Data analytics for greener marine operations: towards a fuel optimization decision support system

Piercey, Caitlin

Data analytics for greener marine operations: towards a fuel optimization decision support system

Caitlin Piercey

***Abstract*—We present an exploratory data analysis (EDA) of data obtained from instrumented Canadian Coast Guard (CCG) vessels. The data is analyzed for the purposes of developing a fuel optimization real-time decision support system (DSS) for on-board implementation. The proposed DSS will assist ship navigators in determining the most fuel-efficient operations for their vessel based on controllable parameters of speed, shaft torque, path, etc., resulting in greener operations. We show how the collected data can be integrated with ice data, and how operational modes can be extracted from the data as a new feature. A preliminary predictive model for fuel flow rate using machine learning provides the baseline predictive accuracy of the data for future implementation as a fuel optimisation DSS.**

I. INTRODUCTION

Predicting the fuel consumption of large marine vessels using machine learning methods has been investigated by researchers seeking to reduce the environmental impact and fuel cost of the marine industry. However, a major limitation of such work has been access to complete and extensive data encompassing the performance of working vessels. To date, work has largely focused on fuel predictions for ferries, cargo ships, or other vessels with defined routes. With a novel dataset obtained from the instrumented Canadian Coast Guard vessel, this work will be used to inform the development of a marine decision support system (DSS) with the goal of reducing the fuel consumed by Canada's marine fleet. Processing and exploring the data is a step towards developing the DSS while ensuring that the components of the model are explainable and interpretable.

In this analysis, we consider several research questions in terms of the preparation of this data for future work. We first consider whether seasonality is present in this dataset or if patterns can be accounted for in terms of ship operations. We also consider how to manage redundancy from multicollinearity due to the mathematical relationship between several features. Finally, we consider the impact of engineering new features on the performance of this dataset in a basic predictive model. Future work will expand on these questions as development of the DSS continues.

The remainder of this paper is organized as follows: Section 2 briefly explores relevant previous work to identify some gaps and limitations. Section 3 outlines the methodology for collecting, processing, and expanding the data through an exploratory data analysis (EDA). The purpose of this EDA is to highlight key relationships between features in the dataset, identify anomalies and errors, and integrate various data sources. This section also details the development of a preliminary prediction model for the purposes of comparing prediction improvements with the inclusion of the engineered features. Section 4 presents the results from the prediction comparisons and illuminates some conclusions drawn from the EDA. Finally, section 5 concludes by discussing the results in terms of limitations and future work.

II. RELATED WORK

Daily noon reports, encompassing information regarding daily fuel consumption, load condition, ship position, average speed and distance, and some weather conditions, are a source of data for fuel prediction models [1], [2]. While this form of data collection does not require expensive instrumentation, the resulting data set is small and manually-created, meaning the potential for human error is notable. Models trained on this data may not be suited to a real-time DSS as they do not capture ongoing maneuvers and have lower accuracy. Work has been completed on elevating noon report data using data mining techniques [3]. Bialystocki and Konovessis, for example, used noon report data in a statistical analysis of fuel consumption [4]; while a simple, effective model, it was slightly less accurate than models built from continuous datasets with an R^2 value below 0.9.

The variables captured in datasets also pose a gap. Presently, data relating to ice characteristics have not been included in any of the models discussed in this review. Several models use datasets with no environmental data [4]–[7]; Simonsen, for example, does not use environmental data and thus optimizes speed without considering resistance [7]. Environmental data is shown to improve model accuracy, such as in Yoo and Kim, which models the physical disturbances of the ocean and how they alter results [8]; Hu et al. [9] similarly reported improved accuracy with the inclusion of environmental data. Du et al. explored modelling with different combinations of meteorological data, sensor data, voyage report data, and AIS data; in a series of publications, they describe accuracy improvements when using AIS data versus estimated geographical positions [10], using continuous data versus voyage data [11], and with the inclusion of meteorological data [12]

III. METHODOLOGY

For this work, Canadian Coast Guard ship (CCGS) the Sir Wilfrid Laurier was instrumented and data was obtained between November 2021 and October 2022 at a sampling rate of 1 Hz. The Laurier is a light icebreaker of approximately 3812 GT. The features in this dataset include the fuel flow rate (L/h), fuel consumption (L), and temperature (deg. C) of the three main ship engines as well as the electric generator, auxiliary engine, and incinerator; the torque (kN/m), power (kW), and angular velocity (RPM) of port and starboard shafts; the speed of the vessel (knots); the date and time; and the geocoordinates of the vessel (latitude/longitude). During collection of this data, the vessel completed a variety of standard Coast Guard operations such as icebreaking, buoy tending, transiting, search and rescue (SAR), Arctic operations, light stationkeeping, etc., as well as periods of docking. The data is not labelled with the operation corresponding to the time stamp of each datapoint as this information was not available at the time of collection.

A. Processing and exploration

Data processing followed a standard procedure for time series data, including detecting outliers and anomalies, missing values, and errors. An EDA was then conducted to investigate the features for relationships, correlations, and patterns to better shape the dataset to meet research goals.

External data was then integrated with the obtained data, specifically data relating to ice concentration in the Canadian arctic and ice thickness. The purpose of integrating this data is to understand ship performance in ice, as the developing DSS will be used to guide fuel-efficient operations in the Canadian Arctic. This data was obtained from the Canadian Ice Services (CIS) database and was integrated with the CCGS data based on time stamps and latitude/longitude coordinates. A difference in geospatial and temporal resolution was observed between the two datasets. Ice observations were recorded daily, compared to the 1 Hz records in the CCGS data, and therefore the time in the CCGS data was dropped and data was merged based on day, under the assumption that ice observations will be relatively constant over the course of a single day. As well, consecutive ice observations have approximately 2 nautical miles between them, while the CCGS data follows a vessel route at a frequency of 1 Hz, and thus the spatial resolution, while varying, is much smaller.

To manage this, bins were created surrounding the geocoordinates of the ice observations, covering an area with a radius of 2 nautical miles around the observation. The CCGS data was then sorted into this bins and the corresponding ice concentration and thickness was integrated with the ship data based on day and bin.

The data contained missing timestamps for a period from November 2021- April 2022, presumably corresponding with the off-season for the vessel while it was not operating. The relevant features, such as ship speed and power, were filled with zeroes under the assumption that the ship was sitting in dock. Other timestamps were also missing covering shorter periods; the reason for these missing values was not determined. The remaining missing values covered short periods (¡ 5 minutes) and were filled using interpolation as we assume that the relationship between variables over such a relatively short period can be approximated as linear.

The dataset was explored for multicollinearity as several of the features had strong mathematical relationships. For example, power is the factor of torque and angular velocity; therefore, the impact of capturing all three features in the dataset was examined. To identify multicollinearity, the variance inflation factor (VIF), a measure of how the variance of a feature is impacted through relations with other features, is calculated for the features of the dataset. Feature reduction and scaling is applied to explore how multicollinearity can be reduced.

We next investigate this dataset in terms of trend and seasonality. Standard seasonal periods are not assumed as the operations of the vessel vary depending on the dynamic needs of the vessel operators; therefore, we assume variations within the data depending on vessel goals rather than time of year. It is worth noting that this vessel, as an Arctic-transiting vessel, does not operate year-round and therefore there are periods within the data where the values reflect the docked status of the ship. Thus, it is essential to explore the seasonality of this data to better understand the patterns and context of the data. Seasonal decomposition plots were generated to investigate the data for residuals, seasonality, and trends at a daily, weekly, and monthly period. In addition to this, the Kruskal-Wallis test was used to test for stationarity in the data.

B. Operational mode identification

A novel categorical feature is created for this dataset classifying the different operations of the vessel. The different operational modes were provided by ship's navigators and defined in terms of ship speed and shaft rotational velocity. We further defined these operations in terms of heading.

Operations were found by segmenting the data using an offline change point detection algorithm based on [13]. A rolling window was applied across the data and summary statistics were generated for each window. The statistics of each window were compared; a discrepancy value is calculated. After analyzing the entire dataset, the most significant discrepancy values are considered as change points, splitting the data into segments. Therefore, windows without a significant discrepancy between their statistical measures would be considered as part of the same operational segment.

After the operational segments were identified via the detected change points, they were classified based on several factors. First, a segment with a mean speed, power, and RPM of 0 is considered as the anchored or docked condition. These datapoints are labelled and removed from further classification. Next, the remaining segments are divided into ”transiting” and ”maneuvering” based on the residuals of a linear fit. If the segment can be linearly represented, it is considered steady state and therefore is classified as transiting, i.e. when the ship is sailing at a relatively constant speed, direction, and power. If the segment cannot be linearly represented, it is considered maneuvering, represented by directional changes or increases/decreases in speed and power. In the next stage, we explore how separating these operations affects predictive performance in the model.

C. Preliminary predictions

After cleaning and augmenting the data, we applied a simple regression model to assess how the novel features affect the

feature	full dataset	reduced	scaled
Port Shaft RPM	89.133566	13.421017	10.068123
Port Shaft Torque	10.431685	7.627560	6.171160
Port Shaft Power	51.309682	61.382465	45.286446
Stbd Shaft RPM	119.342256	12.435581	10.055625
Stbd Shaft Torque	38.214547	2.191068	1.402386
Stbd Shaft Power	87.905809	43.248216	32.378150
lain Engine 1 Total	184493.829053	3.136821	1.705059
lain Engine 2 Total	170001.005198	NaN	NaN
lain Engine 3 Total	104971.441079	NaN	NaN
Aux Gen Total	48.870725	NaN	NaN
Aux Gen Temp	2.555168	NaN	NaN
Electric Gen Total	422.075028	NaN	NaN
Incinerator Flow	1.150826	NaN	NaN
Incinerator Total	52317.374367	NaN	NaN
GPS Speed	44.863262	NaN	NaN

Fig. 1. VIF scores for full dataset, reduced dataset with combined features, and scaled reduced dataset.

predictive performance. Due to the multicollinearity present in the data, we selected ridge regression as our test algorithm. Future work will focus on determining the optimally-performing algorithm.

Accuracy was assessed using blocked time series cross validation, wherein the data is split into training and validation folds based on rolling windows. The training set is always chronologically in front of the validation set; this prevents data from the "future" leaking into data from the "past", i.e. it prevents assuming an impact from the future on the past. This is useful to avoid potentially introducing the impacts of data drift into the training set.

We therefore use the data in the following configurations to make fuel flow predictions:

- Cleaned and processed CCGS data.
- Cleaned and processed CCGS data, ice data.
- Cleaned and processed CCGS data, ice data, VIF reduced.
- Cleaned and processed CCGS data, ice data, operational categories.

The purpose of comparing these configurations is to identify the optimal configuration in terms of predictive accuracy, as well as to identify any areas where further work is needed to refine the dataset.

IV. Results

The multicollinearity present in the data (for select features) is shown via VIF scores in Figure 1.

Seasonal decomposition plots were generated for daily, weekly, and monthly periods and for different features, including speed, shaft rotational velocity, power, etc. An example of a daily seasonal decomposition plot for ship speed is presented in Figure 2.

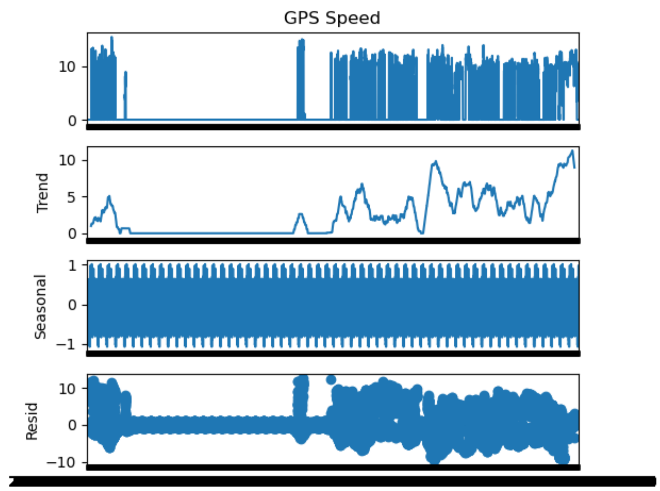

Fig. 2. Seasonal decomposition with daily period for ship speed.

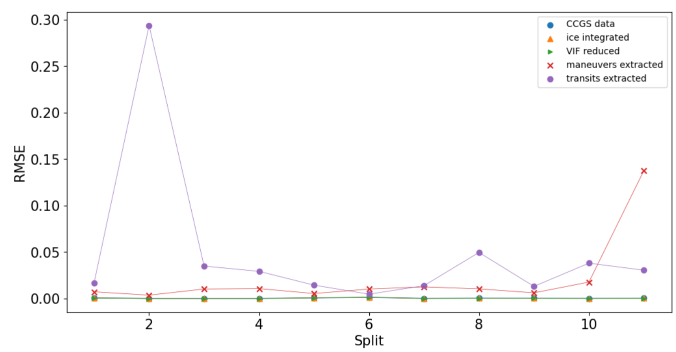

Fig. 3. RMSE scores for 11-fold cross validation.

For the prediction models, the root mean squared error (RMSE) for predictions made on the different datasets is presented in Figure 3.

V. Discussion

Based on the above results, several key conclusions for the future of this project can be determined. As presented in Figure 1, the VIF scores show significant (VIF $\gg$ 10) multicollinearity between features, suggesting that the same information is being captured in multiple features, reducing interpretability of the results i.e. our ability to accurately assess the impact of these variables independently on the target, and indicating increased variance. To test the assumption that angular shaft velocity, torque, and shaft power are such correlated features, power is removed for both the starboard and port shaft as it is represented by the remaining features. This reduced VIF but not below 10, which is still considered severe multicollinearity. High VIF (VIF ¿ ¿ 10) is also observed in the total fuel consumed by the three main engines. This feature is replaced with a summation of the fuel consumed by the three engines and the auxiliary engines, for a total fuel consumed. Further to this, normalizing the data reduces multicollinearity. Removing the individual engine totals and using the new feature of total fuel reduces most of the VIF scores to moderate level (between 5-10); however, a highly multicollinear relationship is still observed between shaft RPM (starboard and port) and ship speed. The relationship between ship speed and shaft RPM is not a simple mathematical expression and both parameters are useful in our DSS as controllable factors for optimisation; therefore, the decision was made to keep these features in the dataset.

This work precedes the development of a fuel optimisation DSS for a working vessel; as the above indicates, the behavior of a working vessel is unique and complex for modelling, as the ship operates under a wide variety of conditions and operational profiles. The lack of predictable seasonality and trend to this data suggests that the resulting model underlying the DSS will need to be adaptive in order to respond to the dynamic data stream. Figure 2 shows the data decomposed into trend, seasonal, and residuals on a daily period. From this plot, we can determine that there is no consistent long term trend in the speed feature. Furthermore, there is significant residual or noise component of the time-series, and the magnitude of the seasonality plot is small relative to the residual and trend plot, suggesting weak seasonality. Potentially a longer-term dataset encompassing years would reveal a yearly or monthly seasonality consistent with the on and off season for the vessel; however, with the current data, we can instead consider the fluctuations in the data to be representative of the dynamic nature of a working vessel, wherein the goals of the vessel are constantly changing and thus so are the vessel parameters of speed, power, etc. As data collection is ongoing, this will be revisited with an extended dataset.

The results of the predictive models provide insight and further questions to answer. Notably, the predictive performance decreases after extracting and separating the different operations into maneuvering and transiting. This may be due to the decreased size of the training set used to train the model for each operation; the transit condition in particular saw a decrease in performance. The transit category was less present in the data and therefore this model was trained on the smallest subset of the data.

As well, without annotated operational data, wherein the segments of the ship's journey have been annotated with their respective operations, we are unable to obtain validation metrics for our segmentation model. As such, the operation segmentation model can be improved in the future with the use of annotated data for validation. Future work will then further refine the operational mode classification by subdividing the "maneuvering" class into specific operations such as icebreaking or buoy tending. These classifications will be made based on descriptions of said operations provided by experts working with the CCG. The integration of this domain knowledge will elevate the data and allow us to better identify how the model performs in different operational conditions.

The results also indicate little difference between the inclusion of ice data; however, this inclusion does not decrease model performance and instead provides important context for the DSS. While this stage of the project focuses on obtaining accurate predictions, future iterations of this project will focus on fuel optimization and recommendations to be supplied by the DSS. A ship operating in ice has specific constraints to consider in order to operate safely and therefore ice as a measurable factor will be present. We can therefore confirm that this data does not decrease performance. Furthermore, obtaining a higher resolution ice dataset may have a stronger impact on the dataset and should be investigated. Finally, as we will further refine our operational mode classifications, data regarding the presence and type of ice will help extract icebreaking maneuvers from the dataset.

There are limitations worth considering to this work. First and foremost, highly accurate wind and wave data was not available for this dataset; future work, however, will focus on an extended dataset which includes these parameters. A second limitation is that there is no annotated data against which to validate our operational mode classifications. This issue will be fixed in a later iteration of this project through instrumenting the vessel with an operational mode selector, so that the navigator will select modes throughout the ship's voyage, created a labelled dataset. Thirdly, the granularity of the ice data is such that it is difficult to visualize and understand the ship's response to ice; the available data indicates an ice concentration in a given region but the actual response of the ship in that region must be inferred by ship behavior. For example, we may make assumptions that a ship operating in an area with moderate ice concentration is actively navigating ice when the ship is moving backwards and then forwards in a breaking motion; without knowing the specific location of ice in relation to the ship, we are unable to definitively determine that this is icebreaking. This issue could be confronted with a higher-resolution dataset of ice images.

With the above assessment, we can proceed with further data preparation on which to train a model as the basis of the DSS; following model development, our work will focus on integrating the model with an interface and evaluating the DSS in terms of both performance and efficacy for human use. The preparation and understanding of this data is the first step in developing a powerful and interpretable DSS.

References

[1] E. Bal Beşikçi, O. Arslan, O. Turan, and A. Ölçer, "An artificial neural network based decision support system for energy efficient ship operations," *Computers Operations Research*, vol. 66, 04 2015.

[2] T. Uyanık, Y. Arslanoğlu, and Kalenderli, "Ship fuel consumption prediction with machine learning," 04 2019.

[3] A. Safaei, H. Ghassemi, and M. Ghiasi, "Correcting and enriching vessel's noon report data using statistical and data mining methods," *European Transport - Trasporti Europei*, 03 2018.

[4] N. Bialystocki and D. Konovessis, "On the estimation of ship's fuel consumption and speed curve: A statistical approach," *Journal of Ocean Engineering and Science*, vol. 1, 03 2016.

[5] T. Uyanık, Y. Yalman, Kalenderli, Y. Arslanoğlu, Y. Terriche, C.-L. Su, and J. Guerrero, "Data-driven approach for estimating power and fuel consumption of ship: A case of container vessel," *Mathematics*, vol. 10, p. 4167, 11 2022.

[6] H. Lee, N. Aydin, Y. Choi, S. Lekhavat, and Z. Irani, "A decision support system for vessel speed decision in maritime logistics using weather archive big data," *Computers Operations Research*, vol. 98, 06 2017.

[7] M. Simonsen, H. J. Walnum, and S. Gössling, "Model for estimation of fuel consumption of cruise ships," *Energies*, vol. 11, 04 2018.

[8] B. Yoo and J. Kim, "Path optimization for marine vehicles in ocean currents using reinforcement learning," *Journal of Marine Science and Technology*, vol. 21, 12 2015.

[9] Z. Hu, Y. Jin, Q. Hu, S. Sen, T. Zhou, and M. T. Osman, "Prediction of fuel consumption for enroute ship based on machine learning," *IEEE Access*, vol. 7, pp. 119497–119505, 2019.

[10] X. Li, Y. Du, Y. Chen, S. Nguyen, W. Zhang, A. Schönborn, and Z. Sun, "Data fusion and machine learning for ship fuel efficiency modeling: Part i – voyage report data and meteorological data," *Communications in Transportation Research*, vol. 2, p. 100074, 2022.

[11] Y. Du, Y. Chen, X. Li, A. Schönborn, and Z. Sun, "Data fusion and machine learning for ship fuel efficiency modeling: Part iii – sensor data and meteorological data," *Communications in Transportation Research*, vol. 2, p. 100072, 2022.

[12] Y. Du, Y. Chen, X. Li, A. Schönborn, and Z. Sun, "Data fusion and machine learning for ship fuel efficiency modeling: Part ii – voyage report data, ais data and meteorological data," *Communications in Transportation Research*, vol. 2, p. 100073, 2022.

[13] C. Truong, L. Oudre, and N. Vayatis, "Selective review of offline change point detection methods," *Signal Processing*, vol. 167, p. 107299, 2020.