

Realistic Curriculum Reinforcement Learning for Autonomous and Sustainable Marine Vessel Navigation

Xiaocai Zhang¹, Zhe Xiao², Maohan Liang³, Tao Liu⁴, Haijiang Li^{5,6}, Wenbin Zhang^{7*}

¹Department of Infrastructure Engineering, The University of Melbourne, Australia

²Institute of High Performance Computing, Agency for Science, Technology and Research (A*STAR), Singapore

³School of Navigation, Wuhan University of Technology, China

⁴College of Transport & Communications, Shanghai Maritime University, China

⁵School of Maritime Economics and Management, Dalian Maritime University, China

⁶Collaborative Innovation Centre for Transport Study, Dalian Maritime University, China

⁷Michigan Technological University, USA

xiaocai.zhang@unimelb.edu.au, xiaoz@ihpc.a-star.edu.sg, mhliang@whut.edu.cn, liutao2@shmtu.edu.cn,

haijiang@dmlu.edu.cn, wenbinzh@mtu.edu

Abstract

Sustainability is becoming increasingly critical in the maritime transport, encompassing both environmental and social impacts, such as Greenhouse Gas (GHG) emissions and navigational safety. Traditional vessel navigation heavily relies on human experience, often lacking autonomy and emission awareness, and is prone to human errors that may compromise safety. In this paper, we propose a Curriculum Reinforcement Learning (CRL) framework integrated with a realistic, data-driven marine simulation environment and a machine learning-based fuel consumption prediction module. The simulation environment is constructed using real-world vessel movement data and enhanced with a Diffusion Model to simulate dynamic maritime conditions. Vessel fuel consumption is estimated using historical operational data and learning-based regression. The surrounding environment is represented as image-based inputs to capture spatial complexity. We design a lightweight, policy-based CRL agent with a comprehensive reward mechanism that considers safety, emissions, timeliness, and goal completion. This framework effectively handles complex tasks progressively while ensuring stable and efficient learning in continuous action spaces. We validate the proposed approach in a sea area of the Indian Ocean, demonstrating its efficacy in enabling sustainable and safe vessel navigation.

Introduction

The maritime transport sector plays a pivotal role in global trade and logistics, as it supports approximately 90% of the global trade volume, being the most cost-efficient transport mode over long distances (Zhang et al. 2024a; Liu et al. 2025). However, it faces growing pressure to align with sustainability goals in response to escalating environmental concerns (Nguyen et al. 2023; Fan et al. 2024; Wei et al. 2025). Greenhouse Gas (GHG) emissions from vessels contribute significantly to global warming. Maritime shipping

produces billions of tons of GHG emissions annually. In response, the International Maritime Organization (IMO) has initiated strategies to reduce shipping emissions by at least 50% by 2050. Despite global efforts, under a business-as-usual scenario, shipping’s carbon dioxide emissions are projected to increase by approximately 10–30% compared to 2008 levels by 2050 (Huang et al. 2025). Figure 1 illustrates the IMO’s projected GHG reduction pathway.




Figure 1: IMO’s GHG reduction pathway.

Safety represents a fundamental aspect of maritime sustainability (Zhou, Li, and Yuen 2023). As emphasized by the IMO, a sustainable maritime transport system must not only prioritize operational efficiency and reduced emissions, but also ensure the protection of human life, cargo, and marine ecosystem by addressing risks such as oil spills, collisions, and fatalities. Failures in safety management can lead to catastrophic outcomes. One notable example is the collision between the SANCHI oil tanker and a bulk carrier in East China Sea, which resulted in the loss of 32 crew members and the discharge or combustion of more than 100,000 tons of petroleum products in the marine environment (Zhang et al. 2022). Investigations concluded that human error was the primary cause of the incident (Wan and Chen 2018).

Optimizing marine operations provides a practical and high-impact approach to reducing GHG emissions within the maritime sector. Unlike fuel-switching or vessel retrofitting, operational measures, such as slow steaming, optimized routing, speed regulation, and Just-In-Time (JIT) arrival strategies (Zhang et al. 2023), can be rapidly adopted

*Corresponding author.

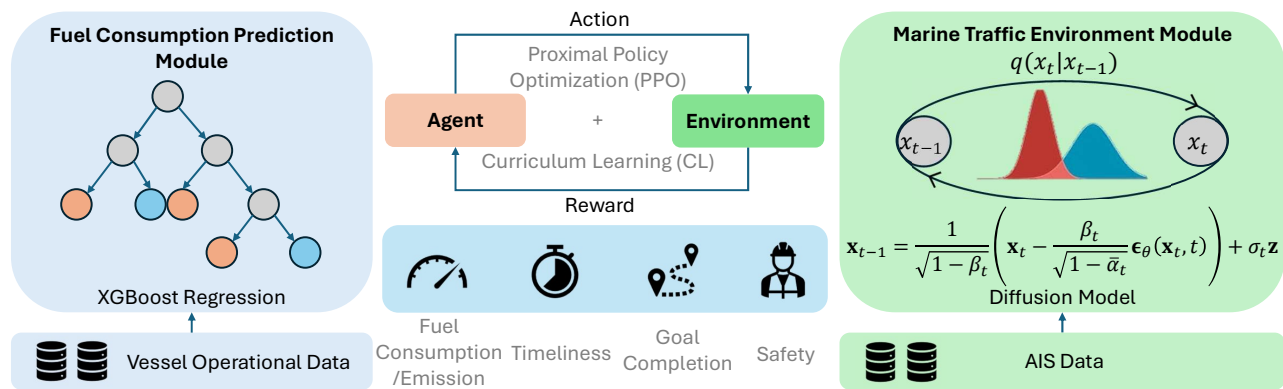


Figure 2: Overview of the methodological workflow in this study.

across the current global fleet with minimal capital expenditure. These strategies not only lower fuel consumption but also improve overall operational efficiency and safety, offering both environmental and economic benefits. According to the IMO, such improvements have the potential to achieve emission reductions of up to 30–40% by 2030 (Comer, Chen, and Rutherford 2018), serving as a crucial interim solution while low- and zero-emission technologies continue to advance. In this context, the application of advanced Artificial Intelligence (AI) techniques to maritime navigation systems emerges as a timely and promising opportunity to revolutionize conventional operational practices.

To support decarbonization in marine vessel operations, recent studies leveraging Deep Reinforcement Learning (DRL) have primarily focused on operational optimization tasks such as routing and scheduling. For example, Moradi et al. (Moradi et al. 2022) proposed a Deep Deterministic Policy Gradient (DDPG) framework to optimize vessel routing with the objective of reducing fuel consumption and CO₂ emissions. Similarly, Wang et al. (Wang et al. 2024a) applied a DDPG-based reinforcement learning approach to optimize transit routes for harbor tugboats in Singapore, taking into account local waterway characteristics, static obstacles, and dynamic environmental factors. Chen et al. (Chen et al. 2025) introduced a hybrid approach that integrates a traditional A* search algorithm (Xiao et al. 2024) with a Double Deep Q-Network (DDQN) to enhance route planning, aiming to minimize sailing time, operational costs, and carbon emissions. Additionally, Hasanvand et al. (Hasanvand et al. 2020) presented an advanced DQN-based strategy for onboard energy scheduling in all-electric ships, formulating a multi-objective optimization problem that balances emission reduction, cost efficiency, and system reliability.

In the context of safe navigation, most existing research focuses on real-time decision-making with the primary objective of collision avoidance. One study proposed an autonomous collision avoidance system for ships using DRL based on a Deep Q-Network (DQN), incorporating a quantitative risk assessment derived from the Closest Point of Approach (CPA) within a maneuverability-aware ship domain that adheres to COLREGs regulations (Wang et al. 2024b). Another work targeted strategy optimization for

Maritime Autonomous Surface Ships (MASS) navigating through complex traffic scenarios. It introduced a risk-aware Safe Reinforcement Learning (SRL) approach built upon the Actor-Critic architecture, which jointly considers collision risk and decision reliability in the policy optimization process (Wang et al. 2024b). Furthermore, Pan et al. (Pan et al. 2025) developed a novel DRL framework utilizing a Dueling Double Deep Q-Network (DDDQN), specifically designed for multi-ship collision avoidance in dynamic and congested environments. This method advances the field by realistically modeling ship interactions and continuously generating safe maneuvering recommendations in real time.

Achieving the multi-objective goals of carbon emission reduction, navigational safety, and timely and successful task completion in autonomous vessel operations remains a significant challenge, even with the support of DRL techniques. Existing reinforcement learning approaches often face limitations in scalability and generalization, particularly under the highly dynamic conditions characteristic of real-world maritime environments. Furthermore, multi-objective optimization in DRL-based vessel autonomous navigation poses significant challenges for effective policy learning and deployment (Li, Zhang, and Wang 2020). To address these limitations, we propose a novel Curriculum Reinforcement Learning (CRL) framework specifically designed for autonomous and sustainable marine vessel navigation. Our approach is grounded in a realistic, data-driven simulation environment constructed from real-world Automatic Identification System (AIS) data (Liang et al. 2024) and augmented with a Diffusion Model (DM) (Ho, Jain, and Abbeel 2020) to simulate dynamic maritime traffic conditions. The real-time traffic conditions are encoded as image-based environmental representations, effectively capturing the spatial complexity of maritime scenarios. In addition, we develop a machine learning-based fuel consumption prediction model trained on a large-scale dataset of real-world operational data including fuel usage records, enabling accurate estimation of fuel consumption. At the core of our DRL solution is a lightweight CRL agent guided by a multi-objective reward structure, facilitating efficient and robust policy training in complex, multi-constraint environments. We evaluate the proposed framework in a region of the Indian Ocean,

demonstrating its effectiveness in achieving safe, efficient, and environmentally sustainable navigation. By integrating curriculum learning, realistic environment simulation, and sustainability-oriented objectives, this work provides a scalable and principled pathway towards advancing AI-powered autonomy in maritime transport.

The main contributions are summarized as follows:

- **A novel CRL framework for marine vessel navigation:** We propose a CRL framework that integrates curriculum learning and reinforcement learning paradigms tailored for autonomous vessel navigation, addressing safety, emissions, timeliness, and goal-oriented performance in complex ocean environments.
- **Realistic data-driven simulation environment:** We construct a high-fidelity simulation environment using real-world AIS-based vessel movement data, enhanced by a DM to simulate diverse and dynamic maritime traffic conditions.
- **Integrated fuel consumption prediction module:** We develop a data-driven fuel estimation model using historical real-world operational records from ocean-going vessels, enabling the agent to receive realistic and accurate fuel consumption feedback and make emission-aware navigation decisions.

Methods

Overview

An overview of the methodological workflow employed in this study is illustrated in Figure 2. The proposed framework consists of three key components: a fuel consumption prediction module, a marine traffic environment module, and CRL-based policy learning modules. The fuel prediction module leverages real-world operational data from vessels to estimate fuel usage accurately. The marine traffic environment module simulates dynamic and realistic maritime conditions using AIS data, providing a high-fidelity training ground for policy learning. The CRL framework is driven by a multi-objective reward function that jointly optimizes fuel and emission reduction, navigational safety, timeliness, and successful task completion.

Fuel Consumption Prediction

Given the direct relationship between fuel usage and GHG emissions—particularly CO_2 , NO_x , and SO_x —robust prediction models are critical for achieving both economic and environmental objectives. These models enable ship operators to estimate fuel requirements for specific voyages based on various dynamic factors, including vessel characteristics, engine specifications, cargo load, weather conditions, and sea state. However, due to limitations in data availability, some of these factors are challenging to obtain consistently.

To address this, we propose a machine learning-based prediction model trained on real-world operational data collected from several hundred international ocean-going vessels over a two-year period. The input features include vessel travel distance, geographical location (latitude and longitude), speed-over-ground (SOG), and vessel particulars such

as Length Overall (LOA), beam, Gross Tonnage (GT), and ship type. Temporal variables, such as month, day, and hour, are also incorporated, as they are known to correlate with varying ocean and sea conditions (Caires and Sterl 2005). Categorical features, including ship type and temporal information, are processed using one-hot encoding. Figure 3 shows the distributions of vessel SOG and ship types.

The model outputs the Fuel Consumption Rate (FCR), measured in metric tons (mt) of fuel per hour. In this study, fuel consumption is measured comprehensively to include usage from the main engine, auxiliary engine, and auxiliary machinery. Four types of fuel are considered: Heavy Fuel Oil (HFO), Low Sulfur Fuel Oil (LSFO), Diesel Oil (DO), and Low Sulfur Gas Oil (LSGO). The total fuel consumption represents the aggregate usage of all four fuel types.

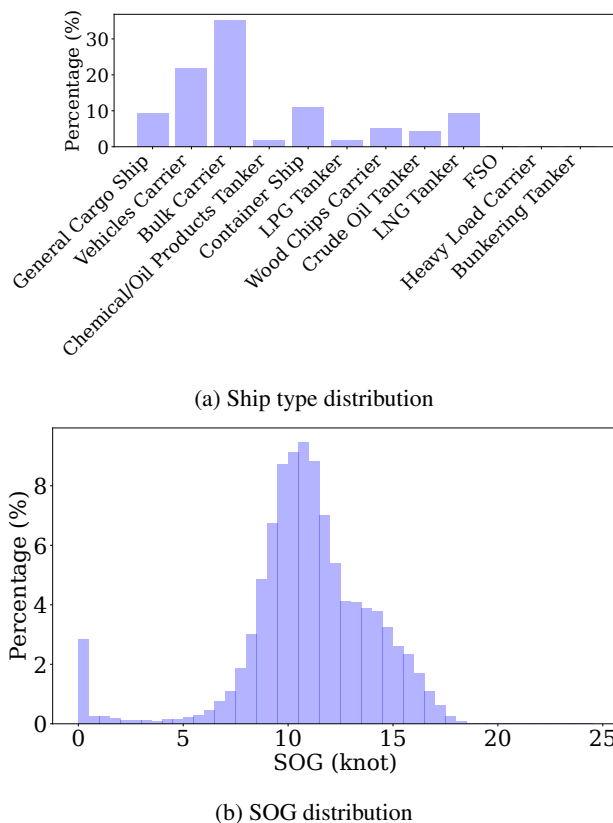


Figure 3: Distribution of vessel features

Following insights from a recent literature review (Nguyen et al. 2025), we adopt eXtreme Gradient Boosting (XGBoost) to predict the FCR using the selected input features. Let the input feature vector be denoted by $\mathbf{x} \in \mathbb{R}^d$, where $d = 86$ represents the total number of features. The prediction model is formulated as:

$$\hat{y} = f_{\text{xgboost}}(\mathbf{x}) = \sum_{k=1}^K f_k(\mathbf{x}), \quad f_k \in \mathcal{F}, \quad (1)$$

where \hat{y} is the predicted FCR, K is the number of decision trees in the ensemble, and each f_k is a regression tree belonging to the functional space \mathcal{F} .

DM-Enhanced Marine Traffic Environment

To enrich the realism and variability of vessel trajectory data in the simulation environment, we incorporate a generative model (Zhang et al. 2024b,a) to generate synthetic AIS trajectories that mimic real-world movement patterns. Diffusion models, originally developed for image generation, have shown promising results in sequential data generation tasks due to their ability to learn complex distributions through iterative denoising.

Let a vessel trajectory be represented as a sequence of positions:

$$\mathbf{x}_0 = ((\phi_{t,1}, \lambda_{t,1}, v_{t,1}), \dots, (\phi_{t,T}, \lambda_{t,T}, v_{t,T})), \mathbf{x}_0 \in \mathbb{R}^{T \times 3}, \quad (2)$$

where \mathbf{x}_0 is the original trajectory and T is the number of time steps. ϕ_t , λ_t , and v_t are latitude, longitude, and SOG, respectively. The forward diffusion process gradually adds Gaussian noise to the data over N steps:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}), \quad t = 1, \dots, N \quad (3)$$

where β_t is a variance schedule controlling the noise level.

During training, the model learns to reverse this process by estimating the added noise ϵ , using a neural network ϵ_θ :

$$\mathcal{L}_{\text{DM}} = \mathbb{E}_{\mathbf{x}_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2] \quad (4)$$

To sample a new trajectory, we initialize with Gaussian noise $\mathbf{x}_N \sim \mathcal{N}(0, \mathbf{I})$, and iteratively apply the learned reverse process, as shown by

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{1 - \beta_t}} \left(\mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}, \quad \mathbf{z} \sim \mathcal{N}(0, \mathbf{I}), \quad (5)$$

where $\bar{\alpha}_t = \prod_{s=1}^t (1 - \beta_s)$ and σ_t is the noise scale.

By leveraging the DM, we are able to generate plausible and diverse vessel trajectories that align with realistic movement patterns observed in AIS data. These synthetic trajectories enrich the simulated environment, thereby improving the quality of downstream DRL model training.

CRL Framework

We formulate the vessel navigation task as a Markov Decision Process (MDP), which is typically defined by four key components: state, action, state transition, and reward (Ni et al. 2023). The curriculum learning introduces a shrinking distance-to-goal threshold that decreases with training episode.

The state representation at time step t consists of two components: the self-vessel state \mathbf{s}_1^t and the environmental state \mathbf{s}_2^t . The self-vessel state \mathbf{s}_1^t captures the vessel's status and contextual information, such as current speed and direction, at its current position. \mathbf{s}_1^t is defined by

$$\mathbf{s}_1^t = [\phi_t, \lambda_t, \phi_t^d, \lambda_t^d, \psi_t, \beta_t, \vartheta_t], \quad \mathbf{s}_1^t \in \mathbb{R}^9, \quad (6)$$

where ϕ_t, λ_t denote the geographical coordinates (latitude and longitude) at time t , and ϕ_t^d, λ_t^d represent the destination coordinates. ψ_t is the vessel's heading angle, β_t is the direction of the ocean current, and ϑ_t is the vessel's current speed. All angle-related features (e.g., heading and current

direction) are encoded using a two-dimensional representation (e.g., sine, cosine encoding), resulting in a total state dimension of 9.

The environmental state \mathbf{s}_2^t is encoded as a three-channel image tensor representing the vessel's surrounding traffic conditions, as formulated by

$$\mathbf{s}_2^t \in \mathbb{R}^{64 \times 64 \times 3}. \quad (7)$$

Each channel corresponds to a 64×64 grid centered on the self-vessel, where each pixel represents a discretized spatial cell in the local coordinate frame, covering a fixed area (e.g., $10 \text{ nm} \times 10 \text{ nm}$). The three channels are defined as follows:

- **Channel 0: Occupancy** — a binary indicator of whether a vessel is present in the cell.
- **Channel 1: SOG** — SOG of the vessel in the cell.
- **Channel 2: Course-Over-Ground (COG)** — COG of the vessel in the cell.

The action space is defined as a continuous two-dimensional space comprising the heading increment and Speed Through Water (STW). At each time step t , the action is represented by

$$\mathbf{a}_t = [\Delta\psi_t, v_t], \quad -\hat{\psi} \leq \Delta\psi_t \leq \hat{\psi}, \quad v_l \leq v_t \leq v_u, \quad (8)$$

$$\psi_{t+1} = \psi_t + \Delta\psi_t, \quad (9)$$

where $\Delta\psi_t$ denotes the change in heading angle, and v_t is the commanded STW. The intervals $[-\hat{\psi}, \hat{\psi}]$ and $[v_l, v_u]$ define the allowable ranges for heading adjustment and speed, respectively. The vessel's COG and SOG are then computed by factoring in the effects of ocean currents on the vessel's intended motion. The subsequent position of the vessel is predicted through simulation based on the applied action and environmental conditions.

The reward function is designed to capture multiple objectives, including safety, fuel consumption, timeliness, and goal completion. Formally, the reward at time step t is demonstrated by

$$r_t = \begin{cases} 30 + 1.5g_t - f_t - s_t, & \text{if } d_{\text{cur}} < \omega(e) \\ 1.5g_t - f_t - s_t - 1.0, & \text{if } d_{\text{cur}} > d_{\text{pre}} \\ 1.5g_t - f_t - s_t - 0.1 \cdot d_{\text{cur}}, & \text{if late and } d_{\text{cur}} \geq \omega(e) \\ 1.5g_t - f_t - s_t, & \text{otherwise} \end{cases}, \quad (10)$$

where g_t , f_t , and s_t denote the reward components for goal completion, fuel consumption, and safety, respectively. The terms d_{cur} and d_{pre} represent the distances from the current and previous vessel positions to the destination, respectively. In the curriculum learning paradigm, we introduce a threshold parameter $\omega(e)$ that progressively decreases over training episodes, transitioning from more challenging to simpler tasks to facilitate gradual learning. The parameter $\omega(e)$ is defined by

$$\omega(e) = \omega_0 \cdot \left(1 - \frac{e}{N_e}\right) + \omega_f \cdot \min\left(\frac{e}{N_e}, 1\right), \quad (11)$$

where e denotes the episode index, and ω_0 is the initial threshold, set to 5 nautical miles (nm) in this study. The goal completion reward g_t is calculated by

$$g_t = d_{\text{pre}} - d_{\text{cur}}. \quad (12)$$

The fuel consumption reward f_t is computed by

$$f_t = \alpha \cdot \frac{f_{XGBoost}(\mathbf{x}_t)}{GT}, \quad (13)$$

The safety reward s_t is computed based on the Distance at Closest Point of Approach (DCPA) and the Time to Closest Point of Approach (TCPA), defined as follows:

$$s_t = \frac{1}{N} \sum_{i=1}^N \left[\text{clip} \left(1 - \frac{DCPA_i}{d_{safe,i}}, 0, 1 \right) \cdot \text{clip} \left(1 - \frac{|TCPA_i|}{t_{max}}, 0, 1 \right) \right], \quad (14)$$

$$d_{safe} = \max \left(\tau \cdot \frac{L_s + B_s + L_t + B_t}{2 \times 1852}, 0.5 \right), \quad (15)$$

where L and B means LOA and beam of a vessel, respectively. Subscripts s and t denote the self-vessel and target vessel involved in a potential collision scenario. To define the safety thresholds, we introduce a buffer multiplier τ , set to 4 in this study. Additionally, the time threshold for collision risk assessment is capped by a maximum allowable time to closest approach, and $t_{max} = 15$ minutes.

The DRL framework in this study is built upon the Proximal Policy Optimization (PPO) algorithm, which comprises two components: an actor network and a critic network. The architecture of the actor network is depicted in Figure 4. The network receives two types of input states, each processed through distinct feature extraction pathways. The environmental state, represented as an image-based tensor, is processed using a lightweight convolutional backbone composed of two separable convolutional layers, allowing for efficient spatial feature extraction. In parallel, the self-vessel state, represented as a structured vector, is processed through dense layers. The extracted features from both branches are then concatenated into a joint latent representation, which is used to compute the action distribution. The critic network adopts a similar architecture to the actor, sharing the same input processing structure, but outputs a scalar value representing the estimated value of the current state.

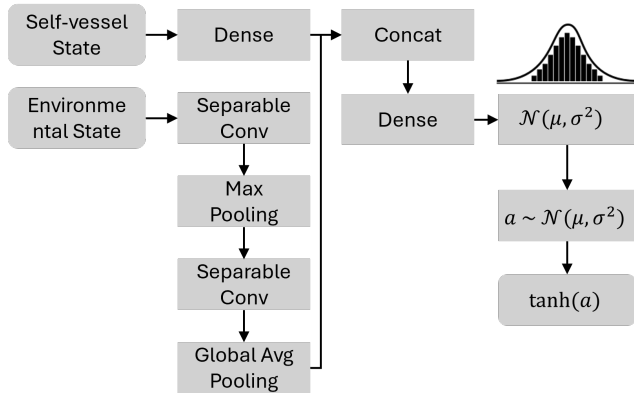


Figure 4: Architecture of the actor network in CRL.

Experiments & Results

The experiments are conducted through simulation within a selected sea area in Indian Ocean, chosen for its high density

of commercial vessel traffic and diverse environmental conditions. While open AIS data for this region is limited, and comprehensive operational datasets for vessel fuel consumption are typically confidential and not publicly accessible, we utilize proprietary data obtained through industry collaboration. To evaluate performance under varying conditions, we design three representative navigation instances. A moderate traffic scenario is simulated, comprising hundreds of vessels distributed across the region to emulate realistic maritime operational environments. To further enhance environmental fidelity, monthly-varying ocean currents are modeled based on historical distributions of current direction and speed. These currents introduce dynamic disturbances that influence vessel trajectories and are integrated into the simulation as time-varying environmental factors. The combined modeling of traffic density and dynamic current flows enables a comprehensive evaluation of the proposed navigation strategy’s robustness and adaptability across a range of realistic and challenging operational conditions.

Fuel Consumption Prediction Performance

The prediction performance is evaluated using an independent test dataset. Three commonly used metrics—Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2)—are employed to assess regression model accuracy. The proposed XGBoost model is compared against several widely adopted machine learning algorithms, including Support Vector Regression (SVR), Multi-Layer Perceptron (MLP), and Random Forest (RF). Overall, XGBoost and RF demonstrate superior performance among the evaluated models. XGBoost achieves the highest R^2 score of 86.10%, indicating strong predictive capability. Furthermore, XGBoost slightly outperforms RF in terms of both RMSE and R^2 , confirming its effectiveness for fuel consumption prediction in maritime applications.

Method	MAE	RMSE	R^2 (%)
SVR	0.4529	0.7050	45.01
MLP	0.2440	0.4916	77.06
ET	0.2116	0.4033	84.56
LightGBM	0.2015	0.3895	0.8560
RF	0.1752	0.3832	86.06
XGBoost	0.1802	0.3827	86.10

Table 1: Performance comparison between fuel consumption prediction models.

Instance Case 1

This instance simulates a vessel navigation task with a time constraint of 21 hours, spanning from 2024-02-08 00:00:00 UTC to 2024-02-08 21:00:00 UTC, requiring the vessel to travel from a specified starting point to its destination. Table 2 presents the performance comparison of various control models based on three key metrics: Accumulated Reward (AR), Accumulated Fuel Consumption (AFC), and Accumulated Safety Score (ASS), where a higher ASS indicates greater navigational risk. The baseline models include CL-ABDDQN, CL-A2C, and DDPG (Wang et al.

2024a; Moradi et al. 2022). Both CL-ABDDQN and CL-A2C adopt the curriculum learning paradigm within the Advantage Actor-Critic (A2C) framework. All baselines are trained and evaluated under the same simulated environment and reward structure as the proposed CRL framework, with the only difference being the underlying DRL architecture. The ABDDQN model builds upon the DDQN with an action-branching architecture, which effectively handles high-dimensional discrete action spaces by decoupling action components. This branching structure enhances learning stability and improves exploration efficiency (Tavakoli, Pardo, and Kormushev 2018). Among all models, the proposed CRL framework demonstrates the best overall performance across all three evaluation metrics. Notably, CL-A2C fails to complete the navigation task, resulting in a negative AR. The DDPG model underperforms, yielding lower AR and higher AFC and ASS values. While CL-ABDDQN demonstrates competitive performance, it still falls short of CRL in terms of fuel efficiency and safety. The navigation trajectory for this scenario is illustrated in Figure 5, where the blue marker indicates the starting point and the red marker denotes the destination.

Method	AR	AFC	ASS
CL-ABDDQN	155.527	20.509	1.466
CL-A2C	-82.034	13.092	7.501
DDPG	150.424	33.533	1.84
CRL	154.018	18.015	0.888

Table 2: Performance comparison of different control models in Instance Case 1.




Figure 5: Vessel navigation trajectory for Instance Case 1.

Instance Case 2

This instance simulates a vessel navigation task over a 46-hour period, from 2024-02-08 00:00:00 UTC to 2024-02-09 22:00:00 UTC, requiring the vessel to travel from a specified starting point to its destination. Table 3 presents the performance of various control models. Due to the longer duration compared to Instance Case 1, all three metrics exhibit higher values across the models. The proposed CRL framework achieves the lowest AFC while maintaining a balanced

AR and ASS. CL-ABDDQN obtains the highest reward, whereas CL-A2C yields the lowest safety score, indicating a safer trajectory. Both models incorporate the Curriculum Learning (CL) framework, highlighting its positive impact on DRL performance. Overall, CRL demonstrates the most balanced performance across all three indicators. The navigation trajectory for this instance case is visualized in Figure 6, where the blue marker indicates the starting point and the red marker marks the destination.

Method	AR	AFC	ASS
CL-ABDDQN	298.026	36.229	5.543
CL-A2C	287.175	21.487	3.952
DDPG	279.439	20.974	5.063
CRL	294.148	19.963	4.754

Table 3: Performance comparison of different control models in Instance Case 2.



Figure 6: Vessel navigation trajectory for Instance Case 2.

Instance Case 3

This instance simulates a vessel navigation task with a 21-hour time window, from 2024-02-08 00:00:00 UTC to 2024-02-08 21:00:00 UTC, requiring the vessel to travel from a designated starting point to its destination. Table 4 summarizes the performance of various control models. In this case, the CRL framework achieves the lowest ASS, indicating the safest navigation behavior, although it incurs a slightly higher AFC compared to the most fuel-efficient models. Conversely, the CL-A2C model records a negative AR (-100.699) along with high AFC and ASS values, reflecting failed goal completions and divergence in learning. The DDPG model demonstrates the best fuel efficiency, achieving the lowest AFC, but at the expense of safety, as indicated by the highest ASS. Overall, while CRL does not significantly outperform all baseline models, it shows a clear advantage over CL-A2C. Its performance is comparable to that of CL-ABDDQN, with both models achieving a balanced trade-off among reward, safety, and fuel consumption. Figure 7 illustrates the vessel's trajectory for this scenario, where the blue marker indicates the starting point and the red marker represents the destination.

Method	AR	AFC	ASS
CL-ABDDQN	155.035	20.901	2.332
CL-A2C	-100.699	35.154	7.258
DDPG	148.615	9.002	2.710
CRL	145.015	28.618	2.143

Table 4: Performance comparison of different control models in Instance Case 3.

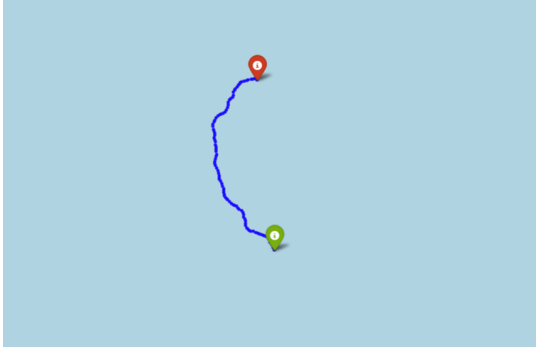


Figure 7: Vessel navigation trajectory for Instance Case 3.

Ablation Analysis

To evaluate the impact of the proposed curriculum learning paradigm within the DRL framework, we perform an ablation analysis by comparing the full model (CRL) with a variant that excludes the CL component (CRL w/o CL). The CRL model incorporates curriculum learning into the training process, while CRL w/o CL shares the same architecture but is trained without the CL mechanism. The results are summarized in Table 5. In Instance Case 2, the CRL w/o CL model obtains a negative AR, indicating failure to complete the navigation task and highlighting the critical role of curriculum learning in guiding the agent toward successful goal completion under challenging conditions. In Instance Cases 1 and 3, the full CRL model outperforms its counterpart across all evaluation metrics, with the exception of fuel consumption in Instance Case 3, where CRL w/o CL shows a slight advantage in fuel efficiency. These results confirm the effectiveness and necessity of curriculum learning in enhancing training stability, task success rate, and overall performance in complex maritime navigation scenarios.

Instance	Method	AR	AFC	ASS
Case 1	w/o CL	139.629	21.893	1.069
	CRL	154.018	18.015	0.888
Case 2	w/o CL	-708.765	48.951	3.274
	CRL	294.148	19.963	4.754
Case 3	w/o CL	149.950	19.557	5.901
	CRL	145.015	28.618	2.143

Table 5: Ablation analysis comparing CRL with and without CL across three instance cases.

Figure 8 illustrates the training reward trends of the CRL model compared to its variant without the curriculum learn-

ing paradigm. The comparison reveals that CRL leads to faster convergence and greater training stability. With CRL, the agent rapidly achieves higher accumulated rewards in the early episodes and maintains stable performance throughout training, indicating efficient policy learning. In contrast, the model without curriculum learning shows slower initial improvement and more pronounced fluctuations, including periods of significant performance degradation, suggesting difficulty in exploring sparse or complex reward spaces. Overall, incorporating curriculum learning improves sample efficiency, stabilizes the learning process, and results in more consistent reward optimization.




Figure 8: Comparison of training reward trends between CRL and CRL without CL.

Conclusion

In this study, we proposed a realistic Curriculum Reinforcement Learning (CRL) framework designed for autonomous and sustainable navigation of marine vessels. The framework incorporates a data-driven simulation environment built from real-world AIS data, a generative diffusion model to capture dynamic maritime conditions, and a machine learning-based fuel consumption prediction module to provide realistic emission feedback. The DRL agent is embedded within the curriculum learning paradigm and is guided by a multi-objective reward function that jointly optimizes navigational safety, fuel efficiency, timeliness, and successful task completion, enabling efficient and stable policy learning in complex continuous control environments. Extensive simulation experiments conducted in a selected sea region of the Indian Ocean validate the effectiveness of the proposed approach. Our CRL framework demonstrates superior performance over state-of-the-art DRL baselines, consistently achieving better trade-offs among cumulative reward, safety, and fuel consumption. Additional ablation studies highlight the critical role of curriculum learning in enhancing training stability and accelerating convergence.

Acknowledgments

Haijiang Li was supported in part by the National Natural Science Foundation of China under grant no. 52501430, and Tao Liu was supported in part by the National Natural Science Foundation of China under grant no. 52571408.

References

- Caires, S.; and Sterl, A. 2005. 100-year return value estimates for ocean wind speed and significant wave height from the ERA-40 data. *Journal of Climate*, 18(7): 1032–1048.
- Chen, X.; Hu, R.; Luo, K.; Wu, H.; Biancardo, S. A.; Zheng, Y.; and Xian, J. 2025. Intelligent ship route planning via an A* search model enhanced double-deep Q-network. *Ocean Engineering*, 327: 120956.
- Comer, B.; Chen, C.; and Rutherford, D. 2018. Relating short-term measures to IMO's minimum 2050 emissions reduction target. *International Council on Clean Transportation*. Retrieved from.
- Fan, Y.; Agand, P.; Chen, M.; Park, E. J.; Kennedy, A.; and Bae, C. 2024. Sequential modeling of complex marine navigation: Case study on a passenger vessel (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 23484–23485.
- Hasanvand, S.; Rafiei, M.; Gheisarnejad, M.; and Khooban, M.-H. 2020. Reliable power scheduling of an emission-free ship: Multiobjective deep reinforcement learning. *IEEE Transactions on Transportation Electrification*, 6(2): 832–843.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Huang, G.; He, Z.; Zhao, P.; and Zhang, C. 2025. A review of climate-related disasters impact on global shipping carbon emissions. *Transportation Research Part D: Transport and Environment*, 139: 104553.
- Li, K.; Zhang, T.; and Wang, R. 2020. Deep reinforcement learning for multiobjective optimization. *IEEE Transactions on Cybernetics*, 51(6): 3103–3114.
- Liang, M.; Liu, R. W.; Gao, R.; Xiao, Z.; Zhang, X.; and Wang, H. 2024. A survey of distance-based vessel trajectory clustering: Data pre-processing, methodologies, applications, and experimental evaluation. *arXiv preprint arXiv:2407.11084*.
- Liu, T.; Wang, P.; Zhang, Z.; Lei, Z.; Ye, Y.; Huo, Y.; Zhang, X.; He, G.; and Wu, H. 2025. An Approach to Multi-AAV Ship Detection Based on Mobile Edge Computing Scenarios. *IEEE Transactions on Intelligent Transportation Systems*, 26(10): 16869–16886.
- Moradi, M. H.; Brutsche, M.; Wenig, M.; Wagner, U.; and Koch, T. 2022. Marine route optimization using reinforcement learning approach to reduce fuel consumption and consequently minimize CO2 emissions. *Ocean Engineering*, 259: 111882.
- Nguyen, S.; Gadel, M.; Wang, K.; Li, J.; Zhang, X.; Kong, S.-C.; Fu, X.; and Qin, Z. 2025. Maritime decarbonization through machine learning: A critical systematic review of fuel and power prediction models. *Cleaner Logistics and Supply Chain*, 14: 100210.
- Nguyen, S.; Leman, A.; Xiao, Z.; Fu, X.; Zhang, X.; Wei, X.; Zhang, W.; Li, N.; Zhang, W.; and Qin, Z. 2023. Blockchain-powered incentive system for jit arrival operations and decarbonization in maritime shipping. *Sustainability*, 15(22): 15686.
- Ni, C.; Duan, Y.; Dahleh, M.; Wang, M.; and Zhang, A. R. 2023. Learning good state and action representations for Markov decision process via tensor decomposition. *Journal of Machine Learning Research*, 24(115): 1–53.
- Pan, R.; Zhang, W.; Wang, S.; and Kang, S. 2025. Deep reinforcement learning model for Multi-Ship collision avoidance decision making design implementation and performance analysis. *Scientific Reports*, 15(1): 21250.
- Tavakoli, A.; Pardo, F.; and Kormushev, P. 2018. Action branching architectures for deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Wan, Z.; and Chen, J. 2018. Human errors are behind most oil-tanker spills. *Nature*, 560(7717): 161–163.
- Wang, S.; Zhang, X.; Li, J.; Wei, X.; Lau, H. C.; Dai, B. T.; Huang, B.; Xiao, Z.; Fu, X.; and Qin, Z. 2024a. Fuel-saving route planning with data-driven and learning-based approaches: A systematic solution for harbor tugs. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*.
- Wang, Y.; Xu, H.; Feng, H.; He, J.; Yang, H.; Li, F.; and Yang, Z. 2024b. Deep reinforcement learning based collision avoidance system for autonomous ships. *Ocean Engineering*, 292: 116527.
- Wei, X.; Lau, H. C.; Xiao, Z.; Fu, X.; Zhang, X.; and Qin, Z. 2025. Bi-objective dynamic tugboat scheduling with speed optimization under stochastic and time-varying service demands. *Transportation Research Part E: Logistics and Transportation Review*, 193: 103876.
- Xiao, Z.; Zhang, X.; Fu, X.; Zhang, L.; Xu, H.; Liu, R. W.; Chong, C. S.; and Qin, Z. 2024. Innovating Waterway Route Planning as a Service for Marine Traffic Applications. *IEEE Transactions on Intelligent Vehicles*.
- Zhang, X.; Fu, X.; Xiao, Z.; Xu, H.; and Qin, Z. 2022. Vessel trajectory prediction in maritime transportation: Current approaches and beyond. *IEEE Transactions on Intelligent Transportation Systems*, 23(11): 19980–19998.
- Zhang, X.; Fu, X.; Xiao, Z.; Xu, H.; Wei, X.; Koh, J.; Ogawa, D.; and Qin, Z. 2023. Prediction of vessel arrival time to pilotage area using multi-data fusion and deep learning. In *Proceedings of 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, 2268–2268. IEEE.
- Zhang, X.; Fu, X.; Xiao, Z.; Xu, H.; Zhang, W.; Koh, J.; and Qin, Z. 2024a. A dynamic context-aware approach for vessel trajectory prediction based on multi-stage deep learning. *IEEE Transactions on Intelligent Vehicles*, 9(11): 7193–7207.
- Zhang, X.; Xiao, Z.; Fu, X.; Wei, X.; Liu, T.; Yan, R.; Qin, Z.; and Zhang, J. 2024b. A Viewpoint Adaptation Ensemble Contrastive Learning framework for vessel type recognition with limited data. *Expert Systems with Applications*, 238: 122191.
- Zhou, Y.; Li, X.; and Yuen, K. F. 2023. Sustainable shipping: A critical review for a unified framework and future research agenda. *Marine Policy*, 148: 105478.