Assessing Stationary Maritime Traffic Density in Territorial Dispute Zones Using Open-Source Sentinel-2 Satellite Imagery

Paul Menhart
Munich University of Applied Sciences
paul@menhart.net

Florian Eiselbrecher
Munich University of Applied Sciences
florianeiselbrecher@gmail.com

Chiara Perocco
Munich University of Applied Sciences
chiara.perocco@hm.edu

Mila Stillman
Munich University of Applied Sciences
mila.stillman@hm.edu

Anna Kruspe
Munich University of Applied Sciences
anna.kruspe@hm.edu

ABSTRACT

Although Sentinel-2 images are seldom employed for object detection due to their lower resolution, their accessibility and cost-effectiveness make them a valuable tool for continuous global maritime monitoring. This paper proposes using publicly available satellite imagery to detect ships in the maritime conflict zones of the South China Sea, aiming to assess the intensity of conflict through a fine-tuned YOLO object detection model. Ship density is estimated by aggregating detected ships over each area, demonstrating that Sentinel-2 is a valuable data source for this task. Additionally, the correlation between the number of ships detected in a given week and the number of articles reporting on conflicts in the relevant maritime zones during the same time period was calculated. Preliminary results suggest that this approach could serve as a valuable indicator for evaluating military tensions in the region, enhancing conflict transparency and supporting risk reduction efforts.

Keywords

Ship Detection, Military Conflict, Satellite Imagery, South China Sea

INTRODUCTION

Transparency is crucial for effective action in conflict zones and holds the potential to discourage aggression. While satellite imagery's contribution to straightforward global transparency is debated (Witjes & Olbrich, 2017), analyzing open satellite imagery to detect conflict intensity and escalations offers a promising avenue for stabilizing transparency in these regions. Open-source material makes such analyses feasible and accessible to a wide range of parties contributing to improved responses and potential mitigation of aggressive actions.

In this study, object detection models were employed to perform a qualitative analysis of disputed maritime areas. Using open-source Sentinel-2 satellite imagery, a cross-regional comparison of ongoing territorial disputes in the South China Sea was performed, examining the correlation between the presence of ships and conflict intensity. An example of the used type of data in comparison with commercial high-resolution data is shown in figure 1.

The South China Sea conflict was selected due to its prolonged nature, maritime context, and the clear distinction between conflict-related activities and routine maritime operations in the region. It is assumed that vessels operating

(a) **(b)**

Figure 1. Sentinel-2 L2A image (10m resolution) of Namyit Island on May 17, 2024 (a) (European Union, 2025) and MAXAR image (3-5m resolution) of Namyit Island on May 16, 2024 (b) (Center for Strategic & International Studies (CSIS), 2025b)

in the vicinity of islands and reefs claimed by various states have either been granted access by the controlling state or else violate the territorial integrity of its claimed waters. The South China Sea conflict has persisted since the 1980s, involving several regional actors and organizations, including the People's Republic of China (PRC), Malaysia, Taiwan, Vietnam, the Philippines and the Association of Southeast Asian Nations (ASEAN). Due to the significant political, strategic and economic importance of the region, trans-regional actors such as the United States and its allies have also been actively involved. More information on the area is visible in figure 2.

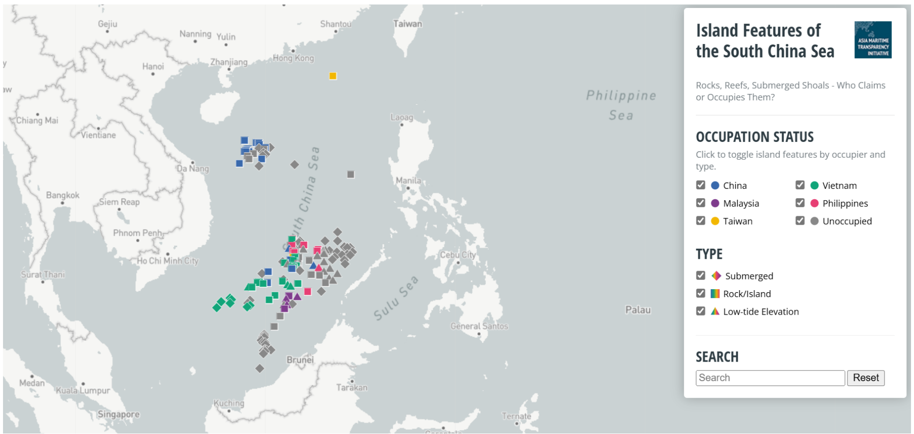

Figure 2. Island Features of the South China Sea (Center for Strategic & International Studies (CSIS), 2025c)

In territorial disputes, states often aim to assert boundaries based on nationalistic grounds, driven by concerns over reputation damage, the risk of losing their monopoly on the use of force if challenged, and the fear that compromise may be perceived as a sign of weakness (Huth, 1996). In the case of the South China Sea, competing sovereignty claims have led to territorial tensions, highlighted by incidents involving military vessels and efforts to reinforce claims through the construction of military facilities on islands, islets, and reefs. (Center for Strategic & International Studies (CSIS), 2025c) In addition to ideological motivations for engaging in conflict, states also consider economic factors (Walter, 2009). The importance of the South China Sea arises from its strategic location near major maritime trade routes and ports, such as the Strait of Malacca, as well as its abundant natural resources, including an estimated 190 trillion cubic feet of natural gas, 11 billion barrels of oil, and rich fishing grounds (Center for Strategic & International Studies (CSIS), 2025c).

Peace and conflict research explores the causes, course, and consequences of conflicts. This analysis focuses on the second aspect, the progression of a conflict. Over time, a conflict can evolve in terms of its intensity, the scale of organized violence. The threat or use of political violence to achieve operational objectives diminishes human security not only for the parties directly involved but also for neutral actors who rely on safe passage through the disputed region for trade. (Glick & Taylor, 2010; UNDP (United Nations Development Programme), 2022)

Studies using open source material have the potential to improve conflict transparency, reduce the risk of unwanted escalation, and support informed incident response by providing a better understanding of the evolution of conflict intensity in a remote and contested maritime region through the analysis of open access optical Sentinel-2 satellite imagery. Our paper seeks to demonstrate that the usage of such material for the task at hand is feasible.

The main contributions of this paper include:

- A self-labeled dataset of Sentinel-2 images relevant to ship detection in the South China Sea
- A YOLO model fine-tuned for the task
- A first analysis of the relationship between detected ship densities and newspaper mentions of the conflict

The data and the model developed are available under https://huggingface.co/Punsi/ship-detection, and example results are showcased at https://www.maritime-conflict-ship-detection.eu.

The rest of the paper is structured as follows: It begins by providing an overview of related research on ship detection with other means and in other areas, and on the usage of Sentinel data in conflicts. This is followed by a description of the data collection and refinement process. The proposed approach, including model training and post-processing techniques, is then introduced, along with a discussion of the obtained results. An analysis of the correlation between detected ships and conflict-related newspaper articles is presented next. Finally, the study concludes with a conclusion and suggestions for future research.

RELATED WORK

Ship detection is a well-established field, especially in the analysis of civilian maritime activities, such as the study of recreational maritime traffic utilizing Sentinel-2 imagery, as demonstrated by Mayrä et al. (Mäyrä et al., 2024), as well as other shipping activities that vary in ship size and regional scale (Bakirci, 2025). These differ from our research in terms of the region, the South China Sea, and the underlying scenario, which involves a conflict zone. It is also utilized as a tool to enhance transparency regarding environmentally harmful or illegal shipping operations by identifying hidden fishing fleets (Park et al., 2020) and other clandestine maritime activities. (Byrne et al., 2021) However, most of this research relies on expensive, restricted, high-resolution imagery, whether from civilian optical sensors or synthetic aperture radar (SAR) systems. SAR data provides the benefit of coverage during the day and night in most weather conditions. Sentinel-1 SAR is an exception, as it is openly available and used as demonstrated by Grover et al. (Grover et al., 2019). However, the imaging is not fully accessible worldwide including the South China Sea.

In the current scientific discourse surrounding conflict research, the analysis of Sentinel-2 imagery focuses primarily on the visible impacts of ground-based armed conflict. A key area of focus is the assessment of structural damage, similar to the study by Fakhri and Gkanatsios (Fakhri & Gkanatsios, 2021), which analyzed the destruction of buildings and infrastructure resulting from the fight against ISIS in and around Mosul, Iraq, in 2019.

Another area under consideration is Ukraine due to the Ukraine-Russia war. Building damage assessment in Kyiv, Ukraine is analyzed by Aimaiti, Y. et al using Sentinel-2 data (Aimaiti et al., 2022).

Another important aspect is analyzing changes in land use due to armed conflict, such as the abandonment of cropland in central Mali in 2019, as demonstrated by Boudinaud and Orenstein (Boudinaud & Orenstein, 2021).

In terms of the ship detection method used, several object detection models were evaluated for ship detection tasks based on their documented performance in literature, including YOLO (Redmon et al., 2016), Faster R-CNN (Ren et al., 2016) and VFNet (Zhang et al., 2021). YOLO was selected due to its proven effectiveness and successful application in previous object detection studies, particularly in maritime contexts (Wang et al., 2022; Yang et al., 2024; Yu et al., 2024).

DATA

In this section, we discuss our considerations of the region of interest as well as data collection steps. We also describe two dataset used to fine-tune the model for ship detection.

Region of interest

In order to define the area of interest, a list of all the islands, reefs and other points of interests in the South China Sea was compiled based on information from the Asia Maritime Transparency Initiative. To scan a larger area for ships, practical limitations on image size required the division of the region into tiles. The tiling should have seamless coverage of our area of interest while also preventing overlap between the images.
The satellite images have a size of 1100 x 1100 pixels with a resolution of 10 meters per pixel, meaning each image depicts an area of 121 square kilometers. Each degree of latitude corresponds to a constant distance of roughly 110 kilometers. This is not the case for the degrees of longitude, as one degree of longitude corresponds to a larger distance near the equator than near the poles. To address this, the area was divided into bands of equal latitude, with tiles stacked horizontally starting from a longitude of 0 degrees. Every coordinate gathered in the initial step was used to mark all tiles within a one-tile radius. In Figure 3 and Figure 4, the resulting tiles are plotted in blue, while the source coordinates are marked in red.

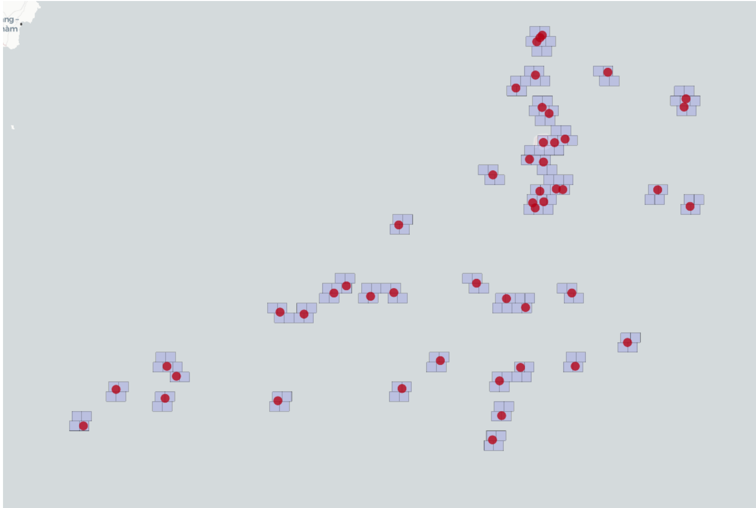

Figure 3. Detailed Visualization of Search Space Tilings.

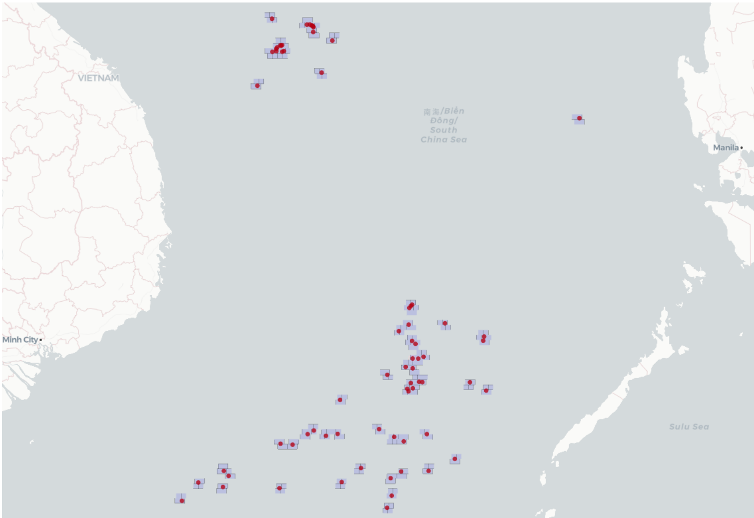

Figure 4. Comprehensive Overview of the Search Space.

Data collection

Google Earth Engine (Gorelick et al., 2017) was used to retrieve images from Sentinel-2. A total area of 26,015 square kilometers was scanned and split into 215 tiles. Each tile was scanned once per week in the time window between 2019-01-01 and 2024-12-01, which equates to 309 scans per tile. The total number of images scanned is 215 x 309 = 66,435 scans, which took approximately 6 hours. The limiting factor was the network bandwidth, as each image needs to be downloaded twice (once as an RGB-image and once as a gray-image for the cloud maps), adding up to 132,870 images downloaded.

Open-source Sentinel-2 RGB imagery has a relatively low resolution (10m) compared to commercial high-resolution Earth observation constellation providers such as Planet or Maxar (3-5m) (an example can be seen in figure 5). Ships vary in size depending on their type and use, from 5m recreational craft to 300+m container ships in the maritime trade. Consequently, smaller ships are harder to detect in moderate-resolution Sentinel-2 imagery.

In order to improve computational performance and establish a connection between islands and respective actors in the conflict, the analysis was restricted to the maritime regions around occupied and claimed islands. A total of 71 islands and reefs and their close surroundings were considered, including their longitude/latitude position. The data also include the state affiliation of each island (PRC: 29, Vietnam: 27, Philippines: 9, Malaysia: 5, ROK: 1), but this information was not used as part of this work. (Center for Strategic & International Studies (CSIS), 2025c)

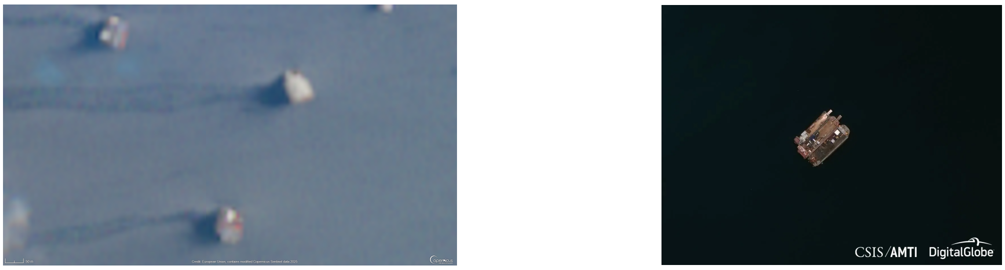

(a) **(b)**

Figure 5. Sentinel-2 Level 2A (10m resolution) image of moored vessels at Mischief Reef, captured on 2024-05-17 (a) (European Union, 2025) and MAXAR image (3–5m) of moored vessels at Mischief Reef, captured on 2018-03-04 (b) (Center for Strategic & International Studies (CSIS), 2025a).

Available Sentinel-2 imagery was manually analyzed with regards to its temporal frequency as this is an important factor for continuous conflict analysis. Mischief Island was chosen as an example because it has been one of the main points of tension in the region since its conversion into a large military air and naval base.

The analysis showed that Mischief Reef images have been available since the first quarter of 2017. During the first year and a half of data availability, the area image refresh rate fluctuated around 15 days, with a maximum of 35 days. The analysis showed a relatively stable refresh rate every 5 days since the third quarter of 2018. Cloud cover was not taken into account in this initial analysis.

Model training data

To enable vessel detection, two datasets were considered for model fine-tuning: first, an existing dataset for vessel detection, ”Vessel Detection from Sentinel-2 Raw” (VDS2RAW) (Prete & Meoni, 2024), and second, a self-labeled dataset.

Fine-tuning the basic YOLOv8 model and the pretrained YOLOv8 model by Mayra et al. (Mäyrä et al., 2024) on the VDS2RAW dataset led to poor ship detection performance on the test portion of the self-labeled South China Sea dataset, as demonstrated by the comparison between Figure 6 and Figure 8. As a result, only the self-labeled dataset was selected for the subsequent fine-tuning tasks. As shown in Figure 8, the pre-trained object detection model (Mäyrä et al., 2024), trained on the Finish coastline only, also performed poorly in the selected region.

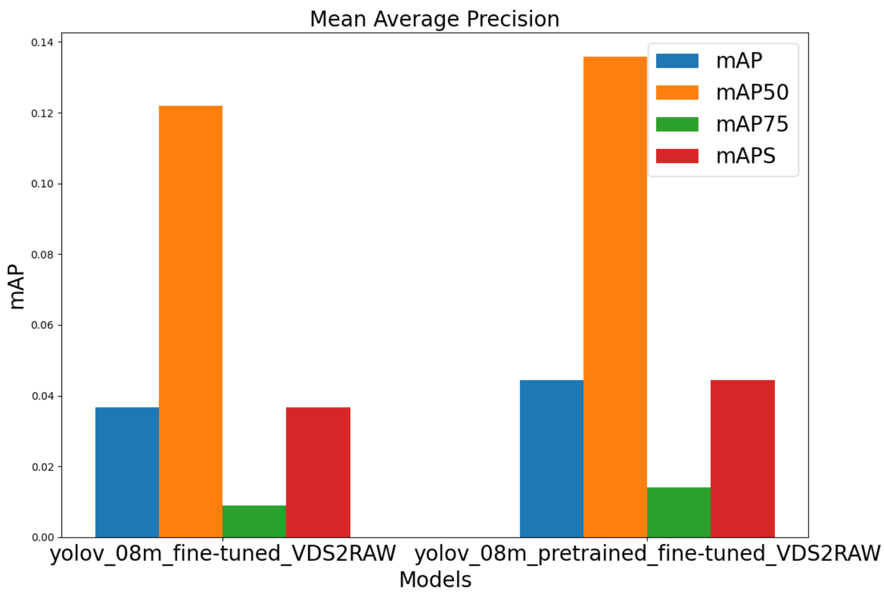

Figure 6. The performance of the VDS2RAW dataset finetuning was evaluated on the test portion of the self-labeled South China Sea dataset for both the basic YOLOv8 model and the pretrained YOLOv8 model by Mayra et al. (Mäyrä et al., 2024). Model precision was assessed using Mean Average Precision (mAP) at various Intersection over Union (IoU) thresholds, such as mAP50 at IoU = 0.50.

The performance difference may be attributed to a domain gap, particularly due to environmental differences between the source domain (the Baltic (Mäyrä et al., 2024) and North Seas (Prete et al., 2024)) and the target domain (the South China Sea). (Kornfein et al., 2023) These differences could include variations in water color, brightness, and contrast, which are influenced by different phytoplankton (algae) concentrations due to different distances of the oceans to the equator. (Sun et al., 2023) Additionally, ships may vary in size, color, brightness, and contrast, depending on regional shipping characteristics, such as the types of fishing (Watch, 2025), construction, or naval vessels present. (Center for Strategic & International Studies (CSIS), 2025a)

1st Dataset: VDS2RAW

The first dataset considered was the "Vessel Detection from Sentinel-2 Raw" (VDS2RAW) dataset. This dataset was created for a vessel detection system using VFNet with ResNet-18 as the backbone for bounding box prediction. Unlike the area of interest, the dataset focused on shipping activity along the Danish coastline. It contains Sentinel-2 L1C images acquired in 2019. Out of 390 L0 granules, a subset of 166 labeled raw granules, including bands B02, B03, and B04, were divided into training (105), validation (27), and test (34) sets. The downloaded data is provided as TIF files for bands B02, B03, and B04, along with .json files containing bounding box data (483, 119, and 93 annotations for each section) (Prete & Meoni, 2024) (Prete et al., 2024). To use the data, both the TIF and JSON files required pre-processing. Since the objective was to detect ships on RGB images, the first step was to merge the frequency bands for the Red (B04), Green (B03), and Blue (B02) channels into a single RGB image. However, misalignment in the band data caused a red-green-blue mirroring effect for ships in some images. Cross-correlation was used to align the bands as accurately as possible. An alternative approach involving feature extraction for image alignment was also tested but resulted in artifacts in some images and was subsequently discarded. Finally, the images were converted to the appropriate datatype (from uint16 to uint8) and saved as PNG files for fine-tuning object detection models, such as YOLO. An example of the pre-processing process is shown in figure 7. A second step was needed to convert the JSON files, which contained the bounding boxes, into the model-specific format (e.g., YOLO: x-center, y-center, width, height) and save them as text files. These steps resulted in a dataset comprising labeled training, test, and validation data.

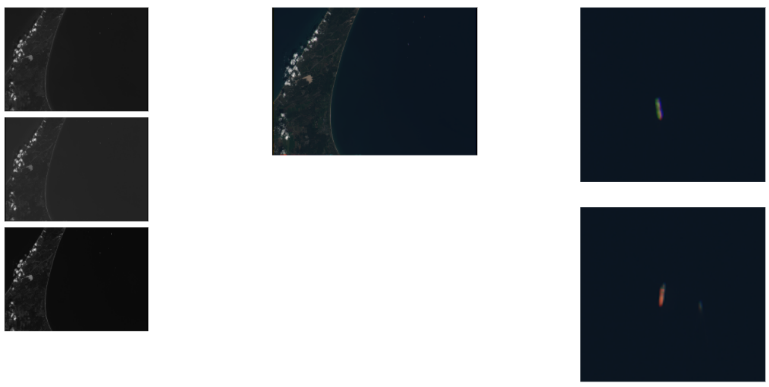

Individual red [B04] (top), green [B03] (center), blue [B02] (bottom) band images

Merged RGB Image

Band alignment issue: Unaligned image (top) and aligned image (bottom) (small deviations even after cross-correlation remain)

Figure 7. Pre-processing of VDS2RAW dataset

2nd Dataset: Self-Annotated

A dataset of images collected from the South China Sea was also considered. To this end, we hand-labeled 524 images with 1,198 ship objects, and divided the resulting dataset into training (80%), test (10%), and validation (10%) sets. The selected images were captured at the start of each quarter in and around Singapore, selected for its high density of ships of all sizes due to its location in the Malacca Strait and proximity to the South China Sea. Each image was annotated to identify the ships and determine their positions. A unique identifier (0) was assigned to each ship, as only one class (ship) was considered. The position and size of the ships were represented by bounding boxes in the format of x-center, y-center, width, and height. The dataset is available under https://huggingface.co/Punsi/ship-detection.

PROPOSED APPROACH

In this section, we first describe the fine-tuning process for our detection model and provide evaluation results. We then discuss the steps taken for post-processing the results.

Object Detection Models

As described in the current research section, a pre-trained YOLOv8 model (YOLOv8, 2024), based on the standard YOLOv8 architecture, was used. This model incorporates a variant of the Cross Stage Partial Darknet (CSP Darknet) as its backbone within the convolutional neural network (CNN), allowing for efficient feature extraction across multiple resolutions. It refines and combines features from various feature maps to improve object detection at different scales.

For multi-scale object detection, the model employs a Feature Pyramid Network (FPN), which upsamples and integrates low- and high-resolution layers. Additionally, it incorporates a Path Aggregation Network (PAN) to improve contextual understanding by reducing the information gap between lower spatial features and higher semantic layers. The model's head includes multiple detection modules for bounding box prediction, a critical component for our application.

We compare a standard YOLOv8 model (YOLOv8, 2024), trained for general-purpose object detection on the COCO dataset, with the model from (Mäyrä et al., 2024), which is specifically trained on Sentinel-2 imagery to detect recreational vessels near coastlines and quantify disturbances caused by small recreational craft. This setup allowed us to compare the original pre-trained model with a version fine-tuned on a limited custom dataset.

The standard COCO-trained YOLOv8 model performed poorly on the sample images, yielding an mAP50 of 0.0, since it was not trained on satellite imagery. In contrast, the YOLOv8 model trained specifically for ship detection was able to identify ships in test images from the South China Sea, though its performance was lower, achieving an mAP50 of 0.24 compared to the mAP50 of 0.89 in its original training region (Mäyrä et al., 2024).

Since this model was originally developed to detect small recreational craft along the Finnish coastline, we hypothesized that fine-tuning it with a self-labeled dataset (described in section "Model training data") from the target region would enhance its performance. After adjusting the model parameters and fine-tuning, the mAP50 increases from 0.24 to 0.71. The full results are shown in figure 8. We make the fine-tuned model available under https://huggingface.co/Punsi/ship-detection.

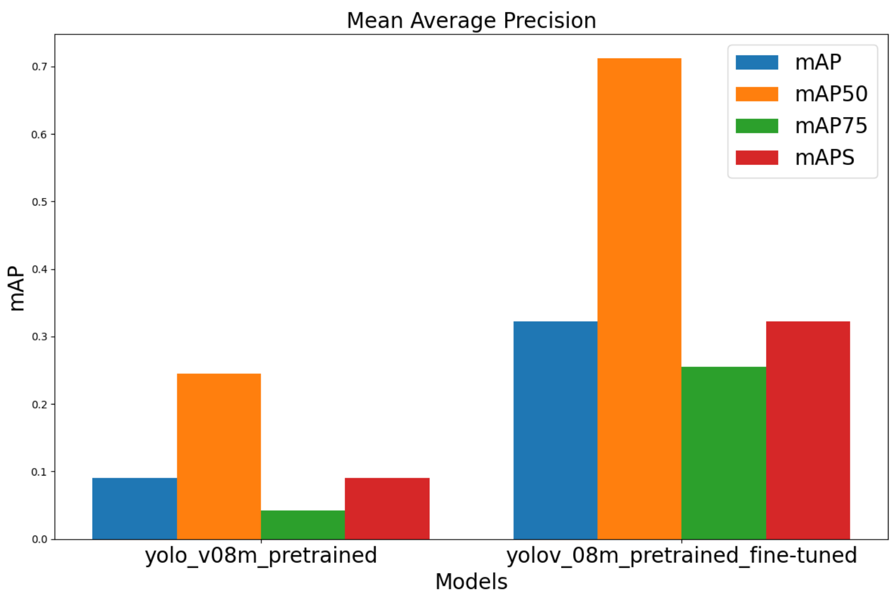

Figure 8. Performance metrics were evaluated across successive rounds of model testing. Mean Average Precision (mAP) was used to assess model precision at various Intersection over Union (IoU) thresholds (e.g., mAP50 at IoU = 0.50).

Post-processing

For post-processing, two filters were applied to the satellite imagery: a cloud mask and a coastline filter.

Cloud Mask Filter

The cloud mask was generated using the Google Earth Engine API, where Sentinel-2 cloud probability maps were downloaded at the same scale and with the same image ID as the RGB bands. This ensures that the cloud probability mask aligns with the original image. However, we noticed that cloud segmentation varies significantly across different probability thresholds. At lower probabilities, clearly visible clouds are often missed, while at higher probabilities, the algorithm struggles with detecting transparent cloud layers and distinguishing clouds from the wake of moving ships. An example is shown in figure 9.

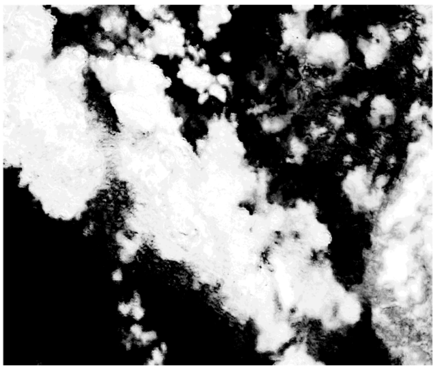

Figure 9. Sentinel-2 cloud mask (80-100%)

Land Filter

For the land surface data, a shapefile containing the latest OpenStreetMap data, with coastlines represented as polygons, was used (OpenStreetMap contributors, 2019). It is shown in figure 10. The surface filter checks whether the detected bounding box edge points intersect with any land surface areas. If an intersection is found, the detection is considered a false positive. However, since the shapefile contains millions of polygons, naively computing intersections between all detected bounding boxes and each polygon was computationally expensive. An R-tree could be employed to efficiently index and query the polygons, thereby significantly reducing the computational cost.

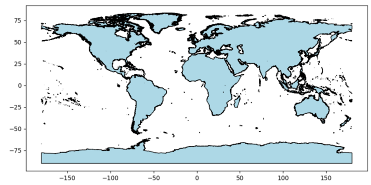

Figure 10. OpenStreetMap Global Coastlines

DETECTION RESULTS AND EVALUATION

With access to a server equipped with two A100 GPUs and 64 CPU cores, the detection model, along with pre- and post-processing, was executed in 24 parallel threads. In total, 39,735 objects were detected. Manual validation revealed that, out of 200 detected objects 53 were true positives, 102 were false positives, and 45 were false negatives, meaning that clearly visible ships were not detected. This results in an accuracy of 26.5%. Approximately 90% of the false positives were tiny clouds that were not filtered by our cloud detection filter, the rest of the false positives were primarily rock formations or waves. Vessel size was identified as the most significant factor affecting detection rates. Up to 80% of larger ships were correctly detected, while only around 30% of small boats were successfully identified.

The detection results for December and February 2024 are presented in figure 11. Both example images show the fine-tuned models ability to detect shipping around islands involved in territorial disputes such as Mischief reef.

An additional analysis of weekly ship density via a heat-map and the detection results of vessels with their respective geocoordinates are generated using the Folium library, as demonstrated in figure 12 for the week of February 6th, 2024. (Story, 2025) This highlights the framework's capability to detect vessels across large maritime sectors in the

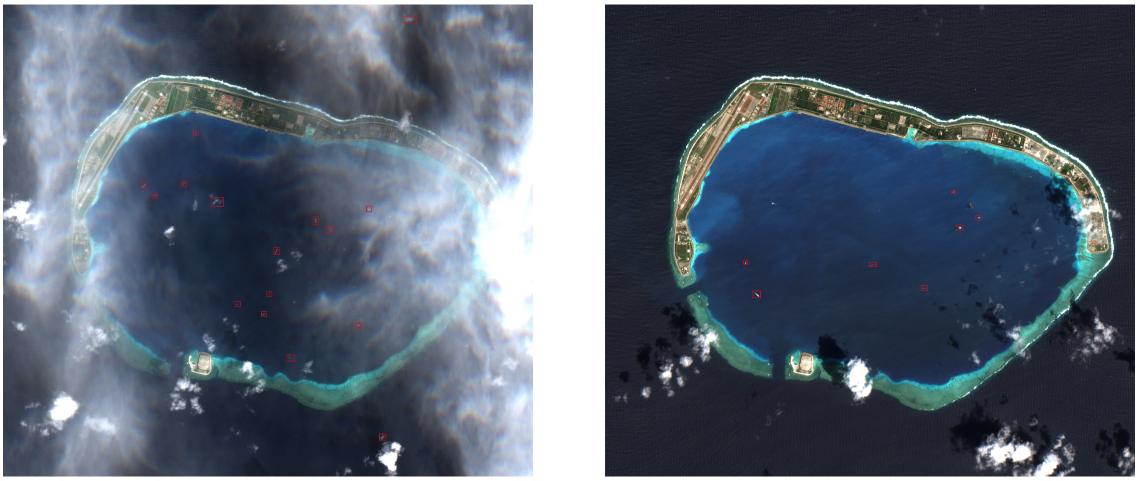

(a) (b)

Figure 11. Mischief Reef Ship Vessel Detection Results for December 2023 (a) and for February 2024 (b)

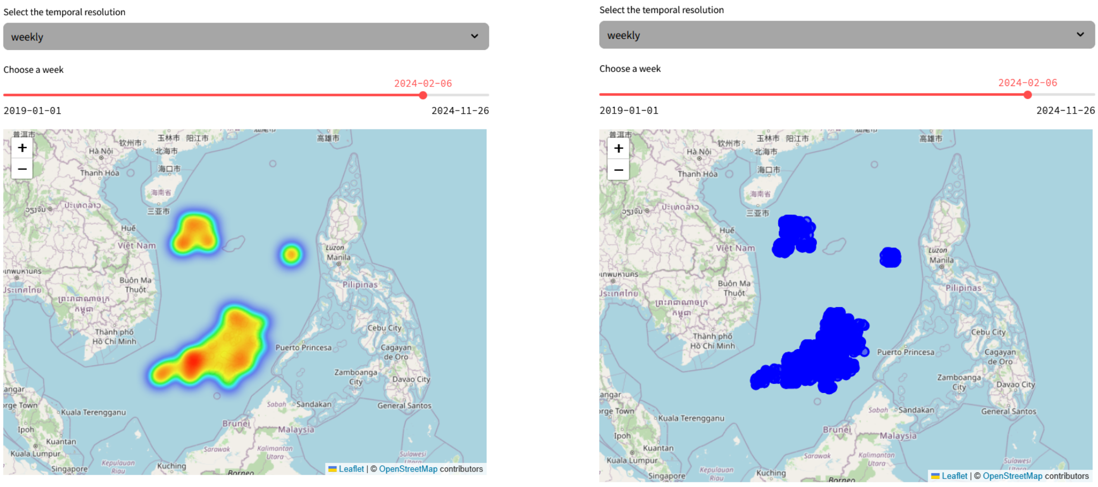

(a) (b)

Figure 12. Weekly Ship Density Heatmap (a) and Vessel Detection Results (b) for February 6, 2024

South China Sea at various time points, creating a tool that can support a transparent and comprehensive analysis of maritime conflict.

More examples and detailed insights are available under https://www.maritime-conflict-ship-detection.eu.

CORRELATION WITH NEWSPAPER ARTICLES

In an additional experiment, we calculated the Pearson Correlation between the number of detected ships within a given week and the number of articles published on the topic of the South China Sea during the same period across different datasets in order to determine possible correlations and examine the linear relationship between articles and number of detected ships. A strong positive correlation would suggest that an increase in the number of ships at a specific time is associated with a corresponding increase in the number of articles. This could lead to the conclusion that when significant or multiple events occur in the South China Sea, it results in both a high number of ships and an increased volume of media coverage.

Datasets The first dataset used is the "The Star" dataset from Kaggle (Azraimohamad, 2025). The Star is a well-known Malaysian news website and newspaper, one of the most widely circulated English-language outlets in the country. It provides coverage of various topics including national news, politics, business, sports, entertainment, lifestyle, and more. The dataset contains news articles' titles, text, sections, categories, authors, published dates, keywords, summaries, URLs, and top images. It covers a variety of topics such as nation, technology, and crime, with articles ranging from short news briefs to long-form news. The dataset includes data from December 2023 to December 2024.

The second dataset was created by the South China Sea Data Initiative and includes data from 2009 to 2023 (South China Sea Data Initiative, 2025). This dataset contains geo-located events related to ongoing territorial disputes in the South China Sea. It includes the columns 'Country, Year, Month, Day, Latitude, Longitude, Level and Notes'. These events were compiled from publicly available news sources and are therefore by no means a comprehensive account of all clashes that occurred in the westernmost Pacific during the decade in question. Interviews with journalists in several littoral states also indicate that a significant number of incidents go unreported, whether due to political pressure, lack of capacity, or insufficient verifiable details of the events.

Correlations: "The Star" In order to specifically identify relevant articles in the "The Star" dataset, a keyword pair search was carried out using the Python library NLTK (Natural Language Toolkit)[1]. The two keywords had to occur in combination, regardless of whether the combination occurred within the same column or across several columns. The search was carried out in the text-based columns `['title', 'text', 'keywords', 'summary']`.

To facilitate correlation analysis with the number of ships, the number of articles with the relevant keywords was aggregated on a weekly basis. January 1, 2019 was used as the starting point for calculating the weeks, because this is also when ship detection began.

Different keyword combinations were tested. Then, the correlations between the number of articles and the number of detected ships for every keyword combination were calculated, except for the combinations "military conflict", "South China Sea" and "sea event", "South China Sea" due to a low number of occurrences.

The results are as follows:

1. "conflict", "South China Sea"; amount: 66; correlation: 0.1805 (Figure 13)
2. "military conflict", "South China Sea"; amount: 3
3. "sea event", "South China Sea"; amount: 0
4. "military exercise", "South China Sea"; amount: 17; correlation: 0.1187
5. "escalation", "South China Sea"; amount: 18; correlation: -0.0686
6. "territorial disputes", "South China Sea"; amount: 11; correlation: 0.0721

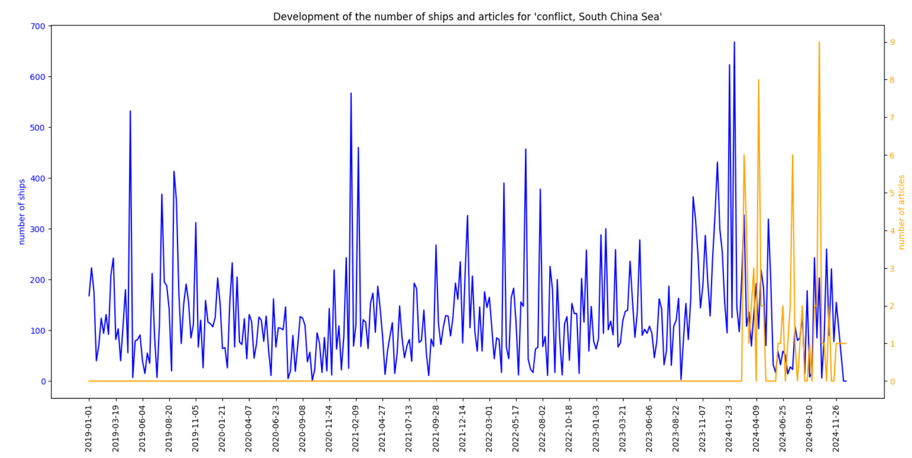

Figure 13. Development of the number of ships and articles in "The Star" with a combination of the keywords "conflict" and "South China Sea"

[1] https://www.nltk.org

Particularly striking are two peaks for the calculated correlation including the keywords ”conflict” and ”South China Sea” (Figure 14) in the periods March and October 2024 (orange), which coincide with the peaks in the number of ships recorded in these periods (blue). This observation led to a deeper analysis of the articles published in the corresponding weeks. The first peak falls in the week of March 5, 2024, when six articles containing the keywords ”conflict” and ”South China Sea” were published. The second peak is in the week of October 8, 2024 and comprises nine articles.

However, an analysis of this week’s articles from 2024-10-06 to 2024-10-12 revealed that not all articles with the keywords ”conflict” and ”South China Sea” directly describe specific events in the South China Sea. This means that these articles do not necessarily provide precise data or information on specific events in the region at that time. Nevertheless, all of these articles discuss approaches to conflict resolution in the South China Sea, suggesting that there may be increased activity or tensions there. The simultaneous accumulation of articles and registered vessels could indicate that tensions in the South China Sea are increasing, which would explain the increased number of recorded vessels. It should be noted, though, that the ship detection introduces artifacts as described above, and is therefore not 100% reliable.

Correlations: South China Sea Data Initiative The South China Sea Data Initiative dataset was used to analyze ship movements over a longer period from 2019 to 2023. No keyword filtering was required here, as the dataset only contains events related to territorial disputes in the South China Sea. In addition, the coordinates of the events were specified so that all articles relating to the area under consideration were taken into account (latitude: 7.3 to 17.09; longitude: 109.6 to 117.76). The counts are shown in figure 14.

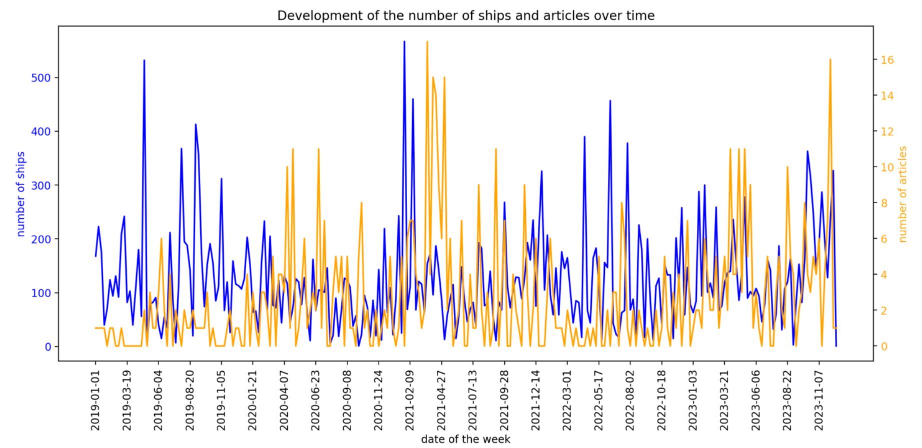

Figure 14. Development of the number of ships and articles in the South China Sea Data Initiative dataset

Even if no explicit correlation was found in this analysis (correlation: 0.0299), interesting patterns can be identified. The number of ships recorded shows clear peaks during the weeks 2021-01-26 and 2021-02-16, and comparable peaks also occur in the number of articles counted during the weeks 2021-03-23, 2021-04-06, 2021-04-13 and 2021-05-04. This raises the question of how quickly news reports on events are published. For this reason, an analysis of the articles in the interim periods was carried out. The National University of Singapore provides a summary of reports on events in the South China Sea (Centre for International Law (CIL), 2025). Although the articles do not provide exact dates as to when the events took place, it is known that they did take place. There is currently no evidence to refute the assumption that the articles refer to the detected ships, but since the reports only appear at a later date, they do not correlate with the number of ships.

DISCUSSION

The proposed fine-tuned YOLO model enables the detection of shipping activity in the South China Sea, as illustrated in Figure 11. As shown in Figure 12, the density of detected vessels varies over time and across different regions.

Despite the absence of a strong statistical correlation (correlation coefficient: 0.0299), the delayed release of news articles, such as those summarized by the National University of Singapore, may account for the lack of synchronization with ship counts. Nonetheless, both datasets exhibit noticeable peaks, suggesting a potential underlying relationship worth further exploration.

One limitation of the current model is the presence of false positives, particularly clouds mistakenly identified as ships, which introduces significant noise into the dataset. Minimizing these inaccuracies is essential for verifying any meaningful connection between detection results and media reporting.

We demonstrated the independent functionality of two key components: the ship detection model focused on the South China Sea, and the use of conflict-related media data as a potential validation tool. However, to perform a robust correlation analysis, both components need to be scaled up and refined. This includes expanding the fine-tuning dataset beyond a single category and increasing the emphasis on conflict-related content in the media dataset.

FUTURE WORK

Building upon the findings presented in this paper, several areas of ongoing research will further enhance our understanding of the relationship between maritime activity and conflicts.

Expanding News Data Sources The limitations of the Malaysian news data set used in this study require the exploration of larger and more comprehensive data sets, such as the Global Database of Events, Language, and Tone (GDELT, 2022; Leetaru & Schrodt, 2013). The use of such resources could lead to more refined insights between maritime activity and conflict reporting.

Refining Conflict Intensity Metrics News article counts serve as an initial proxy of the intensity of the conflict. Future research will explore alternative metrics, including casualty figures (where available) and established conflict indices, to provide a more nuanced assessment.

Accounting for Temporal Trends To isolate conflict-related influences on maritime activity, future analyses will incorporate techniques to deconvolve and remove cyclical or seasonal patterns from ship traffic data before correlation analysis.

Improving Ship Detection Enhancements to ship detection accuracy are crucial. This includes: (1) Expanding the training dataset for fine-tuning the detection model; (2) Improving cloud filtering algorithms; and (3) Transitioning from raw ship counts to ship density (ships/km^2) metrics, accounting for average cloud coverage.

Regional Specificity Spatial analysis at a subregional level may reveal localized correlations obscured by aggregate data. Negative correlations in specific subregions could, for example, indicate suppressed economic activity due to conflict.

Tool Development and Transferability A key objective is to extend the applicability of this methodology to other regions of interest. This involves developing a tool that enables users to define custom geographic areas and keywords. The tool would then generate reports on the correlation between ship activity and keyword frequency, providing a conflict intensity score with associated confidence intervals. This requires addressing engineering challenges related to: (1) Efficiently storing and querying ship density data indexed by location and timestamp; (2) Developing a user-friendly interface; and (3) Establishing a robust hosting infrastructure.

CONCLUSION

This work establishes a foundational framework for enhancing transparency in maritime conflict zones using openly available Sentinel-2 imagery. It combines machine learning-based ship detection over large maritime regions with conflict perception analysis through Natural Language Processing techniques. This work demonstrates that even minor modifications to pre-trained models significantly improve object detection in the specific maritime conflict region.

A preliminary analysis of the relationship between ship traffic and the frequency of conflict-related news coverage in a regional news source suggests a potential correlation between ship presence and media reporting on conflicts. Further work is needed to refine and enhance these initial conclusions.

Our ongoing work can be accessed at: https://www.maritime-conflict-ship-detection.eu/

REFERENCES

Aimaiti, Y., Sanon, C., Koch, M., Baise, L. G., & Moaveni, B. (2022). War related building damage assessment in kyiv, ukraine, using sentinel-1 radar and sentinel-2 optical images. *Remote Sensing*, *14*(24), 6239. https://doi.org/10.3390/rs14246239

Azraimohamad, A. (2025). *News article (weekly updated)*. Retrieved February 7, 2025, from https://www.kaggle.com/datasets/azraimohamad/news-article-weekly-updated

Bakirci, M. (2025). Advanced ship detection and ocean monitoring with satellite imagery and deep learning for marine science applications. *Regional Studies in Marine Science*, *81*, 103975. https://doi.org/10.1016/j.rsma.2024.103975

Boudinaud, L., & Orenstein, S. A. (2021). Assessing cropland abandonment from violent conflict in central mali with sentinel-2 and google earth engine. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *XLVI-4/W2-2021*, 9–15. https://doi.org/10.5194/isprs-archives-XLVI-4-W2-2021-9-2021

Byrne, J., Byrne, J., Kuo, L., & Sung, L. (2021, March 22). Project sandstone special report: Black gold: Exposing north korea's oil procurement networks.

Center for Strategic & International Studies (CSIS). (2025a). *Mischief reef* [Asia maritime transparency initiative]. Retrieved February 7, 2025, from https://amti.csis.org/mischief-reef/

Center for Strategic & International Studies (CSIS). (2025b). *Namyit island* [Asia maritime transparency initiative]. Retrieved February 7, 2025, from https://amti.csis.org/namyit-island/

Center for Strategic & International Studies (CSIS). (2025c). *South china sea features* [Asia maritime transparency initiative]. Retrieved February 7, 2025, from https://amti.csis.org/scs-features-map/

Centre for International Law (CIL). (2025). *March 2021 south china sea and ASEAN news updates now available* [Centre for international law]. Retrieved February 7, 2025, from http://cil.nus.edu.sg/activities/march-2021-south-china-sea-and-asean-news-updates-now-available/

European Union. (2025). *Copernicus browser* [Copernicus browser]. Retrieved February 27, 2025, from https://browser.dataspace.copernicus.eu/

Fakhri, F., & Gkanatsios, I. (2021). Integration of sentinel-1 and sentinel-2 data for change detection: A case study in a war conflict area of mosul city. *Remote Sensing Applications: Society and Environment*, *22*, 100505. https://doi.org/10.1016/j.rsase.2021.100505

GDELT. (2022). *GDELT project*. Retrieved February 17, 2025, from https://www.gdeltproject.org/

Glick, R., & Taylor, A. M. (2010). Collateral damage: Trade disruption and the economic impact of war. *Review of Economics and Statistics*, *92*(1), 102–127. https://doi.org/10.1162/rest.2009.12023

Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., & Moore, R. (2017). Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*. https://doi.org/10.1016/j.rse.2017.06.031

Grover, A., Kumar, S., & Kumar, A. (2019). Ship detection using sentinel-1 sar data. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*.

Huth, P. K. (1996). *Standing your ground: Territorial disputes and international conflict*. University of Michigan Press.

Kornfein, C., Willard, F., Tang, C., Long, Y., Jain, S., Malof, J., & Bradbury, K. (2023). Closing the domain gap: Blended synthetic imagery for climate object detection. *Environmental Data Science*, *2*, e39. https://doi.org/10.1017/eds.2023.33

Leetaru, K., & Schrodt, P. A. (2013). GDELT: Global data on events, location, and tone. *ISA Annual Convention*.

Mäyrä, J., Virtanen, E., Jokinen, A.-P., Koskikala, J., Väkevä, S., & Attila, J. (2024, May 14). Mapping recreational marine traffic from sentinel-2 imagery with yolov8. https://doi.org/10.2139/ssrn.4827287

OpenStreetMap contributors. (2019). *OpenStreetMap*. Retrieved February 20, 2025, from https://www.openstreetmap.org/#map=8/47.714/13.349

Park, J., Lee, J., Seto, K., Hochberg, T., Wong, B. A., Miller, N. A., Takasaki, K., Kubota, H., Oozeki, Y., Doshi, S., Midzik, M., Hanich, Q., Sullivan, B., Woods, P., & Kroodsma, D. A. (2020). Illuminating dark fishing fleets in north korea. *Science Advances*, *6*(30), eabb1197. https://doi.org/10.1126/sciadv.abb1197

Prete, R. D., & Meoni, G. (2024). VDS2raw - VFNet ResNet-18 FPN (weights). https://doi.org/10.5281/zenodo.10637982

Prete, R. D., Salvoldi, M., Barretta, D., Longépé, N., Meoni, G., Karnieli, A., Graziano, M. D., & Renga, A. (2024, November 5). Enhancing maritime situational awareness through end-to-end onboard raw data analysis. https://doi.org/10.48550/arXiv.2411.03403

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016, May 9). You only look once: Unified, real-time object detection. https://doi.org/10.48550/arXiv.1506.02640

Ren, S., He, K., Girshick, R., & Sun, J. (2016, January 6). Faster r-CNN: Towards real-time object detection with region proposal networks. https://doi.org/10.48550/arXiv.1506.01497

South China Sea Data Initiative. (2025). *Data - south china sea data 2012-2023* [South china sea data initiative]. Retrieved February 7, 2025, from https://scsdi.org/data/

Story, R. (2025). *Folium - folium 0.19.4 documentation* [Python visualization]. Retrieved February 20, 2025, from https://python-visualization.github.io/folium/latest/

Sun, X., Brewin, R. J., Sathyendranath, S., Dall'Olmo, G., Airs, R., Barlow, R., Bracher, A., Brotas, V., Kheireddine, M., Lamont, T., Marañón, E., Morán, X. A. G., Raitsos, D. E., Shen, F., & Tilstone, G. H. (2023). Coupling ecological concepts with an ocean-colour model: Phytoplankton size structure. *Remote Sensing of Environment*, *285*, 113415. https://doi.org/https://doi.org/10.1016/j.rse.2022.113415

UNDP (United Nations Development Programme). (2022). *2022 special report on human security*. New York.

Walter, B. F. (2009). *Reputation and civil war: Why separatist conflicts are so violent*. Cambridge University Press.

Wang, N., Er, M. J., Chen, J., & Wu, J. G. (2022). Marine object detection based on improved YOLOv5. *2022 5th International Conference on Intelligent Autonomous Systems (ICoIAS)*, 43–47. https://doi.org/10.1109/ICoIAS56028.2022.9931205

Watch, G. F. (2025). Datasets and code: Vessel identity.

Witjes, N., & Olbrich, P. (2017). A fragile transparency: Satellite imagery analysis, non-state actors, and visual representations of security. *Science and Public Policy*, *44*(4), 524–534. https://doi.org/10.1093/scipol/scw079

Yang, D., Solihin, M. I., Ardiyanto, I., et al. (2024). A streamlined approach for intelligent ship object detection using EL-YOLO algorithm. *Scientific Reports*, *14*, 15254. https://doi.org/10.1038/s41598-024-64225-y

YOLOv8. (2024, January 1). *YOLOv8: Object detection algorithm for accurate recognition*. Retrieved February 7, 2025, from https://yolov8.org/

Yu, C., Yin, H., Rong, C., Zhao, J., Liang, X., Li, R., & Mo, X. (2024). YOLO-MRS: An efficient deep learning-based maritime object detection method for unmanned surface vehicles. *Applied Ocean Research*, *153*, 104240. https://doi.org/https://doi.org/10.1016/j.apor.2024.104240

Zhang, H., Wang, Y., Dayoub, F., & Sünderhauf, N. (2021, March 4). VarifocalNet: An IoU-aware dense object detector. https://doi.org/10.48550/arXiv.2008.13367