GIST: Gear Type Identification by Spatiotemporal Trajectory Transformation for Monitoring Fisheries

Amir Yaghoubi Shahir[1], **Tilemachos Charalampous**[1], **Mahsa Keramati**[1], **Fatemeh Movafagh**[1], **Uwe Glässer**[1], **Hans Wehn**[2]

[1]School of Computing Science, Simon Fraser University, Burnaby, BC, Canada
[2]MDA Space Ltd., Richmond, BC, Canada
{sayaghou, tcharala, mkeramat, fma44, glaesser}@sfu.ca, hans.wehn@mda.space

Abstract

Illegal, Unreported, and Unregulated (IUU) fishing aggravates the global crisis caused by overfishing, threatening the sustainability of marine ecosystems and fisheries worldwide. Distinctive operational characteristics of fishing vessels result in unique footprints on marine environments and socio-economic structures, depending on their fishing method and gear type such as trawlers with non-selective gear that disrupts the seabed, purse seiners using Fish Aggregating Devices (FADs), and longliners notorious for high bycatch rates. As these vessels play an essential role in commercial fishing and the industry, effective monitoring, regulation, and enforcement are critical to mitigate the devastating consequences of overfishing and promote sustainable fishing practices. To this end, this paper introduces a novel multi-stage method for Gear type Identification by Spatiotemporal trajectory Transformation (**GIST**). This method proposes a *data-centric* approach that employs domain knowledge to facilitate the deployment of an efficient and accurate analysis of operational patterns of fishing vessels derived from Automatic Identification System (AIS) data. Our method first extracts fishing patterns from vessel trajectories to refine data integrity and isolate only the most relevant activities, thereby ensuring a more accurate result. Next, it encapsulates the distributional insights of fishing activities into fixed-sized "images" as actionable input for a multi-class CNN-based classifier. Utilizing **GIST** bypasses complicated linear analyses of time series data for rendering lengthy trajectories, advancing an efficient gear type identification with 97% accuracy. To the best of our knowledge, **GIST** is the first to use a multi-stage pipelining method to distinguish three principal gear types widely used globally. Our experiments confirm **GIST**'s practicability and effectiveness, marking a significant advancement towards stricter enforcement of regulations in the fight against IUU fishing.

1 Introduction

IUU fishing is a global threat to marine ecosystems and fisheries, and is estimated to account for up to 30% of all fishing worldwide, causing economic damage of approximately $23 billion (USD) per year (Food and Agriculture Organization of the United Nations 2020). Effective monitoring and enforcement of fishing regulations is crucial to mitigate the devastating consequences of overfishing and promote sustainable fishing practices (Phelps Bondaroff, Reitano, and van der Werf 2015). With the enormous commercial fishing vessel fleet operating across vast coastal waters and open oceans that cover over 70% of our planet's surface, maritime regulatory authorities face extreme challenges in their efforts to enforce compliance of fisheries with fishing laws and regulations (United Nations Conference on Trade and Development 2018; Global Fishing Watch 2017). In addition, although the most prevalent fishing vessel types play a central role in commercial fisheries, they are frequently involved in IUU fishing. Combating IUU fishing, therefore, necessitates automated detection of fishing activities that deviate from the expected operational characteristics of these vessels. The operational characteristics of fishing vessels are influenced by various factors, including the type of fishing gear, vessel capacity, fishing licenses, and seasonal or geographical restrictions. Among these, the type of fishing gear is the most distinctive, as each gear type targets specific fish species and leaves unique footprints on marine environments and socio-economic structures.

Accurate gear type identification is a fundamental step toward detecting potential IUU fishing, as it can indicate various illegal practices (Deng et al. 2005; Kiparissis et al. 2011), such as targeting restricted species or using unauthorized gear types (Inter-American Tropical Tuna Commission 2021). Once the gear type is identified, the amount of fishing gear of a specific type used on fishing grounds over a given period can be determined, allowing for the estimation of the vessel's fishing effort and the potential overfishing risks associated with it (Food and Agriculture Organization of the United Nations 2001). Additionally, understanding the type of gear and the specifics of each fishing method provides insights into the threats they pose to the sustainability of marine ecosystems, including bycatch and seabed disturbance.

Using rich and reliable Vessel Monitoring System (VMS) data has shown promising results for identifying fishing gear types (Marzuki et al. 2017). However, the limited accessibility of VMS data hinders the timely and broad application of these methods. Therefore, Automatic Identification System (AIS) data—transmitted from marine vessels—is widely used as a compelling alternative for developing gear type identification methods (Kim and Lee 2020; Rodriguez-Albala et al. 2024; Carlos et al. 2021). Although AIS data is

publicly available, it is prone to noise, inconsistencies, and corrupted or missing data, posing challenges for AIS-based gear identification methods. Spatiotemporal information is the most reliable attribute in AIS data, and the operational movement characteristics of fishing vessels are closely tied to their gear type (Xing et al. 2023). As a result, developing practical and efficient methods using trajectory data to identify fishing gear types has become a challenge for law enforcement and fisheries monitoring agencies. Vessel trajectories are often lengthy and vary widely in duration, which complicates the ability of existing methods to maintain consistent performance across different gear types and trajectory lengths, limiting their generalizability to all major gear types. Therefore, adopting a data-centric approach by systematically engineering the trajectory data paves the way for a practical and robust implementation of an efficient high-performance model (Ng 2021).

To address these challenges, we propose Gear type Identification by Spatiotemporal trajectory Transformation (**GIST**), a novel multi-stage pipelining method that combines trajectory data transformation and efficient model design for robust and practical gear type identification. **GIST** is designed to classify three major gear types widely used globally: *trawlers*, *purse seiners*, and *longliners*. First, we use **TripTracker** (Shahir et al. 2021) to extract routine fishing activities from vessel trajectories, enabling a focused analysis of fishing patterns. **TripTracker** is an unsupervised learning framework developed in our previous research, validated by domain experts and currently used in industry, that accurately isolates fishing segments, or *fishing islands*. In the second stage, **GIST** uses the distributions of vessel speed and direction changes during fishing activities to transform fishing islands into fixed-size, image-like representations. This novel transformation encodes key distinctive behavioral information into an interpretable format and allows for consistent handling of the diverse and unstructured spatiotemporal data used for marine monitoring. The inherently skewed distribution of fishing islands across gear types calls for data augmentation techniques to balance the dataset. Finally, the curated image-like data is fed into a 2D convolutional neural network (CNN), which is parameter-efficient, fast, and more interpretable than sequential models like LSTM, as shown in our experiments. This CNN effectively captures spatiotemporal patterns from the image-like representations. By fine-tuning the CNN-based classifier with a hyperparameter optimization framework, we achieve 97% gear type identification accuracy across all three major gear types. Our experiments validate the performance and efficiency of **GIST** through an extensive comparative study, underscoring the importance of each stage in the proposed method.

2 Role of Gear Type in Sustainable Fisheries

Accurate gear type identification is crucial for implementing effective conservation measures—such as marine protected areas and gear-specific regulations—that are essential for mitigating the environmental impacts of fishing methods and promoting sustainability (De Souza et al. 2016; Tassetti, Ferrà, and Fabi 2019). Understanding gear types allows authorities to strategically target their efforts against overfishing, habitat destruction, and marine ecosystem degradation (Gilman, Passfield, and Nakamura 2012).

As discussed earlier, leveraging spatiotemporal information from AIS data is a practical approach for identifying vessel gear types (Kim and Lee 2020). AIS is a self-reporting system that periodically transmits data about maritime vessels via an onboard transceiver (MarineTraffic 2021). Required by the SOLAS international maritime treaty, AIS ensures safety compliance by providing both static information (e.g., Maritime Mobile Service Identity (MMSI), a unique vessel identifier) and dynamic data, such as GPS coordinates, Speed Over Ground (SOG), and Course Over Ground (COG), at regular intervals (typically every minute) (Kim and Lee 2018). Different gear types are designed to target specific fish species or groups of species (NOAA Fisheries 2024) and are used for fishing methods with distinct operational characteristics. Therefore, vessel movement patterns are closely associated with fishing gear types, and detailed dynamic information from SOG and COG provides valuable insights for reliable gear identification.

AIS-based gear type identification approaches take advantage of machine learning and deep learning model training. AIS is mostly used to train various discriminative models for diverse classification goals (e.g., fishing vs. non-fishing activity) other than gear identification (Chuaysi and Kiattisin 2020; Chen et al. 2020; Sturgis et al. 2024; Sans and Rodriguez 2023; Ferreira et al. 2022; Wang et al. 2021), while there are few methods addressing gear type identification (Kim and Lee 2020; Rodriguez-Albala et al. 2024; Carlos et al. 2021). Feeding trajectory data to stacked 1D CNNs (Kim and Lee 2020) achieves high performance for some gear types; however, the model's performance varies across them. Rodriguez et al. (Rodriguez-Albala et al. 2024) empirically show that extracting local and global features from the online signature verification domain improves general gear classification accuracy. Moreover, Carlos et al. (Carlos et al. 2021) propose a hybrid approach that combines supervised autoencoder dimensional reduction with traditional classifiers to identify specific fishing gear behaviors from AIS data.

These methods demonstrate the significance of spatiotemporal analysis of AIS dynamic information in monitoring and mitigating IUU fishing. However, their performance varies across gear types, requires costly model training, and focuses on gear types not widely used in the industry, limiting their practical deployment. Our proposed method addresses these challenges with a multi-stage design that transforms trajectory data into interpretable, image-like representations and employs an efficient CNN-based architecture to effectively learn spatiotemporal patterns. Additionally, **GIST** focuses on the three gear types most commonly used in the industry, achieving high classification performance across all three gear types: trawlers, purse seiners, and longliners.

Transforming trajectory time series data into images is a common practice in maritime applications, facilitating the

exploration and identification of underlying patterns in vessel navigation for maritime traffic surveillance. Several studies transform AIS data into trajectory images (Zhang et al. 2022; Chen et al. 2020) or use density maps, space-time cubes, relation graphs, Kernel Density Estimation (KDE), and stacked 3D trajectory bands (Liu et al. 2021). These techniques convert trajectory data into visually interpretable forms; however, they are primarily suited for maritime activity modeling, limiting their applicability for detecting IUU fishing activities.

To elucidate the significance of gear type classification, the next section delves into the details of the three fishing gear types predominantly used by industrial fishing vessels, highlighting their distinctive characteristics and operational implications.

2.1 Operational Fishing Vessel Characteristics

Here, we detail the operational characteristics of the three major fishing vessel types—trawlers, purse seiners, and longliners—focusing on their fishing methods, movement patterns, and gear types. Figure 1 highlights the unique operational features and implications of the three vessel types by comparing their movement patterns during their fishing activity extracted by **`TripTracker`**.

Trawlers drag nets either along the sea floor (bottom trawling) or in midwater (pelagic trawling), operating at slow, steady speeds between 2 and 6 knots during fishing operations, which typically last 3 to 5 hours. Their movement patterns generally consist of straight lines with occasional turns to adjust to seabed topography or initiate new paths. When traveling to or from fishing grounds, trawlers cruise at speeds of 8 to 10 knots (labeled as "Steaming" in Figure 1). They target various species, including cod, haddock, sardines, and anchovies. Trawling is not a selective fishing method and often results in the unintended capture of non-target species (TMT and Network 2023).

Purse Seiners move around schools of fish to deploy their nets in circular or arc patterns. Once the fish are surrounded, they slow down to about 2 knots or less to encircle the school and close the bottom of the net (pursing) before hauling in the catch. These nets can be extraordinarily large, reaching up to 3000 meters in length, 300 meters in depth, and weighing around 6 tons. Purse seiners primarily target tuna and other species that aggregate in large schools, often utilizing Fish Aggregating Devices (FADs) to improve operational efficiency. Approximately 66% of tuna caught globally each year is captured by purse seiners, which often operate far from shore where tuna species are abundant (V. Restrepo and Koehler 2024; Tracking and Network 2021b).

Longliners deploy lines that can extend up to tens of kilometers, equipped with hundreds to thousands of baited hooks. These lines are set at various depths and left in the water for hours, with minimal vessel movement during soak times. Longliners typically move slowly at speeds of 2 to 6 knots in straight lines while setting and hauling lines, as depicted in Figure 1. The fishing trip duration varies significantly: coastal vessels operate shorter trips lasting less than

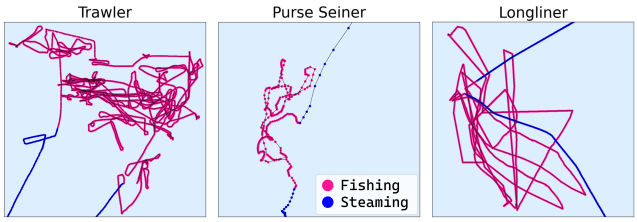

Figure 1: Trajectories of Major Fishing Vessel Types

20 days, while large, distant-water longliners can remain at sea for 140 to 180 days, covering vast oceanic regions. Longliners commonly target deep-sea species such as tuna (albacore, bigeye, and yellowfin), swordfish, sharks, and halibut, which are attracted to the bait (De Souza et al. 2016; Tracking and Network 2021a).

In the following section, we formally define the problem of gear type identification based on spatiotemporal vessel trajectories.

2.2 Problem Statement and Notations

A **vessel** $\mathrm{v} \in \mathbb{V}$ is a marine vessel capable of engaging in fishing activities and is equipped with AIS to continuously transmit AIS reports. Each **report**, $r_{\mathtt{t}}^{\mathrm{v}}$, provides detailed static and dynamic information about v at time $\mathtt{t}$.

A **trajectory** of a fishing vessel, $\mathcal{T}_{\mathtt{T}}^{\mathrm{v}}$, is a temporally ordered set of reports received from v during time period $\mathtt{T}$:

$$\mathcal{T}_{\mathtt{T}}^{\mathrm{v}} = \{r_{\mathtt{t}}^{\mathrm{v}} \mid \mathtt{t} \in \mathtt{T}\}$$

Each fishing vessel is designed for a specific fishing method and gear type to efficiently target certain fish species. The **gear type** associated with v is denoted as g^{v}.

> Suppose a fishing vessel v is engaged in fishing activity at some point along its trajectory. The objective is to learn a function $\gamma(\cdot)$ that accurately classifies the gear type of v based on the information in $\mathcal{T}^{\mathrm{v}}$.

$$\gamma(\mathcal{T}^{\mathrm{v}}) = \mathrm{g}^{\mathrm{v}}$$

3 Methodology

Figure 2 illustrates the multi-stage outline of **`GIST`**, starting with a fishing vessel trajectory (a). First, **`TripTracker`** extracts the fishing island(s) from the trajectory (b). The extracted fishing island is then divided into 128 equal segments (c), with Speed Over Ground (SOG) and changes in Course Over Ground (ΔCOG) values binned for each segment (d). These binned segments are stacked to form the image-like representation (e), which is ultimately fed into a CNN-based classifier to identify the gear type (f). Each stage is explained in detail in the following sections.

3.1 Fishing Islands

Fishing vessels typically engage in various activities along their trajectories, with stationary, steaming, and fishing being the routine ones. Stationary patterns exhibit similar

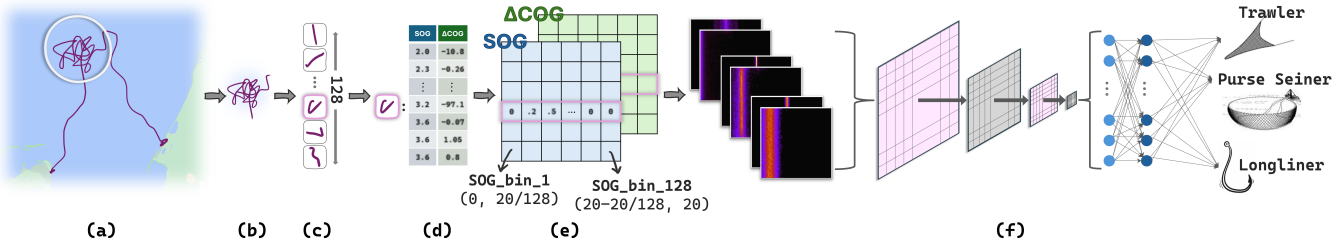

Figure 2: **GIST** Outline: From Vessel Trajectory to Gear Type Identification

movement across different vessel types, and steaming patterns are also not significantly distinct. However, the operational characteristics of fishing activities, as described in Section 2.1, are indicative of each vessel type. Using the entire trajectory of a vessel for gear type identification can degrade accuracy, as the similarity of non-fishing activities across vessel types dilutes the distinctiveness of fishing patterns. Therefore, precisely isolating fishing activity from the rest of the trajectory data is crucial for accurate and robust gear type identification (Zha et al. 2023; Chung et al. 2019).

We use **TripTracker** to extract fishing islands for our gear type identification method. **TripTracker** is an unsupervised learning framework—developed in our previous research, evaluated by domain experts, and currently used in industry—that effectively detects routine fishing vessel activities. Figure 3 compares the Speed Over Ground (SOG) of fishing and steaming activities extracted by **TripTracker** for the three different vessel types, confirming the robustness of its activity segmentation and aligning with the operational movement characteristics reported in the literature.

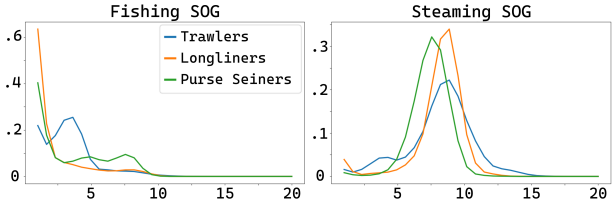

Figure 3: Distribution of Speed Over Ground for Vessel Types During Fishing and Steaming Activities

Trawlers exhibit slow, consistent speeds (up to 4 knots) during fishing activity, with a distinct peak at 8–10 knots during steaming activity (TMT and Network 2023). Purse seiners generally operate at very low speeds (0–3 knots), while actively encircling fish schools and show moderate peaks (6–9 knots) during steaming (Tracking and Network 2021b). Longliners usually maintain slow speeds (up to 6 knots) during their fishing operation, with a peak range of 8–10 knots when steaming between different fishing grounds (Tracking and Network 2021a). It is worth noting that, while this figure does not provide sufficient detail for direct gear type identification, it confirms that **TripTracker** accurately segments fishing and steaming activities, as its outputs align with established movement patterns in the literature.

A **fishing island**, $f_{\mathtt{T}}^{\mathtt{v}}$, extracted using **TripTracker**, is a subset of the trajectory of fishing vessel $\mathtt{v}$, where $\mathtt{v}$ is continuously engaged in fishing activity during $\mathtt{T}$. Formally:

$$f_{\mathtt{T}}^{\mathtt{v}} = \{r_{\mathtt{t}}^{\mathtt{v}} \mid \forall \mathtt{t} \in \mathtt{T},\ r \text{ indicates a fishing pattern}\}$$

3.2 Image-like Representations of Fishing Islands

Fishing vessels, when engaged in fishing activities, typically move at slow speeds and make consecutive turns, with patterns varying based on the fishing method and gear type. Therefore, as described in Section 2, to classify fishing vessels, we use only features from AIS data that represent vessel movement.

▷ **Extracting all fishing islands.** As discussed in Section 3.1, we focus solely on fishing islands within vessel trajectories, as these segments contain the most informative patterns for distinguishing between various fishing methods and gear types (Figure 2 (a) and (b)). By enhancing the quality of input data through targeted extraction, such as isolating fishing activity, we improve the robustness and accuracy of gear type classification. After all, a model can only be as good as the data it is fed (Zha et al. 2023). Formally:

$$\forall \mathcal{T}^{\mathtt{v}}, \exists \{f_i^{\mathtt{v}}\} \quad \text{such that} \quad f_i^{\mathtt{v}} \subseteq \mathcal{T}^{\mathtt{v}} \text{ and } |\{f_i^{\mathtt{v}}\}| \geq 0$$

▷ **Partitioning fishing islands into 128 equal segments.** Segments within a fishing island are non-overlapping and equal in size, although segments from different fishing islands may vary due to significant differences in fishing island sizes. Fishing vessel trajectories exhibit complex yet comparable patterns during fishing activities. Vessels of the same type follow similar actions in their fishing methods, from deploying the gear to hauling it, with each action dictating specific movement behaviors. Variations in the total length of fishing activities using the same gear arise from differences in the duration and repetition of these actions. Trawlers maintain low speeds with minimal changes in course, purse seiners exhibit faster, circular patterns around schools of fish, and longliners show higher speeds during line setting and minimal movement during soak times—each reflecting distinct fishing methods. Given the consistency of behavior within each fishing method, we expect to observe closely related patterns in similar segments across different fishing islands of the same vessel type or recurring patterns specific to that gear type (Figure 2 (c)).

A fishing island qualifies for transformation into an image-like representation only if its size exceeds 128, meaning the minimum duration of fishing activity is longer than

two hours. This step ensures that smaller fishing islands, which may not adequately represent a fishing method, are filtered out.

▷ **Binning features of segments within fishing islands.** As described earlier, only SOG and ΔCOG from AIS data are used to represent the movement of a fishing vessel, and each feature has a distinct range. For clarity, SOG can range from 0 to 20 knots, while ΔCOG can vary from -180 to 180 degrees. Within the range of each feature, 128 non-overlapping, equal-width sub-intervals (bins) are defined, with each bin represented by distinct edges. Formally:

$$\text{SOG\_bin}_{\,i} = \left((i-1)\cdot\frac{20}{128}, i\cdot\frac{20}{128}\right) \quad \text{where} \quad 1 \le i \le 128$$

$$\Delta\text{COG\_bin}_{\,i} = \left(-180+(i-1)\cdot\frac{360}{128}, -180+i\cdot\frac{360}{128}\right) \quad \text{where} \quad 1 \le i \le 128$$

Next, for each segment within a fishing island, histograms are calculated for both features based on their respective bins. In particular, the values of each feature within a segment are distributed across the 128 bins, with each bin counting the number of values falling within its defined range.

▷ **Creating image-like representations from segment histograms.** Each fishing island is divided into 128 segments, with each segment representing its corresponding row in the image-like structure (*i*-th segment corresponds to the *i*-th row). Similarly, each of the 128 bins corresponds to a column in the image-like structure (*j*-th bin corresponds to the *j*-th column). Thus, in this 128×128 image-like structure, $\text{pixel}_{i,\,j}$ represents bin_j in the *i*-th segment. Finally, the values of all bins for each feature are normalized, and the normalized values—ranging in $[0,\ 1]$—are treated as the intensity of their corresponding pixels, creating a two-channel image (SOG and ΔCOG) for each fishing island (Figure 2 (d) and (e)).

An **image-like representation**, $\mathtt{I}_f$, is a transformation of the fishing island f within a given $\mathcal{T}^{\text{v}}$, into a 128×128 image. It comprises two channels: one for the vessel's speed (SOG), and one for changes in the vessel's course (ΔCOG).

Actionable Insights in Image-likes

In production scenarios, data and models must be continuously updated to maintain consistent and reliable results. Proper data maintenance is crucial and begins with a thorough understanding of the data. Real-world data is massive and complex, making its analysis challenging. A concise, low-dimensional yet informative representation enhances the practicality of analysis, particularly since human perception is typically limited to two or three dimensions (Zha et al. 2023). Transforming a fishing activity time series into a fixed-size "image", regardless of its length, while preserving key information about vessel movement characteristics, provides actionable insights. This approach facilitates efficient

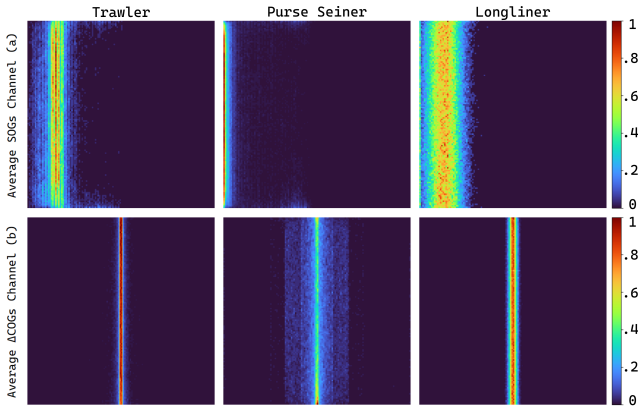

Figure 4: Avg. Image-like Representations of Vessel Types

fine-tuning or retraining of the model and enables targeted analysis of edge cases or misclassified samples.

SOG Channel: When vessel speed is low, bins corresponding to higher speeds have zero values. As a result, the left side of the SOG channel is brighter, with higher pixel intensity indicating low-speed fishing activities. Figure 4 (a) illustrates the average SOG channel for each vessel type. For example, trawlers exhibit a bright vertical line on the left side with relatively low diversity, reflecting their slow and steady speeds during trawling operations.

ΔCOG Channel: The ΔCOG channel ranges between -180 and 180, with the middle column representing a 0° change in course. During fishing, vessels either move straight or turn gradually, resulting in a bright middle column with a gradient on either side. Figure 4 (b) shows the average ΔCOG channel for each vessel type. For instance, purse seiners display the widest gradient around the middle column, capturing their characteristic circular or arc-like movements around schools of fish.

3.3 Gear Type Identification Using Image-likes

To classify the image-like representations derived from fishing islands, each with the shape (128, 128, 2), we design a convolutional neural network (CNN)-based classifier comprising a convolutional backbone and a classifier. The backbone stacks several convolutional blocks, and the classifier outputs a multi-class probability distribution. The gear type of different fishing vessels—trawler, purse seiner, or longliner—serves as the three output classes.

$$\forall f^{\text{v}}, \text{Classifier}_{\text{CNN}}(f^{\text{v}}) = \text{g}^{\text{v}}$$

Each convolutional block includes a Conv2D (2D Convolution), a MaxPooling2D, and a Dropout layer. The Conv2D layers extract in-depth features from the input, employing ReLU or Tanh activation functions to introduce non-linearity. This is followed by MaxPooling2D layers, which downsample the feature maps to emphasize the most salient features. To prevent overfitting, Dropout layers randomly omit subsets of features during training, ensuring the model generalizes well to new data. In the classifier, fully connected (Dense) layers process the extracted features and

classify the input into one of the three gear types (Figure 2 (f)).

This network architecture ensures robust feature extraction followed by accurate and efficient classification, culminating in a Dense output layer that uses a Softmax activation function to generate the probability distribution across the classes.

4 Experiments

In this section, we assess the effectiveness of our proposed method for identifying vessels' gear types. We first elaborate on the characteristics of the datasets and detail the process of addressing skewed data distribution through augmentation with synthetic samples. Next, we provide a brief overview of the implementation setup, followed by a comprehensive analysis of our experiments to demonstrate the efficacy of our approach.

4.1 Data Characteristics

For the experimental analysis, we utilize AIS data collected by the U.S. Coast Guard to monitor and track marine vessels operating in coastal waters around North America (NOAA Office for Coastal Management 2017). AIS reports are filtered to one report per minute for each vessel and stored in geodatabases organized by Universal Transverse Mercator (UTM) Zones. This study focuses on the West Coast (UTM Zones 1 to 11) for the calendar year 2017.

AIS data provide valuable information about vessel trajectories; however, it is prone to manipulation, missing values, and noise (Harati-Mokhtari et al. 2007). Therefore, as a primary step in data preprocessing, we linearly interpolate AIS data using the great circle method (Oceanic and NOAA) to reconstruct continuous vessel trajectories. Among the various features in AIS reports, we focus on Speed Over Ground (SOG) and change in Course Over Ground (ΔCOG), as they contain the most relevant information for analyzing vessel movements. Erroneous values are inevitable in AIS data, but while latitude and longitude are considered the most reliable features, we take an additional preprocessing step: recalculating both speed and course at each data point using kinematic and trigonometric equations based on vessel coordinates. This step enhances the reliability of the data.

Next, we label vessel trajectories with their corresponding gear types by mapping the MMSI numbers of vessels to the Vessel Identity dataset provided by Global Fishing Watch (GFW) (Global Fishing Watch 2023). This dataset contains pairs of MMSI numbers and their associated gear types, categorizing them into five vessel types: 'trawlers', 'fixed_gear', 'other_fishing', 'purse_seines', and 'drifting_longlines'. For this research, we exclude 'fixed_gear' and 'other_fishing' categories, as each encompasses multiple fishing methods rather than a single type of vessel or gear. According to the dataset description, 'fixed_gear' includes three distinct categories: pots and traps, set longlines, and set gillnets, each employing different fishing methods.

We analyze AIS data from over 1,200 vessel trajectories, of which 892 vessels have matching MMSI numbers in the Vessel Identity dataset, allowing their types to be classified. Specifically, these include 349 trawlers, 100 purse seiners, and 66 longliners, while the excluded categories comprise 241 fixed gear vessels and 136 other fishing vessels.

4.2 Data Augmentation

As mentioned in Section 4.1, trajectories with no fishing islands or those shorter than 128 minutes are filtered out, resulting in 5,773 fishing islands—extracted from 892 vessel trajectories—with a skewed distribution across gear types. Specifically, our transformation method produces 5,067 image-like representations for trawler vessels, 663 for purse seiners, and only 43 for longliners. To mitigate the learning bias toward the majority class, we generate synthetic samples to augment the minority class and balance the dataset.

To maintain the model's generalization capability for real-world scenarios, we maximize the use of real samples in the training dataset and minimize reliance on synthetic samples. Consequently, we only augment the longliner class, ensuring each class has 663 image-like representations.

Data Augmentations in Different Modalities

We use Generative Adversarial Networks (GANs) to generate synthetic samples for the longliner class. GANs are widely recognized for producing high-quality, realistic samples and are commonly used to augment imbalanced datasets (Chatziagapi et al. 2019; Fawakherji et al. 2020; Motamed, Rogalla, and Khalvati 2021; Meor Yahaya and Teo 2023). In our approach, data augmentation is performed using two modalities: first, by augmenting the sequential trajectory data before transforming them into image-like representations; and second, by directly augmenting the image-like representations. We utilize the Synthetic Data Vault (SDV) Python library (Patki, Wedge, and Veeramachaneni 2020), which supports synthetic data generation across various modalities, including sequential data. The following sections elaborate on the details of our data augmentation process using SDV for both sequential and image-like data.

Augmentation of Sequential Data. The PAR Synthesizer within the SDV library is employed to generate synthetic sequences for the longliner class by learning the SOG and ΔCOG patterns from the time series data of real fishing islands. The synthesizer is trained for over 500 epochs, generating 750 synthetic sequences. High-quality sequences longer than 128 time steps are then selected, resulting in 663 viable samples for the longliner class.

Augmentation of Image-like Data. To augment the longliner class in the image-like modality, we use the Conditional Tabular GAN (CTGAN) from the SDV library. CTGAN is designed to generate synthetic tabular data that aligns with the structure of our image-like representations. After training for 750 epochs, CTGAN generates 620 synthetic image-like representations for the longliner class.

Comparative Analysis of Synthetic Data

Figure 5 compares the SOG channel of the image-like representations of a real longliner sample with synthetic counterparts generated in both sequential and image-like modalities.

Modality\Metric	Per gear type									Overall						
	Trawler			Purse Seiner			Longliner			Macro Average			Weighted Average			Acc.
	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	
Sequential	0.98	0.97	0.98	0.98	0.97	0.97	0.96	0.97	0.97	0.97	0.97	0.97	0.97	0.97	0.97	0.97
Image-like	0.97	0.95	0.96	0.96	0.98	0.96	0.99	0.99	0.99	0.96	0.97	0.97	0.97	0.96	0.96	0.97

Table 1: Comparative analysis of data augmentation in sequential modality and image-like modality

The top row of Figure 5 demonstrates that synthetic samples generated in the sequential modality exhibit greater diversity in speed, while those generated in the image-like modality are more similar to the real samples.

Table 1 shows that augmenting data with synthetic samples generated in the image-like modality yields classification performance comparable to the sequential modality, with a slight improvement for the longliner class.

To evaluate the synthetic samples, we mapped pixel values to a logarithmic scale to emphasize near-zero values and highlight discrepancies between synthetic and real samples, as shown in the bottom row of Figure 5. These discrepancies arise because transforming a fishing island into an image-like representation typically results in many zero-valued pixels, as fishing vessels move at low speeds during fishing activity, leaving bins representing faster speeds with values of zero. However, due to the probabilistic nature of generative models, these zero-valued pixels are often approximated with values close to zero. Consequently, synthetic samples in the image-like modality inevitably differ from real samples.

This discrepancy can mislead the model into learning spurious features, as the approximated near-zero values can be interpreted as defining characteristics of the longliner class. Such spurious features can overshadow the true discriminative features, harming the model's generalizability. Empirical observations confirm that this issue is particularly evident in both the SOG and ΔCOG channels.

While our experiments indicate that data augmentation in the image-like modality achieves slightly better results, we ultimately use synthetic samples generated in the sequential modality to preserve the model's generalizability.

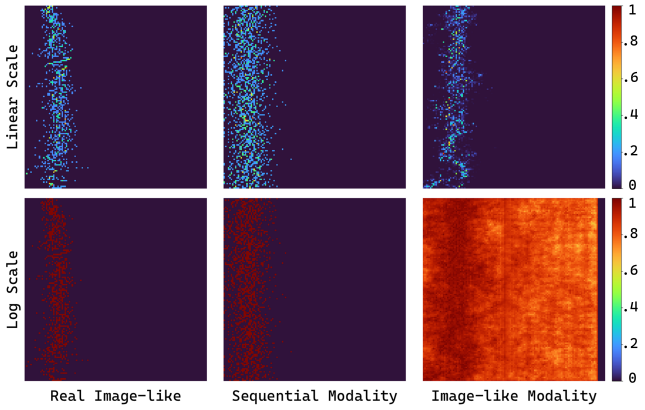

Figure 5: Comparison of SOG Channel for Real and Synthetic Samples: Sequential vs. Image-like Modalities

4.3 Architecture Design

Our gear type classification model, introduced in Section 3.3, is a convolutional neural network (CNN)-based classifier comprising a convolutional backbone and a classifier. The backbone stacks several convolutional blocks to extract robust features from the input, while the classifier outputs a multi-class probability distribution for the gear types—trawler, purse seiner, or longliner.

To optimize the hyperparameters of this architecture for maximum performance, we employ KerasTuner (O'Malley and et al. 2021), a scalable hyperparameter optimization framework that systematically explores a wide range of configurations. KerasTuner evaluates different model architectures across 100 iterations, searching for the configuration that maximizes our model's classification accuracy.

During this process, we explore up to eight convolutional blocks, varying filter sizes (16, 32, 64, 128), kernel sizes (3×3, 5×5, 7×7), padding strategies (valid, same), and activation functions (ReLU, Tanh). Additional components like MaxPooling and Dropout layers are also tested to enhance the model's robustness and prevent overfitting.

The optimal architecture consists of four convolutional blocks. Each block includes a Conv2D layer, followed by a MaxPooling2D layer and a Dropout layer. The first Conv2D layer uses a 3×3 kernel, 16 filters, ReLU activation, and a dropout rate of 0.1. The final Conv2D layer employs a 3×3 kernel, 32 filters, ReLU activation, and a dropout rate of 0.5 for regularization.

The classifier is composed of two fully connected (Dense) layers. The first Dense layer employs a Tanh activation function, while the final Dense layer uses a Softmax activation function to output the probability distribution across the three gear type classes.

4.4 Experimental Setup

We split the dataset into 60% for training, 20% for validation, and 20% for testing. This split results in 398 samples per class for training, 133 samples per class for validation, and an additional 133 samples per class for testing. To evaluate our experiments under consistent conditions, all trainings are conducted on an NVIDIA GeForce RTX 2080 Ti GPU. The models are trained for 100 epochs using a batch size of 16 and an initial learning rate of 1e-3.

4.5 Empirical Analysis

Baselines. To empirically evaluate the efficacy of our proposed method, we compare its classification performance against several baselines: LSTM-seq, an LSTM network trained on sequential time series data of fishing islands in

Method\Metric	Per Gear Type									Overall						Acc.
	Trawler			Purse Seiner			Longliner			Macro Average			Weighted Average			
	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1	
LSTM-seq	0.91	0.45	0.60	0.89	0.93	0.91	0.66	0.96	0.78	0.82	0.78	0.76	0.82	0.78	0.76	0.78
LSTM-img	**0.98**	0.89	0.93	0.96	**0.99**	**0.98**	0.92	**0.97**	0.95	0.96	0.95	0.95	0.95	0.95	0.95	0.95
GIST-w/oDA	0.95	0.96	0.95	**0.98**	0.98	**0.98**	0.56	0.50	0.53	0.83	0.81	0.82	0.95	0.95	0.95	0.95
GIST	**0.98**	**0.97**	**0.98**	**0.98**	0.97	0.97	**0.96**	**0.97**	**0.97**	**0.97**	**0.97**	**0.97**	**0.97**	**0.97**	**0.97**	**0.97**
GIST *vs.* LSTM-seq	7%	52%	38%	9%	4%	6%	30%	1%	19%	15%	19%	21%	15%	19%	21%	19%
GIST *vs.* LSTM-img	0%	8%	5%	2%	-2%	-1%	4%	0%	2%	1%	2%	2%	2%	2%	2%	2%
GIST *vs.* **GIST**-w/oDA	3%	1%	3%	0%	-1%	-1%	40%	47%	44%	14%	16%	15%	2%	2%	2%	2%

Table 2: Performance evaluation of our proposed method compared to baselines. Best performance per metric is in **bold**, and the performance gains of our method over baselines are highlighted in blue.

vessel trajectories, using SOG and ΔCOG features; LSTM-img, an LSTM trained on flattened image-like representations treated as sequential data; **GIST**-w/oDA, our method without data augmentation, allowing us to evaluate the effectiveness of the method and the data augmentation independently; and **GIST**, our method with data augmentation in the sequential modality. It is worth noting that LSTM-seq and LSTM-img are trained using the same augmentation modality as **GIST**.

Gear Type Detection Performance. Table 2 presents the gear type classification results in terms of precision, recall, F1-score, and accuracy. Our method, **GIST**, consistently outperforms LSTM-seq, demonstrating its effectiveness. Moreover, **GIST** achieves superior performance compared to **GIST**-w/oDA, particularly for the longliner class, underscoring the positive impact of data augmentation on the model's generalization. While the performance of **GIST** and LSTM-img is comparable—expected since both models are fed essentially similar input data and leverage powerful deep learning architectures—**GIST** offers significant practical advantages. For this comparison, both models were trained for 400 epochs with a batch size of 16. However, as shown in Table 3, **GIST** has four times fewer parameters and five times lower FLOPs (Floating Point Operations per second), a measure of computational complexity quantifying the operations needed to process input data (Sovrasov 2018-2024). Lower FLOPs translate to faster runtime (142 ms vs. 951 ms), making **GIST** significantly more efficient.

Additionally, our experiments confirm that producing image-like representations enhances categorical contrasts, avoiding the challenges of using sequential models like LSTMs for vessel trajectories with extreme length variations. This establishes **GIST** as an efficient, fast, and high-performing solution, particularly well-suited for real-world applications.

Model	# Params	FLOPs	Runtime (ms)
LSTM-img	214,851	558,985	951 ± 8.02
GIST	54,195	107,850	142 ± 1.84

Table 3: Complexity and Efficiency: CNN vs. LSTM

5 Conclusions

Identifying the types of gear used by commercial fishing vessels provides actionable insights at the operational and tactical levels to combat illegal fishing and is crucial to mitigating significant threats to global fish stocks and marine biodiversity. A practical gear identification method must prioritize robust performance, efficiency, and data accessibility for timely intervention against illegal fishing activity.

Our proposed analytical framework, **GIST**, transforms publicly available spatiotemporal marine traffic data into sensible input for an efficient CNN-based classifier, achieving an impressive 97% accuracy in distinguishing three major commercial gear types: trawlers, purse seiners, and longliners. By focusing solely on the most informative trajectory segments—fishing activities—**GIST** eliminates irrelevant patterns, enhancing classification performance. Furthermore, transforming time series trajectories into interpretable image-like representations provides valuable insights, while bypassing the complexity of sequential processing, significantly improving efficiency. Extensive experiments confirm **GIST** as a flexible and robust foundation to advance the monitoring of commercial fisheries, a pivotal step in enforcing fishing regulations and promoting sustainable marine practices.

Acknowledgments

This research evolved from several collaborative research projects on maritime domain awareness and security, collectively funded by the Natural Sciences and Engineering Research Council of Canada; Innovation, Science and Economic Development Canada; and MDA Space Ltd., BC. The authors sincerely thank all four anonymous reviewers for their detailed and insightful comments that helped shape the final version of this paper.

References

Carlos, H.; Aranda, R.; Rivera-De Velasco, M.; Rodriguez-Gonzalez, A. Y.; and Méndez-López, M. E. 2021. Fishing gear pattern recognition by including supervised autoencoder dimensional reduction. *IEEE Geoscience and Remote Sensing Letters*, 19: 1–5.

Chatziagapi, A.; Paraskevopoulos, G.; Sgouropoulos, D.; Pantazopoulos, G.; Nikandrou, M.; Giannakopoulos, T.;

Katsamanis, A.; Potamianos, A.; and Narayanan, S. 2019. Data Augmentation Using GANs for Speech Emotion Recognition. In *Interspeech*, 171–175.

Chen, X.; Liu, Y.; Achuthan, K.; and Zhang, X. 2020. A ship movement classification based on Automatic Identification System (AIS) data using Convolutional Neural Network. *Ocean Engineering*, 218: 108182.

Chuaysi, B.; and Kiattisin, S. 2020. Fishing vessels behavior identification for combating IUU fishing: Enable traceability at sea. *Wireless Personal Communications*, 115: 2971–2993.

Chung, Y.; Kraska, T.; Polyzotis, N.; Tae, K. H.; and Whang, S. E. 2019. Slice finder: Automated data slicing for model validation. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, 1550–1553. IEEE.

De Souza, E. N.; Boerder, K.; Matwin, S.; and Worm, B. 2016. Improving fishing pattern detection from satellite AIS using data mining and machine learning. *PloS one*, 11(7): e0158248.

Deng, R.; Dichmont, C.; Milton, D.; Haywood, M.; Vance, D.; Hall, N.; and Die, D. 2005. Can vessel monitoring system data also be used to study trawling intensity and population depletion? The example of Australia's northern prawn fishery. *Canadian Journal of Fisheries and Aquatic Sciences*, 62(3): 611–622.

Fawakherji, M.; Potena, C.; Prevedello, I.; Pretto, A.; Bloisi, D. D.; and Nardi, D. 2020. Data augmentation using gans for crop/weed segmentation in precision farming. In *2020 IEEE conference on control technology and applications (CCTA)*, 279–284. IEEE.

Ferreira, M. D.; Spadon, G.; Soares, A.; and Matwin, S. 2022. A semi-supervised methodology for fishing activity detection using the geometry behind the trajectory of multiple vessels. *Sensors*, 22(16): 6063.

Food and Agriculture Organization of the United Nations. 2001. Review of the State of World Marine Fishery Resources. https://www.fao.org/3/y3427e/y3427e0a.pdf.

Food and Agriculture Organization of the United Nations. 2020. The State of World Fisheries and Aquaculture 2020. Sustainability in Action. https://doi.org/10.4060/ca9229en.

Gilman, E.; Passfield, K.; and Nakamura, K. 2012. *Performance assessment of bycatch and discards governance by regional fisheries management organizations*. Iucn.

Global Fishing Watch. 2017. Transshipment: The Invisible Activity Putting Ocean Fish Stocks at Risk. Technical report, Global Fishing Watch.

Global Fishing Watch. 2023. Datasets and Code: Vessel Identity. https://globalfishingwatch.org/datasets-and-code-vessel-identity/. Accessed: 2023-01-01.

Harati-Mokhtari, A.; Wall, A.; Brooks, P.; and Wang, J. 2007. Automatic Identification System (AIS): Data Reliability and Human Error Implications. *Journal of Navigation*, 60(3): 373–389.

Inter-American Tropical Tuna Commission. 2021. IUU Vessel List. https://www.iattc.org/GetAttachment/7478a141-eeda-4594-b77c-a88326e693aa/C-19-02%20IUU%20Vessel%20List. Accessed: 2021-12-01.

Kim, K.-I.; and Lee, K. M. 2018. Deep learning-based caution area traffic prediction with automatic identification system sensor data. *Sensors*, 18(9): 3172.

Kim, K.-i.; and Lee, K. M. 2020. Convolutional neural network-based gear type identification from automatic identification system trajectory data. *Applied Sciences*, 10(11): 4010.

Kiparissis, S.; Fakiris, E.; Papatheodorou, G.; Geraga, M.; Kornaros, M.; Kapareliotis, A.; and Ferentinos, G. 2011. Illegal trawling and induced invasive algal spread as collaborative factors in a Posidonia oceanica meadow degradation. *Biological Invasions*, 13: 669–678.

Liu, H.; Chen, X.; Wang, Y.; Zhang, B.; Chen, Y.; Zhao, Y.; and Zhou, F. 2021. Visualization and visual analysis of vessel trajectory data: A survey. *Visual Informatics*, 5(4): 1–10.

MarineTraffic. 2021. What is the Automatic Identification System (AIS)? https://help.marinetraffic.com/hc/en-us/articles/204581828-What-is-the-Automatic-Identification-System-AIS. Accessed: 2022-08-01.

Marzuki, M. I.; Gaspar, P.; Garello, R.; Kerbaol, V.; and Fablet, R. 2017. Fishing gear identification from vessel-monitoring-system-based fishing vessel trajectories. *IEEE Journal of Oceanic Engineering*, 43(3): 689–699.

Meor Yahaya, M. S.; and Teo, J. 2023. Data augmentation using generative adversarial networks for images and biomarkers in medicine and neuroscience. *Frontiers in Applied Mathematics and Statistics*, 9: 1162760.

Motamed, S.; Rogalla, P.; and Khalvati, F. 2021. Data augmentation using Generative Adversarial Networks (GANs) for GAN-based detection of Pneumonia and COVID-19 in chest X-ray images. *Informatics in medicine unlocked*, 27: 100779.

Ng, A. 2021. Data-centric AI Resource Hub. https://snorkel.ai/. Accessed: 2023-02-01.

NOAA Fisheries. 2024. Fishing Gear and Risks to Protected Species. https://www.fisheries.noaa.gov/national/bycatch/fishing-gear-and-risks-protected-species. Accessed: 2024-01-01.

NOAA Office for Coastal Management. 2017. AIS Vessel Tracks 2017. https://www.fisheries.noaa.gov/inport/item/53161. Accessed: 2022-08-01.

Oceanic, N.; and (NOAA), A. A. 2024. Great Circle Distance and Bearing. https://www.nhc.noaa.gov/gccalc.shtml.

O'Malley, T.; and et al. 2021. Keras Tuner. https://keras.io/keras_tuner/. Accessed: 2023-12-01.

Patki, N.; Wedge, R.; and Veeramachaneni, K. 2020. SDV: The Synthetic Data Vault. https://github.com/sdv-dev/SDV. Python library.

Phelps Bondaroff, T. N.; Reitano, T.; and van der Werf, W. 2015. The Illegal Fishing and Organized Crime Nexus: Illegal Fishing as Transnational Organized Crime. https://globalinitiative.net/wp-content/uploads/2015/04/the-illegal-fishing-and-organised-crime-nexus-1.pdf.

Rodriguez-Albala, J. M.; Peña, A.; Melzi, P.; Morales, A.; Tolosana, R.; Fierrez, J.; Vera-Rodriguez, R.; and Ortega-Garcia, J. 2024. Spatio-temporal trajectory data modeling for fishing gear classification. *Pattern Analysis and Applications*, 27(2): 42.

Sans, M.; and Rodriguez, J. 2023. Development and application of a machine-learning and geocomputation workflow for assessing the gear effort of gillnetters operating in the Bay of Biscay. Details pending publication.

Shahir, A. Y.; Charalampous, T.; Tayebi, M. A.; Glässer, U.; and Wehn, H. 2021. TripTracker: unsupervised learning of fishing vessel routine activity patterns. In *2021 IEEE International Conference on Big Data (Big Data)*, 1928–1939. IEEE.

Sovrasov, V. 2018-2024. ptflops: a flops counting tool for neural networks in pytorch framework. https://github.com/sovrasov/flops-counter.pytorch.

Sturgis, R.; Emiya, V.; Couëtoux, B.; and Garreau, P. 2024. Beyond geofencing: Behavior detection using AIS. *Ocean Engineering*, 293: 116630.

Tassetti, A.; Ferrà, C.; and Fabi, G. 2019. Rating the effectiveness of fishery-regulated areas with AIS data. *Ocean & coastal management*, 175: 90–97.

TMT; and Network, I. M. 2023. A MCS Practitioners Introductory Guide To: Trawl Fishing.

Tracking, T. M.; and Network, I. 2021a. MCS Practitioners Introductory Guide to: Longline Fishing.

Tracking, T. M.; and Network, I. 2021b. A MCS Practitioners Introductory Guide to: Purse Seine Fishing.

United Nations Conference on Trade and Development. 2018. *Review of Maritime Transport 2018*. New York and Geneva: United Nations.

V. Restrepo, A. J., H. Murua; and Koehler, H. 2024. Tuna Fisheries' Impacts on Non-Tuna Species and Other Environmental Aspects: 2024 Summary. Technical Report 2024-03, International Seafood Sustainability Foundation, Pittsburgh, PA, USA.

Wang, Y.; Yang, L.; Song, X.; Chen, Q.; and Yan, Z. 2021. A Multi-Feature Ensemble Learning Classification Method for Ship Classification with Space-Based AIS Data. *Applied Sciences*, 11(21): 10336.

Xing, B.; Zhang, L.; Liu, Z.; Sheng, H.; Bi, F.; and Xu, J. 2023. The Study of Fishing Vessel Behavior Identification Based on AIS Data: A Case Study of the East China Sea. *Journal of Marine Science and Engineering*, 11(5): 1093.

Zha, D.; Bhat, Z. P.; Lai, K.-H.; Yang, F.; Jiang, Z.; Zhong, S.; and Hu, X. 2023. Data-centric artificial intelligence: A survey. arXiv:2303.10158.

Zhang, Z.; Huang, L.; Peng, X.; Wen, Y.; and Song, L. 2022. Loitering behavior detection and classification of vessel movements based on trajectory shape and Convolutional Neural Networks. *Ocean Engineering*, 258: 111852.