pauly. 2020. "Illicit Trade in Marine Fish Catch and Its Effects on Ecosystems and People Worldwide." *Science Advances* 6 (9): eaaz3801. https://doi.org/10.1126/sciadv.aaz3801.

Temple, Andrew J., Daniel J. Skerritt, Philippa E. C. Howarth, John Pearce, and Stephen C. Mangi. 2022. "Illegal, Unregulated and Unreported Fishing Impacts: A Systematic Review of Evidence and Proposed Future Agenda." *Marine Policy* 139 (May): 105033. https://doi.org/10.1016/j.marpol.2022.105033.

Ting, Liu, Zhou Baijun, Zhao Yongsheng, and Yan Shun. 2021. "Ship Detection Algorithm Based on Improved YOLO V5." In *2021 6th International Conference on Automation, Control and Robotics Engineering (CACRE)*, 483–87. https://doi.org/10.1109/CACRE52464.2021.9501331.

Wang, Yuanyuan, Chao Wang, Hong Zhang, Yingbo Dong, and Sisi Wei. 2019a. "A SAR Dataset of Ship Detection for Deep Learning under Complex Backgrounds." *Remote Sensing* 11 (7): 765. https://doi.org/10.3390/rs11070765.

Wang, Yuanyuan, Chao Wang, Hong Zhang, Yingbo Dong, and Sisi Wei. 2019b. "Automatic Ship Detection Based on RetinaNet Using Multi-Resolution Gaofen-3 Imagery." *Remote Sensing* 11 (5): 531. https://doi.org/10.3390/rs11050531.

Wei, Shunjun, Xiangfeng Zeng, Qizhe Qu, Mou Wang, Hao Su, and Jun Shi. 2020. "HRSID: A High-Resolution SAR Images Dataset for Ship Detection and Instance Segmentation." *IEEE Access* 8: 120234–54. https://doi.org/10.1109/ACCESS.2020.3005861.

Yang, Rong, Robert Wang, Yunkai Deng, Xiaoxue Jia, and Heng Zhang. 2021. "Rethinking the Random Cropping Data Augmentation Method Used in the Training of CNN-Based SAR Image Ship Detector." *Remote Sensing* 13 (1): 34. https://doi.org/10.3390/rs13010034.

Yasir, Muhammad, Wan Jianhua, Xu Mingming, Sheng Hui, Zeng Zhe, Liu Shanwei, Arife Tugsan Isiacik Colak, and Md Sakaouth Hossain. 2023. "Ship Detection Based on Deep Learning Using SAR Imagery: A Systematic Literature Review." *Soft Computing* 27 (1): 63–84. https://doi.org/10.1007/s00500-022-07522-w.

Young, Darrell L. 2019. "Deep Nets Spotlight Illegal, Unreported, Unregulated (IUU) Fishing." In *2019 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*, 1–7. https://doi.org/10.1109/AIPR47015.2019.9174577.

Zhang, Tianwen, Xiaoling Zhang, Xiao Ke, Xu Zhan, Jun Shi, Shunjun Wei, Dece Pan, et al. 2020. "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. https://doi.org/10.3390/rs12182997.

Zhao, Congxia, Xiongjun Fu, Jian Dong, Rui Qin, Jiayun Chang, and Ping Lang. 2022. "SAR Ship Detection Based on End-to-End Morphological Feature Pyramid Network." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 15: 4599–4611. https://doi.org/10.1109/JSTARS.2022.3150910.