Maritime Behaviour Anomaly Detection with Seasonal Context

Travis Rybicki[1], Martin Masek[2], Chiou Peng Lam[2]

[1] School of Science, Edith Cowan University, Joondalup, Western Australia – trybicki@our.ecu.edu.au
[2] School of Science, Edith Cowan University, Joondalup, Western Australia – (m.masek, c.lam)@ecu.edu.au

Keywords: Maritime Traffic Patterns, Seasonal Context, Trajectory Clustering, Anomaly Detection, Spatiotemporal Data.

Abstract

Monitoring maritime traffic has become an important task for ensuring the safety of vessels, as well as the goods, and persons that they may be transporting. An active area of research is the modelling of expected normal vessel behaviour so as to detect subsequent anomalies in new data. Anomalies indicate that a vessel is not behaving in an expected manner and their detection can be flagged for further investigation to identify whether the vessel needs assistance or intervention. An important factor for some vessels in determining normal behaviour is seasonal context. However, current approaches typically do not incorporate seasonality into the model.
In this paper, an approach is presented where seasonal context is incorporated into the behaviour model. Seasonal context is first incorporated into vessel trajectory data by encoding the month of year into a historic dataset. Following this, a model of normal behaviour is generated using a clustering approach, with DBSCAN used in this paper. Details of setting the DBSCAN parameters appropriately for vessel trajectory data are provided and four distance metrics explored. Resulting cluster models are evaluated in the context of using the model to classify previously unseen data as either fitting the model or constituting an anomaly. The experiments focus on using fishing vessels within two identified seasons to build the normal model, which is evaluated with a mixture of in season and out of season fishing and non-fishing vessel behaviour.

1. Introduction

The monitoring of maritime vessel traffic is essential so that goods, animals, and people onboard a vessel are kept safe. Consequently, the International Maritime Organization (IMO) enforces relevant vessels to carry onboard tracking devices which broadcast a vessels navigation trajectory and identity through the Automatic Identification System (AIS) (IMO., 2015; Riveiro et al., 2018). The AIS data allows the operators of vessels to maintain an awareness of other vessels around them and allows the monitoring of shipping channels to help prevent and respond to incidents.

Besides real-time use of AIS, the information is collected and curated by several organisations that make the data accessible for maritime data analysis applications. An emerging use of the large archive of historical AIS data is to model the expected behaviours of certain vessel types and subsequently use this information to detect abnormal behaviour in newly observed traffic (Zhao and Shi, 2019; Masek et al., 2021). Such modelling has the potential to provide human operators with tools to better manage the large quantities of real-time data from busy shipping channels.

To date, most studies on maritime behaviour modelling focus on the kinematic characteristics of vessel trajectories but ignore the seasonal context of vessel behaviour. This can affect the accuracy of models by limiting their sensitivity when a vessel type displays different 'normal' behaviour in different seasons. A clear example is that of fishing vessels. As shown by Guan et al. (2021), different sub-types of fishing vessels, and thus different fishing behaviour, might be dominant in different seasons. Nguyen et al. (2021), in their work on automated model development without taking season into consideration acknowledge that this lowered performance of their model on fishing boats and proposed that in future work this could be addressed by having separate models for each season.

In this paper, an approach is presented that integrates kinematic characteristic with seasonal characteristics in a unified model for the detection of anomalies in maritime vessel behaviour. Clustering of historical data using spatial and temporal information is used to discover behavioural clusters. The resultant model is then tested using by testing its proximity to the nearest spatiotemporal clusters and determining whether it matches the behaviour expected or whether it is anomalous.

In evaluating the approach, we explore four alternatives to measuring distance between trajectories to form the cluster models. Evaluation is performed using the fishing boat class within two defined seasons to define normal behaviour and tested on a combination of out of season fishing and cargo, tanker, and passenger shipping to represent abnormal behaviour. The remainder of this article is organized as follows: Section 2 presents related work. Section 3 presents details of the proposed approach. Section 4 evaluates the effectiveness of the proposed approach through a set of experiments, and Section 5 draws conclusions.

2. Related Work

Various approaches are found in the literature to modelling of normal maritime vessel traffic. Many of these techniques are based on clustering approaches. Among other similar approaches, Olesen et al. (2023) proposed a two-step modelling approach, first clustering based on location, then further refining by other kinematic variables. In evaluating their work on a dataset from Danish waters, it is reported that their approach competes with deep learning neural networks with an Area Under the Curve (AUC) value of 0.79.

https://doi.org/10.5194/isprs-annals-X-4-2024-295-2024

In Xie et al. (2024), maritime vessel trajectory modelling has also been explored using clustering, followed by using a probabilistic attention-based transformer for anomaly detection by analyzing the clustered data made up of the latitude, longitude, speed, and course-over-ground, applying the method to an inland river system.

Zhao and Shi (2019) modelled cargo and tanker trajectories navigating the Beilun-Zhoushan port, China. After this, the remaining clustered trajectories were used to train a Long Short-Term Memory (LSTM) neural network for detecting anomalous trajectories not following these defined routes.

3. The Approach

The proposed approach consists of two steps. First, a model of expected vessel behaviour in a particular region of interest of the ocean is built through the spatiotemporal clustering of historical vessel data. Second, new vessel trajectories are classified as either fitting the model or representing anomalous behaviour. In this section we overview the data representation used in the proposed approach in 3.1, followed by details of the model development using a clustering approach is discussed in 3.2 and details of the classification step in 3.3.

3.1 Data Representation

In this work, the trajectory for a particular vessel is represented as a time series of multi-dimensional points. Each point within a trajectory consists of a set of values, each normalized in the range [0, 1], corresponding to kinematic parameters: Latitude, Longitude, and Speed, and a temporal parameter to enable seasonality awareness: Month-of-Year. Latitude and Longitude are each normalized using min-max normalization of the region of interest boundaries. Speed values are normalised through division by the maximum observed speed in the historical dataset from the area of interest, which in our experiments was 50 knots. The raw Month-of-Year (MoY) variable repetitively cycles between 1 and 12, therefore a representation is needed to model that, for example, month 12 (December) is as close to month 11 (November) as it is to month 1 (January). This is achieved by Equations (1) and (2), transforming the raw MoY values using cosine and sine transformation (Petneházi, 2019) and scaling the result between 0 and 1.

$$Cos_x = 0.5\left(cos\left(\frac{2\times \pi \times x}{\max(x)}\right) + 1\right) \quad (1)$$

$$Sin_x = 0.5\left(sin\left(\frac{2\times \pi \times x}{\max(x)}\right) + 1\right) \quad (2)$$

3.2 Modelling through Spatiotemporal Clustering

The proposed approach builds a model of expected vessel behaviour by taking historical trajectories of vessels and separating them into distinct clusters. The aim is to group vessel trajectories by similarity in terms of their kinematic properties and the season in which the behaviour occurs.

Density-Based Spatial Clustering of Applications with Noise (DBSCAN) (Ester et al., 1996) was chosen as the clustering algorithm in the proposed approach. DBSCAN does not need the number of clusters to be known a priori and has seen use previously in the maritime domain (Zhao and Shi, 2019; Olesen et al., 2023; Xie et al. 2024).

The DBSCAN algorithm groups a set of points into clusters using two parameters, ε, and n, and also needs a metric to define distance between a pair of data points. The selection of ε and n will be discussed next in 3.2.1, followed by a discussion on distance metrics explored in 3.2.2.

3.2.1 DBSCAN Parameter Setting: As noted in previous work (Zhao and Shi, 2019), the wide range of vessel behaviours presents a challenge in selecting appropriate values for the DBSCAN parameters. Han et al. (2021) proposed a heuristic for estimating the ε parameter using the kth nearest neighbours distance distribution of the dataset, setting ε as 1.5 times the interquartile range plus the third quartile (referred to as the *IQR rule*). Though their promising results are not based on seasonal-aware data, utilising only kinematic parameters (Latitude, Longitude, Speed Over Ground, Course Over Ground, and Heading), we have adopted their approach for our data which includes Month of Year to model seasonality.

The use of the IQR rule still relies on the setting of one parameter, *k*, which we have tuned in our experiments by examining results for various *k* values. From these results, the k that corresponds to the minimum outliers in the clustered data is chosen, similar to the approach in Han et al. (2021). The chosen k is also used as the parameter *n* for DBSCAN and used to calculate ε using the IQR rule.

3.2.2 Calculating Trajectory Similarity: To measure distance between two trajectories, a metric is needed that is applicable to multi-dimensional time series data. Previous works have explored several measures that are suitable for maritime trajectories that include kinematic parameters such as geographic coordinates, speed and heading (Li et al., 2017; Han et al., 2021). We explore some of the established measures with our dataset, which has the addition of the temporal Month of Year parameter to each data point in the trajectory time series.

For this article, four measures are investigated, three of these are based on Dynamic Time Warping (DTW) (Sakoe and Chiba, 1978), with the original DTW distance and two varieties that employ separate forms of trajectory length normalisation. The fourth distance is the Hausdorff distance (Huttenlocher and Kedem, 1990; Huttenlocher et al., 1993). Both DTW and Hausdorff metrics have been employed in maritime trajectory comparisons (e.g., Olesen et al. 2023 among others), though they both have challenges. The Hausdorff distance does not consider direction due to being a min-max function (Shen et al., 2022). And DTW is known to have issues with trajectory lengths where extreme differences lead to a higher distance value regardless how similar (Shen and Chi, 2017).

Distance normalization seeks to remove the influence of differing trajectory lengths. In this way, a dataset can have a relative distance value for each trajectory pair. This article employs two techniques, one from Shen and Chi (2017) shown as Equation (4) and another from Tao et al. (2021) shown as Equation (5).

$$DTW_{norm1}(A,B) = \frac{DTW(A,B)}{N+M} \quad (4)$$

where DTW(A,B) = the DTW value of two trajectories A and B, which have N and M points respectively,

https://doi.org/10.5194/isprs-annals-X-4-2024-295-2024 | © Author(s) 2024. CC BY 4.0 License.

$$DTW_{norm2}(A,B) = \sqrt{\frac{DTW(A,B)}{\max(N,M)}} \tag{5}$$

3.3 Classification of New Vessel Trajectories

Once the model of normal vessel behaviour is built, new trajectories can be classified as either fitting the model or as anomalies. For this, a k-Nearest Neighbour (kNN) classifier was used to determine the cluster that is closest to the trajectory being classified. A threshold, determined from statistical properties of the nearest cluster, is then used to determine whether the trajectory belongs to that cluster or is an outlier.

The following outlines the classifier approach. Let D denote the dataset of vessel trajectories, where each trajectory is assigned a cluster label C_i, $i \in [1...n]$, where n is the number of clusters. Let p be an unseen trajectory to be evaluated.

Given new trajectory p, the distance between p and a trajectory $o \in D$ is denoted by $dist(p,o)$. Whereby $dist(:,:)$ defines some similarity measure. It is reasonable to assume that the best candidate for a similarity measure for the classification step is the particular measure that was used in the clustering step. Thus, we explored here the same four similarity metrics used in the clustering approach (i.e., Hausdorff distance and the three DTW-based metrics).

Given the trajectory p, and positive integer k. The candidate cluster label C_i for p is assigned by using a majority vote from the set of cluster labels of the k-nearest neighbouring trajectories of p. If a tie occurs, k is increased by one until no tie occurs.

In determining the distance threshold, T, to decide whether the trajectory p does belong to its candidate cluster C_i the respective ε value from clustering was used. If the distance between trajectory p and the closest trajectory in Ci > T then p is an outlier with respect to C_i.

4. Experiments, Results and Discussion

Evaluation of the proposed approach was performed in the context of creating models for fishing boat trajectories within two identified fishing seasons and evaluating the models by using them to classify trajectories from fishing vessels out of season and of non-fishing vessel types. An accurate model would be expected to lead to a normal classification for previously unseen fishing vessels within the seasons. Abnormal classification would be expected for out of season fishing (i.e., a seasonal anomaly) and non-fishing vessel types (a mixture of seasonal and spatial anomalies due to their different behaviour).

A dataset of vessel trajectories from a section of the Western Australian coastline was used in the evaluation. The basis of our dataset was the Open Maritime Traffic Analysis Dataset (OMTAD) (Masek et al., 2021), which consists of curated and pre-processed AIS data sourced from the Australian Maritime Safety Authority (AMSA) (Australian Maritime Safety Authority, n.d.). The region of interest, along with examples of one month of cargo vessel trajectories are shown in Figure *1*.

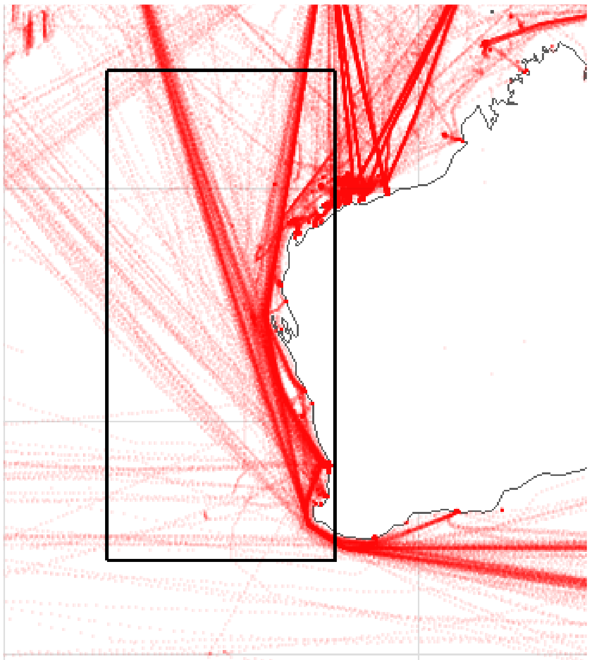

Figure 1. This image shows the location of interest, where the red lines are those of one months of cargo trajectories during January 2019 operating off the coast of Western Australia. Figure adapted from: Masek et al. (2021).

To provide more data evaluation, the original OMTAD dataset, which covers the years 2018-2020 was extended, using the same selection and pre-processing rules to also include fishing vessel trajectories across 2021-2023. From the expanded dataset, the fishing vessel class was isolated and inspected manually, revealing two main periods of fishing activity within each year – February, and May through to July. From these two seasons, 176 trajectories formed a set of candidates for modelling normal behaviour through the clustering approach, the trajectories are depicted in Figure 2(a).

Testing data for the classification step consisted of normal fishing vessel data, where four-fold cross validation was performed, using 132, (i.e., three quarters of the 176) seasonal fishing trajectories to form clusters and the remaining quarter (44 trajectories) as a test set, repeated four times. Two test sets of abnormal data were used. The first test set consisted of 44 fishing vessel trajectories from outside of the two identified seasons, designed to represent seasonal anomalies in otherwise normal spatial behaviour. The second abnormal test set also included the out of season fishing set and in addition included 44 trajectories from each of the other vessel classes: cargo, passenger and tanker randomly selected from the OMTAD dataset to represent a mix of spatial and spatiotemporal anomalies (with respect to in-season fishing vessel behaviour) – i.e. 176 abnormal trajectories in total.

The details of the clustering experiment and results are presented in Section 4.1, with experiments and results of the classification stage presented in Section 4.2.

4.1 Trajectory Clustering Experiments

For these experiments, the DBSCAN parameter ε was set using the IQR rule, described in 3.2.1, through which the optimal value for k, and therefore also the DBSCAN parameter *n,* was determined to be 2. A baseline result of clustering all 176 fishing trajectories was undertaken using each of the four candidate distance metrics. In addition, clustering was repeated for each fold in the four, for further analysis and for subsequent use in classifier evaluation.

Table *1* summarises the clustering results for the complete data set baseline and each of the four-folds. The non-normalised DTW produced a lower number of clusters than other metrics, the clusters upon inspection had spanned larger subregions. The non-normalised DTW algorithm was the only measure that assigned trajectories from separate seasons to a single cluster. This occurred in two of the experiments, indicated in orange shading in the table. Again, this can be related to the relative distances of trajectory pairs within a cluster and their lengths causing wider distance distributions. The normalised versions of DTW and Hausdorff distance clustered trajectories without any seasonal mixing in the clusters. The Hausdorff distance was the most consistent for the number of clusters and outliers between the baseline and each of the four-folds’ results.

As an example of clusters produced, Figure *2*(b) shows the cluster locations of the complete 176 trajectory dataset when using the normalized DTW distance by Tao et al. (2021). As can be seen multiple clusters are found across the region. Manual inspection showed that the clustering was successful in separating trajectories from the two identified seasons. Ten of the trajectories were not assigned to a cluster, indicated in red colour.

When examining the trajectories not assigned to clusters by DBSCAN, there was mostly consistency across the metrics in which trajectories were not clustered. An example of a notable exception was a relatively long trajectory that navigates from the lower port regions on the map and travels toward the top before returning to the source, unlike any other trajectory in the dataset. This trajectory was successfully identified as being too different to be assigned to any cluster by the non-normalised DTW metric, however the normalised DTW – Tao et al. (2021) grouped it into a cluster. The variety of results indicates that there are different strengths and weaknesses in various measures and in normalisation. It may be beneficial in building a model of normality to use a number of metrics to generate separate models to capture the overall behaviour.

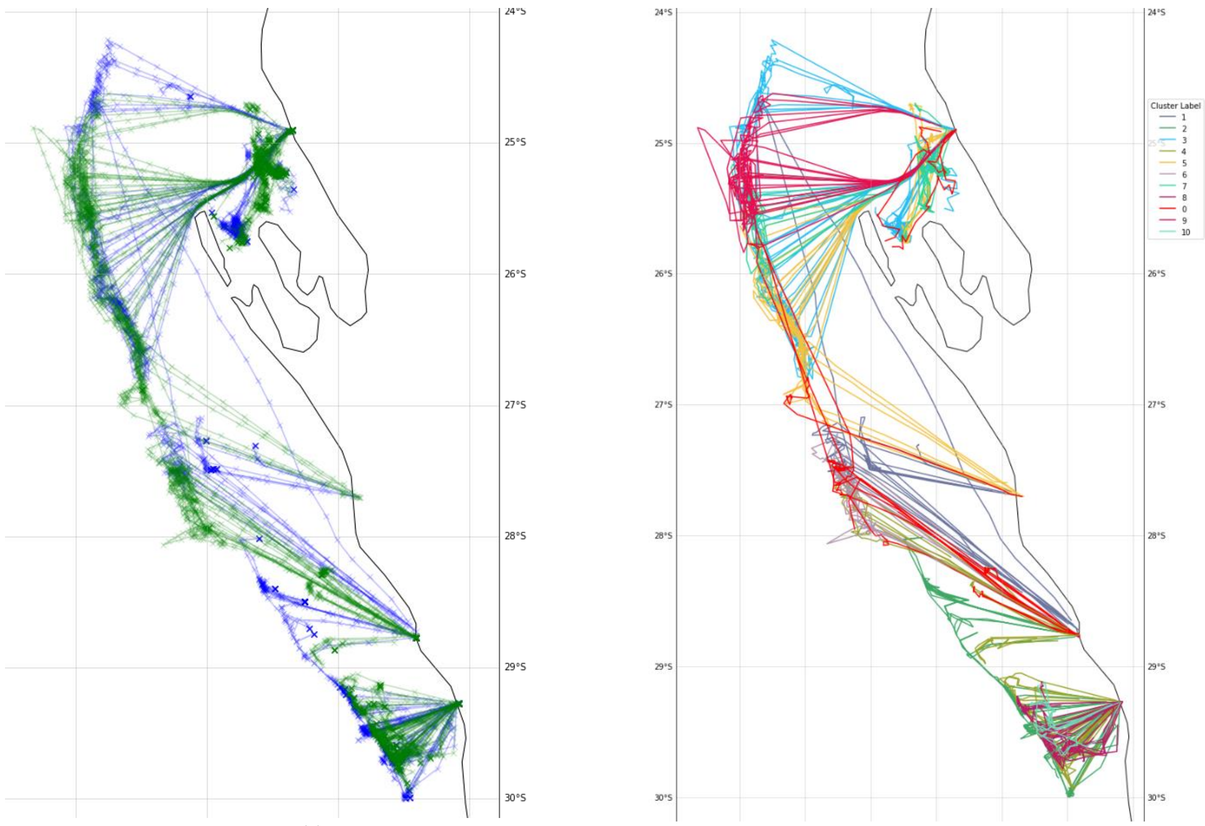

(a) (b)

Figure 2. (a) Locations of the seasonal fishing vessel trajectories, blue details those of season one, and green of season two. (b) Clustering results of all 176 fishing trajectories when minimum points set to 2 using DTW - Tao et al. (2021). Ten unique clusters are found, and 10 trajectories are deemed as not belonging to any cluster (identified by the red trajectory lines).

Similarity Measure	Baseline	1st fold	2nd fold	3rd fold	4th fold
DTW – Non Normalised	**5** (11)	**5** (6)	**4** (10)	**3** (12)	**5** (14)
DTW - Shen and Chi (2017) Normalisation	**13** (16)	**11** (6)	**7** (4)	**10** (9)	**7** (9)
DTW - Tao et al. (2021) Normalisation	**10** (10)	**9** (3)	**8** (4)	**9** (11)	**8** (5)
Hausdorff	**11** (8)	**12** (6)	**10** (7)	**9** (4)	**10** (8)

Table 1. Results of clustering the entire 176 fishing vessels trajectory dataset and each of the 4-fold subsets. The number of resultant clusters is shown, with the number of trajectories not assigned to a cluster in brackets. Experiments where trajectories from different seasons were grouped into the same cluster are indicated in orange shading.

4.2 Trajectory Classification Experiments

The cluster models from each of the four distance metrics were each evaluated for accuracy in classifying unseen trajectories. This was done using a kNN method, as outlined in 3.2.2. The experiments included an investigation of the effect of the kNN parameter *k* on the classifier accuracy.

4.2.1 Temporal Anomaly Detection: A set of experiments on purely temporal anomalies was conducted by using the normal fishing vessel test dataset (the quarter of the trajectories held out for each fold) and a set of fishing vessel trajectories from outside of the two defined seasons. The results of this temporal anomalies experiment are shown in Figure 3 (a), in terms of classifier accuracy as a function of *k*. Each data point represents the average accuracy for the four folds of normal data used in the clustering. It can be seen from these results that the DTW distance normalised using the approach in Tao et al. (2021) produced the highest accuracy. The non-normalised DTW performed significantly worse than the other metrics.

Going into more detail on the temporal anomaly experiment, Table *2* shows the confusion matrices for the 1st Fold, with *k* set to 3. This shows similar performance in correctly classifying normal and abnormal trajectories for all but the non-normalised DTW metric where the only difference in vessel activity was seasonal context. The non-normalised DTW metric results show that the poor overall performance results from miss-classification of the abnormal data (fishing out of season) with normal trajectories being classified on-par with the other metrics.

4.2.2 Spatiotemporal Anomaly Detection: The spatiotemporal anomaly detection experiments were run similarly to the temporal anomaly experiments but using the 176 abnormal trajectories test set described earlier (in addition to the normal trajectories).

Figure 3 (b) shows the spatiotemporal experiment results in terms of kNN classifier accuracy for each of the four cluster models as a function of k. Similarly to the temporal anomaly experiments, DTW-Tao et al. (2021) produced the highest overall accuracy, with the non-normalised version of the DTW producing the lowest accuracy.

Overall, higher accuracy was observed in the spatiotemporal experiments than the temporal experiments. An examination of the results showed that this was due to the high performance on correctly classifying non-fishing vessel types as abnormal. Performance on normal fishing and out of season fishing was identical to the temporal experiments, as the same models and settings were being used. For the non-fishing vessels, classification using the Hausdorff and both normalised versions of DTW models correctly classified all cargo, passenger and tanker trajectories as abnormal. The non-normalised DTW cluster model resulted in accuracies for cargo, passenger and tanker of 93%, 80% and 100% respectively.

DTW – Tao et al. (2021)		Actual	
k = 3, 1_Fold_Temporal		Normal	Abnormal
Predicted	Normal	41	2
	Abnormal	3	42

DTW-Shen-and-Chi-(2017)		Actual	
k = 3, 1_Fold_Temporal		Normal	Abnormal
Predicted	Normal	39	2
	Abnormal	5	42

DTW-Non-Normalised		Actual	
k = 3, 1_Fold_Temporal		Normal	Abnormal
Predicted	Normal	40	17
	Abnormal	4	27

Hausdorff		Actual	
k= 3, 1_Fold_Temporal		Normal	Abnormal
Predicted	Normal	38	2
	Abnormal	6	42

Table 2. Confusion matrices examining the 1st fold Temporal Experiment dataset with k set to 3.

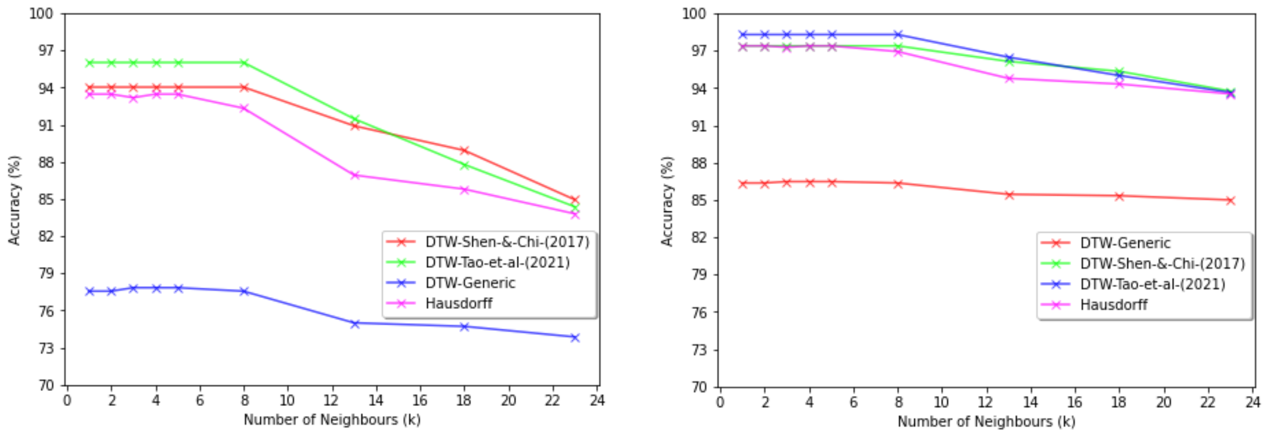

(a) (b)

Figure 3. Classification results: (a) Temporal experiments for seasonal, and non-seasonal fishing trajectories, and (b) Spatiotemporal experiments where the inliers are the respective k-folds seasonal fishing vessels and outliers are off-season fishing, cargo, passenger, and tanker vessels.

4.2.3 Discussion: It is interesting to note that though the different metrics used to produce clusters resulted in different numbers of clusters, they all showed promise in producing a model that could be used successfully in anomaly detection. Whilst the non-normalised DTW model showed lower overall performance, this should not discount the prospect of its use, alongside the other models, in a more comprehensive approach. For example, though not the focus of this study, it may also be considered anomalous is the length of a trajectory overly large compared to the expected trajectory length. In that case, the length-normalised distance measures would hide such a discrepancy in trajectory lengths.

In terms of the number of neighbours, *k*, for the kNN approach, the results show that this is not such a crucial in the dataset that was investigated. It should be noted that as the DBSCAN parameter of *n* was set to 2, which may naturally bias the optimal *k* to be close to 2. Results however showed similar performance up to a value for *k* or 5-8. Beyond *k* of 8, performance steadily decreased. A factor in this may be the low number of data points in each cluster, which could be mediated by using a larger historical dataset for training. Some support for this is seen as performance in the non-normalised DTW model did not see as large a drop in performance with larger *k*. This measure resulted in a lower number of clusters, each with more data points, so more possible neighbours to match to a new trajectory being classified.

5. Conclusion

In conclusion, this article presented an approach for modelling normal behaviour of vessels with seasonal context integrated into the model. The approach uses a clustering technique on a historical dataset of normal trajectories to build the model. Seasonality was incorporated by taking vessel trajectories and adding month of year information to each data point in the trajectory time series to augment the kinematic data. Four distance metrics were investigated within the clustering approach, three were versions of dynamic time warping and the fourth was Hausdorff distance.

Models were evaluated for their ability to discern between temporally abnormal behaviour (i.e., not occurring in the correct season). Further evaluation was performed with a mixture of spatial and temporal abnormalities by using trajectories of vessel types that did not correspond to the behaviour the model was trained with.

Results showed that in the context of a model of fishing vessel behaviour for specific seasons, out of season fishing and vessels not engaged in fishing could be identified based on their spatiotemporal behaviour. These results are promising, leading to prospects of further research in terms of scaling up the approach to a wider range of vessel types and a larger dataset. In particular, the different characteristics of clusters resulting from the various metrics raise the prospect of a multi-model approach and an investigation of which type of model is better suited to specific sub-types of spatiotemporal anomalies.

Acknowledgements

This research is supported by the Commonwealth of Australia as represented by the Defence Science and Technology Group (DSTG) of the Department of Defence. We would like to especially acknowledge Damion Glassborow of DSTG for his assistance in earlier works.

References

Australian Maritime Safety Authority., n.d. Digital Data. https://www.operations.amsa.gov.au/Spatial/DataServices/DigitalData (26 April 2022)

Ester, M., Kriegel, H. P., Sander, J., & Xu, X., 1996: A density-based algorithm for discovering clusters in large spatial databases with noise. *In Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining*, *E. Simoudis, J. Han, and U. Fayyad, Eds. AAAI Press*, 226–231.

Guan, Y., Zhang, J., Zhang, X., Li, Z., Meng, J., Liu, G., Bao, M., & Cao, C., 2021: Identification of Fishing Vessel Types and Analysis of Seasonal Activities in the Northern South China Sea Based on AIS Data: A Case Study of 2018. *Remote Sensing*, *13*(10).

Han, X., Armenakis, C., & Jadidi, M., 2021: Modeling Vessel Behaviours by Clustering AIS Data Using Optimized DBSCAN. *Sustainability*, *13*(15). https://doi.org/10.3390/su13158162

Huttenlocher, D. P., & Kedem, K., 1990. Computing the minimum Hausdorff distance for point sets under translation. In *Proceedings of the sixth annual symposium on Computational geometry*, 340-349.

Huttenlocher, D. P., Klanderman, G. A., & Rucklidge, W. A., 1993. Comparing images using the hausdorff distance. *Ieee Transactions on Pattern Analysis & Machine Intelligence, 15*(9), 850–863. https://doi.org/10.1109/34.232073

IMO., 2015. Revised Guidelines for the onboard operational use of shipborne automatic identification systems (AIS), Resolution A. 1106 (29).

Li, H., Liu, J., Liu, R. W., Xiong, N., Wu, K., & Kim, T.-H., 2017. A Dimensionality Reduction-Based Multi-Step Clustering Method for Robust Vessel Trajectory Analysis. *Sensors (Basel, Switzerland)*, *17*(8), 1792. https://doi.org/10.3390/s17081792

Masek, M., Lam, P., Rybicki, T., Snell, J., Wheat, D., Kelly, L., Glassborow, D., Smith-Gander, C., 2021. The Open Maritime Traffic Analysis Dataset. *MODSIM2021, 24th International Congress on Modelling and Simulation*. Modelling and Simulation Society of Australia and New Zealand, 981-987. https://doi.org/https://doi.org/10.36334/modsim.2021.M9.masek

Nguyen, D., Vadaine, R., Hajduch, G., Garello, R., & Fablet, R., 2021. Geotracknet—a maritime anomaly detector using probabilistic neural network representation of ais tracks and a contrario detection. *Ieee Transactions on Intelligent Transportation Systems*. https://doi.org/10.1109/TITS.2021.3055614c

Olesen, K. V., Boubekki, A., Kampffmeyer, M. C., Jenssen, R., Christensen, A. N., Hørlück, S., & Clemmensen, L. H., 2023. A Contextually Supported Abnormality Detector for Maritime Trajectories. Journal of Marine Science and Engineering, 11(11).

Petneházi, G., 2019. Recurrent neural networks for time series forecasting. arXiv preprint arXiv:1901.00069.

Riveiro, M., Pallotta, G., & Vespe, M., 2018. Maritime anomaly detection: A review. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 8, e1266. https://doi.org/10.1002/widm.1266

Sakoe, H., & Chiba, S., 1978. Dynamic programming algorithm optimization for spoken word recognition. IEEE Transactions on Acoustics, Speech, and Signal Processing, 26(1). https://doi.org/10.1109/TASSP.1978.1163055

Shen, S., & Chi, M., 2017. Clustering student sequential trajectories using dynamic time warping. Proceedings of the 10th International Conference on Educational Data Mining (EDM). (pp. 266-271).

Shen, H., Tang, H., & Yin, Y., 2022. A novel method for ship trajectory clustering. *International Journal of Naval Architecture and Ocean Engineering*, *14*, 100474. https://doi.org/https://doi.org/10.1016/j.ijnaoe.2022.100474

Tao, Y., Both, A., Silveira, R., Buchin, K., Sijben, S., Purves, R., Laube, P., Peng, D., Toohey, K., & Duckham, M., 2021. A comparative analysis of trajectory similarity measures. *GIScience & Remote Sensing*, *58*, 1-27. https://doi.org/10.1080/15481603.2021.1908927

Xie, Z., Bai, X., Xu, X., & Xiao, Y., 2024. An anomaly detection method based on ship behavior trajectory. *Ocean Engineering*, 293, 116640. https://doi.org/https://doi.org/10.1016/j.oceaneng.2023.116640

Zhao, L., & Shi, G., 2019. Maritime Anomaly Detection using Density-based Clustering and Recurrent Neural Network. *Journal of Navigation,* 72(4), 894-916. https://doi.org/10.1017/S0373463319000031

Zhu, R., Ji, X., Yu, D., Tan, Z., Zhao, L., Li, J., & Xia, X., 2020. KNN-Based Approximate Outlier Detection Algorithm Over IoT Streaming Data. *IEEE Access*, 8, 42749-42759. https://doi.org/10.1109/ACCESS.2020.2977114