

Fuel-Saving Route Planning with Data-Driven and Learning-Based Approaches – A Systematic Solution for Harbor Tugs

Shengming Wang¹, Xiaocai Zhang^{1*}, Jing Li^{1,2*}, Xiaoyang Wei¹, Hoong Chuin Lau^{1,3},
Bing Tian Dai³, Binbin Huang⁴, Zhe Xiao¹, Xiuju Fu¹, Zheng Qin¹

¹Institute of High Performance Computing, Agency for Science, Technology and Research, Singapore

²Centre for Frontier AI Research, Agency for Science, Technology and Research, Singapore

³School of Computing and Information Systems, Singapore Management University

⁴School of Computer Science and Technology, Hangzhou Dianzi University

{wang_shengming, zhang_xiaocai, wei_xiaoyang, xiaoz, fuxj, qinz}@ihpc.a-star.edu.sg,
li_jing@cfar.a-star.edu.sg, {hclau, btdai}@smu.edu.sg, huangbinbin@hdu.edu.cn

Abstract

In recent years, there are trends toward cleaner port environments through enforcement by imposed legislation. Transit optimisation of fuel-based port service boats like harbour tugs has emerged as a critical task to reduce fuel consumption and carbon emission. In this paper, an innovative learning-based method, comprising a Reinforcement Learning (RL) model together with a fuel consumption prediction model, was proposed to formulate fuel-saving transit routes. Firstly, an ensemble model is established by combining a Long Short-Term Memory (LSTM) model with a Multilayer Perceptron (MLP) model, predicting fuel use based on tugboat movement and environment factors. Subsequently, an innovative RL based on Deep Deterministic Policy Gradient (DDPG) framework is developed considering the characteristics and obstructions of waterway in Singapore as well as the environmental factors to learn the optimal transit strategy that minimizes fuel consumption. We also demonstrate the efficacy of the solution to generate routes from origin to destination terminals, exhibiting significantly reduced fuel consumption in comparison to real-world transit scenarios.

1 Introduction

Singapore has been one of the busiest ports in the world [Chew *et al.*, 2023]. One-fifth of the world's gross shipping tonnage and one-half of the world's crude oil pass through Singapore's ports. There are approximately 1,000 ships in port at any time. One ship arrives or leaves the ports every two minutes, via the Singapore Strait.

Due to the heavy shipping volume in Singapore's ports, harbour tugs which provide ship handling services play an indispensable and crucial role, including assistance in berthing/unberthing of ships and maneuvering of ships within the crowded confines of the port. In recent years, there have

been trends toward cleaner port environments through enforcement by imposed legislation [Caliskan, 2022]. Figure 1 presents the roadmap of the International Maritime Organization (IMO) on greenhouse gas (GHG) emission reduction strategies. It sets the timelines with goals and promotes green technologies for both ocean-going vessels and harbor crafts. These new legislations would translate to economic pressure on vessel operators, requiring either retrofitting the fleet ready for replacement of cleaner fuels or optimising the operation or both. For example, vessels which are required to meet emission criteria have to be modified. Modifications include installation of scrubbers, selective catalytic reduction systems, humid air motors, or engines which use cleaner fuels. In addition, vessel operators may adopt operational profiles which improve fuel efficiencies. Examples of fuel-efficient operational profiles include imposition of slow steaming or upper limits on vessel speeds, or efficient navigational methods which utilize tide or current conditions to reduce engine demand. Of the energy-saving/emission-reducing methods described above, at the current phase, the use of efficient navigational methods is perhaps the most economical to implement, as these methods do not entail any costly modifications to the vessels. The use of efficient navigational methods is perhaps also most suited to small harbour crafts such as harbour tugs. Harbour tugs operate in home waters where the sea/environmental conditions are well-known and data readily available. By mapping and obtaining optimized navigational profiles, the fuel efficiency of harbour tugs can be improved and operating costs reduced. Higher fuel efficiencies would result in lower emissions, therefore contributing to decarbonization and meeting legislative requirements. In the interim, the approaches implemented can be also applied to greener fuels to reduce the operation cost in the future.

In this work, we conducted a systematic study from the machinery data and tugboat movement traffic data analytics, quantified the determinant factors, and sitting above, built the learning-based models including an ensemble model for fuel consumption estimation/prediction and reinforcement learning (RL) model to generate optimised fuel saving route between origin and destination. In comparison with real-world transit instances in quantitative evaluation, it has been proved

*Corresponding author

that the routes suggested by the proposal solutions allow significant fuel consumption reduction, which in practical exercise will bring in economic benefit for tug operators and contribute towards sustainable green port development.

The **contributions** of this work are highlighted as follows:

- (1) We address an industrial problem statement on sustainability in the maritime sector. Sensor data like fuel flow meter readings, resolution per minute and other engine states are collected for this study with a comprehensive distribution and quality analytics that directly contribute to the downstream learning-based approaches. To the best of our knowledge, this is the first work to formulate such a problem and model in a systematic way with real-world engine sensing data for model training, evaluation and comparison studies.
- (2) By diving into the fuel consumption profiling and distribution, we propose an ensemble model that combines LSTM and MLP which ensures the prediction performance over the full spectrum of navigational states. This work also considers comprehensive categories of fuel consumption relevant features including both the tug navigational states and environmental factors. The feature selection is quantitatively guided by the correlation studies.
- (3) We propose a real-time tug route planning algorithm based on RL. Our solution approach provides a novel definition of waterway distance and incorporates the ensemble-based fuel consumption estimation to define rewards for fuel-saving actions, which is proven to be effective experimentally.
- (4) Our solution approach has been evaluated against historical transit instances of harbour tugs. The efficacy of the solution is validated via case studies covering both regular and irregular cases. The results reveal the promising applicability of the solutions in practical operations to substantially reduce fleet fuel consumption.

The remaining of this paper is organized as follows: Sections 2 and 3 provide a solution overview and data description and analytic insights; The ensemble model for fuel consumption estimation is explained in Section 4, followed by the RL modeling elaborated in Section 5. The performance evaluation is presented in Section 6 along with case studies. We conclude in Section 7 by summarizing the proposed solution and research findings.

2 Overview of Our Approach

In this paper, we introduce an integrated approach, as illustrated in Figure 2, that combines prediction and optimization models to enhance the efficiency of tugboat transit operations. The combined model comprises a stacked Long Short-Term Memory (LSTM) and Multilayer Perceptron (MLP) model, proficient in predicting transit fuel consumption. Concurrently, a Deep Deterministic Policy Gradient (DDPG) model is integrated to design optimal routes for tugboats while optimizing fuel consumption. Significantly, the model is built based on the Singapore waterways, incorporating environmental factors such as water depth, wind and current. In sum-

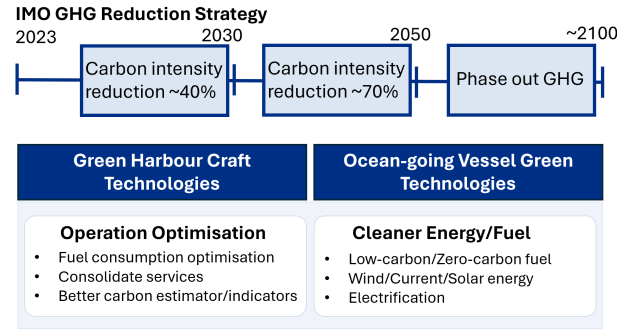


Figure 1: IMO’s GHG reduction pathway and primary strategies for emission reduction.

mary, the solution aims to address the challenges of tugboat transit, offering a versatile approach that incorporates predictive accuracy and route optimization based on real-world conditions.




Figure 2: Structure of the combined model. It consists of a stacked LSTM and MLP model and a DDPG model.

3 Data Preparation

The dataset utilized in this study is comprehensive, encompassing Automatic Identification System (AIS) data, machinery data, and environmental data. AIS data records details of tugboat transit, including navigational states like speed and course attached with timestamp [Xiao *et al.*, 2020; Zhang *et al.*, 2022]. Analysis of the frequency distribution of tugboat speed reveals a range predominantly within 0-16 knots as in Figure 3. Recognizing various operational states such as waiting, transit, and berthing/unberthing, we employ filtering criteria based on speed, distance to the harbor, and proximity to ships to extract the transit state.

From the transit state data, tugboat transit trajectories are extracted and stored in JSON format for each trajectory, providing a structured representation of their movements. The machinery data provides detailed cumulative information on the fuel consumption of both port and starboard side engines of tugboats. Integration of the dataset with trajectory data, based on timestamp alignment with a time interval of approximately 1 minute, facilitates a combined dataset for further analysis. To visualize the relationship between fuel consumption rate and speed, a plot is generated from the combined dataset as shown to Figure 4. The visualization provides an initial insight: when the speed is below 8 knots/s, the mean

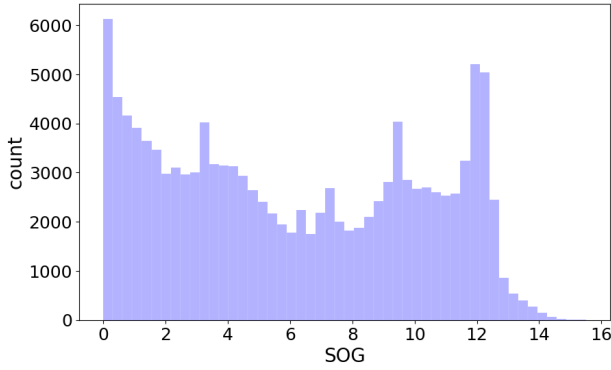


Figure 3: Frequency distribution of speed(knot)

values exhibit a relative consistency, but a noticeable increase is observed after surpassing 8 knots/s. This observation is noteworthy when considering the choice of speed will largely affect the overall fuel consumption. The environmental data,




Figure 4: Distribution of transit fuel consumption rate(Litre/s) versus speed(knot)

including water depth, wind, and current information presented in a grid-based format, is merged with trajectory details based on coordinates. Additionally, to accurately replicate real-world conditions specific to Singapore’s waterways, detailed information about the Singapore waterway pattern is included in the dataset. This comprehensive dataset, enriched with diverse information, forms the basis for our subsequent modeling, enabling an exploration of factors influencing tug-boat transit in the Singapore waterway.

4 Fuel Consumption Estimation Model

Given the collected transit data that contains a number of length-variable trajectories and each of them comprises evenly sampled points (one point per minute) everywhere, we advocate for the development of a time-dependent model functioning as a regressor. It is worth noting that time dependence should be taken into consideration as the fuel consumption rate is influenced by historical features. For example, under the same speed at timestamp t , an increase in speed

typically leads to higher fuel consumption, whereas maintaining a constant speed does not.

Based on this insight, we develop an LSTM model which traverses the trajectories by first encoding the input variables (mainly including vessel operator parameters and environment information) within a history time window w and then projecting the latent representation to the current fuel consumption rate. Specifically, the LSTM model aims to maximize the following likelihood

$$\mathcal{L}_{\text{LSTM}} = \prod_{j=1}^J \prod_{t=w}^{|Tr_j|} p(y_t | x_t, \dots, x_{t-w}), \quad (1)$$

where J is the number of trajectories and $|Tr_j|$ denotes the length of j -th trajectory. The applied LSTM is based on the cell module proposed by [Alahi *et al.*, 2016], which is also used for recent ocean engineering [Yuan *et al.*, 2021]. Denoting gate control signal, input gate, forget gate and output gate by z, z^i, z^f, z^o respectively, which are transformed from last state h_{t-1} and current x_t , we have

$$\begin{aligned} c_t &= z^f \odot c_{t-1} + z^i \odot z \\ h_t &= z^o \odot \tanh(c_t) \\ \hat{y}_t &= \sigma(W h_t), \end{aligned} \quad (2)$$

where the predicted \hat{y}_t is derived by mapping h_t by a linear layer W with Sigmoid as an activation. By assuming that regressor output as a mean of a noisy prediction, characterized by Gaussian distribution, we use Mean Square Error (MSE) loss to train the model

$$\text{MSE}(y_t, \hat{y}_t) = (y_t - \hat{y}_t)^2. \quad (3)$$

However, from Figure 5(a) we can observe that the trained LSTM is not sufficient for real usage; it performs well on the group of data whose true fuel consumption is lower while incurring relatively larger errors on higher true fuel consumption data, known as a lack of subgroup robustness [Liu *et al.*, 2021]. We find this phenomenon can be attributed to the data imbalance, shown in Figure 5(b). It is obvious that error is strongly negatively correlated with data density. Recent research advocates increasing the importance of scarce data. We then conduct oversampling on the minority of training data so that the training data exhibits an approximately uniform distribution, which is proved more stable than its counterpart reweighting strategy by the recent advance [An *et al.*, 2021]. Disappointingly, according to our experiments (histograms w/ oversampling in Figure 5(a)), LSTM cannot benefit much from this workaround, probably because the unreliable time-dependence (e.g., noisy) on limited data is unintentionally amplified by oversampling.

We address this issue by additionally employing an MLP model (also with loss of Eq. (3)) to capture the instant relation between input and output at each timestamp t , which virtually takes advantage of oversampling strategy [Steininger *et al.*, 2021] but not as precise as LSTM on low fuel consumption data. Hence, we combine them by trusting MLP on larger predictions while trusting LSTM on smaller predictions. In practice, we do fuel estimation in the following scheme

$$\hat{y}_t = \begin{cases} \text{MLP}(x_t), & \text{if } \text{MLP}(x_t) > 9 \\ \text{LSTM}(x_{(t-w):t}), & \text{otherwise.} \end{cases} \quad (4)$$

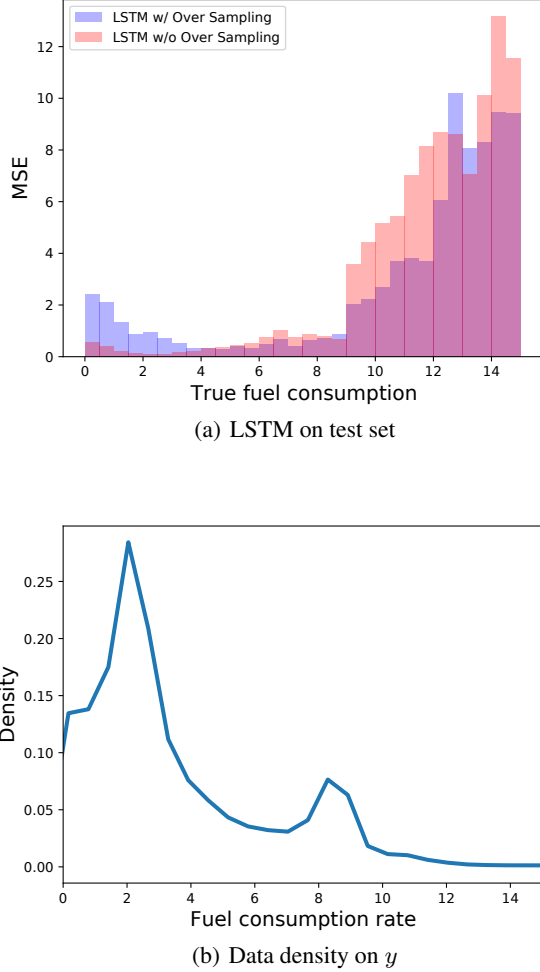


Figure 5: LSTM fails to do precise prediction on scarce data, which cannot be remedied even with the data augmentation technique.

Note that Eq. (4) used for efficacy testing or down-streaming reinforcement learning does not touch true y_t . Such an ensemble is demonstrated as simple and efficient (Refer to Section 6.1).

5 RL-Based Fuel-Saving Route Planning

A state-of-the-art RL framework, termed Deep Deterministic Policy Gradient (DDPG) [Lillicrap *et al.*, 2016] was proposed to simulate the movement and fuel consumption of tugboat in the Singapore port. The following subsections describe the basic principle of DDPG and several important definitions of tugboat optimization based on DDPG.

5.1 DDPG

DDPG is one of the advanced algorithms in deep reinforcement learning, which combines the advantages of deep learning and deterministic policy gradients algorithm with the Actor-Critic structure. DDPG contains an actor network and

a critic network. The actor network (μ) is a policy network that takes the state s_i as input and outputs action a_i , as formulated by

$$a_i = \mu(s_i | \theta^\mu), \quad (5)$$

where θ^μ is the weights of μ . The critic network is a Q-value network that takes both the state s_i and action a_i as input and outputs the Q-value (q_i), as shown by

$$q_i = Q(s_i, a_i | \theta^Q), \quad (6)$$

where θ^Q is the weights of Q . The loss function of the critic network is

$$\mathcal{L}_Q = \frac{1}{N} \sum_i (y_i - Q(s_i, a_i | \theta^Q))^2, \quad (7)$$

and

$$y_i \approx r_i + \gamma Q'(s_{i+1}, \mu'(s_{i+1} | \theta^{\mu'}) | \theta^{Q'}), \quad (8)$$

where y is the target Q-value and μ' and Q' are the target actor and critic networks, which are initialized with the same weights of actor and critic networks, respectively. The objectives of training the critic network is to minimize \mathcal{L}_Q , as indicated by

$$\theta^{Q*} = \arg \min_{\theta^Q} \mathcal{L}_Q, \quad (9)$$

where θ^{Q*} is the optimal weights. The actor network produces the action that can obtain the highest Q-value. Therefore, the objective of μ is to maximize the following function, as defined as

$$\mathcal{L}_\mu = \frac{1}{N} \sum_i Q(s_i, a_i | \theta^Q). \quad (10)$$

The gradient ascent algorithm is adopted by calculating the gradient of \mathcal{L}_μ to the weights θ^μ , as formulated by

$$\nabla \mathcal{L}_\mu(\theta^\mu) = \frac{1}{N} \sum_i \nabla Q(s, a | \theta^Q) (a) \nabla \mu(s | \theta^\mu) (\theta^\mu). \quad (11)$$

The updates for the weights of the target networks are based on the following rules, as shown by

$$\theta^{Q'} = \tau \theta^Q + (1 - \tau) \theta^{\theta^{Q'}}, \quad (12)$$

and

$$\theta^{\mu'} = \tau \theta^\mu + (1 - \tau) \theta^{\theta^{\mu'}}, \quad (13)$$

where τ is a preset soft replacement, and $0 < \tau < 1$.

5.2 Environment and State

The environment is simulated based on real-world navigational conditions in Singapore. It consists of the boundaries of the Singapore waterway as well as the environmental information including water depth, current and wind. Considering the definition of the state, it should take into account that the current state is sufficient to provide all the important information for the agent to predict the next states [Moradi *et al.*, 2022]. For this work, the state comprises the following information:

- Current latitude (φ) of the tugboat;

- Current longitude (λ) of the tugboat;
- Course over ground (COG) at the last state;
- Wind direction and force;
- Current direction and force;
- Water depth.

5.3 Agent and Action

Since DDPG is an algorithm designed specifically for environments with continuous action spaces, we represent the actions by continuous values. In this study, we will define the continuous action space of the agent from two perspectives, i.e., COG (c) and speed over ground (SOG) (v) of the tugboat.

For the SOG action, the selection of speed (v) for the tugboat should satisfy the following constraint referring to the speed frequency distribution, as shown by

$$3 \leq v \leq 16, \quad (14)$$

where the unit of speed is in knot. Another action to generate is the COG, which is in the range of $[0, 360)$. Apart from the magnitude of SOG and COG, another constraint that is very critical in real cases is the tugboat can only travel within the navigable waterway boundary. If the next state generated based on the action is out of the navigable waterway boundary, such as encounter land, pier, restricted area, etc., a value of 2 degrees will be added or deduced to make an iterative adjustment on COG until the agent makes a valid action either by turning clockwise or counter-clockwise.

5.4 Reward

The reward consists of 3 parts. The first part is a waterway distance reward (D). It should be noted that waterway distance is not a straight-line distance, instead it is the common route distance over the extracted regular waterway pattern for harbour tugs. The rationale behind using waterway distance is that straight-line distance will confuse the model in the case of the blockage of the island/land area ahead between the current position of harbour tug to the destination terminal. To implement the waterway distance calculation between any two positions within the port water, an Application programming interface (API) is developed to generate the passage plan over the waterway pattern using an enhanced Theta* algorithm [Daniel *et al.*, 2010]. The API accepts two position coordinates (latitude and longitude) and returns the distance between the two coordinates. Then, the waterway reward is calculated based on the ratio of the waterway distance of the next state to the destination and the waterway distance of the origin to the destination, a negative reward will be given in order to punish the move leading to far away from the destination and to encourage the move closer the destination. The second part is the turning reward (T), as the tugboat is not encouraged to make big turnings. If the turning is within 60 to 90 degrees, half of the magnitude of the turning degree will be given as punishment, and if the turning is more than 90 degrees, there will be a negative reward for the turning degree. The third part is the fuel consumed (F), which is a negative reward based on the prediction from the fuel consumption estimation model. A reward of 3000 is given when

the agent reaches the destination successfully [Muñoz *et al.*, 2019]. The reward function is formulated as

$$r(s_i, a_i) = \begin{cases} +3000 & \text{if succeed} \\ -k_1 \times F_{i,i+1} - k_2 \times D_{i+1} - k_3 \times T_{i,i+1} & \text{otherwise} \end{cases} \quad (15)$$

5.5 Actor-Critic Networks Architecture

Both the actor and critic networks are based on a Multi-Layer Perceptron (MLP), which consists of 2 hidden layers with 512 nodes at each layer. The hidden layers employ ReLU as their activation functions. The output layer of the actor network uses a Tanh activation function to produce actions within -1 to 1 and the critic network's output layer estimates the Q-function and does not have an activation function.

The proposed RL algorithm for fuel-saving route planning for harbor tugs has been illustrated in Algorithm 1. Steps 1 and 2 initialize the weights for actor and critic networks as well as the target actor and target critic networks. Step 3 initializes the replay buffer. Steps 4 to 23 involve training the entire network with each episode in the context of RL. Step 5 initialize the state. Steps 6 to 22 obtain the data for each time step within the episode for training. Steps 7 and 8 get the action and next state. If the obtained next state is within the navigable waterway boundary, reward and network weights will be updated accordingly (from steps 10 to 17). Otherwise, an updated COG will be assigned for the action a_t to repeat the process from step 8 until the derived next state meets the requirement.

Algorithm 1 DDPG for fuel-saving route planning

```

1: Initialize  $\theta^Q$  and  $\theta^\mu$  for critic and actor networks;
2: Initialize  $\theta^{Q'}$  and  $\theta^{\mu'}$  for target networks with  $\theta^{Q'} \leftarrow \theta^Q$ 
   and  $\theta^{\mu'} \leftarrow \theta^\mu$ ;
3: Initialize replay buffer  $R$ ;
4: for episode = 1,  $\dots$ ,  $M$  do
5:   Receive initial state  $s_1$ ;
6:   for  $t = 1, \dots, T$  do
7:      $a_t = \mu(s_t | \theta^\mu) + N_t$ ;
8:     Execute action  $a_t$  and obtain next state  $s_{t+1}$ ;
9:     if  $s_{t+1}$  is in navigable waterway boundary then
10:      Calculate waterway distance;
11:      Get reward  $r_t$ ;
12:      Store transition  $(s_t, a_t, r_t, s_{t+1})$  in  $R$ ;
13:      Sample a minibatch  $(s_i, a_i, r_i, s_{i+1})$  from  $R$ ;
14:      Update  $\theta^Q$  by minimizing  $\mathcal{L}_Q$ ;
15:      Update  $\theta^\mu$  by maximizing  $\mathcal{L}_\mu$ ;
16:      Update  $\theta^{Q'}$  based on Eq. (12);
17:      Update  $\theta^{\mu'}$  based on Eq. (13);
18:     else
19:       Increase/reduce COG in  $a_t$  to get new next
state  $s_{t+1}$ ;
20:       Repeat the process from step 8;
21:     end if
22:   end for
23: end for

```

6 Results & Discussion

The experiment for tugboat route planning is based on trajectory. We randomly select some historical transit trajectories

and train the model based on the real-time environment. A new route is suggested by the model and the total fuel consumption of the trajectories are compared.

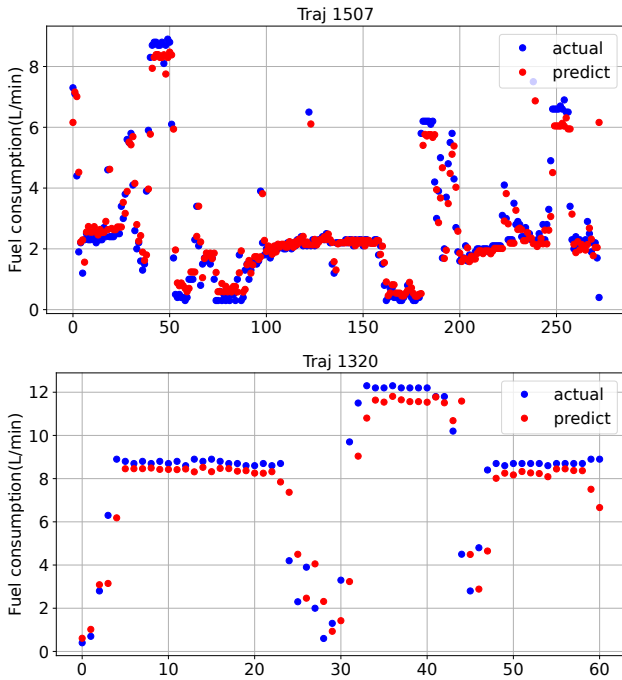


Figure 6: Fuel estimation examples on selected trajectories.

6.1 Estimation Performance

We test the proposed fuel consumption estimation model on a prepared test set to validate its efficacy. Beside MSE, MAE (mean absolute error), R^2 (R Square), and Ratio (Predicted/true value) are used as evaluation metrics. By changing the employed threshold (default as 9) to zero or infinity, we can simply obtain the MLP and LSTM model, whose overall prediction performance is summarized in Table 1. We can see that by leveraging the MLP, the average fuel consumption error can be eventually reduced by 0.34L/min, demonstrating the effectiveness of the proposed ensemble model.

Metric	MSE	MAE	R^2	Ratio (%)
LSTM	3.76	1.23	0.58	106.90
MLP	5.77	1.83	0.51	130.89
Ensemble	2.50	0.89	0.71	102.89

Table 1: Fuel consumption performance on the test set.

Figure 6 also exhibits the predicted results on some randomly selected trajectories. It is evident that given both vessel operations and real-time environmental information, the proposed model can be trained to have the capability to predict the consumed fuel for each time interval while also effectively capturing the underlying trends.



Figure 7: Instance Case 1: Historical human-operated route (top) versus our solution (bottom)

6.2 Instance Case 1

In the figures depicted above, the blue dot denotes the origin, while the orange dot signifies the destination. The historical trajectory reveals a pattern where the tugboat travels a considerable distance away from the Pasir Panjang Terminal before returning, resulting in a prolonged route associated with elevated fuel consumption. In contrast, the optimal route recommended by the model advocates for the tugboat to sail close to the terminal. This strategic adjustment leads to a remarkable reduction, approximately three-fourths, in the overall fuel consumption. The visual representation emphasizes the tangible benefits of the model’s trajectory suggestions, showcasing its potential to significantly enhance fuel efficiency in tugboat transit operations.




Figure 8: Instance Case 2: Historical human-operated route (top) versus our solution (bottom)

6.3 Instance Case 2

In this particular scenario, the tugboat initiates its journey from the terminal and proceeds towards the island. However, in the historical trajectory, an observable deviation occurs as the tugboat follows a curved path on its way to the terminal. Conversely, the model recommends a more direct and efficient route—a straight line guiding the tugboat from the western island to the terminal and then along the terminal to its destination, similar to the actual case. This optimized trajectory, characterized by reduced turns and a more direct path, contributes to a noteworthy reduction in fuel consumption owing to the shorter overall distance covered. The model’s ability to suggest more streamlined routes shows its potential to enhance fuel efficiency and operational effectiveness in tugboat transit scenarios.

6.4 Instance Case 3

In this case, a comparison between the historical transit trajectory and the route suggested by the model reveals distinct path patterns from the origin to the destination. The model, however, proposes a more optimized way of transit, demonstrating the utilization of both environmental factors and speed. Notably, the historical trajectory, with a time interval of 1 minute between each data point, exhibits dense clusters of points, indicative of the tugboat moving at a relatively low speed, while based on the plotting of FO rate versus speed, low speed may not lead to low total fuel consumption, which meets the transit suggestion from the model. Furthermore, this observation indicates the model’s ability to recommend routes that make efficient use of vessel speed and leverage environmental conditions to minimize fuel consumption.

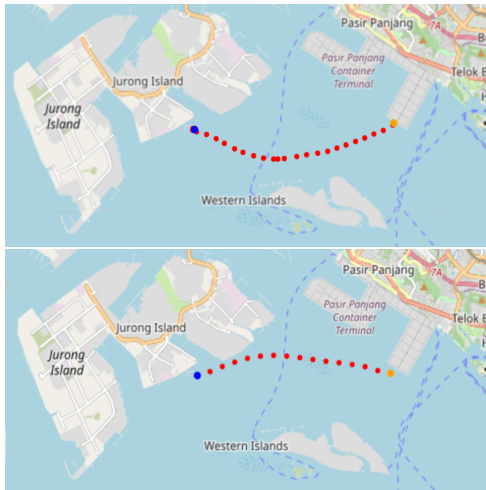


Figure 9: Instance Case 3: Historical human-operated route (top) versus our solution (bottom)

6.5 Fuel Consumption Table

The table below presents the results of historical fuel consumption alongside the fuel consumption suggested by the model for three distinct cases. Notably, across all three cases, there is a substantial reduction in fuel consumption in the




Figure 10: Platform dashboard implemented to facilitate fuel-saving transit training and instruction for tug masters

model-suggested transit compared to historical data. This observation underscores the potential for significant improvements in the way current tugboats navigate, indicating a clear need for optimizing fuel consumption. The discernible reduction in fuel consumption also highlights the efficacy of the proposed model in offering more efficient transit routes.

Case No.	Historical Fuel Consumption (Litre)	Model Fuel Consumption (Litre)
1	306.8 L	79.7394 L
2	174.2 L	148.7517 L
3	181 L	117.3034 L

Table 2: Comparison of historical and model total fuel consumption

6.6 Dashboard Implementation

In accordance with the optimisation models, we have implemented a platform system that can support the retrospective review of harbor tug transits for fuel-saving operations. The dashboard developed is illustrated in Figure 10, which allows directly comparing the optimised transit solutions (through RL-based fuel-saving route generation and speed pattern) with the historical real-world transit and also qualifies the amount of fuel consumption and saving. This helps the tug masters acquire essential insights on how to achieve fuel-efficient transits based on real-world instances.

7 Conclusion

In this paper, we introduce an integrated approach that integrates a fuel consumption prediction model with a Deep Deterministic Policy Gradient (DDPG) model for optimizing tugboat transit operations. The integrated model adeptly suggests optimized transit routes for tugboats, leveraging both speed and environmental factors to enhance operational efficiency. Through trajectory-based analysis, our model consistently demonstrates a substantial reduction in total fuel consumption, indicating significant opportunities for further optimization in tugboat transit trajectories. The results also affirm the model’s efficacy in contributing to fuel savings and underscore its potential for practical implementation in real-world tugboat operations.

Acknowledgments

This work was supported by Maritime AI Research Programme (grant number SMI-2022-MTP-06 funded by Singapore Maritime Institute), and all the helpful supports and valuable domain knowledge from PSA Marine in this project are greatly appreciated.

References

- [Alahi *et al.*, 2016] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.
- [An *et al.*, 2021] Jing An, Lexing Ying, and Yuhua Zhu. Why resampling outperforms reweighting for correcting sampling bias with stochastic gradients. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [Caliskan, 2022] Aylin Caliskan. Seaports participation in enhancing the sustainable development goals. *Journal of Cleaner Production*, 379:134715, 2022.
- [Chew *et al.*, 2023] Lock Yue Chew, Ning Ning Chung, Wen Xuan Sia, Hoai Nguyen Huynh, Glenn Sim, Alvin Chua, and Zhongwen Huang. A data analytic elucidation on the spatial agglomeration of singapore maritime industry. *International Journal on Smart and Sustainable Cities*, page 2340001, 2023.
- [Daniel *et al.*, 2010] Kenny Daniel, Alex Nash, Sven Koenig, and Ariel Felner. Theta*: Any-angle path planning on grids. *J. Artif. Int. Res.*, 39(1):533–579, sep 2010.
- [Lillicrap *et al.*, 2016] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *International Conference on Learning Representations*, 2016.
- [Liu *et al.*, 2021] Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. Just train twice: Improving group robustness without training group information. In *International Conference on Machine Learning*, pages 6781–6792. PMLR, 2021.
- [Moradi *et al.*, 2022] Mohammad Hossein Moradi, Martin Brutsche, Markus Wenig, Uwe Wagner, and Thomas Koch. Marine route optimization using reinforcement learning approach to reduce fuel consumption and consequently minimize co2 emissions. *Ocean Engineering*, 259:111882, 2022.
- [Muñoz *et al.*, 2019] Guillem Muñoz, Cristina Barrado, Ender Çetin, and Esther Salami. Deep reinforcement learning for drone delivery. *Drones*, 3(3):72, 2019.
- [Steininger *et al.*, 2021] Michael Steininger, Konstantin Kobs, Pdraig Davidson, Anna Krause, and Andreas Hotho. Density-based weighting for imbalanced regression. *Machine Learning*, 110:2187–2211, 2021.
- [Xiao *et al.*, 2020] Zhe Xiao, Xiuju Fu, Liye Zhang, Wanbing Zhang, Ryan Wen Liu, Zhao Liu, and Rick Siow Mong Goh. Big data driven vessel trajectory and navigating state prediction with adaptive learning, motion modeling and particle filtering techniques. *IEEE Transactions on Intelligent Transportation Systems*, 23(4):3696–3709, 2020.
- [Yuan *et al.*, 2021] Zhi Yuan, Jingxian Liu, Qian Zhang, Yi Liu, Yuan Yuan, and Zongzhi Li. Prediction and optimisation of fuel consumption for inland ships considering real-time status and environmental factors. *Ocean Engineering*, 221:108530, 2021.
- [Zhang *et al.*, 2022] Xiaocai Zhang, Xiuju Fu, Zhe Xiao, Haiyan Xu, and Zheng Qin. Vessel trajectory prediction in maritime transportation: Current approaches and beyond. *IEEE Transactions on Intelligent Transportation Systems*, 23(11):19980–19998, 2022.