MARITRAC: Maritime trajectory classification using object instance segmentation with model-based generated data augmentation

Enrica d'Afflisio, Leonardo M. Millefiori, Paolo Braca
Research Department
NATO STO-CMRE
La Spezia, Italy
{enrica.d'afflisio, leonardo.millefiori, paolo.braca}@cmre.nato.int

Marco Guerriero
AWS Professional Services, France and EMEA South
Amazon Web Services
Rome, Italy
magerri@amazon.it

***Abstract*—Maritime surveillance, characterized by high-volume data streams, necessitates effective methods for the automatic extraction of meaningful information and accurate classification of vessel patterns. We introduce MARITRAC (maritime trajectory classification), an innovative approach that leverages MASK R-CNN, a state-of-the-art computer vision algorithm, to classify maritime trajectories. The key idea behind MARITRAC is to convert trajectory data into images that capture spatio-temporal patterns. These trajectory images are then used as input to a MASK R-CNN model that is trained on synthetically generated data to classify different types of maritime trajectories. By combining computer vision techniques with trajectory data analysis, MARITRAC provides an effective and automated method for characterizing and distinguishing between different maritime behaviors. To overcome the notable lack of labeled trajectory anomaly datasets, the training is performed with a set of synthetically generated trajectories, created using the piecewise Ornstein-Uhlenbeck dynamic model. The effectiveness of MARITRAC is demonstrated through application and evaluation in two main experiments, involving both synthetic and real-world data. The approach showcases promising performance in classifying maritime trajectories, and the results position MARITRAC as a valuable tool for real-time maritime surveillance.**

***Index Terms*—Maritime anomaly detection, trajectory classification, convolutional neural networks, object detection.**

I. Introduction

In conventional surveillance systems, operators examine the trajectories of monitored targets looking for peculiar patterns to make high-level assessments on their behavior and/or intention. For example, a rectangular trajectory is an indication of a possible reconnoitering operation, and the loitering behavior of a boat is considered as a possible indication of smuggling [1].

In the maritime domain, the analysis of real-world Automatic Identification System (AIS) data shows that most vessels in open seas maneuver rarely, since travelling by the shortest route with a nearly constant speed is advantageous in terms of time and cost [2]. As a result, a vessel's behavior that deviates from the predicted course while underway may indicate actions that are not consistent with the expected behavior. Indeed, depending on the specific circumstances, vessel trajectories that are not compliant with standard navigational patterns may be considered suspicious and worthy of further investigation by the authorities, as possibly correlated with incidents, illegal activities, or even sabotage operations. In such cases, the timely identification of these patterns in vessel trajectories represents an important capability to strengthen Maritime Situational Awareness (MSA) [3], [4].

Recent advancements in deep learning and data-driven methodologies are also pertinent to trajectory classification, and offer promising avenues to support the detection of complex anomalies. These methods are indeed instrumental to scrutinize vessel trajectories automatically and uncover patterns of interest embedded within the large volume of maritime traffic data. Such capabilities are key for extracting spatio-temporal features that are considered signs of abnormal or suspicious behaviors, enabling the detection of patterns or occurrences suggestive of illicit or questionable activities.

Trajectory classification can be seen as an instance of a supervised learning task in data analysis and pattern recognition that involves training models to classify mobility patterns into predefined labels. Using supervised learning approaches, trajectories or their parts can be assigned to some categories, which can be motion modalities, performed activities, or transportation modes. Trajectory classification methods are relevant in quite a number of domains, including vehicle and pedestrian behavior anticipation [5], [6], computational neuroscience [7], sports-analytics [8], animal migration [9]. Each domain is characterized by peculiar challenges, such as the dimensionality of the involved spatio-temporal measurements. Trajectory classification techniques are also crucial within the maritime domain, establishing a rich literature that mainly delve into classifying vessel types [10]–[14], delineating shipping operation areas [15], detecting specific vessel activities such as fishing operations [16]–[20], vessel movements [21], and perform anomaly detection [22]–[26].

Pre-processing is necessary for the majority of the previously mentioned trajectory classification approaches. This step entails understanding and analyzing the data and extracting features that are unique to the trajectories that need to be

This work is supported by NATO Allied Command Transformation (ACT) via project "Data Knowledge Operational Effectiveness" (DKOE).

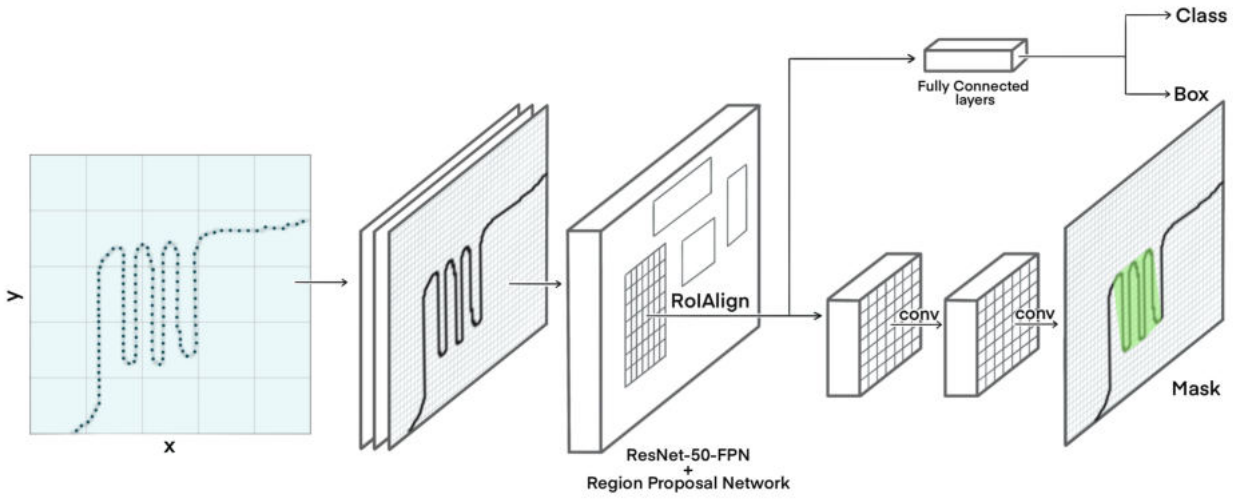

Fig. 1: MARITRAC architecture scheme based on the Mask R-CNN framework for object detection and instance segmentation.

classified. This prevents features that are chosen for one trajectory from being applicable to other patterns. However, the exploration of computer vision methods for trajectory classification has also prompted an intriguing avenue recently. In particular, treating the trajectories as images and availing image recognition techniques, the aforementioned preprocessing stage can be totally avoided, yielding an end-to-end approach for trajectory classification [27].

In the fields of computer vision and image recognition, convolutional neural networks (CNN) are widely used, e.g., to classify sets of images into predefined labels of interest through network layers able to discern different aspects of an image, including shape and color [28]–[31]. However, the use of computer vision approach for the classification of parts of the vessel trajectory has been attempted only recently in [32], where different mobility patterns, indicative of specific vessels' activities, are studied through an image representation procedure and then classified with random forest.

Inspired by the concept presented in [32], here we propose a novel approach called MARITRAC (maritime trajectory classification) to detect and classify parts of maritime trajectories into different categories of patterns of interest. To this aim, the proposed method relies on a mask region-based convolutional neural network (Mask R-CNN) [33] trained on a set of trajectory images that are generated synthetically. Mask R-CNN was originally introduced in [33] as a pixel-level network to identify individual objects, and is a simple, flexible and fast framework for instance segmentation. The task of instance segmentation can be seen as a combination of object detection and semantic segmentation to identify all instances of objects belonging to specific categories in an image. Indeed, each instance is also labeled, i.e., assigned to a category.

For its properties and capabilities, Mask R-CNN is an appealing instance segmentation method for visual perception tasks in various and heterogeneous fields, including the medical sector [34]–[38], remote sensing [39]–[43], the intelligent control of vehicles and pedestrians [44]–[46], and many others. Within the domain of maritime surveillance and situational awareness, to the best knowledge of the authors, the main use of Mask R-CNN has been for ship detection and ship instance segmentation tasks from remote sensing images [47]–[51]. In this work, we focus instead on the detection and instance segmentation of relevant patterns in vessel trajectories and propose MARITRAC, a method based on Mask R-CNN to perform trajectory classification efficiently.

To overcome the lack of labeled vessel trajectory datasets, we generate the training set synthetically by leveraging the on the piecewise Ornstein-Uhlenbeck (OU) dynamic model introduced and validated in [2], [52]–[54]. Indeed, the novelty of MARITRAC lies in the combination of a foundational model together with a peculiar data augmentation approach to overcome such notable lack of labeled vessel trajectory dataset. This design choice is an attempt to take the best of both generative and first principles approaches. Indeed, the Mask R-CNN block can be seen as a foundational model pre-trained against a large generalist dataset, but the training is customized by data augmentation on a smaller dataset of trajectories built according to a statistical motion model, inspired to the model-based deep learning paradigm [55].

Finally, let us remark that, even if the training of MARITRAC is customized against synthetic data, the performance of the so-trained model is assessed not only against synthetic, but also real data. This is another peculiar aspect of the proposed approach, which allows noting a remarkable capability of MARITRAC to perform well also on data that is different from that observed during training.

This paper is organized as follows: the problem formulation is introduced in Sec. II-A, and Sec. II-B provides details of MARITRAC algorithm including the Mask R-CNN architecture and training. In Sec. III, MARITRAC is evaluated through

a series of experiments concerning both synthetic and real-world datasets. Finally, conclusions are presented in Sec. IV.

II. MARITRAC

In this section, our aim is to describe the rationale and the general approach behind MARITRAC and its components.

A. Vessel trajectory classification for anomaly detection

As mentioned in the previous section, a possible indicator of some activities of interest is the deviation of a vessel from its typical underway navigation pattern, where "deviation" is intended in its most general meaning. There are clearly many possible deviations that vessels can perform, and not all of them are necessarily indicators of a specific activity of the vessel. At the same time, there are patterns that can be recognized as possible signs of activities of interest; for example, vessels engaged in fishing operations and/or loitering activities show distinctive trajectories.

MARITRAC is our proposed approach to identify these patterns by automatically analyzing vessel trajectories, e.g., as reported via the AIS. Although the method is pretty general and can be easily extended to an arbitrary number of patterns, this work focuses only on the automatic identification of three specific patterns in vessel trajectories.

- *Drift*: A drifting trajectory pattern is that of a vessel that is not under control and moves due to the combined effects of wind and sea currents. An example of drift pattern is given in the top row of Fig. 2. This type of pattern is traceable to, e.g., trawl fishing activities or to vessels in the dry/wet bulk sector operating in the spot market, when they have not booked the next cargo.
- *Helix*: It refers to a sector search trajectory performed, e.g., during search and rescue operations. This trajectory pattern is made up of several segments that cover a circular area as in the mid row of Fig. 2. The resulting shape resembles a helix with a certain number of blades.
- *Lawnmower*: A back-and-forth pattern of a number of legs, as represented in the bottom row of Fig. 2. This pattern is a possible sign of patrolling, region exploration, inspection and surveying (e.g., for offshore wind farms [56]), and longline fishing activities.

The aforementioned patterns may or may not be interpreted as anomalies depending on the circumstances. For instance, a drifting pattern exhibited by a fishing vessel in a fishing area shall not be considered anomalous, but the same pattern done by another vessel or in a different area may raise an alarm.

Remark. In this work, we focus only and exclusively on the technical aspects of vessel trajectory classification and will not delve into details on how, why and when specific patterns are considered anomalies. For this reason, we will also refer to the specific trajectory patterns of interest as *anomalies* and to the task of automatically identifying them as *anomaly detection*.

B. Methodology

The general idea behind MARITRAC is to leverage recent advances in automatic object detection based on CNNs to identify parts of vessel trajectories that resemble specific patterns of possible interest. Therefore, the first step is to convert vessels trajectories to images. There are several possible ways to perform this transformation, which can also be seen as an embedding of data from the input space to a higher-dimensional space. In this work, we opt for the simplest way, that is converting trajectory by simply plotting it as a continuous black line on a white background, creating for each trajectory one image of a fixed size.

Then, MARITRAC employs a Mask R-CNN architecture [33] to automatically identify trajectory portions that resemble specific patterns of interest, for which attention or further investigation may be needed. The Mask R-CNN is set to detect objects in an image while simultaneously creating a prime segmentation mask for each instance, indeed its framework comprises a first stage where candidate object bounding boxes are proposed, and a second stage that, in parallel to predicting the class and box offset, outputs a binary mask on each predicted bounding box in a pixel-to-pixel correspondence. A pictorial representation of the workflow of the proposed MARITRAC methodology is illustrated in Fig. 1.

C. Synthetic generation of the training set

A notable limitation for the application of data-driven methods to anomaly detection, especially for maritime surveillance, is the lack of labeled datasets of anomalies. To overcome this limitation and take full advantage of powerful object instance segmentation methods such as Mask R-CNN, inspired by the model-based deep learning paradigm [55], we propose the generation of synthetic trajectories for training MARITRAC.

In Sec. III-C, we will introduce three real-world trajectory datasets that we have annotated manually. The process of annotating anomalies is not only very time-consuming and requiring the input of a human operator, but extremely inefficient in the sense that the number of anomalies will be always much smaller (orders of magnitude) than the number of inspected examples. For this reason, for anomaly detection tasks and provided that the underlying data-driven method has sufficient generalization capability, using a synthetic training set can be preferable than manually assembling a labeled training set.

Therefore, we rely on past work on vessel motion modeling [2], [52]–[54] and generate the synthetic training set by simulating several different vessel trajectories according to piecewise OU mean-reverting stochastic processes on the velocity component of the kinematic state. The choice of the OU process is motivated by the fact that it is highly suitable to properly model the behavior of the dynamics of a significant portion of real-world vessels in open sea [52], [53]. Moreover, suitable choices for the parameters of the OU process enable the creation of trajectory patterns that closely resemble those described in Sec. II-A.

Thus, each trajectory in the training set is generated according to the OU model in an area of fixed size and contains a section that resembles one of the patterns of interest. Such pattern does not pertain to the full trajectory, but only to a portion of it; the rest of the trajectory consists of straight

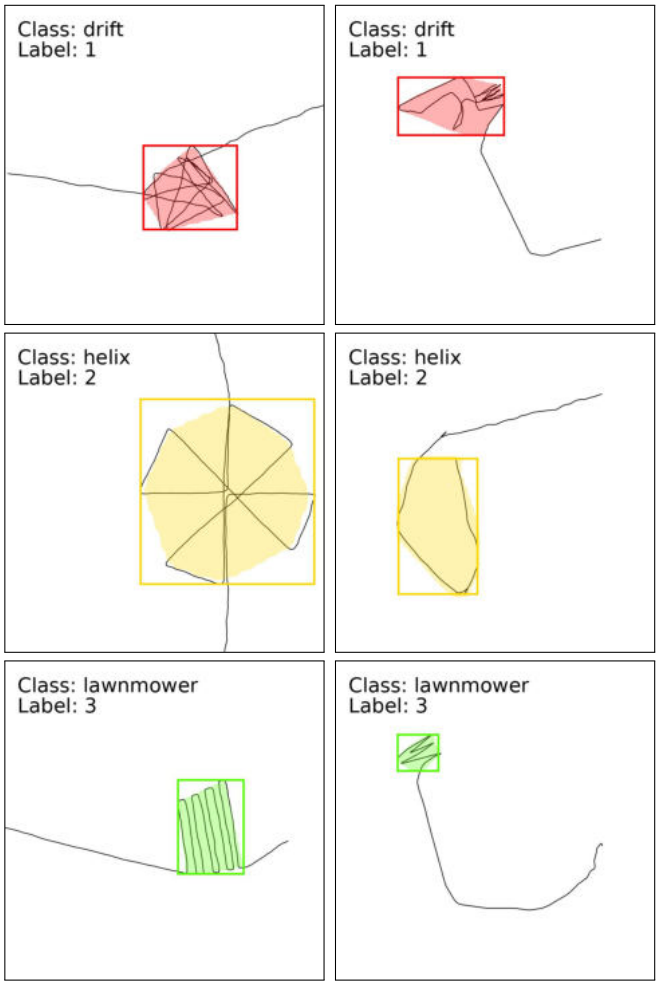

(a) Synthetic patterns (b) Real-world patterns

Fig. 2: Left column: samples of synthetic trajectories generated with the method described in Sec. II-C and used to train MARITRAC. Right column: samples of trajectories from real-world scenario 1; the trajectory images represent the inputs at inference time, while bounding boxes and masks represent the outputs. Each panel contains also a bounding box that contains one of the detected patterns and a mask, which identifies the pixels in the bounding box that are related to the detected pattern. The color is representative of the assigned class: red for the drift pattern, yellow for the helix pattern and green for the lawnmower.

sections representative of a basic non-maneuvering or low-maneuvering behavior of the vessel.

More specifically, trajectories are generated according to a piecewise OU process [53] with a variable number of segments; each segment is characterized by a different set of parameters (most importantly, the duration and the long-run mean velocity parameter) of the underlying OU process, and we refer to "waypoints" as the points where the parameters switch from one set of values to another. Once that the sequence of OU parameters (and waypoints) is generated according to a set of rules that depend on the desired pattern, the synthetic trajectory can be directly simulated as piecewise OU process with such parameters, as in [53]. The specific rules to generate the sequence of waypoints for the three patterns of interest for this work are detailed in the following.

1) Drift pattern: The generic drift trajectory pattern is built starting from a number of concentric circles with increasing radius. A random number of waypoints is generated and randomly located on the circles, each partitioned in fourths. In other words, each waypoint can lie randomly on one of the circles' fourths. An example of drift trajectory generated with this procedure is illustrated in the upper left panel of Fig. 2.

2) Helix pattern: The helix pattern is generated starting from a sequence of waypoints located on a circle of random radius. These points will represent the vertices of a polygon circumscribed by the circle of chosen radius. By suitably connecting the vertices of this polygon in specific order, it is possible to sketch a number of "blades." In other words, by properly connecting the vertices of a regular polygon of 2ℓ edges, it is possible to sketch a ℓ-blade helix. An examples of helix trajectory generated with this procedure is illustrated in the mid row of the left column of Fig. 2.

3) Lawnmower pattern: The lawnmower pattern is essentially built by connecting a variable number of switchbacks. The single switchback section consists of two short horizontal segments and two vertical parallel segments. Consecutive switchbacks can be of increasing size up to a maximum value with respect to the first one. An example of lawnmower trajectory generated with this procedure is illustrated in the bottom row of the left column of Fig. 2.

Trajectories in all the three configurations are also randomly rotated and scaled to ensure that generically oriented patterns with variable size are represented in the training set.

D. Anomaly detection

Consider a trajectory $\{s(t_i)\}_{i=1}^{N}$, composed of N points, where $s(t_j)$ represents the vessel's kinematic state at time t_j.[1] As described in Sec. II-B, the trajectory is transformed into an image $\mathcal{S}$, which represents the input of the Mask R-CNN network. Such a network can be seen, in general, as a function $f_\omega(S)$, which processes the image $\mathcal{S}$ and provides a set of anomalies $\mathcal{A}$:

$$f_\omega(\mathcal{S}) \longrightarrow \mathcal{A} = \{(k_j, p_j, b_j, M_j)\}_{j=1}^{J}, \tag{1}$$

where ω refers to the network parameters, and J is the (variable) number of trajectory anomalies identified by the Mask R-CNN. Each element of $\mathcal{A}$ is composed by: the declared class k_j, its level of confidence $p_j \in [0, 1]$, the bounding box b_j, and the pixel mask M_j. The difference between bounding boxes and masks is easily understood from Fig. 2: bounding boxes are rectangular areas where the Mask R-CNN declares the presence of one of the trajectory patterns of interest, while in masks this indication is reported to the pixel level.

Let us now enforce a decision statistic $\mathcal{T}$ to select only the anomalies whose confidence score is higher than a selected threshold. To this aim, let us define the hypotheses $\mathcal{H}_0$, representing the case of absence of anomaly, and $\mathcal{H}_1$, which represents the case of at least one anomaly present in the input

[1] In this work, we use only the vessel position, but one could also use speed information by encoding it, e.g., in the stroke color.

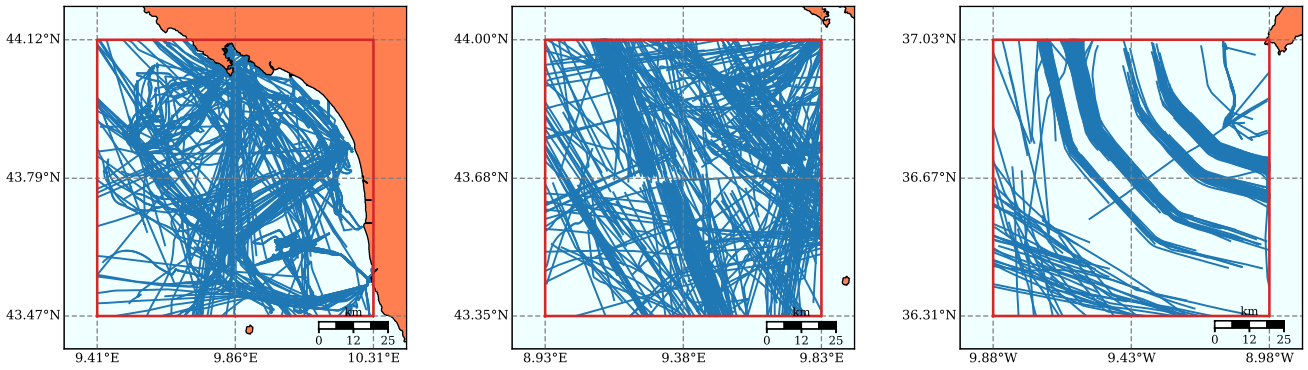

(a) Real-world scenario 1 (b) Real-world scenario 2 (c) Real-world scenario 3

Fig. 3: Real-world AIS datasets used in this work: (a) in the Ligurian Sea close to the coast; (b) in the Ligurian Sea farther from the coast; and (c) in the Atlantic Ocean off Cape St. Vincent.

image. A possible decision statistic, based on the confidence level p_j computed by the Mask R-CNN, is given as follows:

$$\begin{cases} \mathcal{A} = \emptyset \\ p_j \leq \gamma, \, \forall j \end{cases} \quad \mathcal{T}(\mathcal{A}) \text{ declares } \mathcal{H}_0, \\ \text{otherwise} \qquad \mathcal{T}(\mathcal{A}) \text{ declares } \mathcal{H}_1. \tag{2}$$

At this point, we can define the performance metrics of the aforementioned binary test, i.e., the detection probability P_d and false alarm probability P_{fa}, as follows:

$$P_d = \mathcal{P}\left\{\mathcal{T}(\mathcal{A}) \text{ declares } \mathcal{H}_1 \,|\mathcal{H}_1\right\}, \tag{3}$$

$$P_{fa} = \mathcal{P}\left\{\mathcal{T}(\mathcal{A}) \text{ declares } \mathcal{H}_1 \,|\mathcal{H}_0\right\}. \tag{4}$$

It is worth noting that the aforementioned test is not optimal in terms of the Neyman-Pearson criterion [57], as the likelihood functions of the two (composite) hypotheses are not available. Furthermore, MARITRAC is trained by minimizing a combination of classification loss, bounding box regression loss, and mask segmentation loss; this allows the model to learn to simultaneously detect objects, refine their bounding boxes, and produce precise segmentation masks. However, if we interpret p_j as the posterior probability of the existence of the j-th anomaly associated to the given trajectory, then the test $p_j > \gamma$ is basically equivalent to a likelihood ratio test, assuming uniform priors of the two events, anomaly existence and non-existence. This is conceptually similar to comparing the track existence posterior probability to a threshold in a Bayesian multi-target tracking approach to declare/confirm the track, see, e.g., [58]. However, different from the usual assumption of independence among targets, in this context the bounding box of the detected anomalies frequently overlap or are inter-dependents, and for this reason the OR-rule among the different p_j, defined in (2), is quite reasonable but sub-optimal even assuming to have available the exact posterior probabilities. Another quite valid option would be to compare the average of p_j with a threshold, i.e., $\frac{1}{J}\sum_{j=1}^{J} p_j > \gamma$, but this is not investigated further here.

III. Experiments

In this section, we describe the application of MARITRAC to synthetic vessel trajectories first, and then real-world AIS data. To this aim, we have assembled two datasets that serve as benchmarks to assess the performance of the proposed method, specifically in terms of false alarm rate P_{fa} and missed detection rate $1 - P_d$.

A. Training

For our experiments, we have used transfer learning from a Mask R-CNN with a ResNet-50-FPN backbone network [33] whose weights were pre-trained from the COCO dataset [59] and whose head branch was adapted for three output classes, equal to the number of trajectory patterns of interest.

We let the final training of the network run for 42 epochs using a synthetic dataset of 15 000 sample trajectories, i.e., 5000 for each of the three classes, generated as described in Sec. II-A. The trajectories in the training set have significant random variation in terms of orientation, length and characteristics of the anomalous patterns, with an average duration between 12 and 15 hours and average speed of 10 knots. Each trajectory is converted to an image of dimension 768×768 pixels and footprint of approximately 100×100 km^2. Moreover, each image contains exactly one trajectory, which is in part normal and in part anomalous, as in Fig. 2.

B. Synthetic analysis

To assess the performance of MARITRAC, we have generated an independent synthetic dataset of trajectories analogous to the training set. The test set comprises a total of 4000 trajectories for hypothesis $\mathcal{H}_1$, i.e., 1000 images for each of the three classes of patterns of interest, and an additional 1000 images of non-maneuvering trajectories, which represent the hypothesis $\mathcal{H}_0$ and can be used to assess the level of false alarms. Similar to the training set, trajectories in the test set are converted to images of dimension 768×768 pixels and geographical footprint of approximately 100×100 km^2.

The left column of Fig. 2 reports MARITRAC's outputs as evaluated on three input trajectories from the assembled

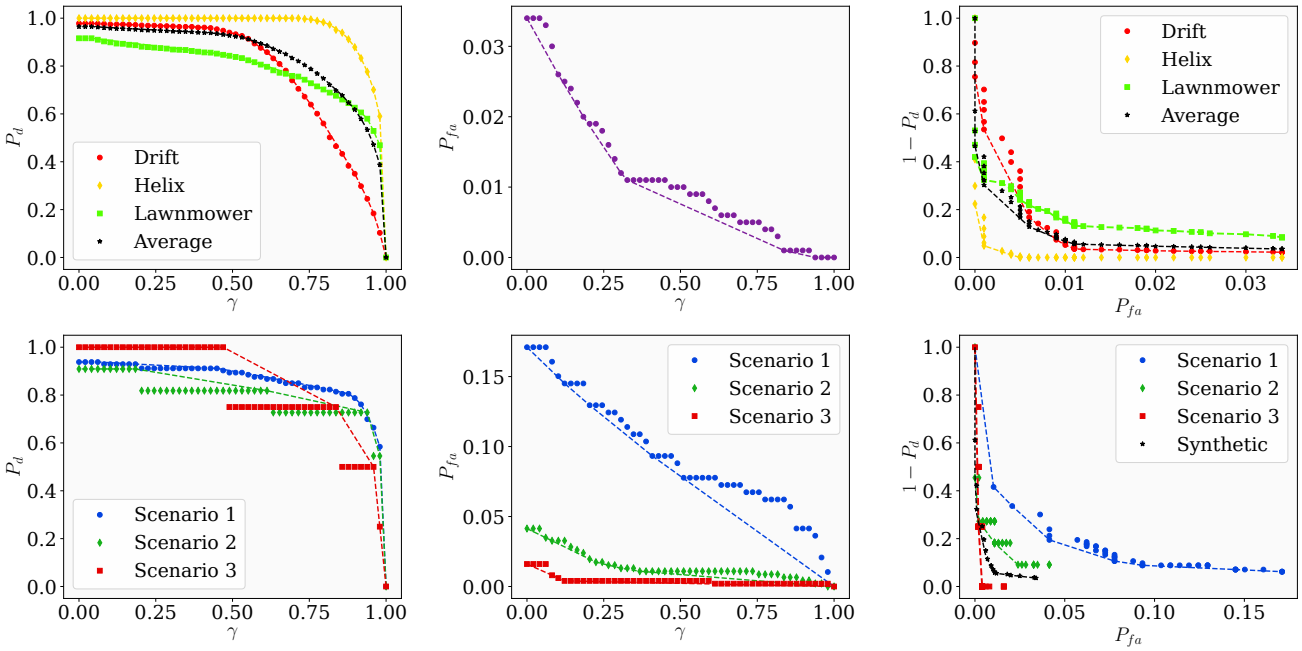

Fig. 4: Performance of MARITRAC with synthetic and real data. The panels in the top row refer to the assessment against the synthetic benchmark dataset described in Sec. III-B, while those in the bottom row refer to the real-world scenarios described in Sec. III-C. The left and central columns illustrate the probability of detection and false alarm, respectively, versus threshold. The right column reports the: probability of missed detection versus false alarm; this is analogous to a ROC curve, as discussed in the text.

synthetic dataset. Specifically, MARITRAC is able to correctly recognize the drift, the helix, and the lawnmower patterns, in red, yellow and green bounding boxes and masks, respectively.

The performance of MARITRAC on synthetic data, in terms of probability of detection and false alarm, is illustrated in the plots in the top row of Fig. 4. The top panel illustrates the probability of detection as a function of the threshold γ in (2) conditioned to each of the pattern classes (color curves and markers); each of these curves is evaluated with data generated according to one specific class of anomaly. In black, we also report the class-unconditioned detection probability, which is equivalent to the average of the class-dependent probabilities when the prior is uniform (as in this case). The mid panel reports the probability of false alarm as a function of the threshold; we see only one curve here as the data under hypothesis $\mathcal{H}_0$ does not depend on the pattern class. The bottom panel, finally, shows the probability of missed detection versus the probability of false alarm under each of the considered pattern classes (color curves and markers). The curves in the right column of Fig. 4 can be viewed as analogous to a receiver operating characteristic curve (ROC) [57] for MARITRAC. In all panels, the dashed curves are the convex hulls computed from the empirical values.

C. Real-world analysis

The datasets used for the real-world analysis in our experimental setup come from AIS data collected in three different geographical areas by NATO STO-CMRE for research purposes as part of its mission. The three areas are all approximately 100×100 km^2 wide and are illustrated in Fig. 3.

We have manually annotated the real-world datasets so as to have a reference ground truth for the evaluation of the performance of the proposed approach. For the scope of this work, and since we are interested in the performance characterization of MARITRAC in terms of probability of detection and false alarm, the annotation is limited only to deciding by visual inspection if a trajectory is anomalous or not; we do not assign anomalous trajectories to specific classes nor do we annotate the bounding box and mask of the anomaly.

1) Scenario 1: The first scenario is located in the Ligurian Sea, in a geographical area that spans from latitude $43.47°$N to $44.12°$N and from longitude $9.41°$E to $10.31°$E. The AIS data of this scenario come from an AIS receiver operated by NATO STO-CMRE in the area of La Spezia. We have isolated AIS tracks from the 3$^{\text{rd}}$ to the 5$^{\text{th}}$ of July, 2023. The AIS data have been organized in continuous trajectories of duration between 10 and 12 hours and minimum cumulative length of 30 nautical miles (nmi). The dataset of this scenario comprises 306 trajectories and is illustrated in Fig. 3a.

2) Scenario 2: The second scenario is close to the first one, but farther from the coast. It spans the geographical area from latitude $43.35°$N to $44.00°$N and from longitude $8.93°$E to $9.83°$E. The AIS data of this scenario also come from the same AIS receiver of the previous scenario. In this scenario, we have isolated AIS tracks from the 3$^{\text{rd}}$ to the 9$^{\text{th}}$ of July, 2023. The AIS data have been organized in continuous trajectories of duration between 20 and 24 hours and minimum cumulative length of 20 nmi. The dataset related to this scenario comprises

470 trajectories and is illustrated in Fig. 3b.

3) Scenario 3: The third scenario is in a completely different geographical area, off the coast of Portugal, and specifically off Cape St. Vincent. The AIS data of this scenario are collected by NATO STO-CMRE from AISHub, a free AIS data sharing network which provides access to real time ship positions for vessel tracking systems [60]. The area of the third scenario spans from latitude $36.31°\mathrm{N}$ to $37.03°\mathrm{N}$ and from longitude $9.88°\mathrm{W}$ to $8.98°\mathrm{W}$. In this scenario, we have isolated AIS tracks from the 3rd to the 9th of July, 2023. The AIS data have been organized in continuous trajectories of duration between 10 and 12 hours and minimum cumulative length of 15 nmi. The dataset related to this scenario comprises 502 trajectories and is illustrated in Fig. 3c.

The performance of MARITRAC against the real-world datasets described above is reported in the bottom row of Fig. 4. For each of the three scenarios: the top panel illustrates the probability of detection P_d as a function of the threshold γ in (2); the mid panel shows the probability of false alarm P_{fa} as a function of the threshold; the bottom panel shows the probability of missed detection versus the probability of false alarm. As before, dashed curves are the convex hulls computed from the empirical values.

As already mentioned, the plot in the right column of Fig. 4 can be interpreted as the analogous of ROC curves. As such, and different from the panels above, they are suitable to be compared directly with the performance evaluated on the synthetic dataset, reported in black also in this plot. These plots also make evident the fact that the actual performance of MARITRAC depends on the specific area under investigation. Indeed, the ROC curve of MARITRAC achieved in scenario 1 is above the other curves, indicating a higher false alarm rate for the same (or approximately same) missed detection rate.

IV. Conclusion

In this study, we presented MARITRAC (maritime trajectory classification), an innovative and effective approach to address the challenges of vessel trajectory classification that merges computer vision and statistical vessel kinematic models. MARITRAC utilizes a mask region-based convolutional neural network (Mask R-CNN) and object instance segmentation to classify maritime trajectories. The unique aspect of this approach lies in the training methodology, which involves using synthetically generated trajectories based on the piecewise Ornstein-Uhlenbeck dynamic model to overcome the lack of labeled trajectory anomaly datasets.

The application and evaluation of MARITRAC in two main experiments, encompassing both synthetic and real-world data, demonstrate its promising performance in classifying maritime trajectories. The results position MARITRAC as a valuable tool for real-time maritime surveillance, offering a robust and efficient method for automatic extraction of meaningful information and accurate classification of vessel patterns. The successful integration of computer vision and statistical kinematic models, along with the use of synthetic data for training, highlights the adaptability and potential of MARITRAC in enhancing the capabilities of maritime surveillance systems. As a result, MARITRAC stands as a noteworthy contribution to the field, addressing a crucial need for advanced methods in handling the complexities of maritime data for improved security and monitoring.

References

[1] V. Krishnamurthy and M. H. Fanaswala, "Intent inference via syntactic tracking," *Digital Signal Processing*, vol. 21, no. 5, pp. 648–659, 2011.

[2] L. M. Millefiori, P. Braca, K. Bryan, and P. Willett, "Modeling vessel kinematics using a stochastic mean-reverting process for long-term prediction," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 5, pp. 2313–2330, Oct. 2016.

[3] G. Soldi, D. Gaglione, S. Raponi, N. Forti, E. d'Afflisio, P. Kowalski, L. M. Millefiori, D. Zissis, P. Braca, P. Willett, A. Maguer, S. Carniel, G. Sembenini, and C. Warner, "Monitoring of critical undersea infrastructures: The Nord Stream and other recent case studies," *IEEE Aerospace and Electronic Systems Magazine*, vol. 38, no. 10, pp. 4–24, 2023.

[4] N. Forti, E. d'Afflisio, P. Braca, L. M. Millefiori, S. Carniel, and P. Willett, "Next-gen intelligent situational awareness systems for maritime surveillance and autonomous navigation [Point of View]," *Proceedings of the IEEE*, vol. 110, no. 10, pp. 1532–1537, 2022.

[5] A. Benterki, M. Boukhnifer, V. Judalet, and C. Maaoui, "Artificial intelligence for vehicle behavior anticipation: Hybrid approach based on maneuver classification and trajectory prediction," *IEEE Access*, vol. 8, pp. 56 992–57 002, 2020.

[6] V. Papathanasopoulou, I. Spyropoulou, H. Perakis, V. Gikas, and E. Andrikopoulou, "A data-driven model for pedestrian behavior classification and trajectory prediction," *IEEE Open Journal of Intelligent Transportation Systems*, vol. 3, pp. 328–339, Apr. 2022.

[7] S. Klibi, M. Vernet, D. Schwartz, and I. R. Farah, "Towards a better multivariate time-series detection of epileptic seizures in electroencephalogram (EEG) using machine learning algorithms," in *2nd International Conference of Smart Systems and Emerging Technologies*, 2022, pp. 142–147.

[8] Z. Ziyi, R. Bunker, K. Takeda, and K. Fujii, "Multi-agent deep-learning based comparative analysis of team sport trajectories," *IEEE Access*, vol. 11, pp. 43 305–43 315, Apr. 2023.

[9] J. De Groeve, N. Van de Weghe, N. Ranc, T. Neutens, L. Ometto, O. Rota-Stabelli, and F. Cagnacci, "Extracting spatio-temporal patterns in animal trajectories: An ecological application of sequence analysis methods (SAM)," *Methods in Ecology and Evolution*, Jul. 2015.

[10] J.-G. Lee, J. Han, X. Li, and H. Gonzalez, "TraClass: trajectory classification using hierarchical region-based and trajectory-based clustering," *Proceedings of The VLDB Endowment*, vol. 1, no. 1, p. 1081–1094, Aug. 2008.

[11] G. K. D. de Vries and M. Someren, "Machine learning for vessel trajectories using compression, alignments and domain knowledge," *Expert Systems with Applications*, vol. 39, Dec. 2012.

[12] K. Chatzikokolakis, D. Zissis, G. Spiliopoulos, and K. Tserpes, "Mining vessel trajectory data for patterns of search and rescue," in *EDBT/ICDT Workshops*, 2018.

[13] T. Yang, X. Wang, and Z. Liu, "Ship type recognition based on ship navigating trajectory and convolutional neural network," *Journal of Marine Science and Engineering*, vol. 10 (1), no. 84, 2022.

[14] A. Pohontu, A. Deliu, and C. Vertan, "Ship type classification: a handwriting signature verification approach for maritime trajectories," in *8th International Symposium on Electrical and Electronics Engineering*, Oct. 2023, pp. 110–115.

[15] F. Mazzarella, M. Vespe, D. Damalas, and G. Osio, "Discovering vessel activities at sea using AIS data: Mapping of fishing footprints," in *17th International Conference on Information Fusion*, Jul. 2014.

[16] F. Natale, M. Gibin, A. Alessandrini, M. Vespe, and A. Paulrud, "Mapping Fishing Effort through AIS Data," *PLoS ONE*, vol. 10, Jun. 2015.

[17] X. Jiang, D. Silver, B. Hu, and E. de Souza, "Fishing activity detection from AIS data using autoencoders," in *Proceedings of the 29th Canadian Conference on Artificial Intelligence on Advances in Artificial Intelligence*, May 2016, pp. 33–39.

[18] E. de Souza, K. Boerder, and B. Worm, "Improving fishing pattern detection from satellite AIS using data mining and machine learning," *PLoS ONE*, vol. 11, Jul. 2016.

[19] I. Kontopoulos, K. Chatzikokolakis, K. Tserpes, and D. Zissis, "Classification of vessel activity in streaming data," in *Proceedings of the 14th ACM International Conference on Distributed and Event-Based Systems*, 2020, p. 153–164.

[20] K.-I. Kim and K. M. Lee, "Convolutional neural network-based gear type identification from automatic identification system trajectory data," *Applied Sciences*, vol. 10, Jun. 2020.

[21] X. Chen, Y. Liu, K. Achuthan, and X. Zhang, "A ship movement classification based on automatic identification system (AIS) data using convolutional neural network," *Ocean Engineering*, vol. 218, 2020.

[22] K. Kowalska and L. Peel, "Maritime anomaly detection using Gaussian process active learning," in *15th International Conference on Information Fusion*, Jul. 2012, pp. 1164–1171.

[23] D. O. D. Handayani, W. Sediono, and A. Shah, "Anomaly detection in vessel tracking using Support Vector Machines (SVMs)," *International Conference on Advanced Computer Science Applications and Technologies*, pp. 213–217, Dec. 2013.

[24] S. Mascaro, A. E. Nicholson, and K. B. Korb, "Anomaly detection in vessel tracks using Bayesian networks," *International Journal of Approximate Reasoning*, vol. 55, no. 1, pp. 84–98, 2014.

[25] Z. Wei, X. Xie, and X. Zhang, "Maritime anomaly detection based on a Support Vector Machine," *Soft Computing*, vol. 26, no. 21, pp. 11 553–11 566, Nov. 2022.

[26] D. Nguyen, R. Vadaine, G. Hajduch, R. Garello, and R. Fablet, "GeoTrackNet—A maritime anomaly detector using probabilistic neural network representation of AIS tracks and a contrario detection," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 5655–5667, 2022.

[27] I. Kontopoulos, A. Makris, and K. Tserpes, "TraClets: A trajectory representation and classification library," *SoftwareX*, vol. 21, p. 101306, 2023.

[28] W. Rawat and Z. Wang, "Deep convolutional neural networks for image classification: A comprehensive review," *Neural Computation*, vol. 29, Jun. 2017.

[29] S. Zhong, Y. Liu, and Y. Liu, "Bilinear deep learning for image classification," *Proceedings of the 2011 ACM Multimedia Conference and Co-Located Workshops*, pp. 343–352, Nov. 2011.

[30] J. Wu, Y. Yu, C. Huang, and K. Yu, "Deep multiple instance learning for image classification and auto-annotation," in *IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2015, pp. 3460–3469.

[31] Y. Endo, H. Toda, K. Nishida, and J. Ikedo, "Classifying spatial trajectories using representation learning," *International Journal of Data Science and Analytics*, vol. 2, pp. 107–117, 2016.

[32] I. Kontopoulos, A. Makris, D. Zissis, and K. Tserpes, "A computer vision approach for trajectory classification," in *22nd IEEE International Conference on Mobile Data Management (MDM)*, 2021, pp. 163–168.

[33] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *IEEE International Conference on Computer Vision*, 2017, pp. 2980–2988.

[34] T. Padma, C. U. Kumari, D. Yamini, K. Pravallika, K. Bhargavi, and M. Nithya, "Image segmentation using Mask R-CNN for tumor detection from medical images," in *International Conference on Electronics and Renewable Systems*, 2022, pp. 1015–1021.

[35] S. Palakvangsa-Na-Ayudhya, T. Sapthamrong, K. Sunthornwutthikrai, and D. Sakiyalak, "GlaucoVIZ: Assisting system for early glaucoma detection using Mask R-CNN," in *17th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology*, 2020, pp. 364–367.

[36] L. mao, T. Yumeng, and C. Lina, "Pneumonia detection in chest X-rays: a deep learning approach based on ensemble RetinaNet and Mask R-CNN," in *Eighth International Conference on Advanced Cloud and Big Data*, 2020, pp. 213–218.

[37] C. Huang, A. Yu, Y. Wang, and H. He, "Skin lesion segmentation based on Mask R-CNN," in *International Conference on Virtual Reality and Visualization*, 2020, pp. 63–67.

[38] E. D. Cherpanath, P. R. Fathima Nasreen, K. Pradeep, M. Menon, and V. S. Jayanthi, "Food image recognition and calorie prediction using Faster R-CNN and Mask R-CNN," in *9th International Conference on Smart Computing and Communications*, 2023, pp. 83–89.

[39] H. Su, S. Wei, M. Yan, C. Wang, J. Shi, and X. Zhang, "Object detection and instance segmentation in remote sensing imagery based on precise Mask R-CNN," in *IEEE International Geoscience and Remote Sensing Symposium*, 2019, pp. 1454–1457.

[40] D. Kumar and X. Zhang, "Improving more instance segmentation and better object detection in remote sensing imagery based on cascade Mask R-CNN," in *IEEE International Geoscience and Remote Sensing Symposium*, 2021, pp. 4672–4675.

[41] Y. Wang, Y. Rao, C. Huang, Y. Yang, Y. Huang, and Q. He, "Using the improved Mask R-CNN and softer-nms for target segmentation of remote sensing image," in *4th International Conference on Pattern Recognition and Artificial Intelligence*, 2021, pp. 91–95.

[42] F. Liu, S. Guan, K. Yu, and H. Gong, "Infrared target detection based on the fusion of Mask R-CNN and image enhancement network," in *China Automation Congress*, 2022, pp. 2011–2016.

[43] Y. Huiming and X. Fuxin, "A remote sensing image target recognition method based on improved Mask-RCNN model," in *IEEE 2nd International Conference on Big Data, Artificial Intelligence and Internet of Things Engineering*, 2021, pp. 436–439.

[44] M. A. Malbog, "MASK R-CNN for pedestrian crosswalk detection and instance segmentation," in *IEEE 6th International Conference on Engineering Technologies and Applied Sciences*, 2019.

[45] C.-H. Lin and Y. Li, "A license plate recognition system for severe tilt angles using Mask R-CNN," in *International Conference on Advanced Mechatronic Systems*, 2019, pp. 229–234.

[46] H. Tahir, M. Shahbaz Khan, and M. Owais Tariq, "Performance analysis and comparison of Faster R-CNN, Mask R-CNN and ResNet50 for the detection and counting of vehicles," in *International Conference on Computing, Communication, and Intelligent Systems*, 2021, pp. 587–594.

[47] Y. Qian, Q. Liu, H. Zhu, H. Fan, B. Du, and S. Liu, "Mask R-CNN for object detection in multitemporal SAR images," in *10th International Workshop on the Analysis of Multitemporal Remote Sensing Images*, 2019, pp. 1–4.

[48] X. Nie, M. Duan, H. Ding, B. Hu, and E. K. Wong, "Attention Mask R-CNN for ship detection and segmentation from remote sensing images," *IEEE Access*, vol. 8, pp. 9325–9334, 2020.

[49] T. Zhang, X. Zhang, J. Li, and J. Shi, "Contextual squeeze-and-excitation Mask R-CNN for SAR ship instance segmentation," in *IEEE Radar Conference*, 2022, pp. 1–6.

[50] Y. Sun, L. Su, Y. Luo, H. Meng, W. Li, Z. Zhang, P. Wang, and W. Zhang, "Global Mask R-CNN for marine ship instance segmentation," *Neurocomputing*, vol. 480, pp. 257–270, 2022.

[51] Y. Liu and Y. Xiao, "Remote sensing warship detection method based on improved Mask R-CNN," in *41st Chinese Control Conference*, 2022, pp. 7100–7105.

[52] L. M. Millefiori, P. Braca, K. Bryan, and P. Willett, "Long-term vessel kinematics prediction exploiting mean-reverting processes," in *19th International Conference on Information Fusion*, 2016, pp. 232–239.

[53] P. Coscia, P. Braca, L. M. Millefiori, F. A. N. Palmieri, and P. Willett, "Multiple Ornstein–Uhlenbeck processes for maritime traffic graph representation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 5, pp. 2158–2170, Oct. 2018.

[54] E. d'Afflisio, P. Braca, and P. Willett, "Malicious AIS spoofing and abnormal stealth deviations: A comprehensive statistical framework for maritime anomaly detection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 4, pp. 2093–2108, 2021.

[55] N. Shlezinger, J. Whang, Y. C. Eldar, and A. G. Dimakis, "Model-based deep learning," *Proceedings of the IEEE*, vol. 111, no. 5, pp. 465–499, 2023.

[56] A. J. I. Foster, M. Gianni, A. Aly, H. Samani, and S. Sharma, "Multi-robot coverage path planning for the inspection of offshore wind farms: A review," *Drones*, vol. 8, no. 1, 2024.

[57] S. M. Kay, *Fundamentals of Statistical Signal Processing, Detection Theory*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1998.

[58] F. Meyer, P. Braca, P. Willett, and F. Hlawatsch, "A scalable algorithm for tracking an unknown number of targets using multiple sensors," *IEEE Transactions on Signal Processing*, vol. 65, no. 13, pp. 3478–3493, 2017.

[59] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*. Springer, 2014, pp. 740–755.

[60] AISHub. [Online]. Available: https://www.aishub.net