The Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-24)

Sequential Modeling of Complex Marine Navigation: Case Study on a Passenger Vessel (Student Abstract)

Yimeng Fan[1*], Pedram Agand[1*], Mo Chen[1], Edward J. Park [1], Allison Kennedy[2], Chanwoo Bae[3]

[1]Simon Fraser University, BC, Canada.
[2]National Research Council Canada, Ontario, Canada.
[3]BC Ferry Inc., BC, Canada.
{yfa56,pagand, ed_park}@sfu.ca, mochen@cs.sfu.ca, allison.kennedy@nrc-cnrc.gc.ca, chanwoo.bae@bcferries.com

Abstract

The maritime industry's continuous commitment to sustainability has led to a dedicated exploration of methods to reduce vessel fuel consumption. This paper undertakes this challenge through a machine learning approach, leveraging a real-world dataset spanning two years of a ferry in west coast Canada. Our focus centers on the creation of a time series forecasting model given the dynamic and static states, actions, and disturbances. This model is designed to predict dynamic states based on the actions provided, subsequently serving as an evaluative tool to assess the proficiency of the ferry's operation under the captain's guidance. Additionally, it lays the foundation for future optimization algorithms, providing valuable feedback on decision-making processes. To facilitate future studies, our code is available at https://github.com/pagand/model_optimze_vessel/tree/AAAI.

Introduction

Numerous researchers explore models for vessel Fuel Consumption (FC) prediction using log-based and sensor-based data. Log-based methods can be associate with human errors and suffer from lower sample frequency. Sensor-based models, follow a process, including data normalization and feature engineering. Nevertheless, a gap exists in the literature as most models lack real-time ship operator input and environmental considerations. Further investigation is needed for feature selection, multicollinearity, non-stationary data handling, and better integration of domain knowledge with data-driven approaches. Agand et al. (2023a) propose a comprehensive approach integrating domain knowledge and data-driven techniques, incorporating physical insights, correlation matrices, and PCA. However, it does not account for data temporality, which may impact performance.

In this work, a reality-based time series forecasting model with the primary objective of auto-regressively predicting states of a ferry is proposed. For this analysis, an operational dataset measured on the ferry every minute over a two year period was used. These records capture the sensor data from a ferry operating in the west coast of Canada for more than 3000 transits along a constant route. This predictive model's pivotal role lies in its ability to generate a precise representation of the ship's states such as location -latitude (LAT) and longitude (LON)-, fuel consumption (FC), and speed over ground (SOG) in future steps. According to Fig. 1, we utilized the previous predicted values in an auto-regressive manner, with information about external disturbance (e.g. wind, current, weather, etc.), and static features (such as direction of movement, docking area, elapsed time of trip, etc). Finally, we open-sourced a reinforcement learning (RL) compatible dataset to D4RL framework in addition to a Gym environment (Fu et al. 2020).

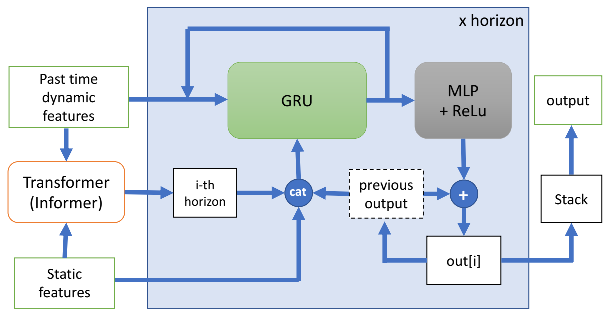

Figure 1: Proposed sequential model architecture

Preprocessing

In this stage, outlier management was addressed using the 1.5 IQR method, treating them as missing data points and employing various imputation techniques for isolated missing values within trips. After devising domain knowledge on draft, cargo, and waves, we developed a clustering method to assign operating modes (mode 1 for autopilot travel and mode 2 for docking regions) based on factors like speed, acceleration, and shaft speed. Wind factors were incorporated through squared relative wind speed and categorized wind direction. Feature selection involved computing the Pearson correlation coefficient and leveraging domain knowledge for decision-making. Feature engineering introduced acceleration and displacement, and we normalized data values to a 0-1 range. Additionally, power transformations were applied to features with skewed distributions, aiming to align them more closely with Gaussian distributions.

[*]These authors contributed equally.

	Metric	NAR	NAR +GRU	AR	AR +GRU
FC	RMSE	0.1876	0.0769	0.0777	**0.0608**
	Std.	0.0856	0.04049	0.0605	**0.0365**
	R^2	-0.023	0.296	0.298	**0.701**
SOG	RMSE	0.1412	0.0322	0.0435	**0.0304**
	Std.	0.0854	**0.0178**	0.0451	0.0179
	R^2	0.091	0.602	0.404	**0.793**
LAT.	RMSE	0.1023	0.0684	0.0613	**0.0385**
	Std.	0.0712	0.0167	**0.0133**	0.0137
	R^2	0.281	0.6193	0.685	**0.874**
LON.	RMSE	0.1370	0.1097	0.0960	**0.0604**
	Std.	0.0842	0.0267	0.0225	**0.0157**
	R^2	0.404	0.671	0.746	**0.899**

Table 1: Prediction results for non-auto-regressive (NAR) and auto-regressive (AR) w/o GRU refinement.

Modeling

As shown in Fig. 1, the framework consists of two components: a pretrained transformer model called informer (Zhou et al. 2021), responsible for executing time series forecasting and feature fusion. Inspired from (Agand et al. 2023b), it is followed by a Gated Recurrent Unit (GRU) module employed to predict the residual, thereby mitigating auto-regressive cumulative errors within the predicted outcomes. The determination of the sequence length necessitates two crucial considerations. Tt must be of sufficient to facilitate accurate forecasting while it should avoid excessive elongation, as it signifies the time interval to wait in each trip before making new predictions. Following trade-offs, a sequence length of 25 minutes, coupled with a prediction horizon of 5 minutes, was deemed the optimal configuration.

Environment

We have open sourced a RL compatible dataset for offline RL setting with a gym environment that can be served as reality-based simulator. We constructed an offline dataset tailored for conventional RL frameworks, comprising current/next observations, actions, rewards, and termination indicators. The actions encompass heading, shaft speed, and mode selections. For observations, we employ both current and next states, encompassing static factors such as trip start hour, direction, dynamic factors including heading rate, resistance (torque/thrust) (Carlton 2019), displacement, and previous outputs, as well as disturbance-related variables like time, weekday, current, season, weather, wind direction, wind force, and water depth. In terms of rewards, we have defined three distinct intermediate rewards. The first penalizes deviations from the top 1% trips of the dataset that uses the least fuels, the second penalizes fuel consumption (FC), and the last comprises sparse rewards, with +1 awarded for on-time arrival at the docking area and a penalty of -0.1 for each minute beyond the schedule. The model outputs four variables: LAT, LON, SOG, and FC.

In the Gym environment, we implemented curriculum learning by defining three reward stages. The first stage focuses on emulating the decisions made by captains in the dataset. In the second stage, rewards are structured to encourage behavior similar to the top 1% of best-performing trips. The third stage emphasizes the minimization of FC while adhering to the designated time schedule. The reinforcement learning process concludes with the issuance of the "done" signal, triggered when the ferry successfully reaches its intended destination. Conversely, if the ferry fails to reach its destination within 25% more time steps than usual, the process times out, resulting in a negative reward.

Conclusion and Future Work

As demonstrated in Table 1, we observe that the root mean square error (RMSE) and coefficient of determination (R^2) exhibit the best performance for all four quantities when employing the AR + GRU approach. The second-best results are associated with the standard deviation (Std.) of SOG and LAT, while FC and LON still exhibit the best performance across all approaches. In general, the AR approach outperforms the others, primarily due to allowing the transformer to learn and mitigate cumulative errors during training.

We outline a step-wise approach to develop a time-series model for a ferry using real data. Additionally, we introduce an offline dataset and a simulator that both can serve as training tools for captains or an environment for machine learning systems. The forecasting model will be instrumental in assessing the optimization model's effectiveness. The Gym environment undergoes three stages of training. Initially, it aims to replicate captain driving behaviors, then shifts to imitating top trips with the lowest FC, and finally strives to surpass captain performance by pursuing to minimize FC while adhering to the time schedule and other limitation. This multi-stage training approach is expected to accelerate the convergence of RL systems. For future direction, we consider leveraging RL to optimize the navigational best practice to maximize fuel efficiency.

References

Agand, P.; Kennedy, A.; Harris, T.; Bae, C.; Chen, M.; and Park, E. J. 2023a. Fuel consumption prediction for a passenger ferry using machine learning and in-service data: A comparative study. *Ocean Engineering*, 284: 115271.

Agand, P.; Mahdavian, M.; Savva, M.; and Chen, M. 2023b. LeTFuser: Light-weight End-to-end Transformer-Based Sensor Fusion for Autonomous Driving with Multi-Task Learning. *arXiv preprint arXiv:2310.13135*.

Carlton, J., ed. 2019. *Marine Propellers and Propulsion (Fourth Edition)*. Butterworth-Heinemann, fourth edition edition. ISBN 978-0-08-100366-4.

Fu, J.; Kumar, A.; Nachum, O.; Tucker, G.; and Levine, S. 2020. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*.

Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; and Zhang, W. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11106–11115.