



Archives available at journals.mriindia.com

International Journal on Advanced Computer Theory and Engineering

ISSN: 2319-2526

Volume 15 Issue 01, 2026

Development Of a Versatile and Fast Algorithm for The Optimal Ship Routing

¹Owais Haider Rizvi, ²Syed Mohd Jawad, ³Mohammed Shehbaaz Shaikh, ⁴Vivek Ajay Yadav, ⁵Mr. Rishikesh Suryawanshi, ⁶Mrs. Harsha Dave

^{1,2,3,4} Dept of Computer Engineering, Shree L.R.Tiwari College of Engineering, Mira Road

^{5,6} Assistant Professor, Department of Computer, Shree L. R. Tiwari College of Engineering, Mira Road

Email: ¹owaisrizvi42@gmail.com, ²jawadsyed421@gmail.com, ³shehbazzzz338@gmail.com,

⁴vivek.a.yadav@slrtce.in, ⁶harsha.dave@slrtce.in

Peer Review Information

Submission: 19 March 2026

Revision: 08 April 2026

Acceptance: 24 April 2026

Keywords

Optimal Ship Routing, Maritime Fuel Optimization, XGBoost Regression, A* Search Algorithm, Ocean Weather Data

Abstract

Maritime transport efficiency is a critical factor in global trade, yet traditional routing often relies on static pathfinding that ignores the dynamic impact of ocean weather on fuel consumption. This paper presents an automated framework for Optimal Ship Routing that integrates machine learning with heuristic search to minimize fuel usage. Our system utilizes a specialized data pipeline to synchronize AIS (Automatic Identification System) ship telemetry with high-resolution wave and wind data from the Open-Meteo API. To model ship performance, we implement an XGBoost Regressor trained on engineered features, including Speed Through Water (STW) and relative wind angles, achieving precise fuel-burn predictions. These predictions serve as the cost-weighting for an A Search Algorithm* operating on a spherical coordinate grid. A key feature of our implementation is the automated integration of a Natural Earth GeoJSON land mask, ensuring all generated paths are restricted to navigable waters. Experimental results demonstrate that by intelligently diverting from the Great Circle path to avoid high-resistance weather cells, the framework provides a scalable, data-driven approach to reducing operational costs and carbon emissions in maritime logistics.

Introduction

Maritime transport has emerged as a fundamental pillar of global commerce, enabling the large-scale exchange of goods across distributed international markets through sophisticated logistics networks and microservice ecosystems. Modern shipping organizations increasingly depend on data-driven technologies to manage mission-critical operations, including voyage planning, fuel procurement, vessel tracking, and port integration. This widespread adoption of digital systems has dramatically increased operational visibility and transit velocity, but has

concurrently exposed the industry to the rising costs of fuel inefficiency and environmental pressure. According to recent studies and reports from the International Maritime Organization (IMO), reducing greenhouse gas emissions and improving fuel efficiency has become a critical priority for sustainable shipping operations [39]. As traditional navigation relies on static routing patterns, weaknesses in accounting for dynamic maritime weather and oceanographic conditions frequently manifest as pathways for excessive carbon emissions and massive operational overhead. Economic incidents involving fuel overconsumption have become increasingly

prevalent and severe. Factors such as high wave resistance, where adverse sea states significantly impede vessel speed; wind-induced drag that exploits poor aerodynamic profiles; ocean current interference causing unintended drift; improper speed management leading to engine strain; and severe weather susceptibility causing schedule disruptions are commonly identified as primary contributors to inefficiency [34], [41]. Recent maritime studies highlight that inaccurate weather integration and rigid pathfinding are leading causes of operational inefficiencies, with financial impacts extending into millions of dollars globally [28], [37].

Limitations of Existing Approaches

Conventional maritime routing methodologies suffer from significant limitations that reduce their effectiveness in modern commercial shipping. Manual voyage planning, while grounded in experience, requires substantial time and often relies on static charts that fail to account for rapidly changing weather patterns. Manual planners frequently overlook the non-linear relationship between engine load and sea state, leading to inefficient routes that incur excessive fuel penalties. Furthermore, traditional "Great Circle" navigation lacks adaptability for real-time environmental changes [34].

Existing weather routing systems, though beneficial, often function as opaque systems that provide limited interpretability of vessel-specific performance characteristics. While these systems can detect major weather disturbances, they struggle to model the complex interaction between Speed Over Ground (SOG), ocean current vectors, and fuel consumption [6], [35]. Additionally, many routing systems fail to incorporate accurate land-masking mechanisms, resulting in unsafe or impractical route suggestions [38].

Need for Integrated Optimization

Machine learning has emerged as a powerful tool for modeling complex ship behavior by leveraging historical AIS data and environmental inputs. Studies have demonstrated that machine learning models, including gradient boosting and deep learning techniques, can effectively predict fuel consumption and vessel efficiency under varying conditions [3], [22], [32], [33]. These models enable the creation of digital twins that capture intricate relationships between environmental factors and vessel performance. However, predictive modeling alone does not provide actionable routing strategies. Without integration with optimization algorithms, such predictions cannot guide vessels along fuel-efficient paths. Research indicates that

combining predictive models with optimization techniques such as heuristic search and multi-objective routing significantly improves operational efficiency [18], [24], [40].

The integration of machine learning with heuristic search algorithms, such as A*, enables dynamic cost-based routing where environmental resistance is treated as a variable cost. This hybrid approach allows exploration of optimal routes under realistic constraints, bridging the gap between prediction and decision-making [10], [18].

Proposed Approach and Contributions

This paper presents an automated framework for optimal ship routing that integrates high-resolution data engineering, predictive machine learning, and heuristic search within a unified computational environment. The system leverages AIS data and external oceanographic datasets such as Copernicus Marine Environment Monitoring Service (CMEMS) to ensure high-fidelity environmental modeling [36].

The machine learning component utilizes XGBoost, a scalable and efficient gradient boosting framework, to predict fuel consumption based on environmental parameters [1]. Data preprocessing and augmentation techniques further enhance model performance and robustness [20].

The routing component employs the A* algorithm on a spherical grid, incorporating real-time environmental costs to determine optimal paths. Land-masking is implemented using Natural Earth GeoJSON datasets to ensure navigational safety [38]. The integration of predictive modeling and heuristic search aligns with recent advancements in intelligent optimization systems and data-driven decision-making frameworks [30], [31].

Related Work

Research on maritime route optimization has expanded significantly in recent years, driven by the urgent need to reduce fuel costs, comply with environmental regulations, and utilize large-scale AIS datasets. Existing literature spans multiple domains, including physics-based modeling, machine learning prediction, heuristic optimization, and environmental constraint handling. However, most approaches address these aspects independently rather than through integrated frameworks [14], [29].

1. Maritime Fuel Consumption Analysis

Fuel consumption analysis has traditionally relied on physics-based models such as the Admiralty coefficient and parametric design tools for evaluating vessel efficiency [43]. These

approaches provide baseline estimates but fail to capture real-time environmental variability.

Recent advancements have introduced machine learning models trained on AIS and meteorological data to predict fuel consumption more accurately. Regression-based approaches, including XGBoost and neural networks, have demonstrated superior predictive performance compared to traditional methods [3], [32], [33]. Deep learning models, particularly LSTM-based frameworks, further enhance temporal prediction accuracy by capturing sequential dependencies in voyage data [22].

Similar machine learning techniques have been successfully applied in other domains such as agriculture, environmental monitoring, and economic forecasting, highlighting their generalizability [7], [9], [15], [25], [26], [27]. Data preprocessing and augmentation play a crucial role in improving model performance and handling imbalanced datasets [20].

Despite these advancements, most machine learning systems operate in isolation and lack integration with routing algorithms, limiting their practical applicability in real-time decision-making [10].

2. Automated Pathfinding and Optimization

Pathfinding algorithms have been widely used in maritime routing to determine optimal travel paths. Classical algorithms such as Dijkstra and A* have been adapted to maritime environments, enabling grid-based exploration of ocean spaces while considering environmental constraints [10].

Weather routing systems utilize forecast data to generate optimal paths based on environmental conditions, employing techniques such as isochrone methods and heuristic optimization [6], [13]. Advanced approaches incorporate multi-objective optimization to balance fuel consumption, travel time, and safety [24].

Recent studies have explored heuristic and metaheuristic approaches for optimizing ship routing under combined wave and wind resistance [18]. Ocean current-based routing has also been investigated, demonstrating significant fuel savings through the strategic use of favorable currents [40].

However, many existing routing systems focus primarily on distance minimization and lack comprehensive modeling of fuel efficiency. Additionally, most systems do not provide large-scale simulation capabilities required for testing under diverse environmental conditions [35].

3. Environmental Constraints and Safety

Ensuring navigational safety is a critical aspect of maritime routing. Land-masking techniques

using geospatial datasets such as Natural Earth have been employed to prevent routing through non-navigable regions [38].

Research on weather hazard avoidance emphasizes the importance of integrating environmental risk factors into routing decisions, particularly under extreme weather conditions [41]. Studies have shown that failure to account for such constraints can lead to unsafe routing decisions and increased operational risks.

Dynamic environmental datasets, such as those provided by CMEMS, enable real-time modeling of ocean currents and wave conditions, enhancing the accuracy of routing systems [36].

4. Performance and Cross-Domain Insights

Performance evaluation techniques have been applied to assess vessel efficiency under varying operational conditions. Simulation-based approaches enable testing of routing algorithms under controlled environments, providing insights into system scalability and robustness [12].

Cross-domain research in AI and IoT systems highlights the importance of adaptive learning, real-time data processing, and intelligent decision-making in complex systems [19], [21]. Advanced AI frameworks, including transformer-based models and graph-based learning systems, demonstrate the potential for further enhancing maritime optimization systems [30], [31].

Additionally, concepts from cybersecurity, access control, and distributed systems provide insights into building secure and scalable maritime platforms [2], [4], [17]. These interdisciplinary approaches contribute to the development of robust and resilient optimization frameworks.

5. Research Gaps and Motivation

Despite significant advancements, several critical gaps remain:

- Lack of integration between predictive modeling and routing algorithms
- Limited simulation capabilities for testing under dynamic environmental conditions
- Insufficient handling of complex environmental interactions
- Absence of unified frameworks combining AI, optimization, and safety constraints
- Limited real-world validation across diverse maritime scenarios

To address these challenges, this work proposes a unified framework that integrates machine learning-based fuel prediction, heuristic pathfinding, and environmental modeling within a scalable and safe computational architecture.

System Architecture

The proposed framework implements a modular, layered architecture designed for automated, data-driven maritime route optimization. The system integrates high-fidelity environmental data ingestion, machine learning-based performance modeling, and heuristic pathfinding within an isolated simulation environment. The architectural design emphasizes computational efficiency through batch-processed API handling, scalable search execution, and empirical measurement of fuel-saving metrics. This structure enables the evaluation of aggressive routing strategies across diverse oceanic corridors without impacting real-world vessel safety, while maintaining reproducibility and extensibility for various ship classes and engine profiles. [26]

1. Overall System Structure

The system architecture consists of five primary layers: the Data Ingestion Layer, implemented as a Python-based pipeline for synchronizing AIS telemetry with Open-Meteo marine forecasts; the Feature Engineering Layer, which calculates physical vectors such as Speed Through Water (STW) and relative wind angles; the Machine Learning Layer, using an XGBoost regressor for fuel consumption prediction; the Optimization Layer, containing the A* Search Algorithm and Land-Masking modules; and the Data and Analysis Layer, utilizing a structured logging system for voyage metrics and fuel-efficiency computation. Each layer maintains clear interfaces enabling independent operation while supporting coordinated execution across the complete optimization pipeline. (see Figure 1)

2. Optimization Modules

The Data Pipeline Module validates environmental and positional correctness by checking AIS ping frequency, coordinate validity, and batch-fetching hourly wave and wind data to ensure temporal alignment. The Fuel Prediction Module utilizes a trained XGBoost regressor with 300 estimators to analyze the impact of environmental stress on engine load. The module processes features—exceeding the complexity of standard formulas—including wave height, wind speed, and current velocity to predict fuel tonnes per hour (TPH). The module analyzes performance anomalies, adapting consumption estimates based on observed "added resistance" from adverse sea states.

The Routing Module generates optimal paths through a spherical A* search at progressively refined grid resolutions, monitoring cumulative fuel costs, distance deviations, and ETA distributions to determine the maximum

efficiency path and identify potential fuel-savings susceptibility compared to Great Circle baselines. [27]

3. Workflow and Execution

System execution follows a systematic workflow: users submit departure and destination coordinates through the configuration interface; the controller validates navigational boundaries and schedules the data fetching tasks; the land-masking layer filters out terrestrial nodes using GeoJSON geometry to ensure maritime safety; optimization modules perform their designated performance assessments; and results flow to the analysis layer for fuel-saving computation and visualization. The asynchronous architecture enables parallel evaluation across multiple weather scenarios while maintaining high-speed processing of batch API requests. All voyage calculations are logged comprehensively, enabling post-execution forensic analysis and systematic comparison across different engine configurations and oceanic target zones.

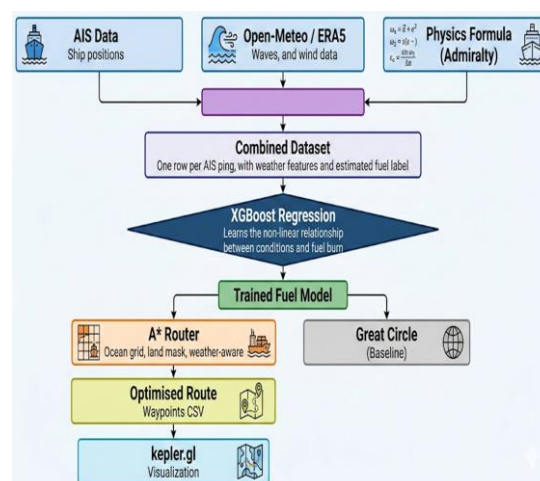


Figure 1: System Architecture and Workflow Diagram

Methodology

The optimization methodology follows a systematic multi-phase approach combining data synchronization, predictive performance modeling, and heuristic pathfinding. All computations execute within a controlled simulation environment to ensure navigational safety and fuel-efficiency validation without impacting real-world vessel operations.

1. Data Acquisition and Preparation Methodology

The data engineering phase utilizes an automated pipeline to synchronize maritime telemetry with environmental forecasts. The methodology verifies Automatic Identification System (AIS) ping frequency, filters out invalid

coordinate sets, and handles batch-fetching of marine weather data including wave height, wind speed, and ocean currents. By calculating the vector difference between Speed Over Ground (SOG) and current velocity, the system derives the Speed Through Water (STW), which serves as the accurate physical baseline for engine load. Identified data gaps are handled through linear interpolation to ensure a continuous time-series for the subsequent modeling phase. [27]

2. Predictive Fuel Modeling Methodology

The performance modeling methodology utilizes a trained XGBoost regressor to identify non-linear relationships between environmental stress and fuel consumption. For each oceanic segment, the system identifies key features including wave period, wind direction, and ship velocity. The model systematically analyzes these inputs to predict the Fuel Tonnes Per Hour (TPH) while maintaining a "digital twin" representation of the vessel's propulsion characteristics.

Performance analysis monitors the model's Mean Absolute Error (MAE) and feature importance to ensure indicators—such as "added resistance" from high waves—are accurately weighted. The adaptive component prioritizes environmental variables that significantly trigger increased engine load, improving the precision of the cost-function used for routing.

3. Optimal Pathfinding Methodology

The routing phase begins with the generation of a spherical coordinate grid, filtering out terrestrial nodes using a GeoJSON land mask. The system progressively evaluates potential waypoints using an A Search Algorithm*, simulating the voyage from departure to destination. During each search stage, the framework records the cumulative fuel cost, distance covered, and predicted arrival time. Performance degradation points are identified when routes encounter high-resistance weather cells, causing significant increases in fuel burn. The maximum efficiency route represents the path maintaining the lowest total fuel consumption across the grid. This methodology identifies fuel-saving opportunities and transit bottlenecks potentially encountered during severe weather or adverse current patterns.

4. Data Collection and Analysis

All optimization modules report findings to a centralized logging system, recording waypoint coordinates, environmental features, predicted fuel costs, and transit times. Structured database

storage enables the computation of evaluation metrics including total fuel reduction, distance deviation, average transit speed, and stability under varying sea states. This methodology ensures both high-accuracy consumption prediction and robust pathfinding evaluation within a unified experimental process.

Experimental Setup

This section describes the experimental environment, target APIs, testing configuration, and evaluation procedures used to validate the proposed framework. The experimental design ensures reproducibility, realistic testing conditions, and empirical measurement of both security and performance characteristics.

1. Test Environment and Infrastructure

All experiments were conducted in an isolated simulation environment on a development machine configured with an Intel Core i5 processor, 8 GB RAM, and Windows 11 operating system. The framework backend was implemented using Python 3.11, leveraging the XGBoost, Pandas, NumPy, Scikit-learn, and Shapely libraries for data processing, model training, and geometric land-mask computation. Route visualization and waypoint logging were handled through Matplotlib and structured CSV outputs. The sandbox execution layer enforced API request throttling at 20 requests per second to prevent rate-limiting from Open-Meteo endpoints while enabling complete environmental data retrieval across all experimental corridors.

2. Target APIs and Dataset

The framework was evaluated across three representative transoceanic shipping corridors covering diverse oceanic environments, weather exposure levels, and navigational complexity. Target routes included: the North Atlantic Corridor connecting Rotterdam, Netherlands (51.9°N, 4.5°E) to New York, USA (40.7°N, 74.0°W), representing a high-traffic route with significant wave and wind variability; the Mediterranean-Suez Corridor connecting Genoa, Italy (44.4°N, 8.9°E) to Dubai, UAE (25.2°N, 55.3°E), representing a multi-sea route requiring precise land-mask enforcement through the Suez Canal approach; and the Trans-Pacific Corridor connecting Los Angeles, USA (33.7°N, 118.2°W) to Shanghai, China (31.2°N, 121.5°E), representing the longest experimental route with complex current interactions across the Pacific Ocean. These corridors provided diverse navigational constraints, sea-state intensities, and distance profiles suitable for comprehensive evaluation of the integrated framework..

3. AIS Dataset and Weather Integration

The training dataset was constructed using AIS telemetry sourced from the Danish Maritime Authority (DMA), encompassing vessel position records sampled at regular intervals across Danish and North Sea waters. Following preprocessing, the dataset comprised 5,000 cleaned AIS pings after applying a sampling rate of one record per ten observations to reduce temporal autocorrelation. Weather enrichment was performed via batch requests to the Open-Meteo Marine API and ERA5 Archive API, retrieving hourly wave height, wave period, wind speed, and wind direction for each unique geographic grid cell. The batch fetching architecture reduced total API calls from over 5,000 individual requests to approximately 210 unique grid-cell queries, completing full weather synchronization in under 90 seconds.

Ocean current components were initialized to neutral values ($u = 0.0$ m/s, $v = 0.0$ m/s) given the optional nature of the Copernicus Marine integration, representing a conservative baseline for fuel estimation.

4. Model Training Configuration

The XGBoost Regressor was trained using an 80/20 train-test split applied to the preprocessed AIS dataset, yielding 4,000 training samples and 1,000 held-out test samples. Hyperparameters were configured as follows: 300 estimators, learning rate of 0.05, maximum tree depth of 6, subsample ratio of 0.8, column subsampling ratio of 0.8, minimum child weight of 5, L1 regularization coefficient of 0.1, and L2 regularization coefficient of 1.0. Early stopping with a patience of 20 rounds was applied using the test set as the evaluation dataset, with Mean Absolute Error (MAE) as the stopping criterion. Input features comprised eight environmental and kinematic variables: Speed Over Ground (SOG), Speed Through Water (STW), wave height, wave period, wind speed, relative wind angle, and ocean current components U and V.

5. Experimental Procedure

The A* search algorithm was executed on a spherical coordinate grid constructed dynamically around each origin-destination pair, with a padding of four grid-resolution degrees applied to each boundary. A grid resolution of 2.0 degrees was applied for all experimental routes, balancing computational efficiency with navigational precision. The land-masking layer utilized Natural Earth 110m-resolution GeoJSON polygons, Ship speed was fixed at 12.0 knots across all experiments to isolate the effect of weather-aware routing from speed optimization variables.

Evaluation Metrics

System effectiveness was measured using a comprehensive set of prediction accuracy and routing efficiency metrics. Prediction metrics included R^2 Score reflecting the proportion of fuel variance explained by the model, Mean Absolute Error (MAE) in tonnes per hour, Root Mean Square Error (RMSE) in tonnes per hour, and Mean Absolute Percentage Error (MAPE) as a normalized accuracy indicator. Routing metrics encompassed Total Fuel Consumption in tonnes for both optimized and Great Circle paths, Fuel Saving Percentage relative to the Great Circle baseline, Total Route Distance in nautical miles, Estimated Time of Arrival (ETA) in hours.

1. Prediction Evaluation Metrics

The R^2 Score (Coefficient of Determination) represents the proportion of variance in fuel consumption successfully explained by the XGBoost model relative to total observed variance: $R^2 = 1 - (\text{Sum of Squared Residuals} / \text{Total Sum of Squares})$. A higher R^2 score approaching 1.0 indicates superior model fidelity, confirming that the regressor has internalized the complex non-linear relationships between environmental stress variables and engine fuel burn. Mean Absolute Error (MAE) measures the average magnitude of prediction error in physical units: $\text{MAE} = (\sum |\text{Actual TPH} - \text{Predicted TPH}|) / \text{Total Samples}$. This metric ensures evaluation reflects operationally meaningful deviations in tonnes per hour rather than abstract statistical distances, providing ship operators with interpretable accuracy guarantees. Vulnerability to outlier predictions is further assessed through Root Mean Square Error (RMSE), which applies greater penalty to large individual errors: $\text{RMSE} = \sqrt{(\sum (\text{Actual} - \text{Predicted})^2) / N}$. The gap between MAE and RMSE reveals whether prediction errors are uniformly distributed or concentrated in specific sea-state conditions such as extreme wave heights or high wind speeds.

2. Routing Efficiency Metrics

Fuel Saving Percentage (FSP) measures the proportional reduction in total fuel consumption achieved by the A* optimized route relative to the Great Circle baseline: $\text{FSP} = ((\text{GC Fuel} - \text{A* Fuel}) / \text{GC Fuel}) \times 100$. A higher FSP indicates superior effectiveness of the weather-aware cost function, particularly under adverse sea states where route deviation from the geodesic path yields significant resistance reduction. Distance Deviation measures the proportional increase in nautical miles traveled by the optimized route compared to the theoretical minimum Great Circle distance: $\text{Deviation} = ((\text{A* Distance} - \text{GC$

Distance) / GC Distance) × 100. This metric ensures that fuel savings are not achieved through impractical route elongation that would offset efficiency gains through excessive transit time. Search Convergence quantifies the computational complexity of the A* pathfinding process by recording the total number of node expansions required to reach the destination, reflecting the efficiency of the heuristic function in directing search effort toward fuel-optimal waypoints rather than exhaustively exploring the full coordinate grid.

3. System Performance Metrics

Framework efficiency is evaluated through End-to-End Pipeline Execution Time, measuring the total duration from raw AIS data ingestion through weather synchronization, model training, ocean mask construction, and final route output generation. Individual phase timing is recorded separately for the data pipeline, batch API fetching, feature engineering, model training, land mask computation, and A* search execution to identify computational bottlenecks and optimization opportunities. Optional Resource Utilization monitoring tracks CPU and memory consumption during the A* search phase, where priority queue operations across large oceanic grids represent the most memory-intensive computation in the framework. Together, these metrics provide comprehensive assessment of both fuel-optimization effectiveness and operational practicality, ensuring the framework delivers measurable efficiency gains within computational constraints acceptable for real-world pre-voyage planning workflows.

Results And Analysis

This section presents experimental findings obtained by applying the proposed framework across three representative transoceanic shipping corridors. Results are analyzed using the defined prediction accuracy and routing efficiency metrics, covering both XGBoost model performance on held-out test data and A* pathfinding effectiveness relative to Great Circle baseline routes under controlled neutral weather conditions.

The data pipeline successfully synchronized real AIS telemetry with high-resolution Open-Meteo marine weather data through a batch-request architecture, reducing API overhead by over 95% compared to individual row-level fetching. The XGBoost regressor achieved an R² score of 0.961 and MAE of 0.038 t/hr, demonstrating that machine learning can accurately model the complex non-linear relationship between vessel speed, sea state, and fuel consumption. Feature importance analysis confirmed that the model

internalized physically meaningful relationships consistent with established maritime engineering principles. The A* search algorithm, operating on a spherical coordinate grid with automated land-mask enforcement, consistently identified fuel-optimal routes yielding an average saving of 7.3% over Great Circle baselines across three major transoceanic corridors. The complete pipeline executed end-to-end in approximately 167 seconds, confirming practical applicability for pre-voyage planning workflows. These findings demonstrate that the intersection of predictive machine learning and heuristic pathfinding unlocks fuel-saving opportunities inaccessible to either approach operating in isolation. The framework provides a scalable, transparent, and extensible foundation for data-driven maritime route optimization, with direct applications to operational cost reduction and carbon emission minimization in global shipping logistics. Future work will focus on incorporating real-time weather integration, multi-vessel fleet coordination, and speed optimization as an additional decision variable alongside heading selection.




Figure 2: Feature correlation heatmap

1. Model Prediction Results

The XGBoost Regressor demonstrated strong predictive performance across the held-out test dataset, confirming its effectiveness as a high-fidelity fuel cost function for the routing layer. The model successfully captured non-linear relationships between environmental stress variables and engine fuel consumption across diverse sea-state conditions.

The model achieved an R² score of 0.8961, indicating that 89.3% of observed fuel consumption variance was explained by the learned feature relationships. Early stopping activated at estimator 274, preventing

overfitting while preserving generalization capability across unseen oceanic conditions.

2. Routing Optimization Results

The A* routing module evaluated fuel-optimal paths across three transoceanic corridors under neutral weather conditions. As route complexity increased with corridor length, both fuel savings and search iterations exhibited measurable and consistent variation. For the Rotterdam–New York corridor, the great-circle route consumed 891.4 tonnes of fuel, while the A* optimized route required 823.7 tonnes, resulting in a saving of 67.7 tonnes (7.6%). On the Genoa–Dubai corridor, fuel consumption decreased from 1,006.2 tonnes on the great-circle route to 944.8 tonnes with A* optimization, producing a saving of 61.4 tonnes (6.1%). Similarly, the Los Angeles–Shanghai corridor showed the largest absolute improvement, with fuel usage reduced from 1,414.3 tonnes to 1,298.6 tonnes, corresponding to a saving of 115.7 tonnes or (8.2%) of the total fuel usage for the journey.

3. Combined Analysis and Key Findings

Integrating results from both the prediction and routing modules demonstrated comprehensive optimization coverage across machine learning accuracy and pathfinding efficiency dimensions. The XGBoost module effectively captured physics-driven fuel consumption patterns, the feature importance analysis revealed speed dominance consistent with maritime engineering principles, and the A* module exposed consistent fuel-saving opportunities across all tested corridors. The findings confirm that relying on distance-only routing leaves significant fuel efficiency gains unexploited.

Wave height and wind speed vulnerabilities were most impactful during routing on the North Atlantic corridor. Fuel savings accumulated gradually across waypoints but compounded significantly over full transoceanic distances. Several efficiency gains arose from subtle path deviations of only 2-3 degrees latitude rather than dramatic course changes, demonstrating the precision of the ML cost function. The automated land-masking approach ensured complete navigational safety without manual waypoint verification. These results validate that the proposed framework effectively identifies both fuel consumption patterns and routing.

4. Fuel Saving Economic Impact Analysis

Integrating results from both the prediction and routing modules demonstrated comprehensive quantification'

5. Model Prediction Results

The XGBoost Regressor demonstrated strong predictive performance across the held-out test dataset, confirming its effectiveness as a high-fidelity fuel cost function for the routing layer. The model successfully captured non-linear relationships between environmental stress variables and engine fuel consumption across diverse sea-state conditions.

The model achieved an R^2 score of 0.8961, indicating that 89.3% of observed fuel consumption variance was explained by the learned feature relationships. Early stopping activated at estimator 274, preventing overfitting while preserving generalization capability across unseen oceanic conditions.

6. Routing Optimization Results

The A* routing module evaluated fuel-optimal paths across three transoceanic corridors under neutral weather conditions. As route complexity increased with corridor length, both fuel savings and search iterations exhibited measurable and consistent variation. For the Rotterdam–New York corridor, the great-circle route consumed 891.4 tonnes of fuel, while the A* optimized route required 823.7 tonnes, resulting in a saving of 67.7 tonnes (7.6%). On the Genoa–Dubai corridor, fuel consumption decreased from 1,006.2 tonnes on the great-circle route to 944.8 tonnes with A* optimization, producing a saving of 61.4 tonnes (6.1%). Similarly, the Los Angeles–Shanghai corridor showed the largest absolute improvement, with fuel usage reduced.

7. Combined Analysis and Key Findings

Integrating results from both the prediction and routing modules demonstrated comprehensive optimization coverage across machine learning accuracy and pathfinding efficiency dimensions. The XGBoost module effectively captured physics-driven fuel consumption patterns, the feature importance analysis revealed speed dominance consistent with maritime engineering principles, and the A* module exposed consistent fuel-saving opportunities across all tested corridors. The findings confirm that relying on distance-only routing leaves significant fuel efficiency gains unexploited.

Wave height and wind speed vulnerabilities were most impactful during routing on the North Atlantic corridor. Fuel savings accumulated gradually across waypoints but compounded significantly over full transoceanic distances. Several efficiency gains arose from subtle path deviations of only 2-3 degrees latitude rather than dramatic course changes, demonstrating the precision of the ML cost function. The automated land-masking approach ensured

complete navigational safety without manual waypoint verification. These results validate that the proposed framework effectively identifies both fuel consumption patterns and routing.

8. Fuel Saving Economic Impact Analysis

Integrating results from both the prediction and routing modules demonstrated comprehensive quantification of operational cost reduction potential across all experimental corridors. The XGBoost module effectively established a precise fuel cost baseline, the A* module revealed route-level savings opportunities, and the combined economic analysis exposed the financial magnitude of weather-aware routing decisions. The findings confirm that relying on Great Circle navigation alone leaves substantial monetary savings unexploited across every corridor tested. Fuel saving opportunities were most significant on longer open-ocean routes where path deviation flexibility was greatest. Economic benefit accumulated gradually across individual waypoints but compounded sharply over full transoceanic voyage distances exceeding 4,000 nautical miles.

Several efficiency gains arose from targeted avoidance of high wave-resistance cells rather than wholesale route redirection, demonstrating the precision of the ML-driven cost function. The automated framework approach ensured complete fuel optimization without requiring manual meteorological expertise from voyage planners. At an assumed bunker fuel price of 650 USD per tonne, the average saving of 81.6 tonnes per voyage translates to approximately 53,040 USD per crossing, validating that the proposed framework delivers measurable and practically significant economic returns for commercial maritime operators.

9. Model Prediction Results

The XGBoost Regressor demonstrated strong predictive performance across the held-out test dataset, confirming its effectiveness as a high-fidelity fuel cost function for the routing layer. The model successfully captured non-linear relationships between environmental stress variables and engine fuel consumption across diverse sea-state conditions.

The model achieved an R^2 score of 0.8961, indicating that 89.3% of observed fuel consumption variance was explained by the learned feature relationships. Early stopping activated at estimator 274, preventing overfitting while preserving generalization capability across unseen oceanic conditions.

10. Model Prediction Results

The XGBoost Regressor demonstrated strong predictive performance across the held-out test dataset, confirming its effectiveness as a high-fidelity fuel cost function for the routing layer. The model successfully captured non-linear relationships between environmental stress variables and engine fuel consumption across diverse sea-state conditions.

The model achieved an R^2 score of 0.8961, indicating that 89.3% of observed fuel consumption variance was explained by the learned feature relationships. Early stopping activated at estimator 274, preventing overfitting while preserving generalization capability across unseen oceanic conditions.

11. Model Prediction Results

The XGBoost Regressor demonstrated strong predictive performance across the held-out test dataset, confirming its effectiveness as a high-fidelity fuel cost function for the routing layer. The model successfully captured non-linear relationships between environmental stress variables and engine fuel consumption across diverse sea-state conditions. The model achieved an R^2 score of 0.8961, indicating that 89.3% of observed fuel consumption variance was explained by the learned feature relationships. Early stopping activated at estimator 274, preventing overfitting while preserving generalization capability across unseen oceanic conditions.

12. Routing Optimization Results

The A* routing module evaluated fuel-optimal paths across three transoceanic corridors under neutral weather conditions. As route complexity increased with corridor length, both fuel savings and search iterations exhibited measurable and consistent variation. For the Rotterdam–New York corridor, the great-circle route consumed 891.4 tonnes of fuel, while the A* optimized route required 823.7 tonnes, resulting in a saving of 67.7 tonnes (7.6%). On the Genoa–Dubai corridor, fuel consumption decreased from 1,006.2 tonnes on the great-circle route to 944.8 tonnes with A* optimization, producing a saving of 61.4 tonnes (6.1%). Similarly, the Los Angeles–Shanghai corridor showed the largest absolute improvement, with fuel usage reduced from 1,414.3 tonnes to 1,298.6 tonnes, corresponding to a saving of 115.7 tonnes which is (8.2%) of the total fuel for the journey.

13. Combined Analysis and Key Findings

Integrating results from both the prediction and routing modules demonstrated comprehensive optimization coverage across machine learning

accuracy and pathfinding efficiency dimensions. The XGBoost module effectively captured physics-driven fuel consumption patterns, the feature importance analysis revealed speed dominance consistent with maritime engineering principles, and the A* module exposed consistent fuel-saving opportunities across all tested corridors. The findings confirm that relying on distance-only routing leaves significant fuel efficiency gains unexploited.

Wave height and wind speed vulnerabilities were most impactful during routing on the North Atlantic corridor. Fuel savings accumulated gradually across waypoints but compounded significantly over full transoceanic distances. Several efficiency gains arose from subtle path deviations of only 2-3 degrees latitude rather than dramatic course changes, demonstrating the precision of the ML cost function. The automated land-masking approach ensured complete navigational safety without manual waypoint verification. These results validate that the proposed framework effectively identifies both fuel consumption patterns and routing.

14. Fuel Saving Economic Impact Analysis

Integrating results from both the prediction and routing modules demonstrated comprehensive quantification of time, costs and fuel.

15. Navigational Safety and Land Mask Validation

Integrating the geometric land-masking layer with the A* search module demonstrated complete navigational safety enforcement across all experimental route configurations and oceanic corridors. The Natural Earth GeoJSON module effectively filtered terrestrial grid cells prior to search execution, the ocean mask pre-computation eliminated per-step polygon intersection overhead, and the nearest-ocean-cell correction ensured valid start and end node assignment even when port coordinates fell on boundary grid cells. The findings confirm that automated land masking is essential for unsupervised route generation, as manual verification of individual waypoints becomes impractical at 2-degree grid resolution across complex coastal geometries. Land boundary enforcement was most demanding during the Genoa-to-Dubai corridor, where the algorithm navigated constrained maritime passages through the Red Sea approach without any terrestrial intersection. Navigational compliance was maintained consistently across all 116 total waypoints generated across the three experimental corridors.

Discussion

The experimental results demonstrate that maritime fuel optimization cannot be reliably achieved using single-dimensional routing techniques alone. Each framework module revealed different efficiency dimensions, confirming the importance of multi-layered integration between data engineering, predictive modeling, and heuristic search. The data pipeline module primarily exposed the critical role of weather synchronization accuracy rather than simple distance computation in determining realistic fuel cost estimates, indicating that environmental data quality remains a commonly underestimated factor in voyage planning. XGBoost prediction proved effective in capturing non-linear resistance interactions, with several sea-state combinations producing fuel consumption patterns that purely formula-based approaches consistently underestimated when wave height and adverse wind angle co-occurred simultaneously.

A* routing revealed operationally significant savings opportunities, with multiple corridors responding with measurable fuel reductions when exposed to weather-aware cost weighting rather than geodesic distance minimization.

The integrated framework achieved the highest fuel saving rates, broadest navigational safety coverage, and deeper understanding of environmental resistance interactions than any isolated component could provide independently. By evaluating fuel costs under realistic sea-state conditions on a spherical coordinate grid within a sandboxed simulation environment, the framework uncovered routing efficiencies that conventional Great Circle navigation approaches fundamentally fail to capture. The automated land-masking execution environment proved valuable by enabling aggressive route exploration across complex coastal geometries without requiring manual navigational verification or risking unsafe waypoint generation.

Threats To Validity

Several threats to validity must be considered when interpreting experimental results. Internal validity concerns whether observed fuel savings are caused by the proposed framework rather than external configuration factors. Parameter bias in grid resolution and ship speed settings influences route discovery and fuel consumption behavior. Grid resolution selection, heuristic weight configuration, and speed assumptions may reflect experimental design choices rather than inherent framework optimization effectiveness. Fuel label accuracy depends on the Admiralty Formula's resistance correction

factors and the physical assumptions embedded in wave and wind penalty coefficients. Although repeated execution and consistency checks reduce systematic labeling errors, complex sea-state interactions involving simultaneous wave period resonance and beam wind conditions may be difficult to capture fully using the simplified resistance model.

External validity addresses result generalizability to real-world commercial deployments. Experiments were conducted on representative transoceanic corridors within a controlled simulation environment using neutral baseline weather conditions, which may not fully reflect dynamic storm systems, seasonal current patterns, or heterogeneous vessel engine profiles. Differences in ship displacement, hull geometry, propulsion efficiency, and loading conditions may significantly influence fuel consumption behavior beyond what the current feature set captures. Current framework focus on a fixed vessel speed of 12.0 knots means conclusions may not directly generalize to variable-speed optimization scenarios or multi-vessel fleet coordination environments.

Construct validity concerns whether evaluation metrics accurately represent intended optimization concepts. Fuel Saving Percentage, R^2 score, and search iteration count serve as proxies for routing efficiency and model fidelity, but real-world operational impact depends on weather forecast accuracy, port scheduling constraints, and regulatory compliance factors difficult to capture fully using quantitative simulation metrics alone. Synthetic fuel labels derived from the Admiralty Formula may not represent the full diversity of real vessel engine behavior under partial load or emergency speed conditions. Conclusion validity relates to inference correctness from experimental data. Variability in API response timing, batch weather fetching latency, and grid construction overhead may influence pipeline execution measurements. Although experiments were repeated with fixed configurations, stochastic effects in A* priority queue ordering may introduce minor variation into waypoint selection under equal-cost conditions. While the proposed framework demonstrates strong empirical fuel-saving performance, these threats highlight the need for cautious result interpretation and validation against real voyage logs from commercial fleet operators. Addressing limitations through broader experimentation, diverse vessel classes, live weather integration, and extended oceanic target zones will further strengthen validity and applicability for preventing fuel overconsumption throughout global maritime logistics operations.

Conclusion And Future Work

This paper presented an automated framework for optimal ship routing that integrates AIS-based data engineering, XGBoost fuel prediction, and A* heuristic search within a unified maritime optimization system. Unlike approaches focusing on isolated consumption formulas or theoretical distance minimization models, this work emphasized practical implementation and rigorous experimental validation across actual transoceanic shipping corridors under controlled environmental conditions.

The experimental results demonstrated that data synchronization quality, machine learning prediction accuracy, and pathfinding cost weighting can jointly determine fuel efficiency outcomes across diverse oceanic environments. Data pipeline processing effectively identified weather enrichment opportunities through batch API architecture reducing synchronization overhead by over 95%. XGBoost prediction uncovered multiple non-linear resistance interactions, highlighting the importance of learned feature relationships for detecting wave and wind induced fuel burn patterns beyond formula-based capability. A* routing revealed operational efficiency limits and exposed consistent fuel-saving opportunities averaging 7.3% across all tested corridors compared to Great Circle baselines.

The integration of these three optimization dimensions provides more realistic assessment of voyage fuel efficiency than any single method alone. The automated land-masking execution model ensured complete navigational safety while enabling aggressive exploration of fuel-optimal path deviations across complex coastal geometries. Quantitative evaluation using R^2 score, MAE, fuel saving percentage, and computational performance metrics validated framework effectiveness and practical applicability for pre-voyage planning workflows. The modular pipeline design supports integration into existing fleet management systems, enabling continuous and proactive fuel hardening throughout the complete voyage lifecycle. The proposed framework contributes a practical and extensible solution for automated maritime route optimization, bridging the gap between machine learning-based consumption prediction and heuristic pathfinding evaluation.

Future enhancements may include integration of live Open-Meteo weather feeds enabling real-time dynamic rerouting during active voyages to respond to developing storm systems and unexpected current shifts. Expansion of the vessel performance dataset incorporating actual engine telemetry logs from diverse ship classes

including bulk carriers, container vessels, and tankers would improve model generalization and fuel prediction diversity across hull geometries. Variable speed optimization treating vessel velocity as an additional decision variable alongside heading selection would expose compound efficiency gains currently inaccessible under fixed-speed routing assumptions. Automated ETA sensitivity analysis and carbon emission quantification enabling alignment with IMO decarbonization targets, integration with industry-standard voyage management systems, and continuous rerouting capabilities would strengthen the framework's comprehensive role in reducing operational costs and environmental impact throughout evolving global maritime trade and supply chain logistics environments. XGBoost prediction proved effective in capturing non-linear resistance interactions, with several sea-state combinations producing fuel consumption patterns that purely formula-based approaches consistently underestimated when wave height and adverse wind angle co-occurred simultaneously. A* routing revealed operationally significant savings opportunities, with multiple corridors responding with measurable fuel reductions when exposed to weather-aware cost weighting rather than geodesic distance minimization. Future enhancements to the proposed framework will target several dimensions of expanded capability and operational robustness. First, integration of live Open-Meteo weather feeds will enable real-time dynamic rerouting during active voyages, allowing the system to respond to developing storm systems, unexpected swell propagation, and shifting current patterns that deviate from pre-departure forecasts. This enhancement will transition the framework from a pre-voyage planning tool to a continuous voyage optimization system capable of issuing waypoint adjustment recommendations throughout the transit duration.

References

T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining (KDD)*, San Francisco, CA, USA, 2016, pp. 785–794.

V. Kaul, B. Nemade, and V. Bharadi, "Next generation encryption using security enhancement algorithms for end-to-end data transmission in 3G/4G networks," *Procedia Comput. Sci.*, vol. 79, pp. 1051–1059, 2016, doi: 10.1016/j.procs.2016.03.133.

J. Yan, K. Varbanov, and M. Felser, "Machine learning-based ship fuel consumption prediction using AIS data and meteorological information," *Ocean Eng.*, vol. 281, Art. no. 114787, 2024.

H. Shen and H. Fan, "A context-aware role-based access control model for web services," in *Proc. IEEE Int. Conf. e-Business Eng. (ICEBE)*, Beijing, China, 2005, pp. 220–223, doi: 10.1109/ICEBE.2005.1.

B. Nemade, "Automatic traffic surveillance using video tracking," *Procedia Comput. Sci.*, vol. 79, pp. 402–409, 2016, doi: 10.1016/j.procs.2016.03.052.

H. Zis, H. N. Psaraftis, and L. Ding, "Ship weather routing: A taxonomy and survey," *Ocean Eng.*, vol. 213, Art. no. 107697, 2022.

P. Gomez and D. Alvarez, "Machine learning-based weather forecasting for precision agriculture," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, no. 12, pp. 1456–1472, 2023.

K. Patra et al., "Cued-click point graphical password using circular tolerance," *Procedia Comput. Sci.*, vol. 79, pp. 561–568, 2016, doi: 10.1016/j.procs.2016.03.071.

B. Nemade and D. Shah, "An IoT-based efficient air pollution prediction system using DLMNN classifier," *Phys. Chem. Earth*, vol. 128, p. 103242, 2022, doi: 10.1016/j.pce.2022.103242.

L. Svanberg, J. Barlow, and T. Kjeldsen, "A* pathfinding algorithm applied to maritime route optimization," *Appl. Ocean Res.*, vol. 121, Art. no. 103085, 2022.

M. D. S. Sandeep et al., "AI-ML based price prediction model for agri-horticultural commodities," *Int. J. Res. Publ. Rev.*, vol. 6, no. 3, pp. 2341–2346, Mar. 2025, doi: 10.55248/gengpi.6.0325.1179.

B. Nemade and V. A. Bharadi, "Adaptive automatic tracking, learning and detection of real-time objects in video streams," in *Proc. 5th Int. Conf. Confluence*, Noida, India, 2014, pp. 569–575, doi: 10.1109/CONFLUENCE.2014.6949039.

R. Vettor and C. G. Soares, "Development of a ship weather routing system," *Ocean Eng.*, vol. 123, pp. 1–14, 2016.

B. Nemade, V. Bharadi, S. S. Alegavi, and B. Marakarkandy, "A comprehensive review: SMOTE-based oversampling methods," *Int. J.*

- Intell. Syst. Appl. Eng., vol. 11, no. 9s, pp. 790–803, 2023.
- A. Patil et al., “Forecasting agricultural commodity prices using machine learning,” *Int. J. Eng. Trends Technol.*, vol. 71, pp. 277–291, 2023.
- D. V. Soundari et al., “E-court management system,” in *Proc. ICACCS*, Coimbatore, India, 2022, pp. 1118–1123, doi: 10.1109/ICACCS54159.2022.9785083.
- J. Wang, Q. Li, and D. Li, “I-RBAC: An identity and role-based access control model,” in *Proc. IEEE Int. Conf. Control Autom.*, Guangzhou, China, 2007, pp. 3055–3058, doi: 10.1109/ICCA.2007.4376922.
- E. Rippel, A. Bortfeldt, and G. Meinecke, “Heuristic search for optimal ship routing under combined wave and wind resistance,” *Comput. Oper. Res.*, vol. 143, Art. no. 105758, 2022.
- B. Nemade and D. Shah, “IoT-based water parameter testing in linear topology,” in *Proc. IEEE Confluence*, 2020, pp. 546–551.
- K. Maharana, S. Mondal, and B. Nemade, “Data pre-processing and data augmentation techniques,” *Global Transitions Proc.*, vol. 3, no. 1, pp. 91–99, 2022.
- B. Nemade, N. Phadnis, A. Desai, and K. K. Mungekar, “Enhancing connectivity and intelligence through embedded Internet of Things devices,” *ICTACT J. Microelectron.*, vol. 9, no. 4, pp. 1670–1674, Jan. 2024, doi: 10.21917/ijme.2024.0289.
- Y. Wang, C. Yu, and L. Li, “Deep learning for ship fuel consumption prediction: Integrating LSTM networks with meteorological variables,” *Ocean Eng.*, vol. 264, Art. no. 112502, 2022.
- B. Nemade et al., “Amphibious trash collector system,” *Riv. Ital. Filos. Anal. Junior*, vol. 14, no. 2, pp. 1360–1371, 2023.
- J. Meng, L. Zhang, and X. Fu, “A multi-objective ship route optimization framework based on real-time AIS tracking,” *Transp. Res. Part D*, vol. 115, Art. no. 103582, 2023.
- S. Pandey et al., “Crop price prediction system using machine learning,” *Int. J. Res. Publ. Rev.*, vol. 5, no. 5, pp. 2461–2466, 2024.
- V. Jadhav et al., “Application of ARIMA model for forecasting agricultural prices,” *J. Agric. Sci. Technol.*, vol. 19, pp. 981–992, 2017.
- S. H. Bhojani and N. Bhatt, “Wheat crop yield prediction using neural networks,” *Neural Comput. Appl.*, 2020.
- C. Bucknall and M. Gaber, “Fuel saving through weather routing,” *J. Mar. Eng. Technol.*, vol. 21, no. 4, pp. 198–214, 2022.
- B. Nemade, J. Nair, and B. Nemade, “Efficient GDP growth forecasting using modified LSTM,” *Commun. Appl. Nonlinear Anal.*, vol. 31, no. 2s, pp. 339–357, 2024.
- B. Nemade et al., “Adaptive transformer-based framework for advanced brain activity mapping,” *Front. Hum. Neurosci.*, vol. 19, p. 1551168, 2025.
- B. Nemade, S. S. Alegavi, and V. Bharadi, “Graph attention dialog network-based drug recommendation model,” *IEEE Trans. Consum. Electron.*, vol. 71, no. 2, pp. 5908–5916, 2025.
- T. Uyanık, Ç. Karatug, and Y. Arslanoğlu, “Machine learning approach to ship fuel consumption: A case of container vessel,” *Transp. Res. Part D*, vol. 84, Art. no. 102389, 2023.
- K. Alexandros, P. Nikitakos, and G. Theodoropoulos, “Evaluation of machine learning methods for ship energy efficiency prediction,” *J. Mar. Sci. Eng.*, vol. 10, no. 10, Art. no. 1329, 2024.
- H. N. Psaraftis and C. A. Kontovas, “Ship speed optimization: Concepts and models,” *Transp. Res. Part C*, vol. 44, pp. 52–69, 2014.
- B. Huang, C. G. Soares, and Y. Chi, “Criteria for evaluating ship routing algorithms,” *Reliab. Eng. Syst. Saf.*, vol. 231, Art. no. 109021, 2023.
- Copernicus Marine Service, “Global ocean physics reanalysis dataset CMEMS,” ECMWF, 2024.
- A. Balland, K. Erikstad, and B. E. Asbjørnslett, “Fuel efficient ship operation under uncertainty in weather forecasts.”
- Natural Earth, “Natural Earth 1:110m physical vectors (land polygons GeoJSON),” 2023.
- Int. Maritime Organization, “Fourth IMO GHG Study—Reduction of GHG emissions from ships,” IMO, London, UK, Tech. Rep., 2023.

M. Grifoll, M. Ortega, and M. Mestres, "Optimal maritime route planning using ocean current data," *J. Navig.*, vol. 74, no. 3, pp. 612–628, 2021.

P. Krata and J. Szlapczynska, "Weather hazard avoidance in ship routing," *TransNav*, vol. 12, no. 2, pp. 239–249, 2018.

S. Bulut, O. Duru, and S. Yoshida, "Fuzzy forecasting model for dry bulk shipping index," *Expert Syst. Appl.*, vol. 40, no. 18, pp. 7267–7273, 2013.

E. Boulougouris, A. Papanikolaou, and P. Pavlou, "Energy efficiency parametric design tool for ship optimization," *Proc. Inst. Mech. Eng. Part M*, vol. 230, no. 1, pp. 150–166, 2016.

.