The 9th International Symposium on Ship Operations, Management & Economics (SOME)
15-16 September 2026, Athens GR

www.sname.org

Fuel Oil Consumption Prediction Model for AI Optimisation Platform

Andreas Kouvaras[1], **Dimitris Stavropoulos**[1], **Aikaterini Karampasi**[1], **Constantine Abazis**[1], **Michael Fan**(FL)[2], **George Margelis**(M)[3]

1. VesselFront, Greece, andreas@vesselfront.com, dimitris@vesselfront.com, katerina@vesselfront.com, costas@vesselfront.com;
2. Tsakos Energy Navigation Ltd., Greece, and University of Strathclyde, Scotland, mfan@tsakoshellas.gr;
3. Tsakos Shipping and Trading S.A., Greece, gmargelis@tsthellas.gr;

The efficient operation of vessels at sea requires a reduction in operating costs, which, in turn, depend heavily on fuel oil consumption. Consequently, accurately predicting fuel oil consumption is important for cost management as well as the environment. Recently, machine learning techniques have been evaluated and customised for fuel oil consumption prediction across a range of vessel types e.g., most notably container ships, bulk carriers, and oil tankers. For instance, the use of deep learning models has led to high predictive accuracy. We assessed the findings of earlier empirical analyses concerning fuel oil consumption prediction, relying on real, large maritime datasets, including millions of position messages emitted from vessels sailing across the globe. Moreover, we performed an ablation study that allows us to appreciate the significance of different feature sets. On one hand, our empirical analysis confirms, in stricter evaluation settings, earlier results on predictive accuracy. On the other hand, our findings indicate that simpler learning models than those proposed in the literature may achieve comparable performance.

KEY WORDS: Fuel Oil Consumption prediction; Vessel's Main Engine; Big Maritime Data; Telemetry Data; Fleet Reports.

INTRODUCTION

Fuel Oil Consumption (FOC) is a critical indicator for the shipping industry, with a direct impact on fleet efficiency, operating costs and environmental impact. FOC prediction has become of great interest as it directly affects the economic and environmental performance of fleets. In weather routing for instance, optimal route selection depends heavily on FOC. Predicting FOC is challenging due to the wide variety of factors affecting consumption such as mechanical, meteorological and sea conditions (Bal Beşikçi et al. 2016; Anh Tran 2021; Bialystocki and Konovessis 2016). Moreover, the variation in fuel type and density adds complexity to the accuracy of the predictions.

Early work on FOC prediction relied on white-box models, i.e., first-principles models developed from the physics governing the operations of the vessel, such as propulsive power, resistance to motion, and the environmental factors affecting such motion. White-box models are interpretable, and their outcomes may be explained in terms of physical principles. On the other hand, these models are highly dependent on parametrization (Lu et al. 2015; Coraddu et al. 2017; Moreno-Gutiérrez et al. 2015). Moreover, such models require comprehensive knowledge of the vessel's mechanics, which can often be complex and hard to acquire (Bialystocki and Konovessis 2016; Gkerekos et al. 2019; Meng et al. 2016; Yao et al. 2012).

The availability of Big Mobility Data offers new opportunities to address these challenges. Such data include real-time and historical observations collected from fleet reports, Automatic Identification System (AIS), and telemetry data from on-board sensors (Artikis and Zissis 2021; Coraddu et al. 2017; Gkerekos and Lazakis 2020; Gkerekos et al. 2019; Anh Tran 2021). Fleet Reports are daily reports compiled by the members of a vessel's crew, including information concerning the vessel's operation, such as the position, speed, course, amount of fuel consumed, and the condition and operation of the engines. Although fleet reports provide a regular update and support decision-making, they are prone to human error as they rely only on the crew. AIS is used to exchange information about the vessels' identity, position, speed, course, draught and other voyage-related information (Bereta et al. 2021). Telemetry data contain information collected from various sensors installed on a vessel, such as data on the position, course, engine, speed and fuel consumption of the vessel, as well as the speed and direction of wind and sea currents, allowing monitoring of the vessels' operation (Gkerekos and Lazakis 2020; Gkerekos et al. 2019).

The use of data analytics techniques can enhance the prediction and optimisation of FOC. For instance, multiple linear regression algorithms using aggregation of fleet reports have been used for predicting FOC (Gkerekos et al. 2019). Recent research investigated techniques such as Artificial Neural Networks (ANNs), Support Vector Machines, Ridge Regression, and ensemble techniques, such as Random Forests and Gradient Boosting Regressors (Bal Beşikçi et al. 2016; Gkerekos et al. 2019; Karagiannidis et al., n.d.; Simonsen et al. 2018; Anh Tran 2021). Gkerekos et al. (2019) used, separately, fleet reports and telemetry data in order to predict FOC. To predict FOC of a vessel's main engine, ANNs seem to outperform other machine learning techniques (Gkerekos et al. 2019; Karagiannidis et al., n.d.; Anh Tran 2021). For instance, Gkerekos and Lazakis (2020) employed an ANN to predict FOC for weather routing.

We examined the efficacy of ANNs to predict a vessel's FOC. We used telemetry data at 15-second intervals and fleet reports from six Aframax tankers covering a period of 4 years. We propose a methodology including data pre-processing techniques, implementation and evaluation of ANNs, and an ablation study in order to examine the importance of each feature subset.

Our contributions may be summarised as follows:

a) Combine telemetry data with fleet reports in order to construct voyages and voyage segments. This way, we perform a thorough evaluation of machine learning models.
b) Propose an automated process for transforming raw data into valuable data, by performing, among others, outlier filtering and feature engineering.
c) Conduct an extensive empirical analysis using 21,445,602 filtered telemetry data points from six vessels operating over a period of four years. The features of our dataset include speed over ground, main engine revolutions per minute (RPM), longitudinal velocity component of sea currents relative to the ship's heading, true wind speed, relative wind direction, distance travelled, slip, mean draft and trim.
d) Present an ablation study which was used to examine the importance of each feature subset. Whilst our empirical analysis confirms, in stricter evaluation settings, earlier results on the accuracy of ANNs for FOC prediction, our findings also indicate that simpler learning models than those proposed in the literature may achieve comparable performance.

The remainder of this paper is structured as follows. It elaborates on the proposed methodology, focusing on aspects of data pre-processing and model training. It presents the employed evaluation metrics. It presents our empirical analysis, and finally, it summarises the presented work and outlines proposals for further work.

METHODOLOGY

Our aim is to predict a vessel's voyage FOC. Towards this goal, we follow a set of pre-processing steps, employ and evaluate a model of ANNs, and perform an ablation study. Figure 1 illustrates our methodology. First, we collected the necessary data for FOC prediction, i.e., fleet reports and telemetry data. We then followed six data pre-processing steps, which we describe in the following sub-sections. We then proceeded to train an ANN per vessel and evaluated the results.

Dataset Acquisition

Table 1 presents an overview of the features of the dataset used to train the ANNs. We focus on crude oil Aframax tankers. Our dataset consists of 6 vessels (V1-V6) covering a period of 4 years. Figure 2 shows the dataset sizes for the raw dataset as well as the dataset sizes after each preprocessing step for vessels V1-V6. We collected telemetry data provided by sensors on-board that contain information about vessel propulsion and performance i.e., speed, RPM and slip, vessel loading condition (i.e., draft and trim), sea and weather conditions (i.e., longitudinal velocity component of sea currents, wind speed and direction), and fuel oil consumption. These data were reported at 15-second intervals. Oil tankers, like other tramp services, do not follow a fixed departure or arrival schedule. Furthermore, telemetry data does not provide information about whether a vessel is at port of when a voyage starts or ends. In order to separate different voyages, we correlated fleet reports with telemetry data.

Table 1. Dataset Features

#	Feature	Description	Units
1	Speed over ground	Speed of vessel in relation to a fixed point on the Earth's surface.	knots
2	Main Engine RPM	Main engine RPM in revolutions per minute	rev/min
3	Longitudinal speed of sea currents	Longitudinal velocity component of sea currents relative to ship's heading and expressed as difference between the speed over ground and speed through water.	knots
4	Wind Speed True	Wind speed experienced by the vessel.	knots
5	Wind Direction Relative	Direction of the wind relative to the vessel's heading.	degrees
6	Distance travelled over ground	Actual distance travelled by the vessel over ground.	nm
7	Slip	Difference between the theoretical distance a propeller would advance based on its pitch and rpm and actual distance travelled by the ship, expressed as fraction of theoretical distance.	%
8	Mean Draft	Average of forward and aft drafts	m
9	Trim	Difference between draft aft and draft fwd.	m

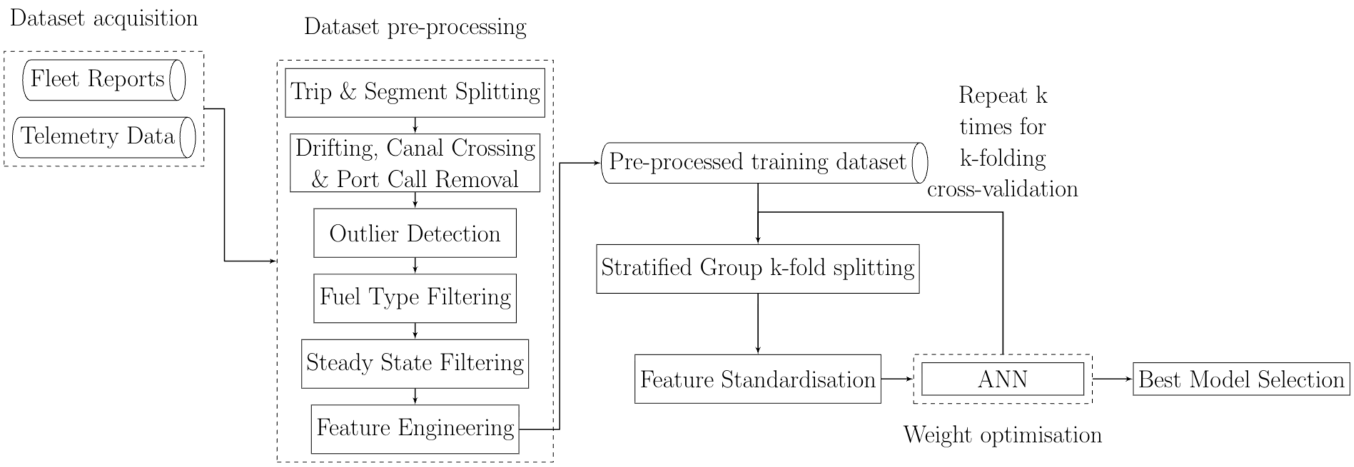

Figure 1. Visual representation of our methodology for understanding FOC prediction. Cylindrical shapes represent datasets, rectangular shapes denote methodology steps, and dashed rectangles indicate overarching processes that encompass multiple steps.

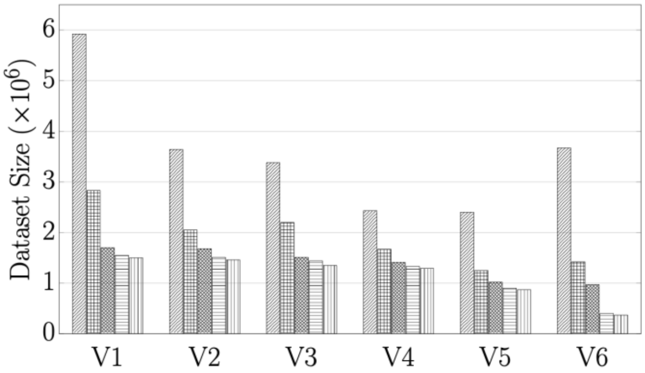

Figure 2. Dataset Sizes for the preprocessing steps for vessels V1-V6. In each vessel, the bars express, respectively, the input data, the data after drifting, canal crossing, and port call removal, the data after outlier detection, the data after fuel type filtering, and the final data after steady state filtering.

Data Pre-processing

Voyage and Segment Splitting. Within port areas and Emission Control Areas (ECAs), a vessel is required to consume a different type of fuel from that which is consumed in open seas outside ECAs. Moreover, during the time a vessel is engaged in loading and discharging operations, only the auxiliary engines and boilers are operating. For these reasons, a uniform approach towards FOC prediction would be complex. Therefore, we predict FOC from port departure to port arrival, excluding their stay at ports, with a combination of telemetry data with fleet reports. To separate the voyages, we used the fleet reports whose type was "Departure" and "Arrival". Each voyage starts from a "Departure" report and ends at the previous report from the next "Departure". In several cases, despite the existence of the "Departure" report, the vessel changed berthing area within the same port or the voyage was eventually cancelled. Therefore, the "Departure" should be followed by a report whose type is "Noon", in order to ensure that the vessel has indeed started its voyage. "Noon" reports have an average duration of 24 hours. Each noon report corresponds to a segment, i.e., a part of the voyage.

Drifting, Canal Crossing, and Port Call Removal. During a Canal Crossing (e.g., Panama or Suez Canal) the vessel often operates at reduced speeds, with frequent changes in shaft power, due to safety and navigational restrictions. In the same vein, during drifting or other manoeuvring operations, the vessel remains stationary or moves minimally with the engine often being in idle mode. These conditions result in fuel consumption patterns that are not indicative of sailing at open-sea. Hence, we remove the data that correspond to fleet reports whose type was "Drifting" or "Canal Crossing". Then, to remove the data corresponding to port calls, we assumed that each report between an "Arrival" report and the last report before the next "Departure" report corresponds to a vessel's stay in port. The data over the periods that the vessels were in idling condition, should not be used for FOC prediction.

Outlier Detection. The raw data may include inconsistent and/or incorrect data entries (e.g., due to logging inconsistencies, human error, or sensor failures) that should be discarded. First, we removed the null values. Second, for each feature, we removed all values that were not in $[\mu - 3\sigma,\ \mu + 3\sigma]$. Third, we removed cases where RPM < 23 and Shaft Power < 0 as they correspond to vessels manoeuvring operations or inconsistent data.

Fuel Type Filtering. Residual Fuel Oil is often the primary fuel for large vessels, during sailing at the open sea. Similarly, Marine Diesel Oil and Marine Gas Oil are commonly used during special conditions such as when sailing at ECAs. This is confirmed by the data we collected. For this reason, we proceeded to fuel type filtering by selecting only those conditions that the vessel consumed Residual Fuel Oil. In this way, our model is not affected by fuel properties such as energy content, density, and efficiency.

Steady State Filtering. In order to identify operating conditions that do not correspond to the usual operational states considered for the pilot-to-pilot FOC prediction model, we need to discard the transient situations. In this respect, M/E Manufacturers provide a minimum engine speed for continuous operation. In this case, we identified in a sliding window the cases on an hourly basis where the minimum and maximum RPM differ from each other by more than 5%. Therefore, to keep only the data points reflecting steady-state operation, we discarded the transient situations i.e., the corresponding minimum RPM values. Figure 2 illustrates the effect of this type of data filtering. Overall, as shown in Figure 2, the valuable data, i.e., the data that can be safely used for FOC prediction, is a fragment of the input data. Moreover, it should be noted that all data pre-processing steps have been automated, i.e., any given input dataset is "automatically" transformed into data ready for FOC prediction.

Feature Engineering. Following Gkerekos et al. (2019) we engineered features and adjusted some other features using various transformations:

(i) The actual distance covered by a vessel was computed. This offers a direct understanding of the energy requirements of a voyage.

(ii) "slip" was computed, given by the difference between the actual distance travelled by a ship and the theoretical distance travelled based on propeller rotation. Slip is an indicator that can provide important information about the performance of the ship. For instance, in the case of a loss of efficiency, fuel consumption can be adversely affected as the vessel has to expend more energy to maintain the desired speed. Slip is expressed as a percentage and can have a negative value if a following sea current or wind exists.

(iii) The effect of the "trim" of the vessel on FOC was computed. By adjusting the trim, the vessel can reduce the resistance, thereby optimising FOC.

(iv) The effect of "mean draft" on FOC was computed.

(v) "Relative wind speed" was transformed to "true wind speed", i.e., a measurement independent of the movement or position of the vessel.

Feature Standardization

All attributes in the dataset are standardized by removing the mean and scaling to unit variance. For an attribute x, the standardized value x' is computed as:

$$x' = \frac{x - \mu}{\sigma} \tag{1}$$

where μ represents the mean of the attribute, and σ is its standard deviation. Standardizing ensures that all attributes contribute equally to the objective function used during model training.

Stratified-Group k-fold Splitting

In order to train a machine learning algorithm to predict a vessel's FOC, we performed stratified k-fold splitting. More precisely, we performed two types of experiments: a) voyage-based splitting, and b) segment-based splitting. The use of voyage-based splitting allows for a stricter evaluation setup, in the sense that each voyage is either at the training set or the test set.

Model Selection

Gkerekos et al. (2019) in their investigation of FOC prediction techniques, revealed that Extra Trees Regressors, Random Forest Regressors, Support Vector Regressors, and ANNs deliver the best performance. Among these, only ANNs and Support Vector Regressors provide the necessary generalization capabilities for predicting FOC. Each vessel has unique technical characteristics. In our case, our dataset consists of six sister vessels, i.e., the vessels have the same dimensions, same Deadweight Tonnage, same engine type, and are built by the same shipyard. However, the maintenance status of each vessel, the trading profile, the sea temperature at the sailing areas, and the idling periods of of each vessel are unique. These parameters significantly affect the FOC performance, thus, we proceeded to train an ANN per vessel. We adopted the optimal hyperparameter values proposed by Gkerekos and Lazakis (2020).

EVALUATION METRICS

Explained Variance (EV)

The explained variance (EV) quantifies the amount of variance a model can capture from a given dataset. Given the true target output y, the estimated target output can be obtained as:

$$\hat{y} = f(x) \tag{2}$$

where $f(\cdot)$ represents any derived model. The explained variance EV is then calculated as:

$$EV(y, \hat{y}) = 1 - \frac{\sigma^2(y - \hat{y})}{\sigma^2(y)} \tag{3}$$

where σ_x denotes the standard deviation of parameter x. The ideal EV score is 1.0, achieved when:

$$\sigma^2(y - \hat{y}) \rightarrow 0 \tag{4}$$

with lower values indicating poorer performance.

Mean Absolute Error (MAE)

The Mean Absolute Error (MAE) measures the expected value of the absolute error (L1 norm) and can be computed as:

$$MAE(y, \hat{y}) = \frac{1}{n} \sum_{i=1}^{n} |y_i - \hat{y}_i| \tag{5}$$

where n represents the number of samples in y, and y_i refers to the i-th sample of y.

Median Absolute Error (Median AE)

The Median Absolute Error (MedAE) is computed as:

$$MedAE(y, \hat{y}) = \text{median}\,(|y_1 - \hat{y}_1|, \ldots, |y_i - \hat{y}_i|) \tag{6}$$

MedAE is particularly robust to outliers because it only considers the median performance, making it less sensitive to extreme values.

Mean Squared Error (MSE)

Following the same formulation as above, the Mean Squared Error (MSE) is computed as:

$$MSE(y, \hat{y}) = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{7}$$

MSE represents the expected value of the squared error. If the 1/n term is omitted, MSE becomes the L_2 loss function. When used as a cost function for optimisation, both MSE and the L_2 loss function produce similar results.

In comparison to MAE, MSE places greater emphasis on larger deviations between the true and estimated targets. For this reason, MAE is generally more robust to outliers.

Coefficient of determination (R^2)

The coefficient of determination (R^2) is computed as:

$$R^2(y, \hat{y}) = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{8}$$

R^2 provided a measure of the quality of future model outputs (predictions), with the best possible R^2 score being 1, and lower values indicating worse performance.

Root Mean Squared Error (RMSE)

The Root Mean Squared Error (RMSE) is calculated as:

$$RMSE(y, \hat{y}) = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{9}$$

RMSE represents the square root of the average of the squared differences between the true values and the predicted values. It provides a measure of the magnitude of the error, with lower values indicating bett4er model performance. The RMSE is particularly sensitive to larger errors, as it penalizes them more due to the squaring of the differences. This makes it a useful metric when large errors are undesirable, but it also makes RMSE more sensitive to outliers compared to other error metrics like MAE.

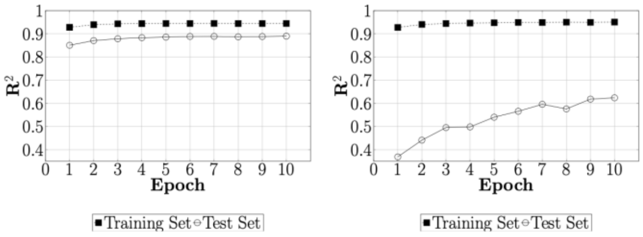

(a) Learning Curves for vessel V3 - R^2

(b) Learning Curves for vessel V6 - R^2

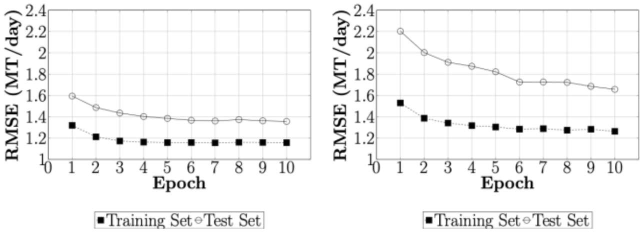

(c) Learning Curves for vessel V3 - RMSE

(d) Learning Curves for vessel V6 - RMSE

Figure 3. Learning curves for voyage-based splitting (vessels V3 and V6).

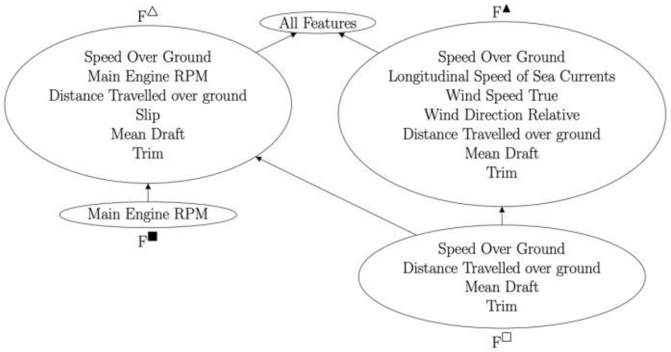

Figure 4. The subset relationships between the feature sets of the ablation study.

EMPIRICAL ANALYSIS

FOC prediction

Table 2 presents the dataset size, number of voyages and segments, average duration of voyages and segments, the number of folds as well as the values of the different metrics with which we evaluated the ANN, which was trained for each vessel in voyage-based and segment-based splitting. Vessels V1-V4 have a comparable dataset size, in contrast to vessels V5 and V6 which have a much smaller dataset size.

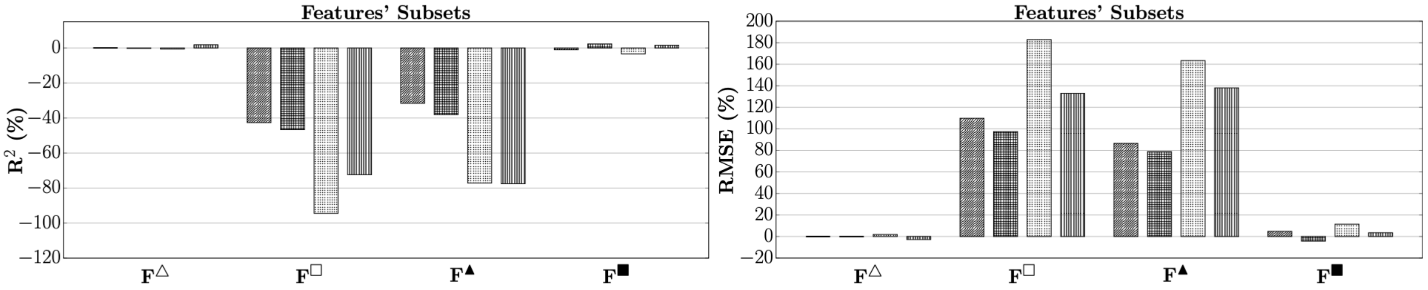

(a) R2 % difference across features' subsets for vessels V1-V4. **(b) RMSE % difference across features' subsets for vessels V1-V4.**

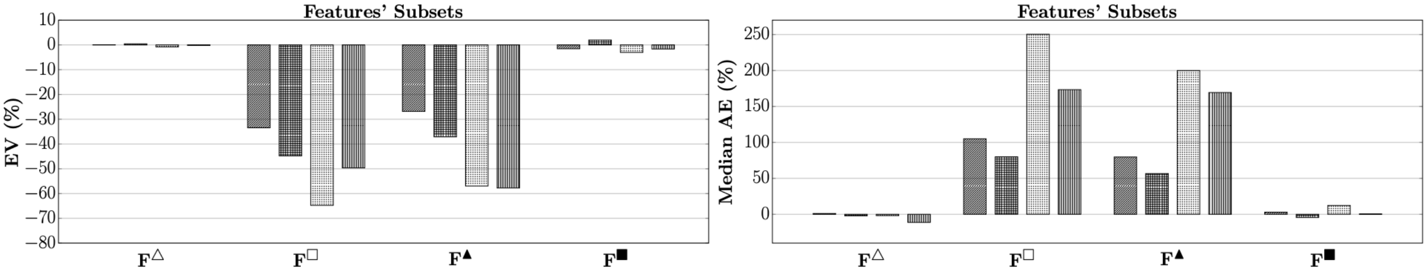

(c) EV % difference across features' subsets for vessels V1-V4. **(d) Median AE % difference across features' subsets for vessels V1-V4.**

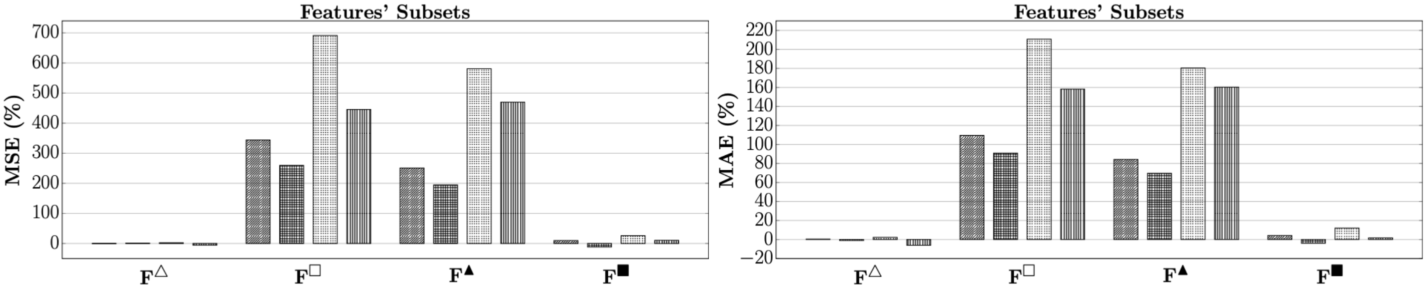

(e) MSE % difference across features' subsets for vessels V1-V4. **(f) MAE % difference across features' subsets for vessels V1-V4.**

Figure 5. Percentage performance difference between the ANN trained on all features and the ANNs trained on feature subsets. In each diagram, the bars express, respectively, the performance difference for V1, V2, V3, and V4.

Table 2. Performance Metrics per Vessel. MT stands for Metric Tons, d for days, and h for hours.

Vessel	Dataset Size	Splitting	#	Avg. Duration	Folds	EV	MAE (MT/d)	Median AE (MT/d)	MSE (MT²/d²)	R²	RMSE (MT/d)
V1	1,501,703	Voyages	36	13.07 d	10	90.95%	1.037	0.852	1.849	89.13%	1.342
		Segments	387	23.86 h	10	92.50%	0.925	0.729	1.518	92.44%	1.230
V2	1,458,960	Voyages	38	8.82 d	10	89.32%	1.645	1.457	4.574	85.57%	2.019
		Segments	339	23.73 h	10	93.45%	1.228	0.953	2.647	93.33%	1.622
V3	1,354,327	Voyages	24	13.01 d	10	90.04%	1.006	0.718	1.922	88.96%	1.357
		Segments	313	23.94 h	10	93.93%	0.873	0.665	1.473	93.89%	1.211
V4	1,290,317	Voyages	18	16.05 d	10	88.52%	1.206	0.926	3.086	84.39%	1.685
		Segments	290	23.91 h	10	94.49%	1.011	0.780	2.281	94.44%	1.501
V5	870,547	Voyages	14	11.62 d	10	74.80%	0.971	0.740	1.897	70.56%	1.329
		Segments	191	23.94 h	10	95.28%	0.809	0.579	1.425	95.15%	1.190
V6	367,582	Voyages	8	11.47 d	5	76.56%	1.241	1.043	3.078	62.38%	1.657
		Segments	92	23.93 h	10	94.90%	0.928	0.694	1.832	94.86%	1.348

The number of voyages for vessels V1-V3 is comparable, while vessels V4-V6 have performed less than half of the voyages of vessels V1-V3. V4 has the longest average voyage duration, while V2 has the shortest. V2 has the largest number of voyages, which means that V2 makes shorter distance voyages. The average duration of the segments is comparable across all V1-V6 vessels. We performed a 10-fold cross-validation in both voyage-based and segment-based splitting. Due to the small number of voyages available for vessel V6, we trained the corresponding ANN with a 5-fold cross-validation. Our results confirm the empirical analysis of Gkerekos et al. (2019) in terms of the accuracy achieved.

Please note that, as compared to previous works (Gkerekos et al. 2019; Gkerekos and Lazakis 2020), we used a bigger volume of data and number of vessels and a larger dataset. As expected, the reported scores in the case of segment-based splitting are higher than those of voyage-based splitting. We note that the effect of dataset size affects the predictive accuracy for vessels V5 and V6 in the case of voyage-based splitting where R^2 decreases significantly, while in the case of segment-based splitting there is no noticeable difference. RMSE is a metric that may explain the significance of the results from the point of view of shipowners who wish to reduce operating costs and thus costs related to FOC. Based on the Ship & Bunker platform, the bunker prices, i.e., the price of fuel, on a worldwide scale is on average 550 U.S. dollars per metric ton (MT). According to Table 2 the average RMSE per vessel in the case of voyage-based splitting is 1.564 MT/day, while the highest reported RMSE is 2.019 MT/day. In Figure 3 we present the learning curves for the vessels V3 and V6, i.e., the best- and worst-performing vessels in the voyage-based splitting case, which is a more challenging task compared to segment-based splitting. For V3 convergence is achieved early, while in the case of V6 predictive accuracy benefits from additional training epochs.

Ablation Study

We conducted an ablation study with four different feature subsets, in the case of voyage-based splitting, for vessels V1-V4 which have the largest dataset size. Figure 4 illustrates each feature subset, while Figure 5 shows the percentage differences of each feature subset experiment compared with the case in which we used all features.

1. The influence of wind direction, wind speed, and speed of sea currents were investigated; thus, we used a feature subset F△ without them. Surprisingly, the weather features (wind direction, wind speed, and speed of sea currents) in our dataset have a very minor effect in predictive accuracy – see Figure 5. This leads us to the conclusion that other weather features, such as waves, which were not examined in this study, are more important. It should be noted that the weather has a direct effect on the speed of a vessel, and the distance travelled, i.e., features that were present in this set of experiments.

2. In weather routing applications, the commonly used vessel-related features are the speed, distance travelled, draft and trim of the vessel. Thus, we trained the ANNs only with those features (see feature subset F□). Our results demonstrate that performance drops significantly when training uses only the aforementioned features.

3. A feature subset without information about RPM was used; i.e., we removed RPM and Slip from the original feature set, in order to examine the influence of RPM in predictive accuracy (see features subset F▲). Our results show that RPM is essential for FOC prediction.

4. The ANNs were trained using only RPM (see feature subset F■). Our results demonstrate that performance is only slightly affected, illustrating the importance of RPM in FOC prediction. Moreover, this result indicates that a simpler model than ANNs may be sufficient for FOC prediction.

CONCLUSIONS

We aimed to understand vessel's fuel oil consumption prediction, by training an ANN-model based on literature findings (Gkerekos and Lazakis 2020; Gkerekos et al. 2019). Our contribution is threefold: an automated pipeline that fuses high-frequency telemetry with fleet reports into voyages and segments. In addition, to the segment-based protocol community reported, and a ablation study, quantifying the marginal contribution of eac feature group. On one hand, our empirical analysis confirmed, in stricter evaluation settings, earlier results on predictive accuracy. On the other hand, our findings indicate that simpler learning models than those proposed in the literature may achieve comparable performance.

Future work will incorporate additional features, such as the effect of waves, ship's fouling condition and fuel oil consumption of auxiliary engines and boilers. Finally, we aim to incorporate domain expert knowledge into training and apply neurosymbolic learning techniques.

REFERENCES

Anh Tran, Tien. 2021. 'Comparative Analysis on the Fuel Consumption Prediction Model for Bulk Carriers from Ship Launching to Current States Based on Sea Trial Data and Machine Learning Technique'. *Journal of Ocean Engineering and Science* 6 (4): 317–39. https://doi.org/10.1016/j.joes.2021.02.005.

Artikis, Alexander, and Dimitris Zissis, eds. 2021. *Guide to Maritime Informatics*. Springer International Publishing. https://doi.org/10.1007/978-3-030-61852-0.

Bal Beşikçi, E., O. Arslan, O. Turan, and A. I. Ölçer. 2016. 'An Artificial Neural Network Based Decision Support System for Energy Efficient Ship Operations'. *Computers & Operations Research* 66 (February): 393–401. https://doi.org/10.1016/j.cor.2015.04.004.

Bereta, Konstantina, Konstantinos Chatzikokolakis, and Dimitris Zissis. 2021. 'Maritime Reporting Systems'. In *Guide to Maritime Informatics*, edited by Alexander Artikis and Dimitris Zissis. Springer International Publishing. https://doi.org/10.1007/978-3-030-61852-0_1.

Bialystocki, Nicolas, and Dimitris Konovessis. 2016. ‘On the Estimation of Ship’s Fuel Consumption and Speed Curve: A Statistical Approach’. *Journal of Ocean Engineering and Science* 1 (2): 157–66. https://doi.org/10.1016/j.joes.2016.02.001.

Coraddu, Andrea, Luca Oneto, Francesco Baldi, and Davide Anguita. 2017. ‘Vessels Fuel Consumption Forecast and Trim Optimisation: A Data Analytics Perspective’. *Ocean Engineering* 130 (January): 351–70. https://doi.org/10.1016/j.oceaneng.2016.11.058.

Gkerekos, Christos, and Iraklis Lazakis. 2020. ‘A Novel, Data-Driven Heuristic Framework for Vessel Weather Routing’. *Ocean Engineering* 197 (February): 106887. https://doi.org/10.1016/j.oceaneng.2019.106887.

Gkerekos, Christos, Iraklis Lazakis, and Gerasimos Theotokatos. 2019. ‘Machine Learning Models for Predicting Ship Main Engine Fuel Oil Consumption: A Comparative Study’. *Ocean Engineering* 188 (September): 106282. https://doi.org/10.1016/j.oceaneng.2019.106282.

Karagiannidis, Pavlos, Nikos Themelis, George Zaraphonitis, and Christos Giordamlis. n.d. *Ship Fuel Consumption Prediction Using Artificial Neural Networks*.

Lu, Ruihua, Osman Turan, Evangelos Boulougouris, Charlotte Banks, and Atilla Incecik. 2015. ‘A Semi-Empirical Ship Operational Performance Prediction Model for Voyage Optimization towards Energy Efficient Shipping’. *Ocean Engineering* 110 (December): 18–28. https://doi.org/10.1016/j.oceaneng.2015.07.042.

Meng, Qiang, Yuquan Du, and Yadong Wang. 2016. ‘Shipping Log Data Based Container Ship Fuel Efficiency Modeling’. *Transportation Research Part B: Methodological* 83 (January): 207–29. https://doi.org/10.1016/j.trb.2015.11.007.

Moreno-Gutiérrez, Juan, Fátima Calderay, Nieves Saborido, Maria Boile, Rafael Rodríguez Valero, and Vanesa Durán-Grados. 2015. ‘Methodologies for Estimating Shipping Emissions and Energy Consumption: A Comparative Analysis of Current Methods’. *Energy* 86 (June): 603–16. https://doi.org/10.1016/j.energy.2015.04.083.

Simonsen, Morten, Hans Jakob Walnum, and Stefan Gössling. 2018. ‘Model for Estimation of Fuel Consumption of Cruise Ships’. *Energies* 11 (5): 1059. https://doi.org/10.3390/en11051059.

Yao, Zhishuang, Szu Hui Ng, and Loo Hay Lee. 2012. ‘A Study on Bunker Fuel Management for the Shipping Liner Services’. *Computers & Operations Research* 39 (5): 1160–72. https://doi.org/10.1016/j.cor.2011.07.012.