DATA-DRIVEN OPTIMIZATION FOR SUSTAINABLE SHIPPING: ADDRESSING WEATHER UNCERTAINTY, ENVIRONMENTAL DYNAMISM, AND HYBRID PROPULSION COMPLEXITY

XI LUO

PhD

The Hong Kong Polytechnic University

2026

The Hong Kong Polytechnic University
Department of Logistics and Maritime Studies

Data-Driven Optimization for Sustainable Shipping: Addressing Weather Uncertainty, Environmental Dynamism, and Hybrid Propulsion Complexity

Xi Luo

A thesis submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy
October 2025

CERTIFICATE OF ORIGINALITY

I hereby declare that this thesis is my own work and that, to the best of my knowledge and belief, it reproduces no material previously published or written, nor material that has been accepted for the award of any other degree or diploma, except where due acknowledgment has been made in the text.

Signature: ____________________

Name of Student: Xi Luo

Abstract

The global maritime industry confronts mounting pressures from stringent emission reduction regulations and escalating fuel costs that constitute a substantial portion of operational expenses. Compared to capital-intensive technological retrofits, operational optimization through speed planning offers a cost-effective and readily implementable pathway for achieving energy efficiency improvements and emission reductions. Ship speed optimization aims to minimize a ship's fuel consumption by determining optimal speed profiles along planned routes, typically following a predict-then-optimize paradigm where fuel consumption is first predicted and then used to optimize speed decisions. Since fuel consumption prediction relies on weather forecasts as critical input variables, and sea and weather conditions exhibit considerable variability during voyages, accurate weather forecasts become essential for effective speed optimization.

However, existing speed optimization methods encounter two challenges that constrain their practical effectiveness. First, current models predominantly rely on deterministic weather forecasts while neglecting inherent forecast uncertainties. These approaches lack scientific frameworks to quantify the operational value of probabilistic forecast information. Second, static pre-voyage planning fails to capitalize on continuously updated weather forecasts available throughout voyages, resulting in suboptimal performance under evolving environmental conditions. Furthermore, while operational optimization offers immediate efficiency gains, achieving deeper

decarbonization targets requires technological innovation. Emerging hybrid electric propulsion systems, which integrate battery storage with conventional engines, introduce greater operational complexity encompassing multi-dimensional coordination challenges in routing, speed selection, and energy management decisions that transcend conventional optimization paradigms.

This thesis first establishes an innovative dual-model evaluation framework that scientifically simulates the planning-execution gap in maritime operations. This framework demonstrates that ensemble weather forecast-based speed optimization achieves approximately 1% additional fuel savings compared to deterministic forecast-based approaches, validated through computational experiments on two nine-day voyages under both ballast and laden conditions. The superior rhumb line-based data fusion method developed for this framework ensures high-quality inputs for the ship's fuel consumption rate predictions. Building upon this foundation, this thesis introduces a rolling horizon optimization approach that enables ship speed replanning based on continuously updated weather forecasts throughout voyages. Computational results on two laden voyages demonstrate that this framework achieves 10.86% fuel savings compared to fixed speed operations and 7.53% additional savings compared to static optimization approaches, unlocking operational efficiency gains beyond conventional pre-voyage planning.

While ship speed optimization offers immediate energy efficiency improvements, achieving deeper decarbonization targets fundamentally requires technological innovation. This thesis further addresses the operational complexity of hybrid propulsion systems by formulating the integrated model for hybrid electric platform supply vessel (PSV) operations. This model simultaneously coordinates routing, speed, charging, and energy management decisions under fuel quota constraints that enforce predetermined emission reduction targets. Using realistic instances from North Sea operations, we benchmark hybrid PSV performance against two diesel-only references representing current industry practices: fixed design-speed operations and speed-optimized oper-

ations. Results demonstrate that operators currently using fixed speeds can achieve substantial emission reductions while simultaneously reducing total costs through combined hybrid propulsion and speed optimization. For operators already employing speed optimization, hybrid PSV operations can achieve moderate emission targets with stable abatement costs of approximately 0.77 NOK/kg CO_2. Our analysis further quantifies the critical role of operational factors in determining the cost-effectiveness of hybrid operations: offshore wind farms prove essential for achieving ambitious emission targets by eliminating costly detours to remote charging stations, while voyage time budget flexibility creates favorable trade-offs between operational efficiency and emission reduction costs. These findings offer actionable insights for maritime operators navigating the transition to hybrid propulsion systems under increasingly stringent environmental regulations.

This thesis systematically addresses weather forecast uncertainty in ship speed optimization and enables adaptive speed replanning with updated weather forecasts throughout voyages. By further extending the optimization framework to integrated hybrid system management, these methods provide practical tools for green shipping, offering viable pathways for the maritime industry to achieve sustainable development under mounting regulatory and economic pressures.

Publications During PhD Study

Arising from the Thesis

1. Luo, X., Yan, R., & Wang, S. (2023). Comparison of deterministic and ensemble weather forecasts on ship sailing speed optimization. *Transportation Research Part D: Transport and Environment*, 121, 103801.

2. Luo, X., Yan, R., & Wang, S. (2024). Ship sailing speed optimization considering dynamic meteorological conditions. *Transportation Research Part C: Emerging Technologies*, 167, 104827.

3. Luo, X., Fagerholt, K., Ormevik, A. B., & Yan, R. (2025). Hybrid electric platform supply vessel routing with speed optimization: Pathways to cost-effective emission reduction. *To be submitted.*

Others

1. Luo, X., Yan, R., Xu, L., & Wang, S. (2024). Accuracy and applicability of ship's fuel consumption prediction models: A comprehensive comparative analysis. *Energy*, 310, 133187.

2. Luo, X., Yan, R., & Wang, S. (2024). After five years' application of the European Union monitoring, reporting, and verification (MRV) mechanism: Review and prospectives. *Journal of Cleaner Production*, 434, 140006.

3. Luo, X., Zhang, M., Han, Y., Yan, R., & Wang, S. (2025). Ship fuel consumption prediction based on transfer learning: Models and applications. *Engineering Applications of Artificial Intelligence*, 141, 109769.

Acknowledgments

As I reach the end of this doctoral journey, I am filled with gratitude for all the people and experiences that have shaped these years into something truly meaningful.

First and foremost, I would like to express my sincere appreciation to The Hong Kong Polytechnic University (PolyU) for granting me a three-year scholarship that not only supported my research but also opened doors to a world of possibilities I had never imagined. This opportunity has truly broadened my horizons and transformed my life.

I am deeply grateful to my supervisor, Professor Shuaian Wang, and my co-supervisor, Dr. Ran Yan, for their invaluable guidance throughout this journey. Their meticulous attention to detail, rigorous academic standards, and patient mentorship have been instrumental in shaping my research direction and refining my work. More than just supervisors, they have been my academic mentors who ignited my passion for scholarly inquiry. Their dedication to excellence has set a standard I will carry with me throughout my career.

A special acknowledgement goes to the PolyU for providing me with a remarkable opportunity to spend three months in Norway. During this exchange, I had the privilege of collaborating with Professor Kjetil Fagerholt and Dr. Andreas Breivik Ormevik on a fascinating research topic. Their insights and collaborative spirit made this experience both enriching and enjoyable, and I am grateful for the chance to learn from such accomplished researchers.

To my friends in Hong Kong, especially my fellow students in Professor Wang's research group: thank you for making this journey so much more enjoyable. Whether we were engaged in deep academic discussions or simply sharing moments of everyday life, every conversation with you has been a treasure. I am grateful to have met each one of you.

My heartfelt gratitude goes to my family, whose unwavering support has been the foundation of my 28-year educational journey. To my father, mother, grandfather, and grandmother: though you may not have had the opportunity for higher education yourselves, you always believed in the value of learning and never stopped supporting my dreams. Without your sacrifices and encouragement, I would not be where I am today. To my sister and brother, your companionship along this long road of life is invaluable and deeply cherished.

To my beloved partner, my dearest friend: these six years we have spent together, learning and growing side by side, have been among the most precious of my life. Thank you for being there through it all.

Finally, I want to thank myself. Through countless ordinary days, joyful moments, and difficult times, I kept reaching upward. I want to remind myself: while people may only accompany us for parts of our journey, I accompany myself from birth until death. Cherish every moment, and live this life beautifully.

This thesis is not just the culmination of research; it is a testament to all the love, support, and belief that surrounded me along the way.

Table of Contents

Abstract **i**

Publications During PhD Study **iv**

Acknowledgments **vi**

List of Figures **xii**

List of Tables **xiv**

Acronyms **1**

1 Introduction **5**

1.1 Background and Motivation 5

1.2 Contributions 9

1.3 Thesis Outline 10

2 Literature Review **11**

2.1 Evolution of Speed Optimization for Conventional Vessels 11

2.1.1 Speed Optimization with Power-Law Models 12

2.1.2 Weather-Informed Speed Optimization 13

2.1.3 Speed Optimization under Uncertainty and Dynamism 15

2.2 Hybrid Electric PSV Operations: Routing and Energy Management Foundations 18

2.2.1 Platform Supply Vessel Routing 18

2.2.2 Electric Vehicle Routing Problem 23
2.3 Conclusion 25

3 Data Foundation **26**
3.1 Ship Operational Data 27
3.1.1 Noon Reports 27
3.1.2 Automatic Identification System Data 28
3.2 Meteorological Data 28
3.2.1 ECMWF Reanalysis V5 28
3.2.2 Ensemble Weather Forecasts 29
3.3 Data Fusion Methods 30
3.3.1 Direct Fusion Method 32
3.3.2 Rhumb Line-based Fusion Method 33
3.4 Conclusion 36

4 Ship Speed Optimization under Weather Forecast Uncertainty: An Evaluation Framework **37**
4.1 Ship Speed Optimization Model: A Predict-then-Optimize Approach 38
4.1.1 Fuel Consumption Prediction Model: Gradient Boosting Regression Trees 38
4.1.2 Modeling the Speed Optimization 40
4.2 Evaluation Framework 47
4.3 Computational Results 54
4.3.1 Experimental Setup 54
4.3.2 Ground Truth Predictor: Comparative Performance of Data Fusion Methods 59
4.3.3 Decision-Input Predictor: Performance Validation 61
4.3.4 Comparative Analysis of Speed Optimization under Different Weather Forecasts 63

4.4 Conclusion 68

5 Ship Speed Optimization under Evolving Meteorological Conditions 69

5.1 A Rolling Horizon Framework for Ship Speed Optimization 70

5.1.1 Problem Description 70

5.1.2 Solution Framework: Rolling Horizon Optimization 71

5.2 Deterministic Subproblem: Time-Distance Network Formulation . . . 78

5.2.1 Speed Optimization Model based on the Time-Distance Network 78

5.2.2 Ship Fuel Consumption Prediction: The Artificial Neural Network Model 84

5.3 Experimental Design 86

5.3.1 Validation Voyages and Model Configuration 87

5.3.2 Comparison Framework 88

5.4 Computational Results 92

5.4.1 Development of the ANN Model 92

5.4.2 Validation of the RHO framework 94

5.5 Conclusion 97

6 Hybrid Electric Platform Supply Vessel Routing with Speed Optimization: Pathways to Cost-Effective Emission Reduction 98

6.1 Problem Description 100

6.2 Modeling the HEPSVRP-SO 103

6.2.1 Time-Space Network 103

6.2.2 Notation 105

6.2.3 Mathematical Formulation 112

6.3 Experimental Design 115

6.3.1 Test Instances 115

6.3.2 Scenario Design 118

6.3.3 Parameter Settings 120

6.4 Results and Analysis 122
6.4.1 Comparative Performance Analysis Across Scenarios 123
6.4.2 Operational Factors Analysis 128
6.5 Conclusion 136

7 Conclusion 137
7.1 Summary of Accomplishments 137
7.2 Recommended Topics for Further Research 140
7.3 Closing 144

Appendices 146
A.1 Algorithms for Rolling Horizon Speed Optimization 146
A.1.1 Rolling Horizon Optimization Framework 146
A.1.2 Time-Distance Network Construction 146
A.2 Detailed Voyage Information 149
A.3 Time-space Network Generation Algorithm 149

List of Figures

3.1 Graphic explanation of the fusion method based on the rhumb line. . 35

4.1 Framework for evaluating speed optimization under different forecast types. 49

5.1 Rolling horizon optimization framework for ship speed optimization. . 76
5.2 Structure of the ANN model for fuel consumption rate prediction. . . 85
5.3 Fuel consumption and sailing speed profiles after optimization. 95

6.1 Illustrative PSV route demonstrating hybrid propulsion operations, charging strategies, and speed optimization. 102
6.2 Time-space network structure with service-transit and battery charging arcs. 105
6.3 An example for temporal sets $\mathcal{T}_{ij}^{S}$, $\mathcal{T}_{ij15}^{SS}$, and $\mathcal{T}_{i15j}^{SE}$. 107
6.4 The battery level dynamics corresponding to the illustrative PSV route. 108
6.5 Geographic distribution of installations in the North Sea region . . . 116
6.6 Total cost of all instances under the design-speed and speed-optimization references. 124
6.7 Average relative cost vs abatement rate under the scenario of adopting the fuel energy quota (a)$E_{\text{base}}^{\text{DS}}$ and (b)$E_{\text{base}}^{\text{SO}}$. 126
6.8 Abatement cost vs abatement rate under the scenario of adopting the fuel energy quota (a) $E_{\text{base}}^{\text{DS}}$ and (b) $E_{\text{base}}^{\text{SO}}$. 128
6.9 Effect of wind farm under the default time budget with the fuel cap parameterized by the speed-optimization reference. 129

6.10 Planned route of instance 4 at abatement rate $r = 0.9$ under (a) scenario 5 and (b) scenario 6. 130
6.11 Battery level of instance 4 at abatement rate $r = 0.9$ under (a) scenario 5 and (b) scenario 6. 131
6.12 Sailing speed and fuel-electric energy distribution for instance 4 with and without wind farm access. 132
6.13 Voyage durations under custom (hatched) and default time budgets when adopting the fuel energy quota (a) $E^{\mathrm{DS}}_{\mathrm{base}}$ and (b) $E^{\mathrm{SO}}_{\mathrm{base}}$. 133
6.14 Percentage cost increase of using a custom time budget versus a default time budget. 135

List of Tables

2.1 Operational PSV routing problems in the literature 21

2.2 Electric vehicle routing problems in the literature 22

3.1 Features in each data source 31

4.1 Notation for speed optimization model formulations 41

4.2 Notation for evaluation framework and dataset construction 50

4.3 Operational and meteorological conditions for Voyage I (9-day ballast voyage) 57

4.4 Operational and meteorological conditions for Voyage II (9-day laden voyage) 58

4.5 Hyperparameters for XGBoost models $\mathcal{F}_{\text{XGB}}^{\text{direct}}$ and $\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$ 60

4.6 Predictive performance comparison of XGBoost models trained based on different data fusion methods 60

4.7 Hyperparameters for GBRT planner model $\mathcal{F}_{\text{GBRT}}^{\text{planner}}$ 62

4.8 Performance of GBRT planner model $\mathcal{F}_{\text{GBRT}}^{\text{planner}}$ on simulated test set . 62

4.9 Speed optimization results for Voyage I (9-day ballast voyage) 64

4.10 Speed optimization results for Voyage II (9-day laden voyage) 65

4.11 Summary of evaluated fuel consumption and savings across optimization strategies 66

5.1 Notation for the rolling horizon optimization framework 73

5.2 Notation for the time-distance network model 77

5.3 Input features and prediction target for ANN fuel consumption prediction model 88
5.4 Hyperparameters and optimal configuration of the ANN regressor . . 93
5.5 Predictive performance of the ANN regressor on test data 93
5.6 Comparative performance of speed optimization schemes on validation voyages 96

6.1 Sets and indices 107
6.2 Parameters 110
6.3 Decision variables 111
6.4 Test instances 117
6.5 Description of scenario design 120
6.6 Parameters for the computational experiment 121
6.7 Cost breakdown for instance 4 under scenario 5 (with wind farm) and scenario 6 (without wind farm) at abatement rate $r = 0.9$ 132

A.1 Detailed operational and meteorological information for Voyage I . . . 150
A.2 Detailed operational and meteorological information for Voyage II . . 151

Acronyms

2mT	2-meter Temperature.
AIS	Automatic Identification System.
ANN	Artificial Neural Network.
CII	Carbon Intensity Indicator.
DM	Disruption Management.
E-LRPTW	Electric Location Routing Problem with Time Windows.
ECMWF	European Centre for Medium-Range Weather Forecasts.
EEDI	Energy Efficiency Design Index.
EEXI	Energy Efficiency Existing Ship Index.
ERA5	ECMWF Reanalysis v5.
ETS	Emissions Trading System.
EU	European Union.
EVRP	Electric Vehicle Routing Problem.
EVRP-RE	EVRP with Renewable Energy.

EVRP-NL-C	EVRP with Nonlinear Charging and Capacitated Stations.
EVRP-NL	EVRP with Nonlinear charging.
EVRP-TW-PR	EVRP with Time Windows and Partial Recharging.
EVRP-TW	Electric VRP with Time Windows.
FCR	Fuel Consumption Rate.
FSR	Forecast Surface Roughness.
GBRT	Gradient Boosted Regression Trees.
GEFS	Global Ensemble Forecast System.
GHG	Greenhouse Gas.
HEPSVRP-SO	HEPSVRP-SO.
HEVRP	Electric Vehicle Routing Problem.
HVRP	Hybrid Vehicle Routing Problem.
IDW	Inverse Distance Weighting.
IMO	International Maritime Organization.
MAE	Mean Absolute Error.
MAPE	Mean Absolute Percentage Error.
MDTS	Mean Direction of Total Swell.
MDWW	Mean Direction of Wind Waves.
ME	Main Engine.
MILP	Mixed-Integer Linear Programming.

MIP	Mixed-Integer Programming.
MM-HEVRP-TW	Multi-Mode Hybrid Electric VRP with Time Windows.
MM-HEVRP	Multi-Mode Hybrid Electric VRP.
MPTS	Mean Period of Total Swell.
MPWW	Mean Period of Wind Waves.
MSE	Mean Squared Error.
MWD	Mean Wave Direction.
MWP	Mean Wave Period.
NOAA	National Oceanic and Atmospheric Administration.
OSP	Order Selection Problem.
PSV	Platform Supply Vessel.
PSVRSP	Platform Supply Vessel Routing and Scheduling Problem.
RHO	Rolling Horizon Optimization.
SEEMP	Ship Energy Efficiency Management Plan.
SHCWWS	Significant Height of Combined Wind Waves and Swell.
SHTS	Significant Height of Total Swell.
SHWW	Significant Height of Wind Waves.
SVPDPCC	Single Vehicle Pickup and Delivery Problem with Capacitated Customers.

SVPP	Supply Vessel Planning Problem.
TDN	Time-Distance Network.
TDVRP-SO	Time-Dependent VRP with Speed Optimization.
VRP	Vehicle Routing Problem.
VRPPD	Vehicle Routing Problem with Pickups and Deliveries.
VRPSPD	Vehicle Routing Problem with Selective Pickups and Deliveries.
WD	10-meter Wind Direction.
WD-SO	Weather-dependent Speed Optimization.
WS	10-meter Wind Speed.
XGBoost	Extreme Gradient Boosting.

Chapter 1

Introduction

1.1 Background and Motivation

Maritime transport serves as the backbone of global trade, carrying more than 80% of world merchandise trade, amounting to 10.9 billion metric tons in 2020 (UNCTAD, 2020). However, the international shipping sector contributes approximately 3% of global greenhouse gas (GHG) emissions (IMO, 2020). In response, the International Maritime Organization (IMO) announced an ambitious target in 2023 requiring the maritime industry to reduce total annual GHG emissions by at least 70% compared to 2008 levels by 2040, ultimately achieving net-zero emissions by around 2050 (IMO, 2023). To operationalize these objectives, the IMO has implemented regulatory initiatives including the Energy Efficiency Design Index (EEDI), the Energy Efficiency Existing Ship Index (EEXI), and the Ship Energy Efficiency Management Plan (SEEMP) (IMO, 2011). Starting from 2023, the Carbon Intensity Indicator (CII) requires ships to calculate and report their annual CII rating (IMO, 2021). Beyond global regulations, regional frameworks such as the European Union's Emissions Trading System (EU ETS) have incorporated shipping into carbon markets (European Commission, 2025; Kotzampasakis, 2025). National governments like Norway have mandated zero-emission technology for platform supply vessels by 2030 (Norwegian Ministry, 2019), addressing the substantial GHG emissions from offshore oil and gas operations that rely heavily on diesel-powered supply vessel services (International Energy Agency, 2023; Thomas et al., 2025; Upasana et al., 2025).

These regulatory pressures coincide with significant economic challenges. Fuel costs constitute over 50% of total operational expenses for many vessels (Wang and Meng, 2012), making energy efficiency both an environmental imperative and an economic necessity. The combination of high fuel prices, volatile energy markets, and increasing carbon costs, as exemplified by the EU ETS, means that reducing emissions has become not only a regulatory obligation but also a critical competitive advantage (Wang et al., 2013). Shipping companies that fail to improve their operational efficiency face direct financial penalties through carbon pricing mechanisms and indirect market disadvantages through poor CII ratings that may limit their access to certain ports and charterers.

In response to these pressures, the maritime sector has explored technical and operational measures (Yan et al., 2021; Yang et al., 2019b; Du et al., 2019). Technical approaches, such as vessel redesign, propulsion system upgrades, and alternative fuel adoption, offer long-term emission reduction potential but require substantial capital investment and extended implementation timelines. For the existing global fleet, such technical retrofits are often economically prohibitive. In contrast, operational measures, such as speed optimization (Beşikçi et al., 2016b), virtual arrival (Jia et al., 2017), and trim optimization (Du et al., 2019), present immediately implementable pathways to enhance ship energy efficiency with minimal capital requirements (Yang et al., 2020; Yan et al., 2020; Fan et al., 2022). Speed optimization has been widely adopted by major shipping companies due to its proven effectiveness and ease of implementation (Maersk, 2021; Hafnia, 2021). The fundamental principle is straightforward: fuel consumption increases exponentially with speed, and by strategically modulating sailing speeds based on operational constraints and environmental conditions, significant fuel savings can be achieved without compromising service reliability.

However, effective implementation of ship speed optimization in real-world maritime operations faces critical limitations in existing research. Speed optimization follows a predict-then-optimize framework: accurate fuel consumption prediction must precede optimal speed profile identification. Since fuel consumption is strongly influenced by meteorological conditions such as wind, waves, and currents, the quality of weather information directly determines optimization effectiveness (Psaraftis and Kontovas, 2014; Yan et al., 2021). However, meteorological forecasts are inherently uncertain due to the chaotic nature of atmospheric and oceanic systems (Lorenz, 1965). While deterministic weather forecasts provide a single predicted scenario representing the most likely future state, ensemble forecasting systems generate multiple plausible scenarios that explicitly characterize prediction uncertainty (World Meteorological Organization, 2012). Despite the operational availability of ensemble forecasts for decades, their potential value for ship speed optimization remains unexplored. More critically, previous studies that attempted to evaluate the benefits of ensemble weather forecasts suffered from a methodological limitation: they lacked rigorous evaluation frameworks that explicitly distinguish prediction capability from physical reality and employ reliable historical data to simulate the planning-execution gap (Hinnenthal and Clauss, 2010; Skoglund et al., 2015). This methodological gap prevents systematic assessment of whether ensemble weather forecasting can deliver tangible operational improvements in ship speed optimization.

Another limitation arises from environmental dynamism during voyage execution. Ocean voyages typically span multiple days or weeks, during which meteorological conditions evolve continuously and weather forecasts are updated regularly. Traditional speed optimization research has predominantly employed static optimization frameworks that determine optimal speed profiles based on forecast information available before departure and maintain these decisions throughout the voyage regardless of forecast updates (Fagerholt et al., 2010; Wang and Meng, 2012; Beşikçi et al., 2016a). This approach fails to exploit the valuable information contained in weather forecast updates and cannot adapt to changing conditions encountered during operations. This can result in suboptimal fuel consumption and compromise the effectiveness of ship speed optimization. Addressing this limitation requires developing adaptive optimization frameworks that can incorporate updated meteorological information and adjust operational decisions accordingly (Wang et al., 2018a; Zis et al., 2020).

While ship speed optimization offers cost-effective emission reductions for conventional vessels with minimal capital investment, achieving the maritime sector's deep decarbonization targets ultimately requires technological transformation. Hybrid diesel-electric propulsion systems have emerged as a particularly promising transitional technology, capable of reducing emissions substantially while maintaining operational flexibility (DNV GL, 2016; Chin et al., 2022; Inal et al., 2022; Paipa-Sanabria et al., 2025). However, the introduction of hybrid propulsion technology creates new operational complexities that significantly expand beyond traditional vessel routing and speed optimization frameworks. Operators must simultaneously optimize charging strategies and energy allocation between diesel and electric power sources while coordinating routes, speeds, and schedules to maximize emission reduction and cost-effectiveness. This operational complexity is further compounded by the limited availability of charging infrastructure, requiring vessels to maintain sufficient energy reserves for potentially long sailing segments between charging locations. The strategic integration of renewable energy infrastructure, such as offshore wind farms that provide high charging rates for ships, adds another dimension to the optimization challenge by offering new charging opportunities that could fundamentally alter optimal routing strategies.

1.2 Contributions

To address these limitations, this thesis systematically develops methodologies that advance from evaluating the uncertainty of weather forecasts to enabling adaptive speed planning and ultimately coordinating multi-dimensional decisions for hybrid propulsion systems, which results in the following three contributions.

First, this thesis develops an evaluation framework for quantifying the operational value of ensemble weather forecasts in ship speed optimization. By explicitly distinguishing prediction capability from physical reality through a novel dual-model architecture, this framework overcomes the impossibility of conducting controlled real-world experiments where the same vessel sails identical routes under different forecast types. The framework enables systematic comparison of deterministic and ensemble forecast-based optimization under realistic conditions.

Second, this thesis extends ship speed optimization from static to adaptive frameworks that can adapt to evolving meteorological conditions during voyage execution. The proposed rolling horizon framework enables vessels to incorporate regularly updated weather forecast information and adjust operational decisions accordingly, addressing the limitation of conventional static speed optimization methods.

Finally, this thesis introduces the Hybrid Electric Platform Supply Vessel Routing Problem with Speed Optimization (HEPSVRP-SO) and develops a mixed-integer programming (MIP) model that simultaneously optimizes route planning, speed selection, charging strategies, and power allocation between diesel and electric propulsion under fuel quota constraints. Computational experiments based on realistic instances from North Sea operations quantify the emission reduction potential of hybrid electric vessel operations and assess the strategic value of offshore wind farms for maritime decarbonization.

1.3 Thesis Outline

The remainder of this thesis is organized as follows. Chapter 2 reviews the relevant literature on ship speed optimization, covering the evolution from foundational models to advanced approaches that incorporate weather forecast uncertainty and operational dynamism, as well as electric vehicle routing problems that provide methodological insights for hybrid vessel operations. Chapter 3 introduces the ship operational data, describes the meteorological data sources, and develops data fusion methods for integrating these heterogeneous data sources into high-quality datasets suitable for predicting ship fuel consumption and optimizing ship speed. Chapter 4 develops the dual-model evaluation framework for comparing deterministic and ensemble weather forecasts in ship speed optimization. Chapter 5 introduces the rolling horizon optimization framework for ship speed optimization under evolving meteorological conditions. Chapter 6 formulates and evaluates the integrated optimization model for hybrid electric vessel routing with speed optimization and energy management. Chapter 7 concludes the thesis and discusses future research directions.

Chapter 2

Literature Review

The review is organized into two research streams. Section 2.1 examines the evolution of ship speed optimization for conventional vessels, tracing the progression from power-law-based speed optimization to weather-dependent speed optimization and ultimately to advanced methods addressing forecast uncertainty and operational dynamism. Section 2.2 examines two research domains that provide the theoretical foundation for hybrid electric vessel optimization: platform supply vessel (PSV) routing in offshore logistics, which establishes the operational context and practical constraints, and electric vehicle routing problems from land transportation, which offer insights for battery management and energy allocation strategies.

2.1 Evolution of Speed Optimization for Conventional Vessels

Ship speed optimization follows a predict-then-optimize framework: first predicting fuel consumption under different speed scenarios, then determining the optimal speed profile through mathematical optimization. This framework has evolved significantly as fuel consumption prediction models have progressed from simplified power-law relationships to sophisticated weather-informed approaches, and as optimization methods have advanced to address forecast uncertainty and operational dynamism. This section reviews this evolution across three stages: optimization based on power-law fuel consumption models, optimization employing weather-informed prediction models, and advanced optimization methods that explicitly address weather forecast uncertainty through ensemble predictions and operational dynamism through adaptive optimization frameworks.

2.1.1 Speed Optimization with Power-Law Models

Early research into speed optimization modelled ship fuel consumption primarily as a function of sailing speed (Ryder and Chappell, 1980). Power-law relationships are employed in these models, where fuel consumption increases with the cube or higher powers of sailing speeds (Fagerholt et al., 2010; Wang and Meng, 2012; Lai et al., 2022; Li et al., 2022a). Fagerholt et al. (2010) proposed a continuous nonlinear programming model to determine optimal speeds along a shipping route, which was discretized and solved using shortest path algorithms. Kim et al. (2016) subsequently developed an exact algorithm for the same problem. Wang and Meng (2012) established speed-fuel consumption relationships from historical operational data and embedded them into a mixed-integer nonlinear programming model to optimize sailing speeds for multiple container ships in a liner shipping network. Yao et al. (2012) developed similar relationships for different container ship sizes and integrated them into an optimization model determining refueling ports, refueling amounts, and sailing speeds to minimize total bunker fuel costs.

While these models successfully demonstrated the economic benefits of speed reduction and established fundamental solution methodologies, they suffered from a critical limitation: ship fuel consumption was modeled solely as a function of speed, implicitly assuming constant environmental conditions. This simplification proved inadequate for practical applications, as actual fuel consumption at identical speeds can vary substantially depending on meteorological conditions. The exclusion of weather factors represents a fundamental limitation that motivated the development of weather-informed optimization approaches.

2.1.2 Weather-Informed Speed Optimization

The recognition that meteorological conditions substantially affect ship fuel consumption leads to a paradigm shift in speed optimization research. Factors such as wind conditions, wave characteristics, ocean currents, trim, and draught significantly influence ship fuel consumption beyond speed alone (Medina et al., 2020; Yu et al., 2021). Early approaches employed physics-based simulation models and optimization methods. Li et al. (2018) developed a simulation model considering wind and irregular wave impacts to calculate bunker consumption rates, then formulated a bi-objective optimization model simultaneously minimizing fuel consumption and maximizing cost savings. This work was extended by Li et al. (2020), who additionally accounted for voluntary speed loss and GHG emissions. Fan et al. (2022) built regression functions relating instantaneous fuel consumption and speed over water to main engine revolutions using operational sensor data, then proposed a multistage decision-making model to minimize overall voyage fuel consumption. Lin et al. (2013) developed a weather routing algorithm based on the composite influence of multiple dynamic environmental elements, optimizing ship trajectories considering wind, waves, and currents. Roh (2013) proposed an economical shipping route determination method accounting for sea state effects, employing isochrone-based techniques to evaluate alternative routes under varying weather conditions.

The emergence of big data technology and artificial intelligence prompts the growing adoption of machine learning algorithms for constructing fuel consumption prediction models (Yan et al., 2021). Several common machine learning algorithms include the tree-based model (Soner et al., 2018; Li et al., 2022b), artificial neural networks (Jeon et al., 2018; Petersen et al., 2012b; Gkerekos and Lazakis, 2020; Kim et al., 2021; Zhu et al., 2021), regularized linear regression (Soner et al., 2019; Wang et al., 2018b), k-nearest neighbors models (Peng et al., 2020; Uyanık et al., 2020), and Gaussian process (Petersen et al., 2012a; Yuan and Nian, 2018). These data-driven approaches learn complex nonlinear patterns directly from operational data without requiring explicit hydrodynamic equations or proprietary vessel specifications. Tarelko and Rudzki (2020) established two artificial neural network models to predict ship fuel consumption rate FCR and instantaneous speed over ground, then formulated a bi-objective optimization model concurrently minimizing fuel consumption and maximizing sailing speed. Du et al. (2019) proposed a two-phase framework for optimizing container ship speed and trim, employing artificial neural networks to predict fuel consumption rates in the first phase, then evaluating alternative optimization strategies in the second phase. Yan et al. (2020) trained a random forest model to estimate fuel consumption under various scenarios using datasets combining noon reports with meteorological data from the European Centre for Medium-Range Weather Forecasts (ECMWF), then developed a mixed-integer linear program optimizing sailing speeds subject to arrival time constraints.

The above-mentioned studies usually adopted ship noon reports in the construction of ship fuel consumption prediction models (Yuan and Nian, 2018; Gkerekos et al., 2018; Tarelko and Rudzki, 2020). However, as the noon report is manually recorded by crews and there is only one record for each day (usually at noon according to the local time), the quality as well as the quantity of noon reports can be low, especially given that meteorological conditions are often volatile (Li et al., 2022b; Du et al., 2022b). Thus, incorporating sea and weather information from the external meteorological database can improve data quality as well as the data quantity, eventually improving the predicting performance of fuel consumption prediction models (Du et al., 2022a).

However, few studies describe in detail a fusion method to combine two or more data sources or discuss the impact of different fusion methods on the prediction accuracy of the fuel consumption prediction model.

Furthermore, despite improvements in predicting ship fuel consumption through the incorporation of meteorological factors, these studies face two limitations. First, these studies obtained weather information from deterministic weather forecasts (Li et al., 2022b; Du et al., 2019; Yan et al., 2020; Lee et al., 2018). However, due to chaos in the atmosphere and ocean and the imperfect depiction of their initial state, the uncertainty is introduced in weather forecasting. Deterministic weather forecasts provide only a single predicted scenario without capturing the range of possible future weather conditions (Lorenz, 1965; Yousuf et al., 2019), which degrades the accuracy of fuel consumption prediction and consequently undermines the effectiveness of speed optimization. Second, forecast accuracy gradually decreases with lead time. However, these studies commit to a complete speed profile before departure without mechanisms to incorporate continuously updated weather forecasts during voyage execution, meaning that speed optimization decisions become progressively suboptimal as more accurate meteorological information becomes available. These deficiencies motivate advanced approaches that explicitly address forecast uncertainty and operational dynamism.

2.1.3 Speed Optimization under Uncertainty and Dynamism

Atmospheric systems exhibit chaotic dynamics wherein small perturbations in initial conditions amplify exponentially, fundamentally limiting deterministic predictability (Lorenz, 1965). Ensemble forecasting systems address this limitation by generating multiple plausible scenarios through repeated model executions with perturbed initial conditions, thereby characterizing the probability distribution of future atmospheric states (World Meteorological Organization, 2012). A typical ensemble forecast consists of a control member generated from the unperturbed initial conditions and multiple perturbed members generated from slightly varied initial conditions. For ship speed optimization, ensemble weather forecasts enable stochastic formulations that minimize expected fuel consumption across all ensemble members, producing more robust solutions that hedge against forecast errors.

While these pioneering studies demonstrated the potential value of ensemble forecasts for maritime operations, closer examination reveals critical methodological limitations. Hinnenthal and Clauss (2010) proposed a Pareto-optimal route optimization method utilizing both deterministic and ensemble weather forecasts with genetic algorithms, introducing robustness as a third optimization objective. However, they did not systematically evaluate actual performance in verified weather conditions, making it difficult to quantify the operational advantages of ensemble weather forecasts. Skoglund et al. (2015) conducted comparisons between deterministic and ensemble-based weather routing using verified weather constructed from consecutive forecast control members, finding that ensemble-optimized routes significantly reduced delay risks and fuel consumption prediction errors. However, they used the same prediction model for both optimization and evaluation, which fails to realistically represent the fundamental distinction in maritime operations between planning based on imperfect predictions and execution governed by true physical laws. Additionally, their verified weather remains weather forecast data rather than actual observations such as reanalysis data. Therefore, a methodological gap remains regarding the need for an evaluation framework that explicitly distinguishes prediction capability from physical reality and uses reliable historical data to realistically simulate the planning-execution gap.

The second limitation concerns operational dynamism. Weather forecasts are updated regularly, typically every six hours, with revisions often differing from earlier predictions, particularly at longer lead times. Static pre-departure optimization cannot adapt to evolving conditions as more accurate information becomes available.

Some studies have attempted to address this limitation by employing dynamic optimization frameworks that adjust speed decisions during voyage execution. Wang et al. (2018a) adopted a dynamic optimization method that determines optimal sailing speeds based on real-time environmental factors provided by sensors at different time steps. Tzortzis and Sakalis (2021) divided the voyage time horizon into smaller segments where sensor-measured weather conditions are considered accurate, solving successive limited-horizon subproblems to determine optimal speeds. However, sensors can only measure meteorological conditions at the vessel's current location and time, without providing forecasts for the remainder of the voyage. Consequently, such methods optimize based on current conditions alone, without considering how the weather will evolve in the future. This can lead to suboptimal decisions if the weather changes significantly along the remaining route.

The evolution from speed-centric models through weather-dependent frameworks to advanced uncertainty-aware and dynamic approaches reflects continuous refinement toward methodologies that better represent maritime operational realities. However, two critical research gaps remain unaddressed. First, while pioneering studies have demonstrated potential benefits of ensemble forecasts, the literature lacks a rigorous evaluation framework that can systematically quantify their operational value by explicitly distinguishing between prediction capability used for planning and physical reality governing actual performance. Second, although dynamic optimization concepts have been explored, existing approaches either rely on sensor measurements without forecast information for future segments or lack systematic frameworks for integrating continuously updated forecasts throughout voyage execution. Addressing these gaps through developing a comprehensive comparison framework for weather forecast uncertainty and a systematic rolling horizon optimization framework for operational dynamism constitutes the central focus of Chapters 4 and 5 of this thesis, respectively.

2.2 Hybrid Electric PSV Operations: Routing and Energy Management Foundations

The preceding section examined the evolution of speed optimization for conventional diesel-powered vessels, revealing critical gaps in addressing weather forecast uncertainty and operational dynamism. While these challenges apply broadly across maritime transportation, this thesis focuses on a more complex operational context involving hybrid electric propulsion systems in offshore supply vessel operations. Hybrid electric vessels integrate dual power sources consisting of diesel engines and battery-electric systems, requiring coordinated optimization of not only speed profiles but also charging strategies and power allocation between propulsion modes. Furthermore, PSVs operate under operational constraints, including sparse charging infrastructure, strict emission regulations, and time-sensitive service requirements for offshore installations. To establish the theoretical foundations for addressing these integrated challenges, this section reviews two research streams: operational routing of PSVs and energy management in electric and hybrid vehicle routing problems.

2.2.1 Platform Supply Vessel Routing

A substantial portion of research on PSV routing and scheduling in the offshore oil and gas industry focuses on the Supply Vessel Planning Problem (SVPP). The SVPP involves determining an optimal fleet of PSVs and their weekly routes and schedules to serve offshore installations that require a specified number of visits. A defining feature of the SVPP is that vessel schedules must be repeatable on a weekly basis, reflecting standard offshore logistics practices. The problem was first studied in a simplified form by Fagerholt and Lindstad (2000) and formally defined by Halvorsen-Weare et al. (2012), who incorporated key practical aspects such as service capacity limits at supply bases and spread departures for voyages to the same installations.

Because the SVPP is a complex combinatorial routing problem, heuristic solution methods are typically employed (Shyshou et al., 2012; Kisialiou et al., 2018; Borthen et al., 2018). More recent studies have introduced additional real-world elements, particularly uncertainty in cargo demand and sailing times, using various modeling approaches and solution methods (Kisialiou et al., 2019; Santos et al., 2022; Cruz et al., 2023).

In contrast to the SVPP, which addresses tactical planning, the HEPSVRP-SO is an operational problem focusing on the detailed planning of the upcoming route for a PSV, involving also decisions about sailing speeds, charging, and the energy management along the route. There are also a few studies dealing with operational routing of PSVs. Aas et al. (2007) developed a mathematical model for routing and scheduling of a single vessel to petroleum installations, considering storage capacity constraints of offshore platforms, to perform both pickup and delivery services. Building on this problem definition, Gribkovskaia et al. (2008) developed a tabu search heuristic designed to solve larger instances of this single-vessel pickup and delivery problem. Subsequent research introduced more complex operational constraints. Cuesta et al. (2017) introduced the Vessel Routing Problem with Selective Pickups and Deliveries, addressing capacity-constrained scenarios where order postponement becomes necessary.

Building on these foundational models that introduced specific operational constraints, other research has focused on developing more comprehensive, integrated decision-making frameworks. de Bittencourt et al. (2021) developed an integrated framework for optimizing cargo assignment, fleet sizing, and delivery planning over fixed, predetermined routes. Silva et al. (2024) addressed an integrated routing and scheduling problem that simultaneously considers multiple vessel trips, pickups and deliveries, and deck space capacity constraints. Stålhane et al. (2019) developed a variable neighborhood search heuristic for offshore supply vessel disruption management, facilitating the re-planning of voyages when vessels deviate from their long-term schedules due to weather or operational disruptions. In addition, speed optimization has emerged as a critical decision variable in operational PSV planning. Ulsrud et al. (2022) considered a time-dependent PSV routing problem with speed optimization that accounts for weather-dependent fuel consumption and variable sailing speeds, while Ormevik et al. (2023) introduced a high-fidelity weather-dependent fuel consumption model for speed optimization on a given, predetermined route to evaluate whether the choice of weather impact model would affect the speed selection strategies.

Despite these advances in operational PSV planning, none of the existing studies address hybrid electric propulsion systems. All previous research focuses exclusively on conventional diesel-powered vessels, without considering the integrated optimization of charging strategies and dynamic energy allocation between diesel and electric power sources that characterize hybrid PSV operations. Table 2.1 summarizes the key characteristics of operational PSV routing studies and positions the HEPSVRP-SO within this literature.

Table 2.1: Operational PSV routing problems in the literature

		Problem Characteristics			**Decision Variables**				
Study	**Problem Name**	**Vessel Type**	**Planning Horizon**	**Fleet**	**Routing**	**Speed**	**Cargo Selection**	**Hybrid Propulsion Management**	**Objective Function**
Aas et al. (2007)	VRPPD	Diesel	3 days	Single	✓				Minimize total travel time
Gribkovskaia et al. (2008)	SVPDPCC	Diesel	3 days	Single	✓				Minimize total travel time
Cuesta et al. (2017)	VRPSPD	Diesel	2-3 days	Single	✓		✓		Minimize travel cost and penalty cost for not handling orders
Stålhane et al. (2019)	VRP for DM	Diesel	2-3 days	Heterogeneous	✓		✓		Minimize sailing cost, chartering cost, and penalty cost for postponed orders and delays
de Bittencourt et al. (2021)	OSP	Diesel	1 voyage	Heterogeneous			✓		Maximize the service level, defined as the percentage of on-time deliveries while penalizing refused requests
Ulsrud et al. (2022)	TDVRP-SO	Diesel	2-3 days	Heterogeneous	✓	✓	✓		Minimize fuel cost, chartering cost, and penalty cost for postponed orders
Ormevik et al. (2023)	WD-SO	Diesel	2-3 days	Single		✓			Minimize total fuel consumption
Silva et al. (2024)	PSVRSP	Diesel	1 voyage	Heterogeneous	✓				Minimize marginal fuel expenditure and vessel utilization time
This study	**HEPSVRP-SO**	**Hybrid**	**2-3 days**	**Single**	✓	✓		✓	**Minimize fuel cost, electricity cost and emission cost**

Note: VRPPD = Vehicle Routing Problems with Pickups and Deliveries; SVPDPCC = Single Vehicle Pickup and Delivery Problem with Capacitated Customers; VRPSPD = VRP with Selective Pickups and Deliveries; DM = Disruption Management; OSP = Order Selection Problem; TDVRP-SO = Time-Dependent VRP with Speed Optimization; WD-SO = Weather-dependent Speed Optimization; PSVRSP = PSV Routing and Scheduling Problem.

Table 2.2: Electric vehicle routing problems in the literature

Study	Problem Name	Problem Characteristics: Vehicle Type	Planning Horizon	Fleet	Decision Variables: Routing	Charging Strategy	Energy Allocation	Objective Function
Schneider et al. (2014)	EVRP-TW	Electric	1 day	Homogeneous	✓	Full charging		Hierarchically minimize vehicle count, then total distance
Keskin and Çatay (2016)	EVRP-TW-PR	Electric	1 day	Homogeneous	✓	Partial charging		Minimize total travel distance
Montoya et al. (2017)	EVRP-NL	Electric	1 day	Homogeneous	✓	Partial charging		Minimize total time of traveling and charging
Mancini (2017)	HVRP	Hybrid	1 day	Homogeneous	✓	Full charging	Mode selection	Minimize total distance and penalty cost for using traditional fuel
Schiffer and Walther (2017)	E-LRPTW	Electric	Multi-day	Homogeneous	✓	Partial charging		Separately minimize multiple different objectives, including total travel distance, number of vehicles used, number of charging stations used, weighted sum of the number of vehicles and charging stations, and total cost encompassing investment and operational expenses
Seyfi et al. (2022)	MM-HEVRP	Hybrid	1 day	Homogeneous	✓	No	Mode selection	Minimize total travel cost
Froger et al. (2022)	EVRP-NL-C	Electric	1 day	Homogeneous	✓	Partial charging		Minimize total time of driving, service, charging, and waiting
Sayarshad (2024)	EVRP-RE	Electric	1 day	Homogenous	✓	Partial charging		Minimize electricity purchase cost and total time of traveling, setup, and charging
Jiang et al. (2025)	MM-HEVRP-TW	Hybrid	1 day	Homogeneous	✓	No	Mode selection	Minimize total travel cost
This study	**HEPSVRP-SO**	**Hybrid**	**2-3 days**	**Single**	✓	**Partial charging**	**Power allocation**	**Minimize fuel cost, electricity cost and emission cost**

Note: EVRP-TW = Electric VRP with Time Windows; EVRP-TW-PR = EVRP with Time Windows and Partial Recharging; EVRP-NL = EVRP with Nonlinear Charging; HVRP = Hybrid VRP; = Electric Location Routing Problem with Time Windows; MM-HEVRP = Multi-Mode Hybrid Electric VRP; EVRP-NL-C = EVRP with Nonlinear Charging and Capacitated Stations; EVRP-RE = EVRP with Renewable Energy. Mode selection refers to discrete choice among predefined operational modes (e.g., pure electric, pure diesel), while power allocation represents continuous proportional distribution between energy sources.

2.2.2 Electric Vehicle Routing Problem

While vessel electrification represents an emerging frontier, the optimization of electric vehicle routing has been extensively studied in land transportation, providing valuable methodological insights for maritime applications. The Electric Vehicle Routing Problem (EVRP) seeks to determine an optimal set of routes for a fleet of electric vehicles to serve a collection of customers. This involves making decisions about not only the sequence of customer visits but also the location and duration of necessary charging stops to accommodate for their limited driving range and battery capacity.

The foundational work by Schneider et al. (2014) established the mathematical formulation for the Electric Vehicle Routing Problem with time windows and recharging stations, introducing the core concepts of range limitations, charging decisions, and battery state management that define this research domain. Key advances in EVRP methodology include the work of Keskin and Çatay (2016), who addressed partial recharging strategies for electric vehicles, demonstrating that this flexibility can lead to significant improvements in routing decisions. Montoya et al. (2017) introduced realistic nonlinear charging functions that more accurately represent battery charging behavior than earlier linear models. Recent developments have focused on charging infrastructure integration and optimization. Schiffer and Walther (2017) studied the simultaneous location-routing problem for electric vehicles with time windows, providing insights for strategic charging infrastructure placement. Froger et al. (2022) addressed capacitated charging stations, explicitly modeling the limited number of chargers available and the potential for queuing. Sayarshad (2024) demonstrated the integration of renewable energy sources with electric vehicle routing, achieving a 42% reduction in battery storage requirements through intelligent charging coordination.

The most relevant extension for this study is the Hybrid Electric Vehicle Routing Problem (HEVRP). Unlike pure electric vehicles, hybrid vehicles can switch to a conventional engine if their battery is depleted, making the electric range a "soft" constraint. Mancini (2017) addressed the hybrid vehicle routing problem, providing mathematical formulations for vehicles that can switch between electric and traditional fuel propulsion and recharge at stations. Seyfi et al. (2022) developed multi-mode hybrid electric vehicle routing models that address four operational modes: pure electric, pure combustion, charging while driving, and boost mode. Building on this, Jiang et al. (2025) extended the problem to include time windows and developed an adaptive large neighborhood search algorithm to solve this more complex variant. Their mathematical formulations for energy management between multiple power sources provide the theoretical foundation for addressing the dual propulsion systems in hybrid electric PSVs.

Despite these advances, research on electric or hybrid vessel routing is still in its early stages. Early work by Kanellos (2013) developed power management methods for all-electric ships without considering routing decisions. More recent studies have begun to integrate these decisions, often using two-stage or sequential optimization frameworks. For example, Wang et al. (2023) proposed a two-stage model that first optimizes ship arrival times at ports based on electricity prices and then optimizes the vessel's speed and energy management for the given voyage. Luo et al. (2025) created a two-stage framework that first optimizes the ship's route considering sea conditions and subsequently optimizes speed and energy management along that predetermined path. Table 2.2 positions the HEPSVRP-SO within the EVRP literature, highlighting the research gap: Existing studies have not developed unified models coordinating hybrid vessel routing, speed optimization, charging strategies, and real-time energy management under fuel quota constraints, particularly in the context of offshore supply operations where charging infrastructure is sparse and energy management is critical for operational feasibility. To address this gap, this study introduces the HEPSVRP-SO along with the first optimization model to integrate the decisions of route planning, sailing speed, charging strategies, and dynamic power allocation between diesel and electric propulsion. The proposed model is then used to systematically evaluate the operational and economic performance of hybrid PSV operations across multiple scenarios, providing quantitative evidence for pathways toward green shipping.

2.3 Conclusion

This literature review reveals a progressive evolution in maritime operational optimization, advancing from simplified power-law-based models to weather-dependent approaches and recent explorations of forecast uncertainty and dynamic replanning. However, critical research gaps remain. For conventional vessel operations, the literature lacks rigorous evaluation frameworks that systematically quantify the operational value of ensemble weather forecasts by distinguishing between prediction capability and actual performance, and lacks systematic approaches for integrating regularly updated weather forecasts throughout voyage execution. For hybrid electric vessel operations, existing PSV routing research focuses exclusively on conventional propulsion, while electric vehicle routing methodologies have not been adapted to maritime contexts. The absence of integrated optimization models simultaneously coordinating routing, speed, charging, and energy allocation represents a significant research gap. Chapters 4, 5, and 6 systematically address these gaps through developing frameworks for evaluating weather forecasts, adaptive speed planning approaches, and integrated decision-making models for hybrid propulsion systems.

Chapter 3

Data Foundation

The empirical analyses presented in Chapters 4 and 5 rely fundamentally on high-quality datasets that integrate vessel operational information with accurate meteorological conditions. This chapter establishes the data foundation for these subsequent investigations by introducing the primary data sources employed in this thesis and detailing the methodological approaches developed to enhance data quality through systematic fusion procedures.

This chapter is organized into three main sections. Section 3.1 introduces the ship operational data obtained from noon reports and the Automatic Identification System (AIS). Section 3.2 presents the meteorological data sources, including the ECMWF Reanalysis v5 (ERA5) dataset and the ensemble weather forecasts provided by the Global Ensemble Forecast System (GEFS). Section 3.3 develops two data fusion methods for integrating low-frequency noon report observations with high-resolution meteorological data. The data foundation and fusion methodologies established in this chapter provide the essential infrastructure for the fuel consumption prediction models and speed optimization frameworks developed in the subsequent chapters.

3.1 Ship Operational Data

The empirical foundation for the research presented in Chapters 4 and 5 is derived from ship operational data. This section introduces the characteristics, coverage, and limitations of the ship operational data employed in this thesis.

3.1.1 Noon Reports

The ship noon report records the sailing conditions of a ship and the sea and weather information around the ship; reports are usually submitted by the captain to shipping companies and shore management parties for the purpose of archiving and planning. The noon report data usually include ship sailing conditions, such as draft, geographic coordinates, loading conditions, and the average sailing speed since the last report. It also records snapshot information about the surrounding weather and sea conditions. For example, the wind speed information is acquired from a single reading of the anemometer by the deck officer when filling out the noon report.

The noon reports employed in this thesis are obtained from a Capesize dry bulk carrier. The reports cover the time span from February 13, 2019 to August 3, 2020, consisting of a total of 331 records. The dataset provides coverage across diverse operational scenarios, including both laden voyages when the vessel carries cargo and ballast voyages when the vessel returns empty after cargo discharge.

The major disadvantage of the ship noon report is that the weather and sea information is recorded manually by crews once a day according to their readings of onboard instruments or subjective visual inspections. Consequently, the meteorological data in the noon report are usually neither adequate nor accurate. To address this issue, more reliable meteorological data should be collected to complement the noon report.

3.1.2 Automatic Identification System Data

The AIS is widely adopted in the shipping industry to exchange sailing information between AIS-equipped terminals, such as between equipped ships and base stations (Yang et al., 2019a). The AIS typically provides static, dynamic, and voyage-related information. Because the direction of wind/waves in the meteorological data is absolute, we use the ship's heading degrees provided by the AIS data to calculate the direction of wind/waves with respect to the ship sailing direction.

3.2 Meteorological Data

As established in Section 3.1.1, the meteorological information recorded in noon reports is usually neither adequate nor accurate. To address this issue, more reliable meteorological data should be collected to complement the noon report. This section introduces the external meteorological data sources employed in this thesis.

3.2.1 ECMWF Reanalysis V5

To mitigate the limitations of the meteorological data provided by the noon reports, we seek to improve the quality of meteorological data by leveraging more reliable external sources. This thesis uses the ERA5 dataset,[1] a reanalysis dataset that is widely considered the most comprehensive and accurate representation of historical weather and climate conditions. ERA5 is a reanalysis dataset provided by the ECMWF; it is regarded as the most complete picture of past weather and climate (ECMWF, 2018). This dataset offers hourly estimations of various meteorological variables from 1979 to the current day, with a horizontal resolution of 0.25° (longitude) × 0.25° (latitude) × 1 h (time) for the atmosphere and 0.5° (longitude) × 0.5° (latitude) × 1 h (time)

[1] https://cds.climate.copernicus.eu/datasets/reanalysis-era5-single-levels.

for ocean waves. Therefore, ERA5 data are regarded as the ground truth of meteorological data, which can be combined with the noon report data for predicting the ship fuel consumption.

Twelve variables from the ERA5 data are retrieved, as shown in Table 3.1. These meteorological variables are combined with the sailing features obtained from the noon reports to develop the fuel consumption prediction model, leading to more reliable and accurate predictions of ship fuel consumption. The absolute wind direction and absolute mean wave direction from the ERA5 data are converted to relative directions with respect to the ship's true heading using the vessel heading derived from AIS data.

3.2.2 Ensemble Weather Forecasts

To improve the ship's operational efficiency, weather information is used for voyage planning before the voyage starts. Due to the inaccessibility of the actual weather data, which are referred to as ERA5 data in this thesis, weather forecast data are used for planning the ship voyage. All weather forecasts employed in this thesis are generated by the GEFS, which is a comprehensive weather forecast model published by the National Oceanic and Atmospheric Administration (NOAA) (NOAA, 2017).

The GEFS generates a range of 21 meteorological forecasts, consisting of a "control member" forecast that is based on the most accurate estimate of the current meteorological conditions and 20 "perturbed member" forecasts that are generated from the perturbed initial meteorological conditions (World Meteorological Organization, 2012). The control member forecast represents the most likely scenario, while the perturbed member forecasts provide a measure of forecast uncertainty. Similar to Hinnenthal and Clauss (2010) and Skoglund et al. (2015), the control member is regarded as a deterministic weather forecast, and the 20 perturbed weather forecasts are used as the ensemble weather forecasts.

The NOAA issues meteorological forecast data four times a day, with the timestamps of 00Z, 06Z, 12Z, and 18Z.[2] The issued dataset produces forecasts every six hours from the initial time, which refers to the starting point from which the forecast data are generated. The forecast data extend up to 384 h in the future (i.e., 16 days) from the initial time and cover various parameters such as temperature, wind speed, precipitation, and atmospheric pressure with a horizontal resolution of 0.5° × 0.5°. Three variables from the NOAA meteorological forecast data are retrieved, namely the 2-meter temperature, the wind speed, and the wind direction. These three variables share identical names in the NOAA weather forecasts and ECMWF reanalysis data. Similarly, the absolute wind directions recorded in the NOAA data are converted to the relative direction with respect to the vessel's true heading.

Table 3.1 provides a summary of features used in this thesis, organized by data source. For each feature, the table specifies the description and the feature processing applied.

3.3 Data Fusion Methods

A voyage in the noon report contains several waypoints that represent the locations traveled through by the ship, and these waypoints can be used to represent a voyage. For example, Figure 3.1(a) shows a voyage containing four waypoints (shown by yellow dots), and the voyage can be represented by (N_1, N_2, N_3, N_4). As shown in Figure 3.1(a), the meteorological data are presented in regular latitude-longitude grids, which are called meteorological grids in this thesis, and only the corner points of a meteorological grid have available meteorological values.

This section introduces two data fusion methods to combine the noon report data with external meteorological data. These methods aim to improve the quality of meteorological information, which is critical for developing accurate fuel consumption prediction models.

[2] The "Z" in "00Z" stands for "Zulu time," which is the same as Coordinated Universal Time.

Table 3.1: Features in each data source

Features	Description	Feature processing
Ship noon report		
Hourly fuel consumption rate of the main engine (ME)	The average fuel consumption per hour of the ship's main engine in unit mt/hour.	The hourly fuel consumption rate is obtained by dividing the overall fuel consumption by the total sailing time.
Average sailing speed	The average sailing speed since the last report in unit knot.	N.A.
Sailing condition	Ballast voyage, if the ship is not loaded; laden voyage, if the ship is loaded.	"1" is used to present "laden voyage" and "0" is used to present "ballast voyage".
Draught	The vertical distance between the waterline and the bottom of the hull (keel) in unit meter.	N.A.
ERA5 reanalysis data		
10-meter wind direction (WD)	This parameter measures the wind direction with respect to the ship's true heading in unit degrees at the height of 10 meters.	The absolute wind direction is transformed to a relative direction to the ship's heading degree from AIS data.
10-meter wind speed (WS)	Wind speed at the height of 10 m in unit m/s.	N.A.
Mean period of wind waves (MPWW)	This parameter is the average time in seconds which is required for two successive wind-sea wave crests to pass a given point on the surface of the sea.	N.A.
Mean direction of wind waves (MDWW)	This parameter is the mean true direction of wind-sea waves with different heights, lengths, and directions in unit degree.	The absolute direction is transformed to a relative direction to the ship's heading degree from AIS data.
Significant height of wind waves (SHWW)	This parameter represents the perpendicular distance from the wind-sea wave trough to the wind-sea wave crest in unit meter.	N.A.
Mean period of total swell (MPTS)	This parameter is the average time in seconds which is required for two successive total swell crests to pass a given point on the surface of the sea.	N.A.
Significant height of total swell (SHTS)	This parameter represents the perpendicular distance from the total swell trough to the total swell wave crest in unit meter.	N.A.
Mean direction of total swell (MDTS)	This parameter is the mean true direction of swells with different heights, lengths, and directions in unit degree.	N.A.
Significant height of combined wind waves and swell (SHCWWS)	This parameter measures the perpendicular distance in meters from the wave trough and the wave crest. Both wind-sea waves and swells are considered.	N.A.
Mean wave direction (MWP)	This parameter is the mean true direction of waves with different heights, lengths, and directions in unit degrees. Both wind-sea waves and swells are considered.	The absolute direction is transformed to a relative direction to the ship's heading degree from AIS data.
Forecast surface roughness (FSR)	This parameter measures the surface resistance depending on the waves over the ocean. It is the aerodynamic roughness length in unit meter.	N.A.
2-meter temperature (2mT)	This parameter represents the highest temperature of the air at two meters above the surface of land, sea, or inland water in unit degrees Celsius (°C).	Temperatures measured in kelvin are converted to °C by subtracting 273.15.
GEFS ensemble forecast		
10-meter wind direction	This parameter measures the wind direction with respect to the ship's true heading in unit degrees at the height of 10 meters.	The absolute direction is converted to a relative direction to the ship's heading degree from AIS data.
10-meter wind speed	This parameter measures the wind speed at the height of 10 m in unit m/s.	N.A.
2-meter temperature	This parameter represents the highest temperature of the air at two meters above the surface of land, sea, or inland water in unit °C.	Temperatures measured in kelvin are converted to °C by subtracting 273.15.

3.3.1 Direct Fusion Method

External meteorological data, which are highly accurate, are incorporated with the original noon report data in this section. Specifically, the first step is to locate the meteorological grid in the external meteorological data in which the waypoint falls, using the geographical coordinates recorded in the noon report. Due to the limited spatial resolution of the meteorological data, there might be no exact meteorological value for the location of the waypoint according to the noon report. To address this issue, the second step is to use the inverse distance weighted (IDW) interpolation method to calculate the meteorological values of the specific voyage point recorded in the noon report (Chen and Liu, 2012). The IDW interpolation method is a deterministic spatial interpolation approach to estimating an unknown value by calculating the weighted average of values of the neighborhood measured points. The weight of a measured point is the inverse distance of the measured point to the unmeasured point.

As shown in Figure 3.1(b), the sea/weather conditions of waypoint N_1 are unknown. As waypoint N_1 is located in the meteorological grid G and the meteorological values of the corner points of G are known, the meteorological values of N_1 can be calculated based on the corner points. Taking the calculation of wind speed at N_1 as an example, the formulas of the IDW method are as follows:

$$Z_{N_1} = \sum_{i=1}^{4} w_i \times Z_{G_i} \tag{3.1}$$

$$w_i = \frac{\frac{1}{d_i}}{\sum_{i=1}^{4} \frac{1}{d_i}}, \tag{3.2}$$

where Z_{N_1} is the unknown wind speed at waypoint N_1; w_i is the weight; Z_{G_i} is the wind speed of measured poin,G_i, i.e., the corner point of the meteorological grid; and d_i is the distance of N_1 to the measured point G_i $(i = 1, 2, 3, 4)$.

3.3.2 Rhumb Line-based Fusion Method

In the fusion method presented in Section 3.3.1, the weather data acquired at a time point cannot accurately represent the weather condition during the whole day. To provide more accurate weather information, it is necessary to interpolate more points traveled on ship's course between two consecutive waypoints of the voyage according to the temporal resolution of the meteorological data. As shown in Figure 3.1(a), 23 points ($P_1, P_2, ..., P_{23}$) between waypoints N_1 and N_2 are interpolated points, and their geographic positions can be estimated by the rhumb line route. Then, the average sea/weather data along the interpolated points can be used as a substitute for the original sea/weather conditions in the noon report.

Specifically, the first key step is to calculate the number of interpolated points according to the temporal resolution of the external meteorological data source and estimate their geographic positions by interpolation. For example, if the noon report is combined with the ERA5 data at a temporal resolution of 1 h, we can interpolate 23 points (as the time difference between two consecutive noon reports is 24 h); if the noon report is combined with the GEFS ensemble forecast data at a temporal resolution of 6 h, we can interpolate 3 points. The whole voyage is divided into a number of segments, on which the ship sails following a rhumb line route (Weintrit and Kopacz, 2011). For example, the voyage in Figure 3.1(a) is segmented into three parts, and the ship moves along the rhumb line on the $N_1 \rightarrow N_2$, $N_2 \rightarrow N_3$, and $N_3 \rightarrow N_4$ segments. The start points of the three segments are N_1, N_2, and N_3, respectively. The latitude and longitude of each interpolated point can be calculated according to the formulas shown below (Williams, 2013):

$$\delta = \frac{d}{R} \tag{3.3}$$

$$\phi_2 = \phi_1 + \delta \times \cos\theta \tag{3.4}$$

$$\Delta\psi = \ln\left(\frac{\tan\left(\frac{\pi}{4}+\frac{\phi_2}{2}\right)}{\tan\left(\frac{\pi}{4}+\frac{\phi_1}{2}\right)}\right) \tag{3.5}$$

$$q = \begin{cases} \cos\varphi, & \text{the ship sails on constant latitude course } \varphi \\ \frac{|\phi_2-\phi_1|}{\Delta\psi}, & \text{otherwise} \end{cases} \tag{3.6}$$

$$\Delta\lambda = \delta \times \frac{\sin\theta}{q} \tag{3.7}$$

$$\lambda_2 = \lambda_1 + \Delta\lambda, \tag{3.8}$$

where δ is the angular distance; d represents the distance between the start point of a segment and the interpolated point on the segment, and is calculated by multiplying the average sailing speed (knots) by sailing time (h); R is the earth's radius; λ_1 and ϕ_1 are the longitude and latitude of the start point, respectively; λ_2 and ϕ_2 are the longitude and latitude of the interpolated point, respectively; ln is the natural log; θ is the ship's heading degree provided by the AIS; and $\Delta\lambda$ is the longitude difference.

When the ship sails on segment $N_1 \to N_2$, the coordinates of the hourly interpolated points $(P_1, P_2, ..., P_{23})$ between waypoints N_1 and N_2 can be obtained using equations (3.3)–(3.8), as shown in Figure 3.1(a). Second, the sea and weather data at each interpolated point can be retrieved by the IDW method. Finally, we average the meteorological values of the interpolated points between the current and previous waypoint, and that of the current waypoint itself. This average is used to replace the original meteorological values of the current waypoint. For example, the wind speeds at the interpolated points are $(Z_{P_1}, Z_{P_2}, ..., Z_{P_{23}})$, and the wind speed at waypoint N_2 is Z_{N_2}. Then the recorded wind speed of N_2 in the noon report is replaced by the average wind speed $(Z_{P_1} + Z_{P_2} + ... + Z_{P_{23}} + Z_{N_2})/24$. Note that the meteorological values at the start point of the first segment in a voyage, i.e., waypoint N_1 of the voyage in Figure 3.1(a), are calculated using the fusion method introduced in Section 3.3.1.

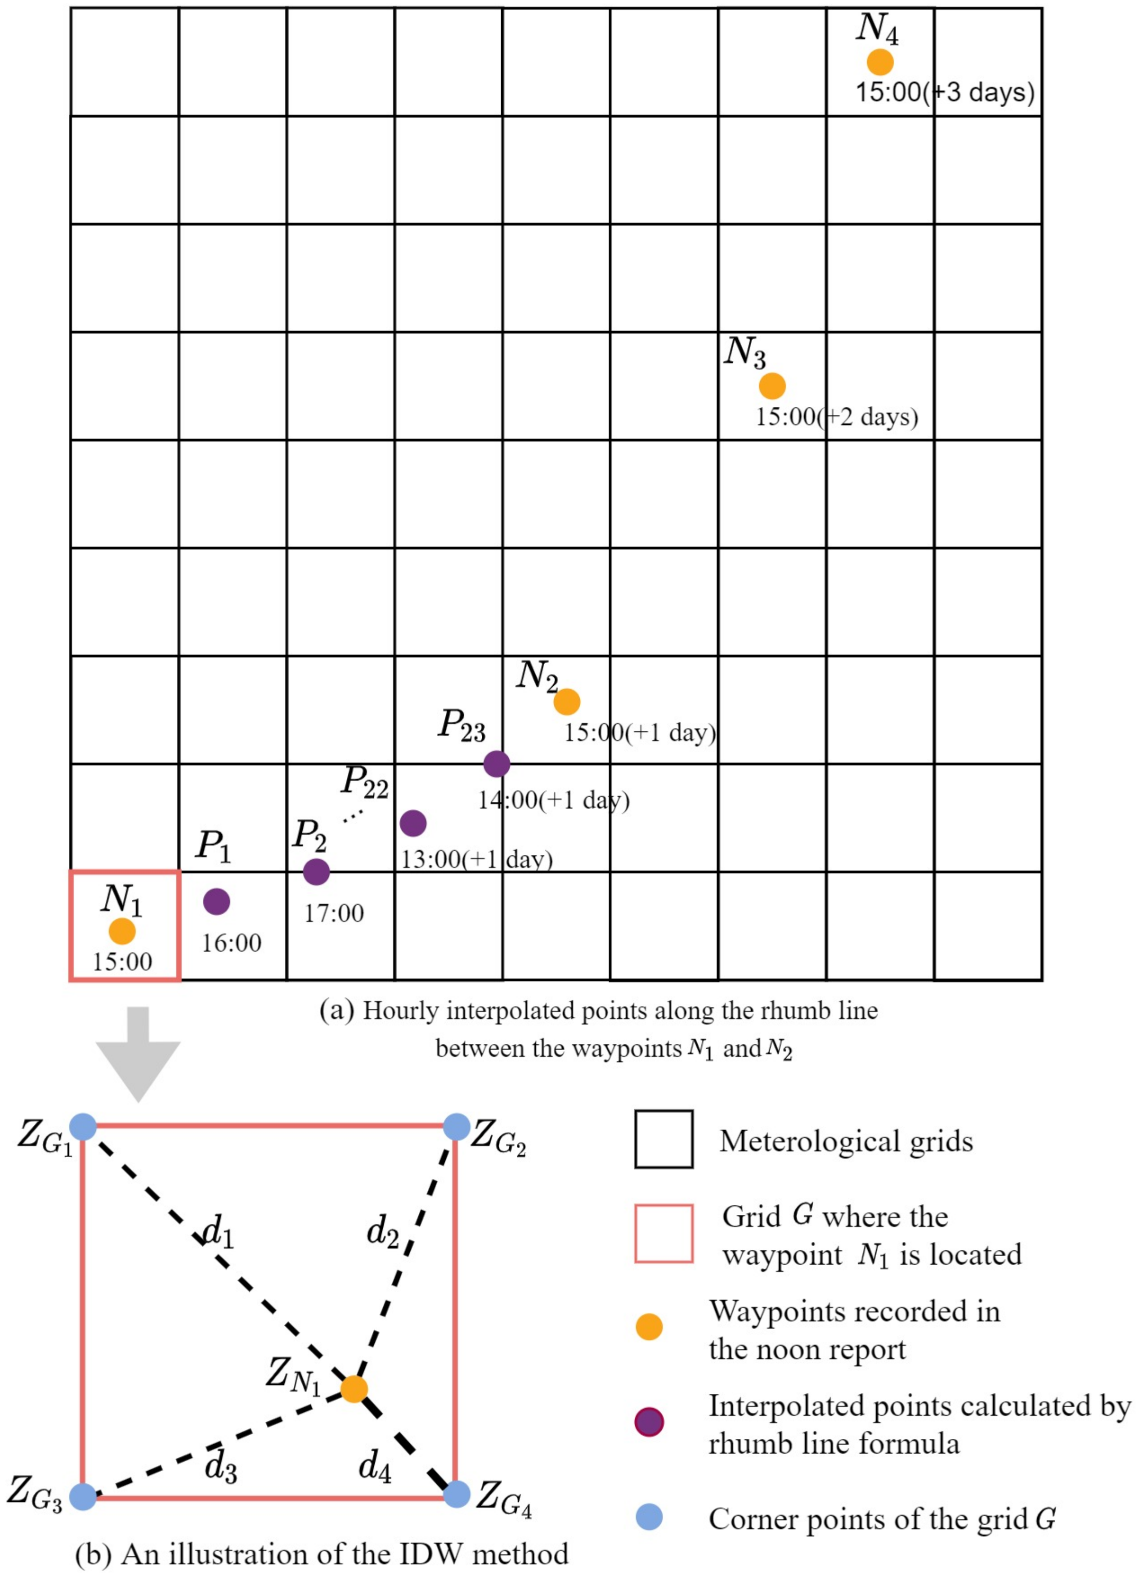

Figure 3.1: Graphic explanation of the fusion method based on the rhumb line.

These two data fusion methods provide alternative approaches to integrating noon report data with external meteorological data. To scientifically determine which fusion method produces more accurate meteorological representations for fuel consumption prediction, their performance will be compared through empirical evaluation in Chapter 4, and the superior method will be selected as the standard data processing approach for all subsequent analyses in this thesis.

3.4 Conclusion

This chapter has established the comprehensive data foundation that underpins the empirical research conducted in Chapters 4 and 5. Three primary data sources have been introduced: ship noon reports providing vessel operational parameters, AIS data supplying vessel heading information, and external meteorological datasets comprising ERA5 reanalysis data for historical conditions and GEFS ensemble forecasts for prospective planning. To address the inadequacy and limited accuracy of meteorological observations in noon reports, two data fusion methods have been developed. The direct fusion method provides a baseline approach through spatial interpolation at waypoint locations, while the rhumb line-based fusion method employs temporal interpolation along vessel trajectories to obtain more representative segment-average meteorological conditions. The comparative performance of these fusion methods will be empirically evaluated in Chapter 4 to determine the optimal data processing approach for subsequent fuel consumption modeling and speed optimization analyses.

Chapter 4

Ship Speed Optimization under Weather Forecast Uncertainty: An Evaluation Framework

Existing ship speed optimization research predominantly relies on deterministic meteorological forecasts, implicitly assuming that a single predicted atmospheric trajectory adequately represents future conditions. While ensemble forecasting systems that characterize prediction uncertainty through multiple plausible scenarios have been operationally available for decades, their potential value for ship speed optimization remains unexplored. As identified in Section 2.1.3 of the literature review, pioneering studies demonstrated potential benefits of ensemble weather forecasts but suffered from a limitation: they lacked evaluation frameworks that explicitly distinguish prediction capability from physical reality and use reliable historical data to realistically simulate the planning-execution gap. Therefore, a research gap remains regarding the evaluation of ensemble weather forecasts in ship speed optimization.

This chapter addresses this methodological gap by developing an innovative computational evaluation framework that simulates the complete planning-execution cycle through a dual-model architecture, overcoming the impossibility of conducting controlled real-world experiments where the same vessel sails identical routes under different forecast types.

The remainder of this chapter is organized as follows. Section 4.1 presents the speed optimization models based on deterministic and ensemble weather forecasts. Section 4.2 develops the dual-model evaluation framework. Section 4.3 reports computational experiments on two validation voyages. Section 4.4 summarizes the findings and identifies the limitations of static speed optimization that motivate the investigation of adaptive ship speed optimization in Chapter 5.

4.1 Ship Speed Optimization Model: A Predict-then-Optimize Approach

This section establishes the speed optimization models for comparison. The models consist of two key components: a fuel consumption prediction model and an optimization model that minimizes the ship's fuel consumption subject to operational constraints. Section 4.1.1 presents the gradient boosting regression trees (GBRT) prediction model, while Section 4.1.2 formulates the optimization problem under both deterministic and ensemble weather forecasts.

4.1.1 Fuel Consumption Prediction Model: Gradient Boosting Regression Trees

The fuel consumption prediction model serves as the foundation for ship speed optimization, providing estimates of fuel consumption rates under varying operational and environmental conditions. We employ GBRT, a powerful ensemble learning method that constructs a strong predictor by sequentially combining multiple weak learners (decision trees) through gradient-based optimization (Friedman, 2001).

The GBRT algorithm operates through an iterative boosting process. At each iteration n, a new decision tree is trained to approximate the negative gradient of the loss function with respect to the current ensemble's predictions. This new tree is then added to the ensemble with an appropriate weight, progressively reducing prediction errors. Algorithm 1 presents the complete construction procedure for a GBRT model. The algorithm initializes with a simple constant prediction (the mean of training targets), then iteratively fits new trees to the residuals (negative gradients) of the current model. Each leaf node of the new tree is assigned an optimal value that minimizes the loss function for samples falling into that region. This process continues for a predefined number of iterations N_{GBRT}, yielding a final additive model that combines predictions from all trees.

Algorithm 1 Construction of a GBRT model

Require: Training set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N}$; Number of decision trees N_{GBRT}
Ensure: A GBRT model $\mathcal{F}_{\text{GBRT}}(\mathbf{x})$

1: Initialize: $f_0(\mathbf{x}) = \arg\min_\rho \sum_{i=1}^{N} L(y_i, \rho)$, where $L(y_i, \rho) = \frac{1}{2}(y_i - \rho)^2$ is half of the squared loss. The value of ρ is estimated by minimizing the loss function, and its optimal value is $\frac{1}{N}\sum_{i=1}^{N} y_i$.
2: **for** $n = 1$ to N_{GBRT} **do**
3: Compute residuals: $r_{in} = -\left[\frac{\partial L(y_i, f(\mathbf{x}_i))}{\partial f(\mathbf{x}_i)}\right]_{f(\mathbf{x}) = f_{n-1}(\mathbf{x})}$ for $i = 1, \ldots, N$.
4: Fit a new tree to residuals: Train the decision tree on $\{(\mathbf{x}_i, r_{in})\}_{i=1}^{N}$ with leaf regions R_{jn} $(j = 1, \ldots, J)$.
5: **for** each leaf region R_{jn} **do**
6: Compute optimal leaf value: $\gamma_{jn} = \arg\min_\gamma \sum_{\mathbf{x}_i \in R_{jn}} L(y_i, f_{n-1}(\mathbf{x}_i) + \gamma)$
7: **end for**
8: Update model: $f_n(\mathbf{x}) = f_{n-1}(\mathbf{x}) + \sum_{j=1}^{J} \gamma_{jn} \mathbb{I}(\mathbf{x} \in R_{jn})$, where $\mathbb{I}(\mathbf{x} \in R_{jn})$ takes the value of 1 if $\mathbf{x} \in R_{jn}$, and 0 otherwise.
9: **end for**
10: **return** $\mathcal{F}_{\text{GBRT}}(\mathbf{x}) = f_0(\mathbf{x}) + \sum_{n=1}^{N_{\text{GBRT}}} \sum_{j=1}^{J} \gamma_{jn} \mathbb{I}(\mathbf{x} \in R_{jn})$

In the context of ship fuel consumption prediction, the GBRT model takes as input a feature vector $\mathbf{x} = [v, \mathbf{w}, \mathbf{u}]$ comprising the sailing speed v, meteorological conditions $\mathbf{w}$, and vessel's operational state $\mathbf{u}$, and outputs the predicted hourly FCR (mt/h). The meteorological conditions $\mathbf{w}$ encapsulate environmental factors such as wind speed, wind direction, wave characteristics, and other relevant atmospheric and oceanic variables that influence ship resistance and propulsion efficiency. The vessel's operational state $\mathbf{u}$ includes attributes such as draft and loading condition (laden or ballast), which affect the vessel's hydrodynamic performance. The model function can be formally expressed as:

$$\mathcal{F}_{\text{GBRT}}^{\text{planner}}(v, \mathbf{w}, \mathbf{u}; \boldsymbol{\theta}_{\text{GBRT}}) \rightarrow \text{Hourly fuel consumption rate (mt/h)}, \tag{4.1}$$

where $\boldsymbol{\theta}_{\text{GBRT}}$ represents the learned parameters of the GBRT model, including the structure and weights of all decision trees. The specific features comprising $\mathbf{w}$ and $\mathbf{u}$, along with the model's training procedure and predictive performance, are detailed in Section 4.3.1. The superscript "planner" distinguishes this model's role within the evaluation framework developed in Section 4.2, where it functions as the decision-maker's prediction tool during the speed optimization phase. This superscript clarifies that the model represents the imperfect predictive capability available to planners when making speed decisions, as opposed to the ground-truth evaluator model introduced later.

4.1.2 Modeling the Speed Optimization

This section formulates the speed optimization problem as a mathematical programming model. The optimization seeks to determine the optimal sailing speed for each segment of the voyage such that total fuel consumption is minimized while satisfying arrival time constraints and operational speed limits. This section first presents the notation and problem description, then introduces the baseline mixed-integer linear programming formulation, and finally demonstrates how different meteorological forecast types lead to distinct optimization objectives. Table 4.1 summarizes the notation employed in the speed optimization model formulations.

Table 4.1: Notation for speed optimization model formulations

Symbol	Description	Unit
Sets and Indices		
$\mathcal{S}$	Set of voyage segments, $\{1, 2, ..., S\}$	-
$s \in \mathcal{S}$	Segment index	-
Parameters		
Voyage Structure		
S	Number of segments in the voyage	-
L_s	Length of segment s	nautical miles
L	Total voyage distance, $L = \sum_{s=1}^{S} L_s$	nautical miles
$T_{\max}$	Maximum allowable sailing time for the voyage	hours
Speed Discretization		
$v_{\min}$	Minimum allowable sailing speed	knots
$v_{\max}$	Maximum allowable sailing speed	knots
M	Number of candidate speeds	-
Ship Fuel Consumption		
$\mathcal{F}_{\text{GBRT}}^{\text{planner}}(\cdot)$	GBRT planner model for the prediction of ship's FCR	metric tons/hours
F_s	Ship's fuel consumption on segment s	metric tons
$\mathbf{w}_s$	Meteorological conditions for segment s	-
$\mathbf{u}_s$	Vessel operational state for segment s	-
Meteorological Forecasts		
$\mathbf{W}^{\text{det}}$	Deterministic weather forecast (GEFS control member)	-
$\mathbf{W}_k^{\text{perturbed}}$	The kth perturbed member of ensemble weather forecasts	-
K	Number of perturbed members in ensemble forecast	-
Decision Variables		
v_s	Sailing speed on segment s (continuous)	knots
v_s^m	The mth candidate speed for segment s (discretized)	knots
y_s^m	Binary variable: 1 if speed v_s^m is selected for segment s, 0 otherwise	-
t_s	The time of arrival to the beginning of segment s	hours
t_{S+1}	The time of arrival at the end of segment S	hours

4.1.2.1 Problem Description

The speed optimization problem addresses a vessel's journey from an origin port to a destination port over a predetermined route. The voyage distance L and maximum allowable sailing time $T_{\max}$ are specified, typically determined by contractual obligations or operational schedules. The ship's loading condition is predetermined and fixed, and the meteorological conditions during the voyage can be acquired in advance from weather forecasts. Because the sea conditions are highly dynamic, the whole route is divided into a number of segments according to the corresponding records in the noon report. We assume that the meteorological conditions and the international conventions the ship must comply with can be viewed as identical within each segment. The minimum and maximum allowable ship sailing speeds are also assumed to be the same in each segment. The departure time of the ship is set to 0, and the estimated time of arrival at destination port should not exceed the latest allowable time of arrival.

4.1.2.2 Mathematical Formulation

The continuous voyage is discretized into S segments, forming the set $\mathcal{S} = \{1, 2, ..., S\}$. Each segment s represents a portion of the route with length L_s, where $\sum_{s=1}^{S} L_s = L$. The decision-maker must select a sailing speed for each segment of the voyage to minimize total fuel consumption while ensuring timely arrival. For each segment s, the meteorological conditions $\mathbf{w}_s$ are extracted from the available weather forecast data, and the vessel operational state $\mathbf{u}_s$ is determined from the operational records. The sailing speed v_s on segment s directly determines both the segment duration $\Delta\tau_s = L_s/v_s$ and the ship's FCR via the prediction model $\mathcal{F}_{\text{GBRT}}^{\text{planner}}(v_s, \mathbf{w}_s, \mathbf{u}_s)$. The auxiliary decision variable is the time of arrival t_s to the beginning of segment s. In addition, t_{S+1} presents the time of arrival at the end of segment S. The optimization objective is to determine the speed profile $\{v_s\}_{s=1}^{S}$ that minimizes the total fuel consumption.

The model is denoted as [**M1**]:

$$F^{\text{total}} = \sum_{s=1}^{S} \mathcal{F}_{\text{GBRT}}^{\text{planner}}(v_s, \mathbf{w}_s, \mathbf{u}_s) \cdot \frac{L_s}{v_s} \tag{4.2}$$

subject to:

$$t_{s+1} = t_s + \frac{L_s}{v_s}, \quad \forall s \in \mathcal{S} \tag{4.3}$$

$$t_1 = 0 \tag{4.4}$$

$$t_{S+1} \leq T_{\max} \tag{4.5}$$

$$v_{\min} \leq v_s \leq v_{\max}, \quad \forall s \in \mathcal{S} \tag{4.6}$$

$$t_s \geq 0, \quad \forall s \in \mathcal{S} \cup \{S+1\}. \tag{4.7}$$

Constraints (4.3) specify the relationship between the time of arrival at the beginning of segment s and at the beginning of segment $s+1$. In this set of constraints, t_s represents the time of arrival at the beginning of segment s, where $s \in \mathcal{S}$. Therefore, as shown in constraints (4.4), $t_1 = 0$ represents the time of arrival at the beginning of segment 1. t_{S+1} represents the time of arrival at the end of segment S. Constraints (4.5) guarantee that the ship's time of arrival at the destination port does not exceed the allowable arrival time. Constraints (4.6) bound the sailing speed in the interval of $[v_{\min}, v_{\max}]$. Constraints (4.7) ensure that the time of arrival at the beginning of every segment is nonnegative.

The optimization problem described above is inherently nonlinear due to the division by v_s in both the objective function and time constraint, as well as the nonlinear relationship captured by the GBRT prediction model. To enable an efficient solution using commercial optimization solvers, we reformulate the problem as a mixed-integer linear programming (MILP) model through the discretization transformation.

The continuous speed range $[v_{\min}, v_{\max}]$ for each segment s is discretized into a finite set of candidate speeds $\{v_s^1, v_s^2, ..., v_s^M\}$, where M denotes the number of speed candidates for each segment. The discretization typically employs a uniform grid (e.g., 0.1-knot intervals) across the feasible speed range. Binary decision variables $y_s^m \in \{0, 1\}$ are introduced, where $y_s^m = 1$ if the mth candidate speed v_s^m is selected for segment s, and $y_s^m = 0$ otherwise. With these transformations, the nonlinear model can be reformulated as the following the MILP model, which we denote as [**M2**]:

$$\min \quad \sum_{s=1}^{S} \sum_{m=1}^{M} y_s^m \cdot \mathcal{F}_{\text{GBRT}}^{\text{planner}}(v_s^m, \mathbf{w}_s, \mathbf{u}_s) \cdot \frac{L_s}{v_s^m} \tag{4.8a}$$

$$\text{s.t.} \quad t_{s+1} = t_s + \sum_{m=1}^{M} \frac{L_s}{v_s^m} \cdot y_s^m, \quad \forall s \in \mathcal{S} \tag{4.8b}$$

$$t_1 = 0 \tag{4.8c}$$

$$t_{S+1} \leq T_{\max} \tag{4.8d}$$

$$\sum_{m=1}^{M} y_s^m = 1, \quad \forall s \in \mathcal{S} \tag{4.8e}$$

$$y_s^m \in \{0, 1\}, \quad \forall s \in \mathcal{S}, m \in \{1, ..., M\} \tag{4.8f}$$

$$t_s \geq 0, \quad \forall s \in \mathcal{S} \cup \{S + 1\}. \tag{4.8g}$$

The objective function (4.8a) minimizes the total fuel consumption across all segments. For segment s, the ship's fuel consumption is computed as the product of the fuel consumption rate (predicted by the GBRT model for the selected speed and conditions) and the sailing duration L_s/v_s^m. Constraint (4.8b) tracks the cumulative sailing time, ensuring temporal consistency across segments. Constraint (4.8c) sets the starting time reference, while constraint (4.8d) enforces the deadline requirement. Constraint (4.8e) ensures that exactly one candidate speed is selected for each segment. Constraints (4.8f) and (4.8g) define the variable domains.

This MILP formulation is linear in the decision variables y_s^m and t_s, and the coefficients $\mathcal{F}_{\text{GBRT}}^{\text{planner}}(v_s^m, \mathbf{w}_s, \mathbf{u}_s)$ are precomputed by evaluating the nonlinear GBRT model at all discretized speed-condition combinations. Consequently, model [**M2**] can be solved efficiently using commercial solvers. The model [**M2**] provides the structural foundation for optimization, but the specific meteorological input $\mathbf{w}_s$ fundamentally shapes the optimization objective.

4.1.2.3 Speed Optimization under Deterministic Weather Forecasts

When deterministic weather forecasts are available, the weather conditions for each segment are uniquely specified. Similar to Skoglund et al. (2015) and Hinnenthal and Clauss (2010), the control member is regarded as the deterministic weather forecast $\mathbf{W}^{\text{det}}$. For each segment s, the meteorological conditions are extracted from the deterministic weather forecast:

$$\mathbf{w}_s^{\text{det}} = \mathcal{W}(\mathbf{W}^{\text{det}}, s), \tag{4.9}$$

where the function $\mathcal{W}(\cdot, s)$ retrieves the weather forecast data corresponding to the starting location and time of segment s from the complete forecast dataset. These conditions representing the meteorological environment for the entire segment. Substituting $\mathbf{w}_s = \mathbf{w}_s^{\text{det}}$ into model [**M2**] yields the deterministic speed optimization model, which we denote as [**M-DET**]:

$$\begin{aligned}
\min \quad & \sum_{s=1}^{S} \sum_{m=1}^{M} y_s^m \cdot \mathcal{F}_{\text{GBRT}}^{\text{planner}}(v_s^m, \mathbf{w}_s^{\text{det}}, \mathbf{u}_s) \cdot \frac{L_s}{v_s^m} && \text{(4.10a)} \\
\text{s.t.} \quad & \text{Constraints (4.8b)} - \text{(4.8g)}. && \text{(4.10b)}
\end{aligned}$$

Solving model [**M-DET**] produces an optimal speed profile $\{v_s^{\text{det}}\}_{s=1}^{S}$ that minimizes the predicted total fuel consumption. The corresponding segment durations are $\{\Delta\tau_s^{\text{det}}\}_{s=1}^{S}$, where $\Delta\tau_s^{\text{det}} = L_s / v_s^{\text{det}}$.

The deterministic forecast-based approach represents the conventional paradigm in ship speed optimization: a single forecast scenario informs the optimization model.

However, the deterministic weather forecast captures only the expected trajectory of the atmospheric state, neglecting the range of possible future conditions. This limitation motivates the development of an alternative optimization approach that explicitly accounts for forecast uncertainty through ensemble weather forecasts.

4.1.2.4 Speed Optimization under Ensemble Weather Forecasts

Ensemble forecasting systems, such as NOAA GEFS, quantify meteorological uncertainty by generating multiple plausible future atmospheric states through perturbations of initial conditions and model physics. The GEFS ensemble comprises $K = 20$ perturbed members $\{\mathbf{W}_k^{\text{perturbed}}\}_{k=1}^{K}$, each representing a distinct but equally probable scenario. These members collectively characterize the probability distribution of future weather conditions. In the context of speed optimization, each ensemble member k implies different meteorological conditions for segment s:

$$\mathbf{w}_s^k = \mathcal{W}(\mathbf{W}_k^{\text{perturbed}}, s), \quad k = 1, ..., K. \tag{4.11}$$

A speed decision v_s on segment s will therefore result in different fuel consumption rates depending on which ensemble member's weather actually materializes. Recognizing that all ensemble members are equally likely under the ensemble forecasting framework, a rational approach to optimization under uncertainty is to minimize the expected ship fuel consumption, where the expectation is taken over all ensemble members. Under the assumption that all ensemble members are equally probable, the expectation operator can be computed as an arithmetic mean over the K members:

$$\begin{aligned} &\mathbb{E}\left[\sum_{s=1}^{S}\sum_{m=1}^{M} y_s^m \cdot \mathcal{F}_{\text{GBRT}}^{\text{planner}}(v_s^m, \mathbf{w}_s, \mathbf{u}_s) \cdot \frac{L_s}{v_s^m}\right] \\ &= \sum_{s=1}^{S}\sum_{m=1}^{M} y_s^m \cdot \left[\frac{1}{K}\sum_{k=1}^{K} \mathcal{F}_{\text{GBRT}}^{\text{planner}}(v_s^m, \mathbf{w}_s^k, \mathbf{u}_s)\right] \cdot \frac{L_s}{v_s^m}. \end{aligned} \tag{4.12}$$

Formally, the ensemble forecast-based optimization model, denoted as [**M-ENS**], seeks to minimize the expected total fuel consumption:

$$\min \quad \sum_{s=1}^{S} \sum_{m=1}^{M} y_s^m \cdot \left[\frac{1}{K} \sum_{k=1}^{K} \mathcal{F}_{\text{GBRT}}^{\text{planner}}(v_s^m, \mathbf{w}_s^k, \mathbf{u}_s) \right] \cdot \frac{L_s}{v_s^m} \tag{4.13a}$$

$$\text{s.t.} \quad \text{Constraints (4.8b)} - \text{(4.8g)}. \tag{4.13b}$$

Solving model [**M-ENS**] produces an optimal speed profile $\{v_s^{\text{ens}}\}_{s=1}^{S}$ that minimizes expected fuel consumption in an uncertain environment. The corresponding segment durations are $\{\Delta\tau_s^{\text{ens}}\}_{s=1}^{S}$, where $\Delta\tau_s^{\text{ens}} = L_s/v_s^{\text{ens}}$.

The ensemble forecast-based model [**M-ENS**] represents a fundamentally different optimization philosophy compared to the deterministic approach [**M-DET**]. Rather than committing to a single forecast scenario and optimizing for that specific case, the ensemble approach acknowledges forecast uncertainty and seeks a solution that performs well on average across the range of plausible scenarios. This strategy aligns with stochastic optimization principles, which prioritize expected performance when future conditions are uncertain.

4.2 Evaluation Framework

To scientifically compare the performance of deterministic versus ensemble weather forecasts in ship speed optimization, we must overcome a methodological challenge identified in the literature review: existing approaches lack a rigorous framework that explicitly distinguishes between prediction capability used for planning and physical reality governing actual performance. In real-world operations, it is impossible to conduct parallel voyages under identical conditions to directly observe which weather forecast type yields superior fuel efficiency. This section develops an innovative evaluation framework to address challenges through a dual-model architecture, enabling performance comparison within a controlled computational environment that realistically simulates the planning-execution gap. The proposed framework simulates the complete decision-making and execution cycle through four interconnected steps, as illustrated in Figure 4.1. The notation employed throughout this section is defined in Table 4.2.

Step 1: Data Fusion for Dataset Construction. The evaluation framework begins with the construction of high-quality datasets by fusing low-frequency noon report data with high-resolution ERA5 reanalysis data. As introduced in Chapter 3, two fusion methods are available: the direct fusion method and the rhumb line-based fusion method. The noon report dataset is partitioned into two disjoint subsets: NR_1 (containing N_1 records) for model training and testing, and NR_2 (containing N_2 records) for weather forecast comparison. By spatiotemporally matching noon report NR_1 with ERA5 reanalysis data through the fusion procedures detailed in Section 3.3, we obtain two candidate datasets: $\mathcal{D}^{\text{direct}} = \{(v_i, \mathbf{w}_i^{\text{ERA5}}, \mathbf{u}_i, y_i) \mid i = 1, ..., N_1\}$ using the direct method, and $\mathcal{D}^{\text{rhumb}} = \{(v_i, \mathbf{w}_i^{\text{ERA5}}, \mathbf{u}_i, y_i) \mid i = 1, ..., N_1\}$ using the rhumb line-based method. Here, for each record i in the noon report, v_i represents the sailing speed, $\mathbf{w}_i^{\text{ERA5}} \in \mathbb{R}^{d_w}$ is the meteorological feature vector from ERA5, $\mathbf{u}_i \in \mathbb{R}^{d_u}$ is the vessel operational state, and y_i is the hourly FCR.

To determine which fusion method produces more accurate meteorological representations, we train two XGBoost models on the respective datasets and compare their predictive performance on held-out test data: $\mathcal{F}_{\text{XGB}}^{\text{direct}}$ trained on $\mathcal{D}^{\text{direct}}$ and $\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$ trained on $\mathcal{D}^{\text{rhumb}}$. The fusion method yielding superior model performance is selected for constructing all subsequent datasets. The corresponding XGBoost model becomes the *evaluator model* $\mathcal{F}_{\text{XGB}}^{\text{evaluator}}$, which will serve as the ground truth predictor in the evaluation phase. Additionally, the noon report dataset NR_2 is fused with ERA5 data using the selected method to produce the evaluation dataset $\mathcal{D}^{\text{eval}} = \{(v_i, \mathbf{w}_i^{\text{ERA5}}, \mathbf{u}_i, y_i) \mid i = 1, ..., N_2\}$, which represents actual meteorological conditions for comparing weather forecast types in the subsequent steps.

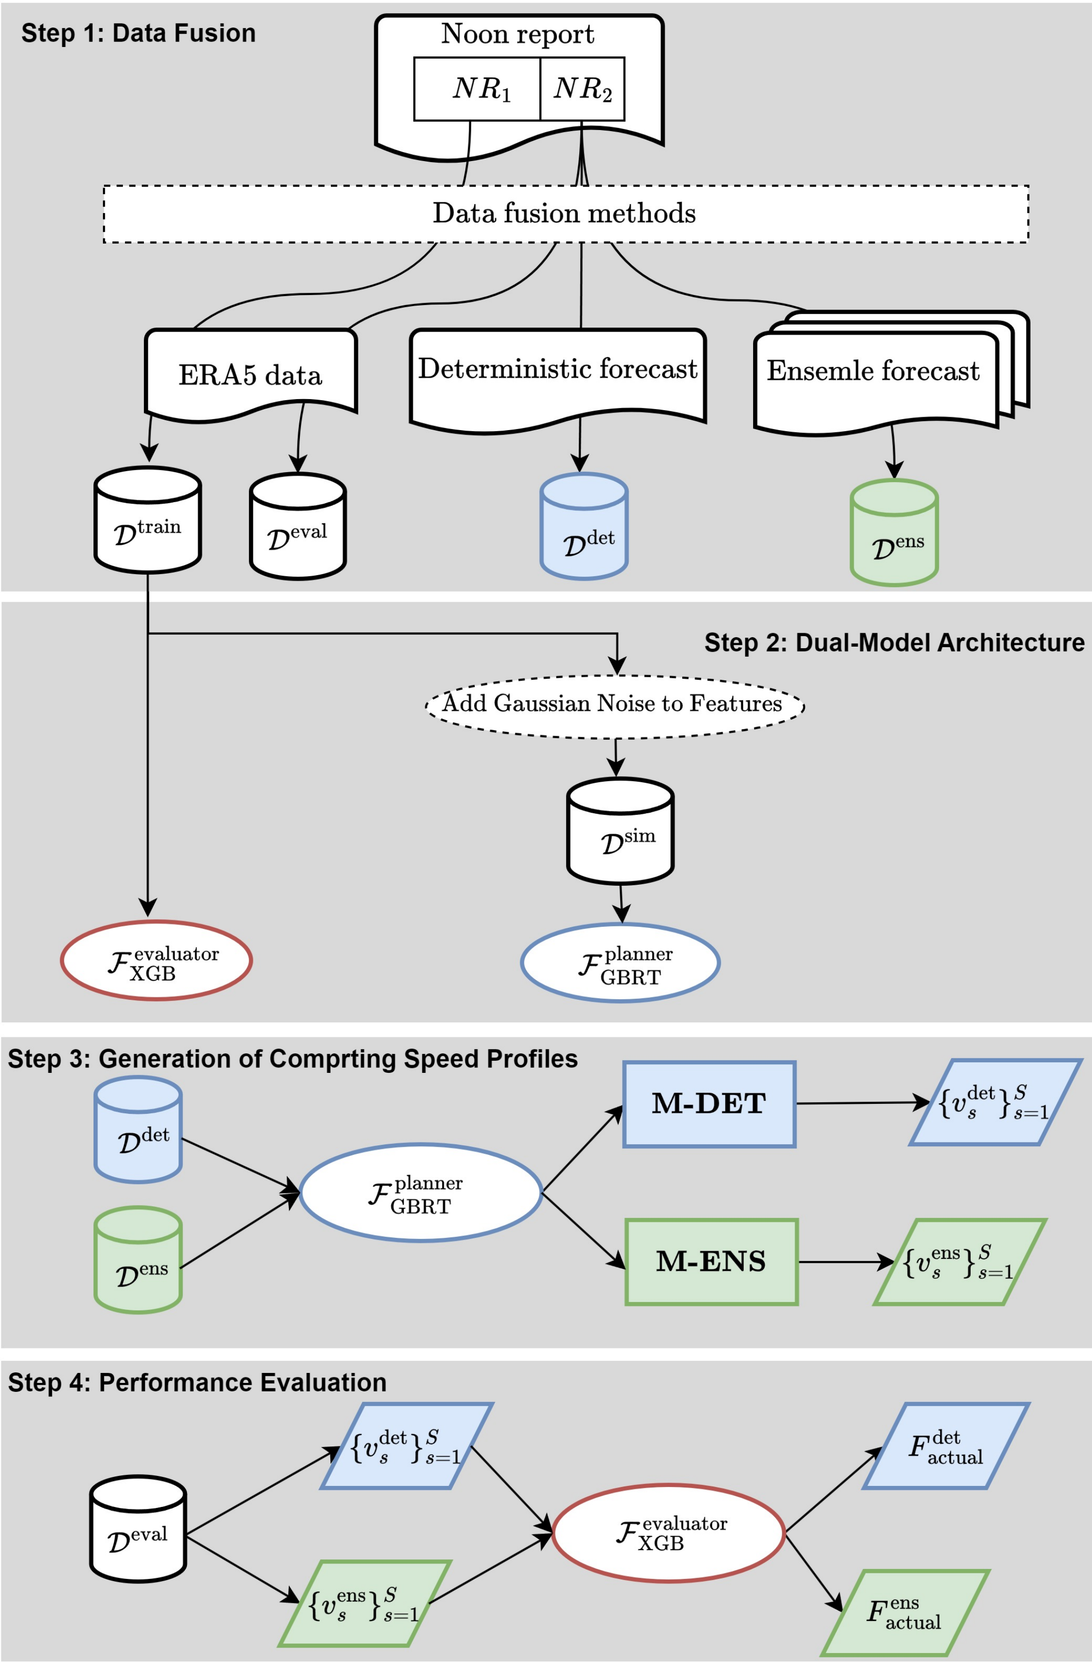

Figure 4.1: Framework for evaluating speed optimization under different forecast types.

Table 4.2: Notation for evaluation framework and dataset construction

Symbol	Description
Dataset Size Parameters	
N_1	Number of records in training/testing dataset NR_1
N_2	Number of voyage segments in validation dataset NR_2
d_w	Dimension of complete meteorological feature vector
$d_{w,f}$	Dimension of meteorological variables provided by the GEFS forecasts (subset of d_w)
d_u	Dimension of vessel operational state vector
Fused Datasets	
$\mathcal{D}^{\text{direct}}$	Dataset from NR_1 fused with ERA5 via direct method: $\{(v_i, \mathbf{w}_i^{\text{ERA5}}, \mathbf{u}_i, y_i) \mid i = 1, ..., N_1\}$
$\mathcal{D}^{\text{rhumb}}$	Dataset from NR_1 fused with ERA5 via rhumb line method: $\{(v_i, \mathbf{w}_i^{\text{ERA5}}, \mathbf{u}_i, y_i) \mid i = 1, ..., N_1\}$
$\mathcal{D}^{\text{eval}}$	Evaluation dataset from NR_2 fused with ERA5: $\{(v_i, \mathbf{w}_i^{\text{ERA5}}, \mathbf{u}_i, y_i) \mid i = 1, ..., N_2\}$
Simulated Dataset (with Noise)	
$\mathcal{D}^{\text{sim}}$	Dataset with Gaussian noise: $\{(v_i^{\text{sim}}, \mathbf{w}_i^{\text{sim}}, \mathbf{u}_i^{\text{sim}}, y_i^{\text{sim}}) \mid i = 1, ..., N_1\}$
α	Noise level coefficient for Gaussian noise injection
$\sigma_v, \boldsymbol{\sigma}_{\mathbf{w}}, \boldsymbol{\sigma}_{\mathbf{u}}$	Empirical standard deviations of features in the base dataset
Forecast-Augmented Datasets	
$\mathcal{D}^{\text{det}}$	Dataset with GEFS control member (deterministic)
$\mathcal{D}_k^{\text{ens}}$	Dataset with the kth perturbed member, where $k = 1, ..., K$
Data Elements and Features	
v_i	Average sailing speed for record i
y_i	Hourly fuel consumption for record i
$\mathbf{u}_i$	Vessel operational state vector for record i
$\mathbf{w}_i^{\text{ERA5}}$	Complete meteorological feature vector from ERA5 reanalysis
$\mathbf{w}_i^{\text{ERA5,other}}$	ERA5 variables not included in GEFS forecasts
$\mathbf{w}_i^{\text{sim}}$	Meteorological features with Gaussian noise: $\mathbf{w}_i^{\text{ERA5}} + \boldsymbol{\epsilon}_{\mathbf{w}}$
$v_i^{\text{sim}}, \mathbf{u}_i^{\text{sim}}$	Speed and vessel state with Gaussian noise
y_i^{sim}	Simulated ship's FCR: $\mathcal{F}_{\text{XGB}}^{\text{evaluator}}(v_i^{\text{sim}}, \mathbf{w}_i^{\text{sim}}, \mathbf{u}_i^{\text{sim}})$
Prediction Models	
$\mathcal{F}_{\text{XGB}}^{\text{direct}}$	XGBoost model trained on $\mathcal{D}^{\text{direct}}$
$\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$	XGBoost model trained on $\mathcal{D}^{\text{rhumb}}$
$\mathcal{F}_{\text{XGB}}^{\text{evaluator}}$	Final XGBoost evaluator model (ground truth predictor)
$\mathcal{F}_{\text{GBRT}}^{\text{planner}}$	GBRT planner model trained on $\mathcal{D}^{\text{sim}}$ (decision-making tool)
Speed Profiles	
$\{v_s^{\text{det}}\}_{s=1}^S$	Deterministic optimization speed profile
$\{v_s^{\text{ens}}\}_{s=1}^S$	Ensemble optimization speed profile
$\{v_s^{\text{orig}}\}_{s=1}^S$	Orignial noon reprot speed profile

Step 2: Dual-Model Architecture Construction. The evaluation framework relies on a dual-model architecture that distinguishes between decision-making and performance assessment. Within this framework, two distinct prediction models are constructed to fulfill different roles: a ground truth predictor (or *evaluator model*) that serves as the evaluator for assessing actual fuel consumption under realized meteorological conditions, and a decision-input predictor (or *planner model*) that simulates the imperfect predictive capability available to decision-makers during voyage planning.

The first component, the *evaluator model* $\mathcal{F}_{\text{XGB}}^{\text{evaluator}}$, is established in Step 1 through training on the high-quality fused dataset. This model serves as the impartial judge in our evaluation framework, computing the actual fuel consumption that would result from any given speed profile when executed under the meteorological conditions that historically occurred along the route. This dataset serves a purely evaluative role, as it does not participate in the optimization process but rather assesses the quality of decisions made by the planner.

The second component, termed the *planner model*, is a GBRT model $\mathcal{F}_{\text{GBRT}}^{\text{planner}}$ that represents the predictive capability available to decision-makers during voyage planning. To realistically simulate the imperfect nature of predictions in operational settings, this model is trained not on the pristine dataset with actual meteorological data, but rather on a simulated dataset $\mathcal{D}^{\text{sim}}$ generated by adding Gaussian noise to the features. Specifically, for each record i in the pristine dataset, independent Gaussian noise is added to all input features:

$$v_i^{\text{sim}} = v_i + \epsilon_v, \quad \epsilon_v \sim \mathcal{N}(0, (\alpha \cdot \sigma_v)^2) \tag{4.14}$$

$$\mathbf{w}_i^{\text{sim}} = \mathbf{w}_i^{\text{ERA5}} + \boldsymbol{\epsilon}_{\mathbf{w}}, \quad \boldsymbol{\epsilon}_{\mathbf{w}} \sim \mathcal{N}(\mathbf{0}, (\alpha \cdot \boldsymbol{\sigma}_{\mathbf{w}})^2) \tag{4.15}$$

$$\mathbf{u}_i^{\text{sim}} = \mathbf{u}_i + \boldsymbol{\epsilon}_{\mathbf{u}}, \quad \boldsymbol{\epsilon}_{\mathbf{u}} \sim \mathcal{N}(\mathbf{0}, (\alpha \cdot \boldsymbol{\sigma}_{\mathbf{u}})^2), \tag{4.16}$$

where α is the noise level coefficient, and $\sigma_v, \boldsymbol{\sigma}_{\mathbf{w}}, \boldsymbol{\sigma}_{\mathbf{u}}$ are the empirical standard deviations of the respective features computed from the pristine dataset. The output labels for the simulated dataset are generated by passing the noisy inputs through the evaluator model: $y_i^{\mathrm{sim}} = \mathcal{F}_{\mathrm{XGB}}^{\mathrm{evaluator}}(v_i^{\mathrm{sim}}, \mathbf{w}_i^{\mathrm{sim}}, \mathbf{u}_i^{\mathrm{sim}})$. This yields the simulated dataset $\mathcal{D}^{\mathrm{sim}} = \{(v_i^{\mathrm{sim}}, \mathbf{w}_i^{\mathrm{sim}}, \mathbf{u}_i^{\mathrm{sim}}, y_i^{\mathrm{sim}}) \mid i = 1, ..., N_1\}$, on which $\mathcal{F}_{\mathrm{GBRT}}^{\mathrm{planner}}$ is trained. This noise injection procedure mimics the weather forecast errors and measurement uncertainties inherent in real-world predictions, ensuring that the planner model exhibits prediction errors similar to those encountered in practice. By separating the planner and evaluator models and training them on data with different quality characteristics, the framework creates a realistic simulation of the planning-execution gap that exists in actual maritime operations.

Step 3: Generation of Competing Speed Profiles. Using the planner model $\mathcal{F}_{\mathrm{GBRT}}^{\mathrm{planner}}$, we generate two speed profiles for the voyage segments in the validation dataset $\mathcal{D}^{\mathrm{eval}}$. For the deterministic approach, the forecast-augmented dataset $\mathcal{D}^{\mathrm{det}}$ combines forecasted variables from the GEFS control member with supplementary variables $\mathbf{w}_i^{\mathrm{ERA5,other}}$ from ERA5. This combination is necessary because GEFS forecasts provide only three meteorological variables (WD, WS, and 2mT), while the fuel consumption prediction model requires additional sea state variables (wave heights, periods, and directions) that are obtained from ERA5 reanalysis data. The planner model predicts ship's FCRs for each segment and candidate speed, populating the objective function of [**M-DET**]. Solving this model yields the deterministic speed profile $\{v_s^{\mathrm{det}}\}_{s=1}^S$ and durations $\{\Delta\tau_i^{\mathrm{det}}\}_{s=1}^S$. For the ensemble approach, we construct K forecast-augmented datasets $\{\mathcal{D}_k^{\mathrm{ens}}\}_{k=1}^K$ using all perturbed members, with each dataset combining the forecasted variables from perturbed member k and the same ERA5 supplementary variables. For each segment and speed, the expected ship's FCR is computed by averaging predictions across all ensemble members. These averaged coefficients populate [**M-ENS**], yielding $\{v_s^{\mathrm{ens}}\}_{s=1}^S$ and $\{\Delta\tau_s^{\mathrm{ens}}\}_{s=1}^S$. Both profiles use identical planner models and optimization frameworks, differing only in meteorological forecast inputs, ensuring a fair comparison.

Step 4: Performance Evaluation. We evaluate the performance of both speed profiles using the evaluator model $\mathcal{F}_{\text{XGB}}^{\text{evaluator}}$ with historical ERA5 reanalysis data, which represents the meteorological conditions that actually occurred. For each speed profile, the total fuel consumption is computed by summing the segment-level consumption:

$$F_{\text{actual}}^{\text{det}} = \sum_{s=1}^{S} \mathcal{F}_{\text{XGB}}^{\text{evaluator}}(v_s^{\text{det}}, \mathbf{w}_s^{\text{ERA5}}, \mathbf{u}_s) \cdot \Delta\tau_s^{\text{det}} \tag{4.17}$$

$$F_{\text{actual}}^{\text{ens}} = \sum_{s=1}^{S} \mathcal{F}_{\text{XGB}}^{\text{evaluator}}(v_s^{\text{ens}}, \mathbf{w}_s^{\text{ERA5}}, \mathbf{u}_s) \cdot \Delta\tau_s^{\text{ens}}. \tag{4.18}$$

The relative fuel savings achieved by the ensemble approach is:

$$\Delta F_{\text{rel}} = \frac{F_{\text{actual}}^{\text{det}} - F_{\text{actual}}^{\text{ens}}}{F_{\text{actual}}^{\text{det}}} \times 100\%. \tag{4.19}$$

A positive value indicates that the ensemble forecast-based optimization achieves lower actual fuel consumption, demonstrating the value of accounting for forecast uncertainty. Furthermore, we evaluate the actual fuel consumption under the original noon report speed profile, which represents the vessel's historical operating strategy without optimization:

$$F_{\text{actual}}^{\text{orig}} = \sum_{s=1}^{S} \mathcal{F}_{\text{XGB}}^{\text{evaluator}}(v_s^{\text{orig}}, \mathbf{w}_s^{\text{ERA5}}, \mathbf{u}_s) \cdot \Delta\tau_s^{\text{orig}}. \tag{4.20}$$

Overall, the comparison yields two insights: (1) comparing $F_{\text{actual}}^{\text{det}}$ versus $F_{\text{actual}}^{\text{ens}}$ quantifies the value of ensemble weather forecasts and (2) comparing both optimized profiles against $F_{\text{actual}}^{\text{orig}}$ demonstrates the overall benefit of forecast-based optimization.

The proposed evaluation framework ensures valid comparison through several methodological strengths. First, the dual-model architecture realistically simulates the planning-execution gap by making decisions with imperfect predictions (planner model with forecast data) while assessing performance with high-quality retrospective data (evaluator model with ERA5 reanalysis data). Second, using the same historical voyages for both approaches ensures that route, realized weather conditions, and operational parameters remain constant, enabling a controlled comparison where observed differences can be attributed solely to the weather forecast type rather than external variability.

4.3 Computational Results

This section presents computational results validating the evaluation framework developed in Section 4.2. Section 4.3.1 describes the experimental configuration, including feature sets for fuel consumption prediction and details of two validation voyages. Section 4.3.2 evaluates and selects between data fusion methods to establish the XGBoost evaluator model. Section 4.3.3 validates the GBRT planner model's predictive capabilities. Finally, Section 4.3.4 quantifies the fuel savings achieved by ensemble forecast-based optimization compared to deterministic approaches under realistic conditions.

4.3.1 Experimental Setup

The fuel consumption prediction models developed in this chapter employ a carefully curated set of features that capture the essential operational and environmental factors influencing vessel's fuel consumption. Table 3.1 presents the complete specification of input features and the prediction target, organized by data source. The input feature set encompasses two categories of variables. *Vessel operational parameters* include sailing speed, loading condition (laden or ballast), and draught, which collectively characterize the vessel's physical state and operational mode. These variables are directly obtained from noon reports and fundamentally determine the vessel's resistance profile and power requirements. *Meteorological conditions* comprise twelve variables extracted from ERA5 reanalysis data that quantify the environmental forces acting on the vessel.

To evaluate the predictive performance of the fuel consumption models developed in this chapter, we employ four standard regression metrics that quantify different aspects of prediction accuracy. Consider a dataset with N samples, where y_i represents the true fuel consumption rate for sample i and $\hat{y}_i$ represents the predicted value. The evaluation metrics are defined as follows:

Mean Squared Error (MSE) measures the average squared difference between predictions and actual values:

$$\text{MSE} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2. \tag{4.21}$$

Mean Absolute Error (MAE) measures the average absolute difference, providing a more interpretable metric in the original units (mt/h):

$$\text{MAE} = \frac{1}{N}\sum_{i=1}^{N}|y_i - \hat{y}_i|. \tag{4.22}$$

Mean Absolute Percentage Error (MAPE) expresses prediction error as a percentage of the true value, facilitating comparison across different scales:

$$\text{MAPE} = \frac{100\%}{N}\sum_{i=1}^{N}\left|\frac{y_i - \hat{y}_i}{y_i}\right|. \tag{4.23}$$

Coefficient of Determination (R^2) quantifies the proportion of variance in the dependent variable explained by the model, with values closer to 1 indicating better fit:

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2}, \tag{4.24}$$

where $\bar{y} = \frac{1}{N}\sum_{i=1}^{N} y_i$ is the mean of the observed values. These metrics provide complementary perspectives on model performance: MSE penalizes large errors more severely, MAE offers robustness to outliers, MAPE enables scale-independent comparison, and R^2 indicates overall explanatory power.

To evaluate the comparative performance of deterministic versus ensemble forecast-based speed optimization, two distinct voyages recorded in dataset $\mathcal{D}^{\text{eval}}$ are deliberately excluded from the training data for both XGBoost and GBRT models, ensuring that the evaluation reflects model performance on genuinely unseen operational scenarios. The two voyages represent complementary loading conditions (ballast and laden), enabling assessment of weather forecast value across different vessel operational states.

Voyage I (Ballast Voyage): The first validation case is a 9-day ballast voyage with a total sailing time of 216 hours (9 days). A ballast voyage refers to a trip where the vessel travels without cargo after delivering its load, sailing at reduced draught but often at higher speeds to minimize non-revenue sailing time. Table 4.3 presents the detailed operational and meteorological conditions for each of the nine voyage segments (corresponding to the nine noon report records). The vessel maintains a constant draught of 8.6 meters throughout this ballast voyage, substantially lighter than its laden condition.

Voyage II (Laden Voyage): The second validation case is a 9-day laden voyage with an identical total sailing time of 216 hours. Table 4.4 presents the operational and environmental profile for this laden voyage. The vessel maintains a draught of 17.5 meters throughout, nearly doubling the ballast draught and representing near-maximum cargo capacity. The sailing speeds are slightly lower than the ballast voyage, ranging from 12.8 to 13.9 knots.

Table 4.3: Operational and meteorological conditions for Voyage I (9-day ballast voyage)

Day	Sailing Hours (h)	Speed (knots)	Ship's FCR (tons/h)	Draught (m)	Actual meteorological data (from ERA5)											
					WD (deg)	WS (m/s)	MPWW (s)	MDWW (deg)	SHWW (m)	MPTS (s)	SHTS (m)	MDTS (deg)	SHCWWS (m)	MWD (s)	FSR (10^{-2}m)	2mT (°C)
1	24	12.63	1.3458	8.6	57.07	8.15	4.98	128.45	1.30	7.14	1.08	150.39	1.81	6.16	0.0181	28.06
2	24	12.58	1.3542	8.6	4.26	3.40	2.84	65.98	0.19	5.57	0.51	60.69	0.57	4.99	0.0045	28.31
3	24	12.92	1.3542	8.6	8.74	1.47	2.55	124.88	0.07	3.75	0.33	90.09	0.34	3.64	0.0032	27.64
4	24	13.29	1.3333	8.6	126.45	0.47	3.46	108.01	0.01	6.39	0.29	77.41	0.29	6.36	0.0048	28.09
5	24	13.83	1.3417	8.6	157.38	0.66	3.14	78.80	0.05	4.68	0.40	5.99	0.41	4.55	0.0036	27.57
6	24	13.50	1.3458	8.6	101.53	3.35	2.53	98.89	0.15	6.93	0.66	64.74	0.66	6.63	8.0360	26.58
7	24	13.29	1.3542	8.6	174.51	3.51	2.57	34.89	0.14	11.86	1.61	48.70	1.62	11.76	0.0025	26.34
8	24	11.67	1.3667	8.6	157.86	5.53	2.95	21.58	0.47	10.91	1.90	57.36	1.96	10.42	0.0060	25.69
9	24	10.17	0.9129	8.6	133.61	4.08	2.14	45.63	0.22	9.24	1.52	63.80	1.53	9.10	0.0024	25.15
Total	216	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–

Table 4.4: Operational and meteorological conditions for Voyage II (9-day laden voyage)

Day	Sailing Hours (h)	Speed (knots)	Ship's FCR (tons/h)	Draught (m)	Actual meteorological data (from ERA5)											
					WD (deg)	WS (m/s)	MPWW (s)	MDWW (deg)	SHWW (m)	MPTS (s)	SHTS (m)	MDTS (deg)	SHCWWS (m)	MWD (s)	FSR (10^{-2}m)	2mT (°C)
1	24	9.17	0.8875	9.2	140.76	5.96	3.11	27.37	0.50	6.67	1.22	78.01	1.34	6.17	0.0067	28.17
2	24	9.96	0.8625	8.6	114.63	4.04	2.49	46.97	0.19	4.43	0.43	76.04	0.49	4.03	0.0033	27.82
3	24	10.96	0.8821	8.6	147.96	1.99	3.23	81.67	0.10	4.47	0.46	19.75	0.49	4.31	0.0036	28.18
4	24	12.00	1.0979	8.6	164.63	3.90	2.07	6.30	0.16	4.57	0.44	9.46	0.48	4.25	0.0028	27.32
5	24	11.20	1.7000	8.6	48.37	1.50	2.91	145.59	0.05	4.31	0.33	132.55	0.34	4.22	0.0029	27.85
6	24	13.35	1.3575	8.6	126.78	3.83	3.06	78.57	0.33	5.46	0.54	43.18	0.66	4.77	0.0053	27.76
7	24	12.83	1.3042	8.6	97.61	4.39	2.90	93.24	0.33	9.14	1.01	26.27	1.08	8.57	0.1158	25.42
8	24	12.71	1.3571	8.6	98.69	7.80	4.17	80.08	1.03	10.58	1.88	10.80	2.17	9.24	0.0001	26.62
9	24	12.00	1.3500	8.6	84.94	5.69	2.96	94.76	0.56	12.07	1.89	52.59	2.01	11.38	0.00007	25.63
Total	216	–	–	–	–	–	–	–	–	–	–	–	–	–	–	–

Note: WD = Wind Direction (relative), WS = Wind Speed, MPWW = Mean Period of Wind Waves, MDWW = Mean Direction of Wind Waves (relative), SHWW = Significant Height of Wind Waves, MPTS = Mean Period of Total Swell, SHTS = Significant Height of Total Swell, MDTS = Mean Direction of Total Swell (relative), SHCWWS = Significant Height of Combined Wind Waves and Swell, MWD = Mean Wave Direction, FSR = Forecast Surface Roughness, 2mT = 2-meter Temperature.

4.3.2 Ground Truth Predictor: Comparative Performance of Data Fusion Methods

Before proceeding with the evaluation of forecast-based speed optimization, we must first determine which of the two data fusion methods yields more accurate meteorological representations for ship's FCR prediction. As described in Step 1 of the evaluation framework, the direct fusion method produces dataset $\mathcal{D}^{\text{direct}}$ and the rhumb line-based fusion method produces dataset $\mathcal{D}^{\text{rhumb}}$, both constructed by fusing the noon report data NR_1 with ERA5 reanalysis data. The comparison is conducted by training two separate XGBoost models, $\mathcal{F}_{\text{XGB}}^{\text{direct}}$ on $\mathcal{D}^{\text{direct}}$ and $\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$ on $\mathcal{D}^{\text{rhumb}}$, and assessing their predictive performance on held-out test data.

Both datasets are randomly partitioned into training and test sets using an 8:2 ratio, with the larger subset used for model training and hyperparameter tuning, and the smaller subset reserved for evaluating generalization performance. Both XGBoost models are implemented using the XGBoost 1.6.1 library and configured with identical training procedures to ensure a fair comparison. The hyperparameters for each model are optimized through Bayesian optimization using the Hyperopt 0.2.2 library with five-fold cross-validation on the training set. This optimization process systematically searches the hyperparameter space to identify configurations that maximize predictive performance while minimizing overfitting risk. Table 4.5 presents the hyperparameters tuned for both models, along with their search spaces and the optimal values selected for $\mathcal{F}_{\text{XGB}}^{\text{direct}}$ and $\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$. Table 4.6 presents the predictive performance of both models on the held-out test set, evaluated using four metrics: MSE, MAE, MAPE, and R^2.

Although $\mathcal{F}_{\text{XGB}}^{\text{direct}}$ achieves slightly lower MSE (0.023 vs 0.024), $\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$ demonstrates superior performance across the three remaining metrics: lower MAE (0.107 vs 0.114 mt/h), lower MAPE (7.706% vs 8.023%), and higher R^2 (0.777 vs 0.721). The higher R^2 value indicates that the rhumb line-based fusion method explains approximately

77.7% of the variance in ship's FCR, compared to 72.1% for the direct method, representing a meaningful improvement in explanatory power. This superior performance across multiple metrics suggests that the rhumb line-based fusion method, which employs hourly interpolation along the vessel's path to obtain representative segment-average meteorological conditions, captures the true environmental exposure more accurately than the direct spatial matching approach.

Table 4.5: Hyperparameters for XGBoost models $\mathcal{F}_{\text{XGB}}^{\text{direct}}$ and $\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$

Hyperparameter	Interpretation	Search Space	$\mathcal{F}_{\text{XGB}}^{\text{direct}}$	$\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$
`max_depth`	Maximum depth of a tree	[2, 17]	15	6
`n_estimators`	Number of boosting trees	[10, 300]	48	83
`learning_rate`	Learning rate shrinks the step size to prevent overfitting	[0.00001, 1]	0.45	0.16
`min_child_weight`	Minimum sum of instance weight required in a child	[0, 10]	5.58	9.54
`colsample_bytree`	Subsample ratio of columns in the construction of each tree	[0.1, 1]	0.85	0.64
`subsample`	Subsample ratio of training instances	[0.4, 1]	0.77	0.94
`reg_alpha`	L1 regularization term on weights	[0, 2]	0.20	0.23
`reg_lambda`	L2 regularization term on weights	[0, 2]	0.61	1.76
`gamma`	Minimum loss reduction for further partitioning on a tree leaf node	[0, 2]	0.29	0.18

Table 4.6: Predictive performance comparison of XGBoost models trained based on different data fusion methods

Model	MSE	MAE (mt/h)	MAPE (%)	R^2
$\mathcal{F}_{\text{XGB}}^{\text{direct}}$	0.023	0.114	8.023	0.721
$\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$	0.024	0.107	7.706	0.777

Based on this comprehensive comparison, the rhumb line-based fusion method is selected as the standard data processing approach for all subsequent analyses in this chapter. Consequently, the XGBoost model $\mathcal{F}_{\text{XGB}}^{\text{rhumb}}$ is designated as the evaluator model $\mathcal{F}_{\text{XGB}}^{\text{evaluator}}$, serving as the ground truth simulator in the evaluation framework. All datasets constructed hereafter, including $\mathcal{D}^{\text{sim}}$, $\mathcal{D}^{\text{det}}$, $\mathcal{D}^{\text{eval}}$, and the ensemble datasets $\{\mathcal{D}_k^{\text{ens}}\}_{k=1}^{K}$, will employ the rhumb line-based fusion procedure to ensure consistency and data quality throughout the evaluation.

4.3.3 Decision-Input Predictor: Performance Validation

Having established the evaluator model $\mathcal{F}_{\text{XGB}}^{\text{evaluator}}$ in Section 4.3.2 as a reliable ground truth simulator with strong predictive performance (MAE $= 0.107$ mt/h, $R^2 = 0.777$), we now develop and validate the planner model $\mathcal{F}_{\text{GBRT}}^{\text{planner}}$ that represents the imperfect predictive capability available to decision-makers during voyage planning. The planner model must exhibit realistic prediction errors to simulate operational conditions, while maintaining sufficient accuracy to support effective optimization decisions.

To create training data reflecting operational uncertainty, we generate the simulated dataset $\mathcal{D}^{\text{sim}}$ by adding Gaussian noise to the features of $\mathcal{D}^{\text{rhumb}}$, as described in Section 4.2. The noise level coefficient is set to $\alpha = 0.05$, meaning that the standard deviation of the added noise equals 5% of each feature's empirical standard deviation. The output labels are generated by passing the noisy inputs through the evaluator model $\mathcal{F}_{\text{XGB}}^{\text{evaluator}}$. The resulting simulated dataset $\mathcal{D}^{\text{sim}}$ is used to train and test the GBRT planner model $\mathcal{F}_{\text{GBRT}}^{\text{planner}}$.

The planner model is implemented using the scikit-learn 1.0.2 library. Dataset $\mathcal{D}^{\text{sim}}$ is randomly partitioned into training and test sets using an 8:2 ratio. Hyperparameters are optimized through Hyperopt with five-fold cross-validation on the training set, with the search spaces and optimal values presented in Table 4.7. Table 4.8 presents the planner model's performance on the held-out test set. The model achieves excellent predictive metrics: MSE of 0.003, MAE of 0.039 mt/h, MAPE of 2.738%, and R^2 of 0.932, demonstrating strong capability to predict ship's FCR under noisy operational conditions.

The planner model's superior metrics relative to the evaluator model reflect the different evaluation contexts rather than contradictory performance: the *planner model* is tested on simulated data with controlled noise characteristics, while the *evaluator model* is tested on real operational data with inherent complexity. Importantly, the planner model's strong performance on simulated data confirms that it has successfully learned patterns within the noisy training environment, exhibiting operationally realistic prediction errors (5% noise level) while maintaining sufficient accuracy for effective optimization. This validation establishes the planner model's suitability as the decision-making tool in the speed optimization models **[M-DET]** and **[M-ENS]**, enabling rigorous evaluation of forecast-based optimization approaches in the following section.

Table 4.7: Hyperparameters for GBRT planner model $\mathcal{F}_{\text{GBRT}}^{\text{planner}}$

Hyperparameter	Interpretation	Search Space	Optimal Value
`max_depth`	Maximum depth of each decision tree	[2, 10]	3
`min_samples_leaf`	Minimum samples required per leaf node	[1, 20]	7
`min_samples_split`	Minimum samples required for splitting an internal node	[2, 20]	17
`max_features`	Number of features considered when searching for best split	[1, 15]	15
`n_estimators`	Number of regression trees	[10, 300]	264
`learning_rate`	Learning rate diminishes the contribution of each tree	[0.00001, 1]	0.06
`subsample`	Fraction of samples for fitting base learners	[0.4, 1]	0.87

Table 4.8: Performance of GBRT planner model $\mathcal{F}_{\text{GBRT}}^{\text{planner}}$ on simulated test set

Model	MSE	MAE (mt/h)	MAPE (%)	R^2
$\mathcal{F}_{\text{GBRT}}^{\text{planner}}$	0.003	0.039	2.738	0.932

4.3.4 Comparative Analysis of Speed Optimization under Different Weather Forecasts

In this section, we evaluate the actual fuel consumption performance of speed profiles generated by the deterministic forecast-based optimization model [**M-DET**] versus the ensemble forecast-based optimization model [**M-ENS**] when executed under historical meteorological conditions on the two validation voyages. Both optimization models are implemented as mixed-integer linear programming problems and solved using the CPLEX optimizer. The speed discretization employs a uniform grid with 0.1-knot intervals across the feasible speed range $[v_{\min}, v_{\max}]$ for each segment. The minimum allowable speed is set to $v_{\min} = 8$ knots based on operational safety considerations, while the maximum speed $v_{\max} = 18$ knots reflects typical service speed constraints for the Capesize vessel class. This discretization yields 100 candidate speeds per segment, resulting in mixed-integer programs with 900 binary variables and 10 continuous time variables for each voyage. The maximum allowable sailing time $T_{\max}$ is set equal to the actual sailing time recorded in the noon reports (216 hours for both voyages), ensuring that optimized solutions meet the original arrival schedule without delays.

For the deterministic approach, the GEFS control member provides forecasted wind speed, wind direction, and 2-meter temperature for each segment, which are combined with ERA5 wave parameters and vessel operational data to form the input dataset $\mathcal{D}^{\text{det}}$. The planner model $\mathcal{F}_{\text{GBRT}}^{\text{planner}}$ predicts ship's FCRs for all candidate speeds under these deterministic forecast conditions, and model [**M-DET**] is solved to obtain the optimal deterministic speed profile $\{v_s^{\text{det}}\}_{s=1}^{9}$.

For the ensemble approach, all $K = 20$ perturbed members provide forecasted meteorological variables, forming the collection of datasets $\{\mathcal{D}_k^{\text{ens}}\}_{k=1}^{20}$. For each segment and candidate speed, the planner model predicts ship's FCRs under each member's conditions, and the average across all 20 predictions forms the expected fuel consumption rate used in model [**M-ENS**]. Solving this optimization problem yields the optimal ensemble forecast-based speed profile $\{v_s^{\text{ens}}\}_{s=1}^{9}$.

Both optimized speed profiles, along with the original noon report speeds, are then evaluated using the ground truth predictor $\mathcal{F}_{\text{XGB}}^{\text{evaluator}}$ with ERA5 meteorological conditions from the evaluation dataset $\mathcal{D}^{\text{eval}}$ to compute actual fuel consumption for each segment and the total voyage.

(1) Results for Voyage I

Table 4.9 presents the optimized speeds, sailing durations, and evaluated actual fuel consumption for each segment of Voyage I under two optimization strategies. The deterministic forecast-based optimization model [**M-DET**] achieves total evaluated fuel consumption of 290.28 mt with a total sailing time of 212.52 hours. The ensemble forecast-based optimization model [**M-ENS**] achieves lower total evaluated fuel consumption of 283.67 mt with a comparable sailing time of 212.54 hours. The ensemble approach achieves 6.60 mt (2.27%) fuel savings compared to [**M-DET**], demonstrating incremental value from explicitly accounting for forecast uncertainty in the ballast voyage context.

Table 4.9: Speed optimization results for Voyage I (9-day ballast voyage)

Segment	Optimized Speed (knots)		Sailing Time (h)		Actual Fuel Consumption (mt)	
	[**M-DET**]	[**M-ENS**]	[**M-DET**]	[**M-ENS**]	[**M-DET**]	[**M-ENS**]
1	13.1	13.1	22.7433	22.7433	30.5624	30.5624
2	13.1	13.1	23.0810	23.0810	31.1340	31.1340
3	13.1	13.1	23.3887	23.3887	31.3853	31.3853
4	13.0	13.0	24.4178	24.4178	31.8067	31.8067
5	10.7	18.0	31.2338	18.5668	40.7133	38.5168
6	13.1	13.1	24.3290	24.3290	33.2213	33.2213
7	10.7	10.7	29.0115	29.0115	32.2260	32.2318
8	18.0	10.7	15.6583	26.3411	33.6074	29.1938
9	13.1	13.1	18.6601	18.6601	25.6203	25.6203
Total	–	–	**212.5235**	**212.5393**	**290.2765**	**283.6723**

(2) Results for Voyage II

Table 4.10 presents the corresponding results for Voyage II, conducted under laden conditions with substantially higher draft. The deterministic approach **[M-DET]** achieves total evaluated fuel consumption of 265.81 mt with a total sailing time of 203.71 hours. The ensemble approach **[M-ENS]** achieves total evaluated fuel consumption of 267.07 mt with a total sailing time of 173.99 hours. In contrast to Voyage I, **[M-ENS]** consumes an additional 1.26 mt (0.47%) compared to **[M-DET]** in this laden voyage.

Table 4.10: Speed optimization results for Voyage II (9-day laden voyage)

Segment	Optimized Speed (knots)		Sailing Time (h)		Actual Fuel Consumption (mt)	
	[M-DET]	**[M-ENS]**	**[M-DET]**	**[M-ENS]**	**[M-DET]**	**[M-ENS]**
1	13.0	13.0	16.9857	16.9857	25.3172	25.3172
2	10.7	10.7	22.2359	22.2359	25.4832	25.4832
3	13.1	13.1	19.9880	19.9880	27.0677	27.0677
4	13.0	13.0	22.1410	22.1410	34.6152	34.6152
5	10.7	10.7	27.2183	27.2183	36.1595	36.1595
6	10.7	13.0	28.4722	23.4348	30.8582	32.1174
7	13.1	13.1	16.6841	16.6841	23.2543	23.2543
8	10.7	10.7	27.0107	27.0107	31.8861	31.8861
9	13.1	13.1	22.9783	22.9783	31.1678	31.1678
Total	–	–	**203.7142**	**173.9864**	**265.8084**	**267.0676**

(3) Aggregate Performance Analysis

Table 4.11 summarizes the evaluated actual fuel consumption under three speed strategies across the two validation voyages: the original noon report speeds without optimization, deterministic forecast-based optimization **[M-DET]**, and ensemble forecast-based optimization **[M-ENS]**. The table presents both absolute fuel consumption values and relative savings across three pairwise comparisons: Original vs. **[M-DET]** quantifies the value of deterministic forecast-based optimization, Original vs. **[M-ENS]** quantifies the value of ensemble forecast-based optimization, and **[M-DET]** vs. **[M-ENS]** isolates the incremental value of using ensemble weather forecasts rather than deterministic weather forecasts.

Table 4.11: Summary of evaluated fuel consumption and savings across optimization strategies

Voyage	Evaluated Fuel Consumption (mt)			Fuel Savings (%)		
	Original	[M-DET]	[M-ENS]	Orig vs. [M-DET]	Orig vs. [M-ENS]	[M-DET] vs. [M-ENS]
Voyage I (Ballast)	304.12	290.28	283.67	4.55	6.72	2.27
Voyage II (Laden)	292.44	265.81	267.07	9.11	8.68	−0.47
Total	596.56	**556.09**	**550.74**	6.79	7.68	0.96

The original noon report speeds result in a total evaluated actual fuel consumption of 596.56 mt across both voyages. Deterministic forecast-based optimization **[M-DET]** reduces this to 556.08 mt, achieving substantial savings of 40.47 mt (6.79%). Ensemble forecast-based optimization **[M-ENS]** further reduces fuel consumption to 550.74 mt, achieving total savings of 45.82 mt (7.68%) relative to the original noon report speeds. The incremental benefit of ensemble weather forecasts over deterministic weather forecasts is 5.34 mt (0.96%) across both voyages combined. This corresponds to an average daily saving of approximately 0.30 mt. Extrapolating this to an annual operation of 200 sailing days, as derived from the ship's noon reports, yields a reduction of roughly 60 mt of fuel and 192 tons of CO_2 (based on an emission factor of 3.2 kg CO_2/kg fuel). Furthermore, given that the ensemble forecast data are publicly available, this improvement represents a net economic gain without increasing operating costs.

However, this aggregate incremental benefit masks heterogeneity across loading conditions. In the ballast voyage (Voyage I), both optimization approaches deliver substantial improvements over the original speeds, with **[M-DET]** saving 4.55% and **[M-ENS]** saving 6.72%. The ensemble approach provides a 2.27% incremental benefit over deterministic optimization in this ballast context. In the laden voyage (Voyage II), **[M-DET]** achieves even larger relative savings of 9.11% over the original speeds, demonstrating strong optimization performance. However, **[M-ENS]** achieves 8.68% savings relative to original speeds while underperforming **[M-DET]** by 0.47%, resulting in a negative incremental benefit from ensemble weather forecasts in the laden condition.

These contrasting results provide insights into the conditional value of ensemble weather forecasts in ship speed optimization. The strong incremental performance in ballast condition (2.27% additional saving) may reflect greater environmental sensitivity due to reduced displacement and increased susceptibility to wind and wave effects, where hedging against forecast uncertainty through ensemble information proves particularly valuable. The negative incremental performance in laden condition (-0.47%) suggests that when vessels operate at higher displacement with more predictable resistance characteristics, ensemble averaging may be less effective, particularly if the deterministic weather forecast provides adequate accuracy.

Overall, the results demonstrate that ship speed optimization based on meteorological forecasts delivers substantial fuel savings relative to baseline operations, with forecast-based optimization reducing consumption by approximately 7% across the two voyages. Ensemble weather forecasts provide net incremental benefits of 0.96% on aggregate, though with context-dependent performance that varies by loading condition. These findings contribute to understanding the operational value of ensemble meteorological forecasts for maritime decision-making, while highlighting that ensemble forecast-based approaches do not universally outperform deterministic weather forecasts in every scenario. The results motivate further investigation into adaptive utilization strategies of weather forecasts and more sophisticated stochastic optimization formulations that better account for operational nonlinearities under different loading states.

4.4 Conclusion

In this chapter, we establish the rhumb line-based data fusion method as an effective approach for integrating low-frequency noon report data with high-resolution ERA5 reanalysis data. Then, we develop an innovative dual-model evaluation framework that explicitly distinguishes prediction capability from physical reality. This framework simulates the planning-execution gap that exists in actual maritime operations, enabling a fair comparison of different forecast types without the confounding factors that plague real-world A/B testing. Finally, applying this framework to compare deterministic versus ensemble weather forecasts reveals that ensemble weather forecasts provide context-dependent operational benefits. Computational experiments on two voyages demonstrate that forecast-based optimization delivers fuel savings of approximately 7% relative to baseline operations, with ensemble weather forecasts providing aggregate incremental benefits of 0.96% compared to deterministic forecast-based optimization. However, this benefit masks heterogeneity: positive performance in ballast voyage (2.27% additional saving) but negative performance in laden voyage (-0.47%), revealing that the operational value of ensemble weather forecasts varies with vessel loading states.

Chapter 5

Ship Speed Optimization under Evolving Meteorological Conditions

Speed optimization is essential for reducing a ship's fuel consumption in maritime transportation. Previous studies have treated speed optimization as a static decision made before departure. The resulting speed profile is then executed without modification throughout the voyage. Thus, static speed optimization assumes that the meteorological forecasts obtained before departure remain valid throughout the voyage. However, meteorological organizations, such as NOAA, typically update weather forecasts four times daily. The static optimization paradigm fails to exploit these regularly updated weather forecasts, which limits fuel savings in real-world operations.

This chapter develops a rolling horizon optimization framework to incorporate continuously updated meteorological forecasts throughout voyages. Section 5.1 presents the problem and the optimization framework. Section 5.2 develops the time-distance network model and ship's FCR prediction model. Section 5.3 describes the experimental design and validation voyages. Section 5.4 reports the results.

5.1 A Rolling Horizon Framework for Ship Speed Optimization

In this section, we present the rolling horizon optimization framework for ship speed optimization under continuously updated meteorological forecasts. Section 5.1.1 characterizes the dynamic nature of the ship speed optimization problem. It identifies three features that distinguish the problem from the static optimization: a multi-stage decision structure, continuously evolving meteorological information, and temporal decision interdependencies. Section 5.1.2 develops the rolling horizon optimization framework, detailing how speed profiles are iteratively re-optimized and executed throughout the voyage using the most current meteorological forecasts available at each decision epoch.

5.1.1 Problem Description

The transoceanic voyage of a vessel fundamentally constitutes a dynamic optimization problem rather than a static one. This dynamic nature arises from three interrelated characteristics that distinguish it from conventional optimization frameworks: sequential decision-making across multiple stages, evolution of meteorological information, and state-dependent decision interdependencies.

Consider a vessel sailing from an origin port to a destination port over a voyage spanning several days or weeks. The decision-maker must determine sailing speeds for successive segments of the voyage. This problem exhibits three defining characteristics that distinguish it from static optimization problems and necessitate an adaptive replanning framework. First, the multi-stage structure means that each speed decision affects not only the ship's immediate fuel consumption but also its arrival time at the next decision point. This point dictates when an updated forecast becomes available. Meteorological agencies issue weather forecasts at regular intervals, creating discrete decision epochs throughout the voyage. Second, information continuously evolves as each forecast update provides revised weather predictions for the entire remaining voyage, with more recent forecasts offering greater accuracy due to reduced lead times. Third, decisions exhibit temporal interdependency: speed choices at any stage constrain the feasible decision space at subsequent stages, as faster speeds early in the voyage may enable the vessel to traverse regions before adverse weather arrives, while conservative choices may delay it into different meteorological regimes.

The dynamic structure above renders static optimization approaches fundamentally inadequate. Static optimization determines the entire speed profile before departure and cannot incorporate weather forecast updates during voyage execution or adapt when actual conditions diverge from initial predictions. The challenge is therefore to develop a solution framework that systematically accommodates multi-stage decision-making, leverages continuously updated meteorological information, and accounts for decision interdependencies while maintaining computational tractability.

Formally, the speed optimization problem seeks to minimize the total fuel consumption accumulated over all segments of the voyage from origin to destination port. This objective must be achieved subject to two critical constraints. First, the vessel must arrive at the destination within the specified maximum allowable sailing time, ensuring schedule reliability for cargo delivery and port operations. Second, sailing speeds in each segment must remain within safe operational bounds, reflecting engine capacity limitations and navigational safety considerations. The decision variables are the sailing speeds assigned to each segment throughout the voyage, which collectively determine both the total fuel consumed and the arrival time at the destination.

5.1.2 Solution Framework: Rolling Horizon Optimization

To address the optimization problem characterized above, we adopt a **rolling horizon optimization (RHO)** framework. The fundamental principle of this framework is to decompose an overarching time-evolving optimization problem into a sequence of more manageable deterministic optimization subproblems that roll forward in time, each solved with the benefit of the most current weather forecasts. At each decision epoch, the framework solves an optimization problem that plans for the remaining voyage, but implements only the immediate next decision before repeating the process with updated weather information. This "plan globally, act locally" paradigm enables the decision-maker to continuously revise speed decisions as new meteorological forecasts become available. Table 5.1 summarizes the key notation introduced for the RHO framework.

In this chapter, we employ NOAA meteorological forecasts, which possess a convenient structural property: the forecast update cycle equals the forecast temporal resolution, both being 6 hours. This alignment, denoted as $\Delta T^{\text{update}} = \Delta T$, ensures that each RHO cycle involves planning over segments of duration ΔT and executing exactly one segment before the next weather forecast becomes available.[3] This one-to-one correspondence between update cycles and segment durations simplifies both the framework presentation and computational implementation. The RHO framework operates through an iterative four-step cycle that repeats at each forecast update point $h \in \mathcal{H}$ throughout the voyage.

Step 1: Forecast Acquisition. When the vessel reaches update point h at time T_h^{start}, the decision-maker obtains the latest meteorological forecast dataset $\mathbf{W}_h$ issued by the weather forecasting agency. Let T_h^{nearest} denote the most recent forecast issue time prior to or at T_h^{start}; this forecast is issued at a time satisfying $T_h^{\text{nearest}} \leq T_h^{\text{start}} < T_h^{\text{nearest}} + \Delta T^{\text{update}}$. The forecast dataset contains weather predictions at discrete time intervals: $\mathbf{W}_h = \{\mathbf{W}_h^0, \mathbf{W}_h^1, \mathbf{W}_h^2, ..., \mathbf{W}_h^{K_{\max}}\}$, where each $\mathbf{W}_h^\kappa$ represents the forecasted meteorological conditions at time $T_h^{\text{nearest}} + \kappa \cdot \Delta T$ for $\kappa \in \{0, 1, ..., K_{\max}\}$.

[3] More generally, if the forecast update cycle ΔT^{update} differs from the segment duration ΔT, the framework would execute segments covering the time interval $[T_h^{\text{start}}, T_h^{\text{start}} + \Delta T^{\text{update}})$ before re-optimizing. The equal-duration case employed here is characteristic of most operational meteorological forecasting systems, where forecast updates and data temporal resolution typically coincide.

For each remaining segment s in the current optimization, the relevant meteorological conditions $\mathbf{W}_h(s)$ are extracted from the appropriate time step in this forecast dataset. The temporal structure determining how the remaining voyage time T_h^{rem} is divided into remaining segments, and how segment durations are aligned with forecast time steps, is detailed in Section 5.1.2.1.

Table 5.1: Notation for the rolling horizon optimization framework

Symbol	Description	Unit
Sets and Indices		
$\mathcal{H}$	Set of weather forecast update points, $\{1, 2, ..., H\}$	-
$h \in \mathcal{H}$	Index of weather forecast update point (optimization iteration)	-
$\mathcal{S}_h$	Set of remaining segments at update h, $\{1, 2, ..., S_h\}$	-
$s \in \mathcal{S}_h$	Segment index (re-numbered at each update)	-
Parameters		
Voyage-related Parameter		
S_h	Number of remaining segments at the hth update	-
H	Number of weather forecast updates	-
Temporal Parameter		
ΔT^{update}	Forecast update cycle	hours
ΔT	Forecast temporal resolution (segment duration)	hours
T_h^{start}	Optimization start time at update h	time
T_h^{nearest}	Most recent forecast issue time before T_h^{start}	time
Δt_h	Time offset: $\Delta t = T_h^{\text{start}} - T_h^{\text{nearest}}$ at update h	hours
$\Delta \tau_h^s$	Duration of segment s at update h	hours
T_h^{rem}	Remaining available sailing time at update h	hours
Meteorological Forecast Parameter		
$K_{\max}$	Maximum forecast horizon in time steps (64 for 16 days)	-
$\mathbf{W}_h$	Complete weather forecast dataset issued at update h	-
$\mathbf{W}_h^{\kappa}$	Weather weather forecast at time step $\kappa \in \{0, 1, ..., K_{\max}\}$ (at time $T_h^{\text{nearest}} + \kappa \cdot \Delta T$)	-
$\mathbf{W}_h(s)$	Weather conditions for segment s at update h	-
Decision Variables		
v_h^s	Sailing speed for segment s at update h	knots

Step 2: Optimization. Using the newly acquired forecast $\mathbf{W}_h$, the framework formulates and solves a deterministic optimization problem to minimize the total fuel consumption over all remaining segments from the vessel's current position to the destination port. This optimization yields a complete speed profile $\{v_h^1, v_h^2, v_h^3, ..., v_h^{S_h}\}$ specifying the optimal sailing speed for each of the S_h remaining segments. The

optimization considers the entire remaining voyage horizon, thereby accounting for the long-term implications of near-term speed decisions.

Step 3: Execution. Despite computing a complete speed profile for all remaining segments, the RHO framework implements only the first element of this solution: the vessel sails the next segment at speed v_h^1 for duration ΔT. The remainder of the optimized profile $\{v_h^2, v_h^3, ..., v_h^{S_h}\}$ is not executed. This selective implementation is the defining characteristic of the rolling horizon approach, preserving flexibility to revise subsequent decisions when more accurate meteorological information becomes available.

Step 4: Rolling Forward. After completing the segment sailed at speed v_h^1, the vessel advances to the next forecast update point $h + 1$. At this new decision epoch, the problem state has evolved: the vessel occupies a new geographic position, one segment has been completed (so $S_{h+1} = S_h - 1$), and the remaining segments are renumbered starting again from 1. The cycle then repeats from Step 1 with the acquisition of the newly available weather forecast. This iterative process continues until the vessel reaches its destination, at which point $h = |\mathcal{H}|$ and only one segment remains.

5.1.2.1 Temporal Structure at Each Update

An implementation detail concerns how the remaining voyage time T_h^{rem} at update point h is divided into S_h segments, each corresponding to a forecast time step. This temporal structure ensures proper alignment between segment boundaries and forecast data points, enabling each segment to utilize distinct meteorological predictions.

The key subtlety arises from the potential misalignment between the vessel's optimization start time T_h^{start} and the forecast issue times. Specifically, T_h^{start} may occur at any point within the interval $[T_h^{\text{nearest}}, T_h^{\text{nearest}} + \Delta T^{\text{update}})$, where T_h^{nearest} denotes the most recent forecast issue time. The time offset $\Delta t_h = T_h^{\text{start}} - T_h^{\text{nearest}}$ represents how far the vessel's decision point lags behind the forecast issuance. This offset affects the segment duration structure, which can be formally expressed as:

$$\Delta\tau_h^s = \begin{cases} \Delta T - \Delta t_h, & s = 1 \\ \Delta T, & s \in \{2, ..., S_h - 1\} \\ T_h^{\text{rem}} - (S_h - 1) \times \Delta T - (\Delta T - \Delta t_h), & s = S_h \end{cases} \tag{5.1}$$

where $\Delta\tau_h^s$ denotes the duration of segment s in the hth optimization. The first segment has a shorter duration $\Delta T - \Delta t_h$ to ensure it concludes precisely at time $T_h^{\text{nearest}} + \Delta T$, coinciding with the next forecast update. Middle segments align perfectly with forecast time steps, each having duration ΔT. The final segment accommodates the remaining time to reach the destination within the voyage deadline. This temporal partitioning ensures that segment boundaries align with forecast update points wherever possible, enabling the vessel to receive new meteorological information precisely when transitioning between segments. Under this structure, each segment s corresponds to exactly one forecast time step.

5.1.2.2 Illustration of the Framework

Figure 5.1 provides a comprehensive illustration of the rolling horizon process across three complementary perspectives. Panel (a) depicts the complete framework over the voyage timeline. The horizontal axis represents time, marked by forecast issue times $T_1^{\text{issue}}, T_2^{\text{issue}}, T_3^{\text{issue}}, ..., T_{|\mathcal{H}|}^{\text{issue}}$ occurring at intervals of $T_{\text{update}} = \Delta T = 6$ hours. Each row corresponds to one RHO iteration at update point h. At $h = 1$ (departure), the vessel obtains weather forecast $\mathbf{W}_1$ and computes the initial speed profile $\{v_1^1, v_1^2, v_1^3, ..., v_1^{S_1}\}$ covering all segments to the destination. However, only v_1^1 (solid color) is implemented, while the remaining planned speeds (lighter shading) are provisional. At update $h = 2$, having completed the first segment, the vessel obtains updated weather forecast $\mathbf{W}_2$ and re-optimizes for the $S_2 = S_1 - 1$ remaining segments, computing $\{v_2^1, v_2^2, ..., v_2^{S_2}\}$ and implementing only v_2^1. This pattern continues until the final segment.

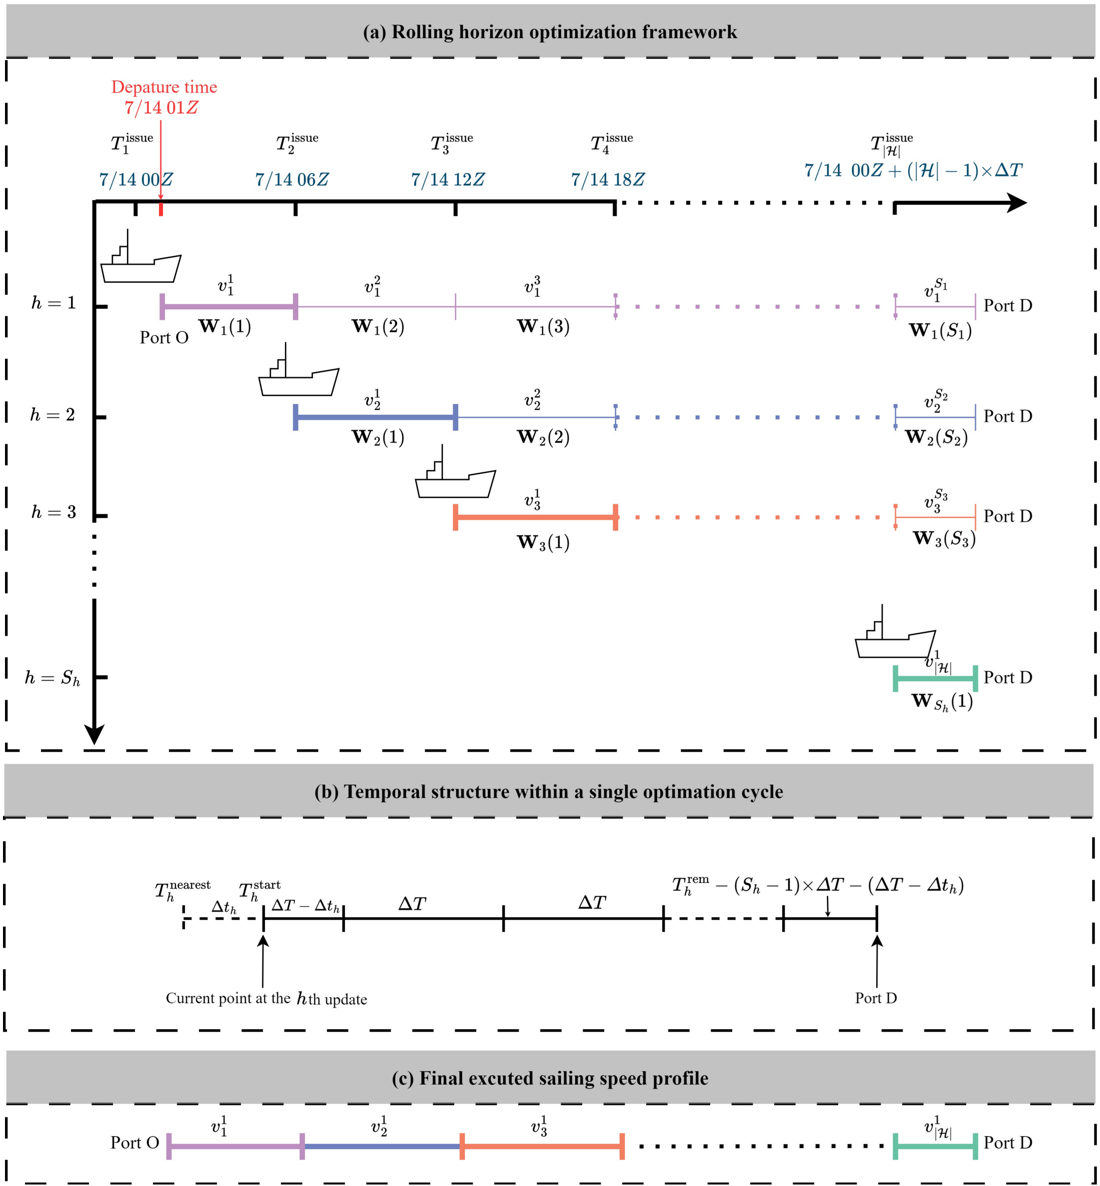

Figure 5.1: Rolling horizon optimization framework for ship speed optimization. Panel (a) illustrates the complete RHO process over the voyage timeline. Panel (b) details the temporal structure within a single optimization cycle at update h. Panel (c) shows the final executed speed profile.

Panel (b) illustrates the temporal structure described above, showing how the time offset affects segment durations and ensures alignment with forecast time steps. Panel

(c) shows the final executed speed profile $\{v_1^1, v_2^1, v_3^1, ..., v_{|\mathcal{H}|}^1\}$, comprising the first-segment speeds from each successive optimization. Each executed speed represents the optimal decision made with the most current meteorological information available at that decision epoch.

The power of the RHO framework lies in its ability to balance foresight with adaptability: by optimizing over the entire remaining voyage at each step, it maintains strategic awareness of how current decisions affect future options and fuel consumption, yet by executing only the immediate decision and re-optimizing, it continuously incorporates updated meteorological information and adapts to evolving conditions. The complete algorithmic procedure for the rolling horizon framework is formally presented in Algorithm 2 in Appendix A.1.2.

Table 5.2: Notation for the time-distance network model

Symbol	Description	Unit
Sets and Indices		
$\mathcal{G}_h$	Time-distance network at update h: $\mathcal{G}_h = (\mathcal{V}_h, \mathcal{A}_h)$	-
$\mathcal{V}_h$	Set of nodes (state space) at update h	-
$\mathcal{A}_h$	Set of arcs (feasible transitions) at update h	-
$\mathcal{T}_h$	Set of discrete time points at update h: $\{\tilde{t}_0, \tilde{t}_1, ..., \tilde{t}_{S_h}\}$	-
$\mathcal{L}_h^s$	Set of possible remaining distances at time point $\tilde{t}_s$	-
$(\tilde{t}_s, \ell)$	Node representing state (time $\tilde{t}_s$, remaining distance ℓ)	-
Parameters		
L_h^{rem}	Remaining voyage distance at update h	nautical miles
Δd	Distance discretization interval	nautical miles
$v_{\min}$	Minimum allowable sailing speed	knots
$v_{\max}$	Maximum allowable sailing speed	knots
Decision Variables and Outputs		
$x^h_{\tilde{t}_s \ell \tilde{t}_{s'} \ell'}$	Binary: 1 if arc from $(\tilde{t}_s, \ell)$ to $(\tilde{t}_{s'}, \ell')$is selected, 0 otherwise	-
$v^h_{\tilde{t}_s \ell \tilde{t}_{s'} \ell'}$	Sailing speed on arc from $(\tilde{t}_s, \ell)$ to $(\tilde{t}_{s'}, \ell')$	knots
$F^h_{\tilde{t}_s \ell \tilde{t}_{s'} \ell'}$	Fuel consumption (weight) on arc from $(\tilde{t}_s, \ell)$ to $(\tilde{t}_{s'}, \ell')$	tons
π_h	Optimal path from source to sink at update h	-

5.2 Deterministic Subproblem: Time-Distance Network Formulation

At each update point h in the rolling horizon framework, the central computational task is to determine the optimal speed profile for the remaining voyage that minimizes total fuel consumption under the currently available meteorological forecast. This section presents a specialized graph-based formulation that we term the **time-distance network (TDN)**. Before presenting the detailed construction of the time-distance network, we first present the notation for the TDN formulation, as shown in Table 5.2.

5.2.1 Speed Optimization Model based on the Time-Distance Network

5.2.1.1 Construction of the time-distance network

The time-distance network $\mathcal{G}_h = (\mathcal{V}_h, \mathcal{A}_h)$ constructed at update point h is a directed acyclic graph where nodes represent feasible spatiotemporal states of the vessel and arcs represent feasible speed decisions connecting these states.

(1) Node definition

Each node $(\tilde{t}_s, \ell) \in \mathcal{V}_h$ represents a state where the vessel is at time point $\tilde{t}_s$ with remaining distance ℓ to the destination. The construction involves two interrelated discretization processes: temporal discretization that defines the time points, and spatial discretization that determines the remaining distances to the destination at each time point. Recall that at update h, the remaining voyage comprises S_h segments with durations $\{\Delta\tau_h^s\}_{s=1}^{S_h}$ given by Equation (5.1). The discrete time points in the network correspond to the boundaries between these segments:

$$\mathcal{T}_h = \left\{\tilde{t}_0, \tilde{t}_1, \tilde{t}_2, ..., \tilde{t}_{S_h}\right\}, \tag{5.2}$$

where $\tilde{t}_0 = T_h^{\text{start}}$ represents the current time (start of optimization), and subsequent time points are defined recursively as:

$$\tilde{t}_s = \tilde{t}_{s-1} + \Delta\tau_h^s, \quad s = 1, 2, ..., S_h \tag{5.3}$$

with the final time point $\tilde{t}_{S_h}$ corresponding to the latest allowable arrival time at the destination. This temporal discretization ensures that each time interval $[\tilde{t}_{s-1}, \tilde{t}_s]$ corresponds to exactly one segment s and aligns with a forecast time step from $\mathbf{W}_h$.

For each time point $\tilde{t}_s \in \mathcal{T}_h$ (corresponding to the completion of segment s), the set of possible remaining distances $\mathcal{L}_h^s$ is determined by the physical constraints of sailing speeds. From the initial state $(\tilde{t}_0, L_h^{\text{rem}})$, the vessel's reachable positions at subsequent time points are bounded by sailing at speeds between $v_{\min}$ and $v_{\max}$. For time point $\tilde{t}_s$, the feasible distance range $[\text{lb}_s, \text{ub}_s]$ is computed as:

$$\text{lb}_s = \max\left(0, L_h^{\text{rem}} - \sum_{i=1}^{s} v_{\max} \cdot \Delta\tau_h^i\right) \tag{5.4}$$

$$\text{ub}_s = \max\left(0, L_h^{\text{rem}} - \sum_{i=1}^{s} v_{\min} \cdot \Delta\tau_h^i\right). \tag{5.5}$$

The physical interpretation of these bounds is as follows. Starting from the initial state with remaining distance L_h^{rem}, the vessel traverses the first s segments. If the vessel sails at maximum speed $v_{\max}$ throughout these segments, it covers the greatest possible distance $\sum_{i=1}^{s} v_{\max} \cdot \Delta\tau_h^i$, leaving the minimum remaining distance to the destination, which defines the lower bound lb_s. Conversely, if the vessel sails at minimum speed $v_{\min}$, it covers the least possible distance $\sum_{i=1}^{s} v_{\min} \cdot \Delta\tau_h^i$, leaving the maximum remaining distance, which defines the upper bound ub_s. The $\max(0, \cdot)$ operator ensures that remaining distances are non-negative, as negative values would indicate positions beyond the destination, which are physically infeasible. If $\text{ub}_s -$

$\mathrm{lb}_s \neq 0$, this range is discretized using interval Δd to generate the discrete set:

$$\mathcal{L}_h^s = \{\mathrm{lb}_s, \mathrm{lb}_s + \Delta d, \mathrm{lb}_s + 2\Delta d, ..., \mathrm{ub}_s\}. \tag{5.6}$$

The number of discrete values is $\lfloor(\mathrm{ub}_s - \mathrm{lb}_s)/\Delta d\rfloor + 1$. Special cases occur at the boundaries: the source node corresponds to the initial state $(\tilde{t}_0, L_h^{\mathrm{rem}})$, comprising a single node with remaining distance equal to the full remaining voyage; the sink node corresponds to arrival $(\tilde{t}_{S_h}, 0)$, representing the destination. The complete node set is:

$$\mathcal{V}_h = \left\{(\tilde{t}_s, \ell) : \tilde{t}_s \in \mathcal{T}_h, \ell \in \mathcal{L}_h^s\right\}. \tag{5.7}$$

(2) Arc definition

A crucial property of the TDN is that arcs exist only between nodes at consecutive time points. An arc from node $(\tilde{t}_{s-1}, \ell)$ to node $(\tilde{t}_s, \ell')$ represents a speed decision for segment s. The existence of an arc is dependent upon the fulfillment of three conditions. First, the nodes must be temporally adjacent, meaning the arc connects consecutive time points $\tilde{t}_{s-1}$ and $\tilde{t}_s = \tilde{t}_{s-1} + \Delta\tau_s$. Second, the vessel must progress toward the destination: $\ell' < \ell$. Third, the implied sailing speed $v^h_{\tilde{t}_{s-1}\ell\tilde{t}_s\ell'} = (\ell - \ell')/(\tilde{t}_s - \tilde{t}_{s-1})$ must satisfy the operational constraints $v_{\min} \leq v^h_{\tilde{t}_{s-1}\ell\tilde{t}_s\ell'} \leq v_{\max}$.

This adjacency constraint ensures that each arc corresponds to a single segment's sailing decision, maintaining the one-to-one correspondence between network arcs and physical segments established in the RHO framework. The arc set can be formally expressed as:

$$\begin{aligned}\mathcal{A}_h = \Big\{&((\tilde{t}_{s-1}, \ell), (\tilde{t}_s, \ell')) : s \in \{1, 2, ..., S_h\}, \ell \in \mathcal{L}_h^{s-1}, \\ &\qquad \ell' \in \mathcal{L}_h^s, \ell' < \ell, v_{\min} \leq \frac{\ell - \ell'}{\tilde{t}_s - \tilde{t}_{s-1}} \leq v_{\max}\Big\}.\end{aligned} \tag{5.8}$$

(3) Arc weight computation

The weight of each arc represents the fuel consumption required to traverse it, computed using the ship's FCR prediction model formulated by the artificial neural network (ANN). The computation requires determining not only the sailing speed but also the meteorological conditions and vessel operational state at the appropriate location and time.

For arc $((\tilde{t}_{s-1}, \ell), (\tilde{t}_s, \ell')) \in \mathcal{A}_h$, we first determine the sailing speed from the distance traveled and time elapsed:

$$v^h_{\tilde{t}_{s-1}\ell\tilde{t}_s\ell'} = \frac{\ell - \ell'}{\tilde{t}_s - \tilde{t}_{s-1}}. \tag{5.9}$$

To obtain the meteorological conditions for this arc, we must identify the geographic location corresponding to the arc's starting point and the time at which the vessel would be at that location. Assuming the vessel follows a rhumb line (constant bearing) route, the geographic coordinates can be extrapolated from the route geometry. Let (λ_O, φ_O) denote the coordinates of the departure port and θ the vessel's true heading. For a node $(\tilde{t}_s, \ell)$ in the network, the distance already traveled from the departure port is $L - \ell$. The geographic coordinate (λ, φ) corresponding to this position is calculated through the following rhumb line navigation formulas (Williams, 2013):

$$\delta = \frac{L - \ell}{R} \tag{5.10a}$$

$$\varphi = \varphi_O + \delta \cos\theta \tag{5.10b}$$

$$\Delta\psi = \ln\left(\frac{\tan(\frac{\pi}{4} + \frac{\varphi}{2})}{\tan(\frac{\pi}{4} + \frac{\varphi_O}{2})}\right) \tag{5.10c}$$

$$q = \begin{cases} \cos\varphi_O, & \text{if sailing on constant latitude course} \\ \frac{|\varphi - \varphi_O|}{\Delta\psi}, & \text{otherwise} \end{cases} \tag{5.10d}$$

$$\Delta\lambda = \frac{\delta \sin\theta}{q} \tag{5.10e}$$

$$\lambda = \lambda_O + \Delta\lambda, \tag{5.10f}$$

where R is the Earth's radius, δ is the angular distance, and $\Delta\lambda$ is the difference in longitude between the departure port and this position. The arrival time at this position is determined by the vessel's progress through previous segments.

With the geographic coordinates and time identified, the meteorological conditions $\mathbf{w}$ for the arc are extracted from the forecast dataset $\mathbf{W}_h$ at the corresponding forecast time step. The vessel operational state $\mathbf{u}$ (draft and loading condition) is obtained from the most recent operational record. The arc weight is then computed as:

$$F^h_{\tilde{t}_{s-1}\ell\tilde{t}_s\ell'} = \mathcal{F}_{\text{ANN}}\left(v^h_{\tilde{t}_{s-1}\ell\tilde{t}_s\ell'}, \mathbf{w}, \mathbf{u}; \boldsymbol{\theta}_{\text{ANN}}\right) \times (\tilde{t}_s - \tilde{t}_{s-1}). \tag{5.11}$$

A special case arises for arcs connecting to the sink node $(\tilde{t}_{S_h}, 0)$. For some nodes in the penultimate stage, the remaining distance ℓ may be less than the minimum distance achievable at speed $v_{\min}$ over the final segment duration, i.e., $\ell < v_{\min} \times \Delta\tau_h^{S_h}$. This indicates that the vessel can reach the destination within the allotted time even when sailing at speeds below $v_{\min}$, or equivalently, at any speed in $[v_{\min}, v_{\max}]$ since the distance is short. Since the optimization objective is to minimize the ship's fuel consumption, we assign such arcs the minimum fuel consumption achievable across all feasible speeds:

$$F^h_{\tilde{t}_{S_h-1}\ell\tilde{t}_{S_h}0} = \min_{v^h_{\tilde{t}_{S_h-1}\ell\tilde{t}_{S_h}0} \in \{v_{\min}, v_{\min}+0.1, \ldots, v_{\max}\}} \left\{\mathcal{F}_{\text{ANN}}(v^h_{\tilde{t}_{S_h-1}\ell\tilde{t}_{S_h}0}, \mathbf{w}, \mathbf{u}; \boldsymbol{\theta}_{\text{ANN}}) \times \Delta\tau_h^{S_h}\right\}, \tag{5.12}$$

where the minimization is performed over discretized speed values with 0.1-knot intervals.

(4) Graph Pruning for Computational Efficiency

Not all nodes in $\mathcal{V}_h$ lie on feasible paths from source to sink. Nodes that cannot reach the destination within the remaining time and speed constraints represent infeasible intermediate states and can be eliminated to reduce computational burden. We employ a backward pruning procedure: starting from the penultimate time stage $\tilde{t}_{S_h-1}$, we iteratively remove nodes with zero out-degree (no outgoing arcs to feasible successor nodes). This process propagates backward through time stages, eliminating all nodes that do not lie on at least one complete path from source to sink. The pruned network retains only nodes participating in feasible solutions, significantly reducing the problem size while preserving solution optimality.

5.2.1.2 Model formulation

With the TDN $\mathcal{G}_h$ fully constructed, the speed optimization problem for update h can be formulated as a minimum-cost path problem. The objective is to find the path from the source node $(\tilde{t}_0, L_h^{\text{rem}})$ to the sink node $(\tilde{t}_{S_h}, 0)$ that minimizes the total fuel consumption. This is equivalent to solving:

$$\text{minimize} \quad \sum_{((\tilde{t}_{s-1},\ell),(\tilde{t}_s,\ell'))\in\mathcal{A}_h} F^h_{\tilde{t}_{s-1}\ell\tilde{t}_s\ell'} \cdot x^h_{\tilde{t}_{s-1}\ell\tilde{t}_s\ell'} \tag{5.13}$$

$$\text{subject to} \quad \sum_{(\tilde{t}_{s-1},\ell):((\tilde{t}_{s-1},\ell),(\tilde{t}_s,\ell'))\in\mathcal{A}_h} x^h_{\tilde{t}_{s-1}\ell\tilde{t}_s\ell'} - \sum_{(\tilde{t}_{s+1},\ell''):((\tilde{t}_s,\ell'),(\tilde{t}_{s+1},\ell''))\in\mathcal{A}_h} x^h_{\tilde{t}_s\ell'\tilde{t}_{s+1}\ell''} = b^h_{\tilde{t}_s\ell'}, \quad \forall(\tilde{t}_s,\ell') \in \mathcal{V}_h \tag{5.14}$$

$$x^h_{\tilde{t}_{s-1}\ell\tilde{t}_s\ell'} \in \{0,1\}, \quad \forall((\tilde{t}_{s-1},\ell),(\tilde{t}_s,\ell')) \in \mathcal{A}_h \tag{5.15}$$

where the flow balance parameter $b^h_{\tilde{t}_s\ell'}$ is defined as:

$$b^h_{\tilde{t}_s\ell'} = \begin{cases} 1, & (\tilde{t}_s,\ell') = (\tilde{t}_0, L_h^{\text{rem}}) \quad \text{(source node)} \\ -1, & (\tilde{t}_s,\ell') = (\tilde{t}_{S_h}, 0) \quad \text{(sink node)} \\ 0, & \text{otherwise (intermediate nodes)} \end{cases} \tag{5.16}$$

Constraint (5.14) enforces flow conservation: for the source node, flow out equals one; for the sink node, flow in equals one; for all intermediate nodes, flow in equals flow out. Constraint (5.15) ensures that each arc is either selected or not.

This formulation represents a classical shortest path problem on a directed acyclic graph, which can be solved efficiently using Dijkstra's algorithm. The algorithm identifies the path π_h that minimizes the objective function (5.13), corresponding to the speed profile with minimum total fuel consumption. The shortest path π_h returned by Dijkstra's algorithm is a sequence of nodes:

$$\pi_h = \left\{(\tilde{t}_0, \ell_0), (\tilde{t}_1, \ell_1), (\tilde{t}_2, \ell_2), ..., (\tilde{t}_{S_h}, \ell_{S_h})\right\}, \tag{5.17}$$

where $\ell_0 = L_h^{\text{rem}}$ and $\ell_{S_h} = 0$. Each consecutive pair of nodes in this path corresponds to one arc in the network, representing the speed decision for one segment. This yields the complete optimal speed profile $\{v_h^1, v_h^2, ..., v_h^{S_h}\}$ for all remaining segments. Consistent with the rolling horizon principle articulated in Section 5.1.2, only the first element v_h^1 of this profile is implemented. After sailing segment 1 at speed v_h^1 for duration $\Delta\tau_h^1$, the vessel reaches the next forecast update point $h+1$, at which time the entire optimization process repeats with updated meteorological information. The remaining speeds $\{v_h^2, v_h^3, ..., v_h^{S_h}\}$ are discarded once new weather information becomes available. Algorithm 3 in Appendix A.1.2 provides the complete pseudocode for the TDN construction described in this section.

5.2.2 Ship Fuel Consumption Prediction: The Artificial Neural Network Model

As established in the previous section, the construction of the time-distance network requires evaluating the ship's fuel consumption associated with each feasible arc, which represents sailing a specific distance at a specific speed under specific meteorological conditions. This evaluation task demands a prediction model capable of accurately estimating ship's FCRs across diverse operational contexts. This chapter employs an ANN to predict a ship's FCRs. ANNs have emerged as powerful tools for ship's fuel consumption modeling due to their ability to capture complex nonlinear relationships between operational parameters and energy consumption without requiring explicit physical equations (Hassoun, 1995; Hastie, 2009). Unlike physics-based resistance models that demand extensive vessel-specific hydrodynamic coefficients and are sensitive to modeling assumptions, ANNs learn energy consumption patterns directly from historical operational data, making them particularly suitable for practical applications where detailed technical specifications may be unavailable.

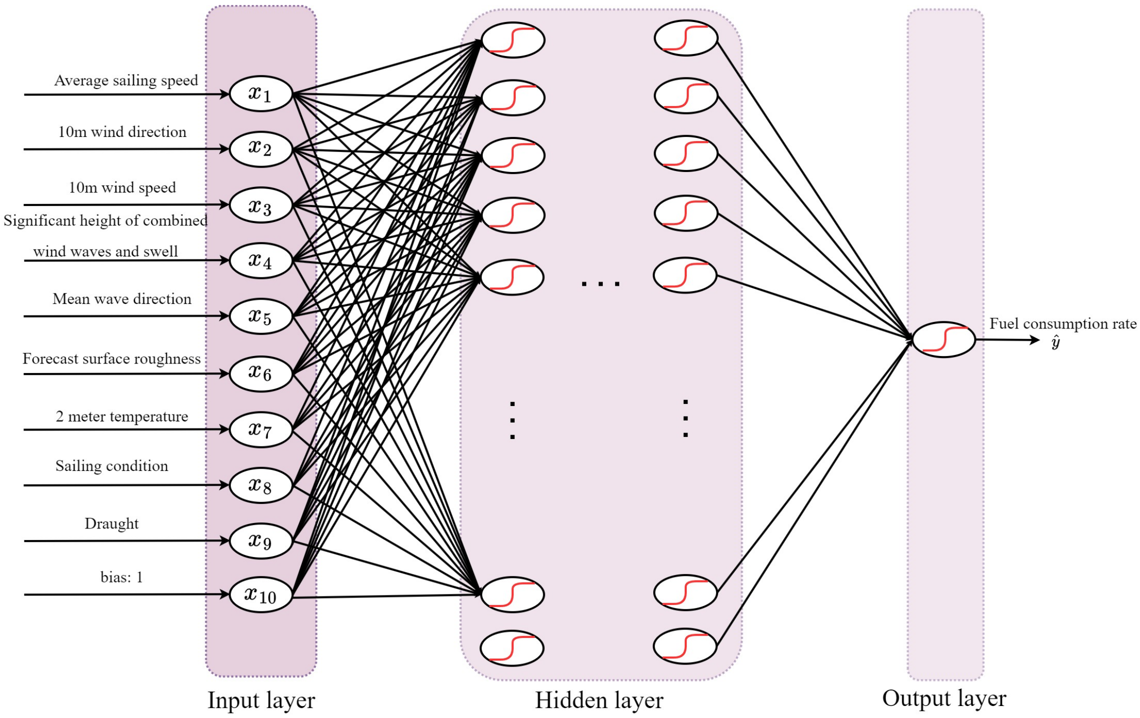

Figure 5.2: Structure of the ANN model for fuel consumption rate prediction. The input layer receives operational variables (speed, meteorological conditions, and vessel operational state), processes them through multiple hidden layers with nonlinear activations, and produces the predicted fuel consumption rate at the output layer.

The structure of the ANN model follows a feedforward multilayer perceptron architecture, as illustrated in Figure 5.2. The input layer receives operational variables that characterize the fuel consumption context: sailing speed v, meteorological conditions $\mathbf{w}$, and vessel operational state parameters $\mathbf{u}$. These inputs are processed through one or more hidden layers, where each neuron applies a nonlinear activation function to a weighted combination of its inputs, enabling the network to learn complex feature interactions. The output layer produces a single value representing the predicted fuel consumption rate, denoted as $\mathcal{F}_{\text{ANN}}(v, \mathbf{w}, \mathbf{u}; \boldsymbol{\theta}_{\text{ANN}})$, where $\boldsymbol{\theta}_{\text{ANN}}$ represents the complete set of learned network parameters (weights and biases).

Formally, the ANN computes the fuel consumption rate through a series of transformations. Let $\mathbf{x} = [v, \mathbf{w}, \mathbf{u}]$ denote the concatenated input vector. For a network with $\mathcal{M}$ hidden layers, the computation proceeds as:

$$\mathbf{h}^{(1)} = f^{(1)}(\mathbf{W}^{(1)}\mathbf{x} + \mathbf{b}^{(1)}) \tag{5.18}$$

$$\mathbf{h}^{(\mu)} = f^{(\mu)}(\mathbf{W}^{(\mu)}\mathbf{h}^{(\mu-1)} + \mathbf{b}^{(\mu)}), \quad \mu = 2, ..., \mathcal{M} \tag{5.19}$$

$$\mathcal{F}_{\text{ANN}} = \mathbf{W}^{(\mathcal{M}+1)}\mathbf{h}^{(\mathcal{M})} + \mathbf{b}^{(\mathcal{M}+1)}, \tag{5.20}$$

where $\mathbf{W}^{(\mu)}$ and $\mathbf{b}^{(\mu)}$ are the weight matrix and bias vector for layer μ, $f^{(\mu)}$ denotes the activation function (typically ReLU for hidden layers), and $\mathbf{h}^{(\mu)}$ represents the hidden layer activations. The parameters $\boldsymbol{\theta}_{\text{ANN}} = \{\mathbf{W}^{(\mu)}, \mathbf{b}^{(\mu)}\}_{\mu=1}^{\mathcal{M}+1}$ are learned from historical voyage data through backpropagation, minimizing the discrepancy between predicted and observed fuel consumption rates. In this chapter, the predictive performance of the trained ANN model is evaluated using three standard regression metrics (MSE, MAE, and MAPE) on held-out test data.

5.3 Experimental Design

This section establishes the experimental framework for validating the proposed rolling horizon optimization method. Section 5.3.1 describes the two validation voyages selected from the noon report, providing detailed operational and meteorological characteristics that represent realistic operational scenarios for evaluating the method's effectiveness. Section 5.3.2 presents the comparison framework, defining the proposed RHO method alongside three baseline schemes (fixed speed, single-shot optimization at departure, and optimization with perfect information). Performance metrics are introduced to quantify fuel savings achieved by adaptive re-optimization and measure how closely the method approaches theoretical optimality.

5.3.1 Validation Voyages and Model Configuration

The ANN model is trained on the comprehensive dataset described in Chapter 3, which comprises 331 noon reports collected from a Capesize dry bulk carrier operating between February 2019 and August 2020. To ensure unbiased evaluation of the rolling horizon optimization framework, we reserve two voyages from this dataset as validation cases that do not participate in any stage of model training. The remaining 307 records (92.7% of the dataset) constitute the training set for developing the ANN model, while the two reserved voyages containing 24 records (7.3%) serve as the test set for both assessing fuel consumption prediction accuracy and validating the performance of the proposed optimization framework.

The ANN model employs ten input features to predict hourly fuel consumption rates, as summarized in Table 5.3. To ensure compatibility with operational meteorological forecasts, the three variables provided by the NOAA GEFS control member (10-meter wind speed, 10-meter wind direction, and 2-meter temperature) are included in the training features. To maintain prediction accuracy, additional meteorological and operational variables are incorporated, including wave characteristics, forecast surface roughness, vessel draft, sailing speed, and loading condition. Chapter 3 provides detailed descriptions of each feature's physical interpretation.

With the model configuration established, we now describe the two voyages selected for evaluating the rolling horizon optimization framework. **Voyage I** is an 11-day laden voyage departing on July 14, 2020, covering a total sailing distance of 2701.56 nautical miles with an overall sailing time of 261 hours when following the original speed profile recorded in the noon reports. The vessel maintains a draft of 17.5 meters throughout the voyage, operating in a laden condition with cargo weight approximately at maximum capacity. **Voyage II** is a 10-day laden voyage departing on February 14, 2019, covering a total sailing distance of 2645.18 nautical miles with a total sailing time of 242 hours under the original speed profile.

Table 5.3: Input features and prediction target for ANN fuel consumption prediction model

Data Source	Feature	Min	Max	Mean
	Prediction Target			
Noon report	Hourly fuel consumption rate (mt/h)	0.434	2.184	1.343
	Input Features			
Noon report	Average sailing speed (knot)	6.730	14.920	11.068
	Sailing condition (laden/ballast)	–	–	–
	Draught (m)	8.600	18.700	13.885
ERA5 reanalysis data	Relative wind direction (degree)	0.150	179.675	93.425
	Wind speed (m/s)	0.421	15.178	5.432
	Mean wave period (s)	2.395	13.380	7.138
	Significant height of combined waves and swell (m)	0.150	3.975	1.435
	Relative mean wave direction (degree)	0.769	179.494	87.659
	Forecast surface roughness (m)	2.42e-05	7.23e-04	0.061
	2-meter temperature (°C)	2.286	29.311	25.039

The detailed operational characteristics and meteorological conditions for both voyages are presented in Tables A.1 and A.2 in Appendix 7.3. These tables provide day-by-day information, including fuel consumption rates, sailing speeds, draft, loading conditions, and comprehensive meteorological data obtained from ERA5 reanalysis data, which serve as ground truth for evaluating actual fuel consumption under optimized speed profiles.

5.3.2 Comparison Framework

We evaluate the RHO method against three baseline speed planning approaches to assess the benefits of adaptive speed re-planning with evolving weather forecasts.

(1) Proposed Method: RHO with Dynamic Forecasts

The proposed RHO framework generates a speed profile that achieves total fuel consumption, F^{RHO}, as described in Section 5.1. At each forecast update point h, the method constructs a time-distance network based on the latest NOAA meteorological forecast $\mathbf{W}_h$, solves for the optimal speed profile $\{v_h^1, v_h^2, ..., v_h^{S_h}\}$ for the remaining voyage, and executes only the first segment speed v_h^1 with duration $\Delta\tau_h^1$ before re-optimizing at the next update point. This process repeats through $|\mathcal{H}|$ optimization iterations until the vessel reaches the destination, resulting in an executed speed profile $\{v_1^1, v_2^1, v_3^1, ..., v_{|\mathcal{H}|}^1\}$ with corresponding durations $\{\Delta\tau_1^1, \Delta\tau_2^1, \Delta\tau_3^1, ..., \Delta\tau_{|\mathcal{H}|}^1\}$. The total fuel consumption is calculated as:

$$F^{\text{RHO}} = \sum_{h=1}^{|\mathcal{H}|} \mathcal{F}_{\text{ANN}}(v_h^1, \mathbf{w}_h^{\text{ERA5}}, \mathbf{u}_h; \boldsymbol{\theta}_{\text{ANN}}) \times \Delta\tau_h^1, \tag{5.21}$$

where v_h^1 is the speed executed in the first segment of the hth optimization iteration, $\mathbf{w}_h^{\text{ERA5}}$ represents the actual meteorological conditions during that segment obtained from ERA5 reanalysis data (which serves as ground truth for evaluation), $\mathbf{u}_h$ denotes the vessel operational state (draft and loading condition), and $\Delta\tau_h^1$ is the duration of that segment. This formulation evaluates the actual fuel consumed when following the RHO-generated speed profile under real-world conditions.

(2) Baseline 1: Fixed Speed Scheme

The fixed speed strategy achieves total fuel consumption denoted as F^{fix}, where the vessel maintains a fixed speed throughout the entire voyage. The fixed speed v^{fix} is determined by discretizing the feasible speed range and selecting the speed that minimizes total fuel consumption based on GEFS control member forecast available at departure. Specifically, the lower bound of feasible constant speeds is $v_{\min}^{\text{fix}} = L/T_{\max}$ (the minimum speed required to complete the voyage within the time limit), and the upper bound is $v_{\max}^{\text{fix}} = v_{\max} = 18.0$ knots. The range $[v_{\min}^{\text{fix}}, v_{\max}^{\text{fix}}]$ is discretized in 0.1-knot intervals, and for each candidate speed, the total fuel consumption is computed using the forecast meteorological data. The speed yielding minimum fuel consumption becomes v^{fix}. The voyage is divided into S^{fix} segments based on forecast temporal resolution, and the actual fuel consumption under this fixed speed scheme is:

$$F^{\text{fix}} = \sum_{s=1}^{S^{\text{fix}}} \mathcal{F}_{\text{ANN}}(v^{\text{fix}}, \mathbf{w}_s^{\text{ERA5}}, \mathbf{u}_s; \boldsymbol{\theta}_{\text{ANN}}) \times \Delta\tau_s^{\text{fix}}, \tag{5.22}$$

where $\Delta\tau_s^{\text{fix}}$ denotes the duration of the sth segment under the fixed speed scheme.

(3) Baseline 2: Single-Shot Optimization

The single-shot optimization baseline achieves total fuel consumption denoted as F^{init}, where the vessel follows a speed profile optimized once at departure using the initial weather forecast, without any subsequent updates. This approach captures the value of optimization while highlighting the limitation of ignoring the evolution of weather forecasts. At departure (i.e., the first optimization iteration with $h = 1$), a time-distance network is constructed using the weather forecast $\mathbf{W}_1$ available at time T_1^{start}, and the optimal speed profile $\{v_1^1, v_1^2, ..., v_1^{S_1}\}$ with corresponding durations $\{\Delta\tau_1^1, \Delta\tau_1^2, ..., \Delta\tau_1^{S_1}\}$ is computed via Dijkstra's algorithm on this single network, where S_1 denotes the total number of segments in the initial optimization. The vessel then executes this pre-determined speed profile throughout the entire voyage without any adjustment despite subsequent weather forecast updates. The actual fuel consumption is:

$$F^{\text{init}} = \sum_{s=1}^{S_1} \mathcal{F}_{\text{ANN}}(v_1^s, \mathbf{w}_1^{s,\text{ERA5}}, \mathbf{u}_1^s; \boldsymbol{\theta}_{\text{ANN}}) \times \Delta\tau_1^s, \tag{5.23}$$

where v_1^s is the speed for the sth segment determined in the initial optimization, $\mathbf{w}_1^{s,\text{ERA5}}$ represents the actual meteorological conditions encountered during segment s from ERA5 reanalysis data, $\mathbf{u}_1^s$ denotes the vessel operational state, and $\Delta\tau_1^s$ is the segment duration as planned in the initial optimization. Comparing F^{RHO} with F^{init} quantifies the fuel-saving benefit of incorporating regularly updated forecasts through the rolling horizon approach.

(4) Baseline 3: Optimization with perfect information

To establish a theoretical upper bound on achievable fuel consumption, we define F^{actual} as the total fuel consumption when the speed profile is optimized using perfect meteorological information (i.e., actual conditions from ERA5) known in advance. This oracle baseline assumes that the complete sequence of true meteorological conditions throughout the voyage is available at departure, allowing a single optimization to determine the globally optimal speed profile with perfect foresight. The voyage is divided into S^{actual} segments based on the actual meteorological forecast update cycle, identical to the segmentation used in the rolling horizon framework. This scheme yields the speed profile $\{v_s^{\text{actual}}\}_{s=1}^{S^{\text{actual}}}$. The fuel consumption is:

$$F^{\text{actual}} = \sum_{s=1}^{S^{\text{actual}}} \mathcal{F}_{\text{ANN}}(v_s^{\text{actual}}, \mathbf{w}_s^{\text{ERA5}}, \mathbf{u}_s) \times \Delta\tau_s^{\text{actual}}, \tag{5.24}$$

where $\Delta\tau_s^{\text{actual}}$ denotes the duration of the sth segment under the oracle baseline. This baseline is unattainable in practice (actual conditions are unknown at departure), but it provides an upper bound for optimization performance and enables assessment of how closely the RHO method approximates the theoretical optimum despite weather forecast uncertainty.

(5) Performance metrics

To quantify the performance of the proposed method relative to the baselines, we define three evaluation metrics. The first metric measures fuel savings relative to the fixed speed scheme:

$$\eta_{\text{fix}} = \frac{F^{\text{fix}} - F^{\text{RHO}}}{F^{\text{fix}}} \times 100\%. \tag{5.25}$$

This quantifies the percentage of fuel reduction achieved by the RHO method compared to maintaining a fixed speed throughout the voyage. The second metric isolates the incremental benefit of adaptive re-optimization with updated weather forecasts:

$$\eta_{\text{init}} = \frac{F^{\text{init}} - F^{\text{RHO}}}{F^{\text{init}}} \times 100\%. \tag{5.26}$$

This measures fuel savings relative to single-shot optimization, thereby isolating the value of incorporating weather forecast updates while controlling for the effect of optimization itself. The third metric quantifies deviation from the perfect information bound:

$$\delta_{\text{perfect}} = \left| \frac{F^{\text{RHO}} - F^{\text{actual}}}{F^{\text{actual}}} \right| \times 100\%. \tag{5.27}$$

This measures how closely the RHO method approaches the theoretical optimum achievable with perfect meteorological foresight. Positive values of η_{fix} and η_{init} indicate fuel savings, with larger values representing greater improvement. Smaller values of δ_{perfect} indicate that the RHO method effectively approximates optimal performance despite weather forecast imperfection.

5.4 Computational Results

This section presents the computational results that validate the effectiveness of the proposed rolling horizon speed optimization method. In Section 5.4.1, we describe the development and performance of the ANN fuel consumption prediction model, which serves as the foundation for optimization. In Section 5.4.2, we analyze the speed optimization results on the two validation voyages, comparing the proposed method against the baseline schemes.

5.4.1 Development of the ANN Model

The ANN model for predicting ship's FCRs is implemented using the MLPRegressor function in the scikit-learn library 0.24.2. The hyperparameters of the ANN model are optimized through grid search with five-fold cross validation on the training dataset. Table 5.4 presents the hyperparameters tuned for the ANN regressor, along with their search spaces and optimal values. The key hyperparameters include the network architecture (`hidden_layer_size`), activation function, solver for weight optimization, and maximum number of training iterations.

Table 5.4: Hyperparameters and optimal configuration of the ANN regressor

Hyperparameter	Interpretation	Search Space	Optimal Value
`hidden_layer_size`	Number of neurons in each hidden layer	Layers: {1, 2, 3, 4, 5, 6, 7} Neurons: {8, 16, 32, 64}	(16, 32, 32, 32, 16, 8, 8)
`activation`	Activation function for hidden layers	{tanh, ReLU, sigmoid}	ReLU
`solver`	Optimizer for weight updates	{lbfgs, sgd, adam}	Adam
`max_iter`	Maximum number of iterations	{1,000, 2,000, 3,000, 4,000, 5,000, 6,000, 7,000}	6,000

The search space for `hidden_layer_size` is constrained by architectural principles: the network structure follows a symmetric pattern where the middle hidden layer contains the largest number of neurons, with neuron counts gradually decreasing toward the input and output layers. Specifically, for a network with $\mathcal{M}$ hidden layers indexed $\mu = 1, 2, ..., \mathcal{M}$, the number of neurons satisfies: neurons(μ) $\leq$ neurons(μ+1) for $\mu < \lceil \mathcal{M}/2 \rceil$ and neurons(μ) $\geq$ neurons($\mu + 1$) for $\mu \geq \lceil \mathcal{M}/2 \rceil$. The optimal configuration identified through cross-validation employs a seven-layer architecture with (16, 32, 32, 32, 16, 8, 8) neurons in successive hidden layers, ReLU activation for nonlinearity, and 6,000 training iterations for convergence.

Table 5.5: Predictive performance of the ANN regressor on test data

Model	MSE	MAE	MAPE
ANN Regressor	0.016	0.081	5.95%

After training the ANN model on the entire training dataset using the optimal hyperparameters, its predictive performance is evaluated on the held-out test set. Table 5.5 summarizes the value of performance metrics. The model achieves a MSE of 0.016, MAE of 0.081 tons/hour, and MAPE of 5.95%. These metrics indicate strong predictive accuracy: the MAPE of 5.95% implies that on average, predictions deviate from actual fuel consumption rates by less than 6%. The low MSE and MAE values further confirm that the model effectively captures the nonlinear relationships between operational variables (speed, meteorological conditions, vessel state) and the ship's FCRs, providing reliable fuel consumption estimates across the diverse operating conditions encountered in the test voyages.

5.4.2 Validation of the RHO framework

All speed optimization experiments are conducted on a laptop computer equipped with an Intel Core i7-12700H CPU (2.30 GHz) and 40.0 GB of memory. The optimization problem at each update point is solved using Dijkstra's algorithm implemented in the NetworkX Python library 2.8. The distance discretization interval is set to $\Delta d = 1$ nautical mile, balancing solution accuracy with computational tractability. Speed constraints are $v_{\min} = 8.0$ knots and $v_{\max} = 18.0$ knots. The meteorological forecast update cycle follows NOAA GEFS issuance at 00Z, 06Z, 12Z, and 18Z, resulting in $\Delta T^{\text{update}} = 6$ hours between updates. For Voyage I, the rolling horizon optimization process takes 146 minutes of total computation time, generating speed decisions for 39 segments. For Voyage II, the process requires 220 minutes and produces decisions for 40 segments.

Figure 5.3 illustrates the optimized sailing speeds and corresponding fuel consumption for each segment of the two validation voyages. The blue bars represent fuel consumption in each segment, calculated by multiplying the segment duration by the fuel consumption rate predicted by the ANN model under actual meteorological conditions from ERA5 reanalysis data. The red line with markers indicates the optimized sailing speed for each segment. It can be seen that the fuel consumption in each segment increases as the ship sailing speed increases. After adopting the RHO framework, for Voyage I, the vessel sails for 6 hours in each of the first 38 segments and 0.97 hours in the final segment, totaling 228.97 hours. For Voyage II, the vessel sails for 6 hours in each of the first 39 segments and 3.26 hours in the final segment, totaling 237.26 hours. These durations confirm that the optimized speed profiles enable on-time arrival at the destination within the maximum allowable sailing time.

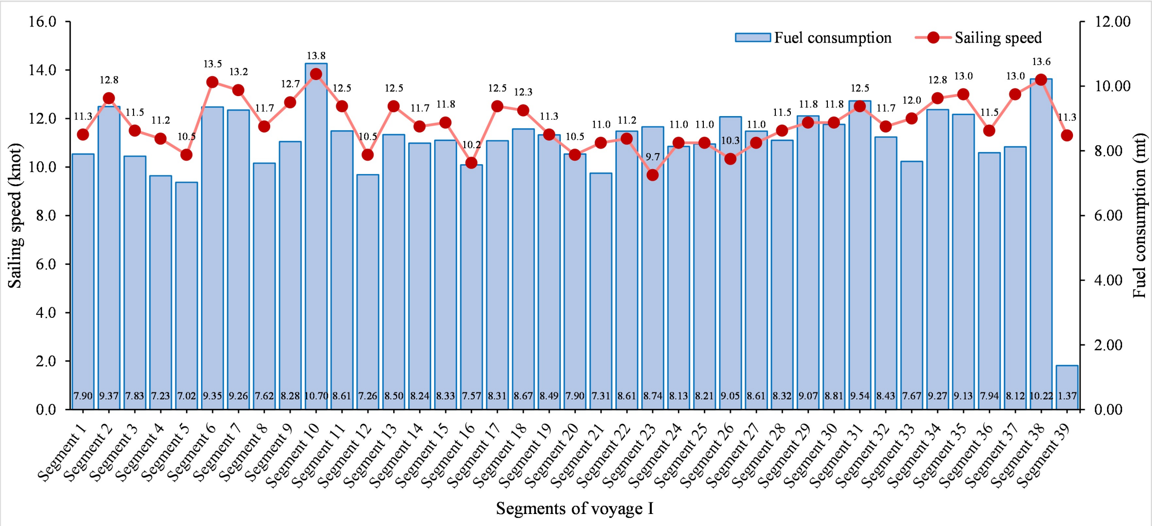

(a) Voyage I

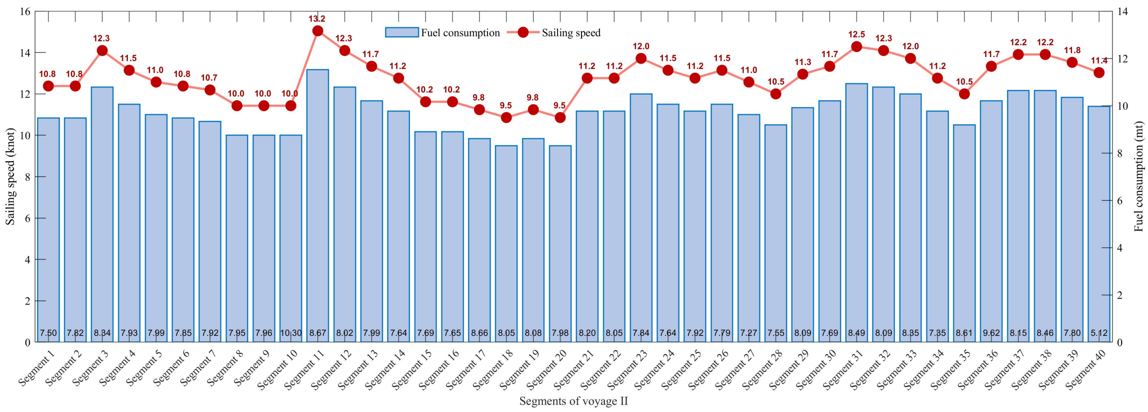

(b) Voyage II

Figure 5.3: Fuel consumption and sailing speed profiles after optimization. (a) Voyage I: 39 segments over 228.97 hours. (b) Voyage II: 40 segments over 237.26 hours.

Table 5.6 presents a comprehensive comparison of fuel consumption under the four speed schemes: the proposed RHO framework (F^{RHO}), fixed speed (F^{fix}), single-shot optimization (F^{init}), and optimization with perfect information (F^{actual}). The results demonstrate substantial fuel savings achieved by the proposed RHO method. Compared to the baseline of maintaining a fixed speed, the RHO method reduces ship's total fuel consumption by 9.46% on Voyage I and 12.22% on Voyage II, corresponding to savings of 33.74 mt and 44.60 mt, respectively. The aggregate fuel savings across both voyages is 10.86% relative to fixed speed operation, translating to 78.34 mt of fuel saved over approximately 21 days of sailing. When compared to single-shot optimization, the RHO method achieves even more impressive gains: 13.50% reduction on Voyage I and 0.60% on Voyage II. The aggregate improvement of 7.53% relative to static optimization underscores the value of incorporating dynamically updated weather forecasts.

Table 5.6: Comparative performance of speed optimization schemes on validation voyages

Voyage	F^{RHO} (mt)	F^{fix} (mt)	F^{init} (mt)	F^{actual} (mt)	η_{fix} (%)	η_{init} (%)	δ_{perfect} (%)
Voyage I	323.01	356.75	373.43	310.29	−9.46	−13.50	4.10
Voyage II	320.08	364.68	322.00	311.22	−12.22	−0.60	2.85

The metric δ_{perfect} measures how closely the RHO method approaches the theoretical optimum achievable with perfect meteorological foresight. For Voyage I, $\delta_{\text{perfect}} = 4.10\%$ indicates that the total fuel consumption achieved by the RHO method is only 4.10% higher than the oracle baseline, demonstrating effective utilization of weather forecasts. For Voyage II, $\delta_{\text{perfect}} = 2.85\%$ shows a smaller gap. These values confirm that the RHO method successfully leverages regularly updated weather forecast information to achieve near-optimal performance.

Overall, the RHO framework consistently outperforms all baseline schemes, validating that adaptive re-optimization with regularly updated weather forecasts yields significant fuel savings over both static optimization and fixed speed operation. This method demonstrates the practical application of regularly updated weather forecasts for ship speed optimization.

5.5 Conclusion

This chapter extends ship speed optimization from a static optimization framework to a rolling horizon optimization framework that exploits continuously updated meteorological forecasts throughout voyages. The rolling horizon optimization approach decomposes the dynamic problem into a sequence of tractable subproblems solved at each forecast update point, with the time-distance network formulation enabling efficient solutions through shortest path algorithms. Validation on two laden voyages demonstrates substantial fuel savings: 10.86% relative to fixed speed operation and 7.53% relative to single-shot optimization using only initial weather forecasts. These results quantify the value of incorporating weather forecast evolution, with the proposed method approaching within 4.10% to 2.85% of the theoretical optimum achievable with perfect foresight.

Chapters 4 and 5 study the speed optimization problem for conventional diesel-powered vessels, establishing approaches for handling weather forecast uncertainty and operational dynamism, respectively. However, future maritime emission reductions require not only operational optimization but also the integration of emerging technologies such as hybrid propulsion systems. Hybrid vessels combine conventional engines with alternative power sources, introducing more complex optimization challenges that extend beyond speed planning to encompass energy management across multiple power sources. Chapter 6 develops an optimization framework for hybrid propulsion to address these complexities inherent in next-generation maritime energy architectures.

Chapter 6

Hybrid Electric Platform Supply Vessel Routing with Speed Optimization: Pathways to Cost-Effective Emission Reduction

The preceding chapters developed advanced operational optimization frameworks for conventional diesel-powered vessels, addressing the challenges of meteorological forecast uncertainty and dynamic environmental conditions. While these operational improvements are valuable for existing fleets, achieving the ambitious emission reduction targets mandated by international regulations requires fundamental technological transformation beyond operational measures alone. Hybrid electric propulsion systems represent one of the most promising technological pathways for maritime decarbonization, combining the operational flexibility of diesel engines with the environmental benefits of electric propulsion. However, this technological advancement introduces unprecedented operational complexity. The optimization problem involves a multi-dimensional coordination challenge encompassing route planning, speed optimization, charging strategies, and power allocation between diesel and electric propulsion systems. This chapter focuses on platform supply vessels (PSVs) operating in the North Sea region. These vessels are responsible for transporting cargo, equipment, and supplies between onshore supply bases and offshore installations, collectively emitting approximately one million tonnes of CO_2 annually. This represents nearly 20% of Norway's total domestic shipping emissions.

This chapter introduces the Hybrid Electric Platform Supply Vessel Routing Problem with Speed Optimization (HEPSVRP-SO). This problem encompasses the comprehensive planning of PSV operations with hybrid diesel-electric propulsion systems over a 2–3 day planning horizon, the typical duration of PSV routes, including route selection, speed optimization along sailing legs, charging decisions at equipped installations, and dynamic energy management between diesel and electric power sources. A distinguishing feature of the HEPSVRP-SO is the incorporation of fuel energy quota constraints that enforce specific emission reduction targets, requiring the optimization model to determine the minimum-cost operational strategy while satisfying predetermined environmental performance standards. We evaluate the performance of hybrid PSV operations through realistic operational scenarios derived from the North Sea operations, benchmarking against two diesel-only baselines representing different operational maturity levels: design-speed operations and speed-optimized operations.

The remainder of this chapter is organized as follows. Section 6.1 presents the detailed problem description, and 6.2 develops an MIP model based on time-space network representation. Section 6.3 describes the experimental design, including instance generation, scenario configurations, and parameter settings. Section 6.4 presents comprehensive computational results and analysis. Section 6.5 concludes the chapter by synthesizing key findings and their implications for hybrid PSV operations.

6.1 Problem Description

The Hybrid Electric Platform Supply Vessel Routing Problem with Speed Optimization (HEPSVRP-SO) addresses the operational planning of a PSV equipped with hybrid diesel-electric propulsion systems servicing offshore installations with predetermined cargo demand that must be satisfied. The hybrid propulsion system enables dynamic power allocation between diesel engines and electric motors. The PSV departs from an onshore depot with a fully charged battery, visits a given set of mandatory offshore installations with cargo demand, and returns to the depot within a specified time horizon. Among these mandatory installations, some of them may be equipped with charging infrastructure, allowing the PSV to recharge its battery packs at these locations. The PSV may also visit other charging-equipped installations without cargo demand to replenish battery power. The PSV's battery state of charge must be maintained between certain limits, typically 20%-80% of the total capacity, to prevent degradation and ensure operational safety.

Beyond the above-mentioned charging-equipped installations, the charging infrastructure also includes a proposed offshore wind farm. Each charging facility has a limited power output capacity, and the amount of charge received depends on both this power output and the charging duration. The inclusion of the wind farm in the network represents a strategic infrastructure investment whose impact on operational efficiency and emission reduction potential requires evaluation. For charging-equipped installations with cargo demand, recharging of the PSV battery takes place in two different ways: 1) automatic service charging that occurs simultaneously with cargo operations, providing battery replenishment without additional time cost; and 2) extended charging where the PSV may choose to remain at the installations beyond service requirements for additional battery replenishment. For charging-equipped installations without cargo demand (such as the offshore wind farm), the PSV might make dedicated charging stops where the entire visit duration is allocated to battery replenishment, enabling strategic detours to meet energy requirements when necessary.

The primary objective of the HEPSVRP-SO is to minimize total operational costs, encompassing fuel costs, electricity costs, and carbon emission costs along the route while ensuring that a predetermined emission reduction target is achieved through adequate electric propulsion. The optimization also incorporates speed decisions for each sailing segment, as different speeds result in varying energy consumption and travel times, creating trade-offs between operational efficiency and energy costs. To achieve emission reduction targets, a specified percentage of the PSV's energy consumption must be supplied by the battery system, effectively limiting diesel fuel usage. Therefore, solving the HEPSVRP-SO provides integrated decisions on route planning, speed selection, and power allocation ratio between diesel and electric propulsion for each sailing segment.

Figure 6.1 provides an illustrative example that demonstrates the key features of the HEPSVRP-SO. A PSV equipped with 20,000 kWh battery departs from the depot and visits installations 10, 17, 8, and 18 sequentially before returning. At installation 10, an offshore wind farm with no service requirements, the PSV performs dedicated charging, replenishing approximately 7,000 kWh at a charging rate of 10 MW. At installation 8, the PSV makes use of flexibility the offshore charging facility provides: 3,000 kWh through automatic charging during service operations, plus an additional 600 kWh from extended charging, totaling 3,600 kWh. The route exhibits speed optimization with sailing speeds varying from seven knots on diesel-dominant segments to 10 knots on battery-powered segments. The allocation of power between diesel and electric propulsion also varies across segments. Specifically, the power allocation is mixed on the initial leg from the depot, transitioning to electric-dominant propulsion after wind farm charging. Subsequently, the power allocation is mixed after charging at installation 8, and finally, it is diesel-dominant on the last leg returning to the depot. This example illustrates the integrated optimization of routing, charging, speed, and energy allocation decisions that characterize the HEPSVRP-SO.

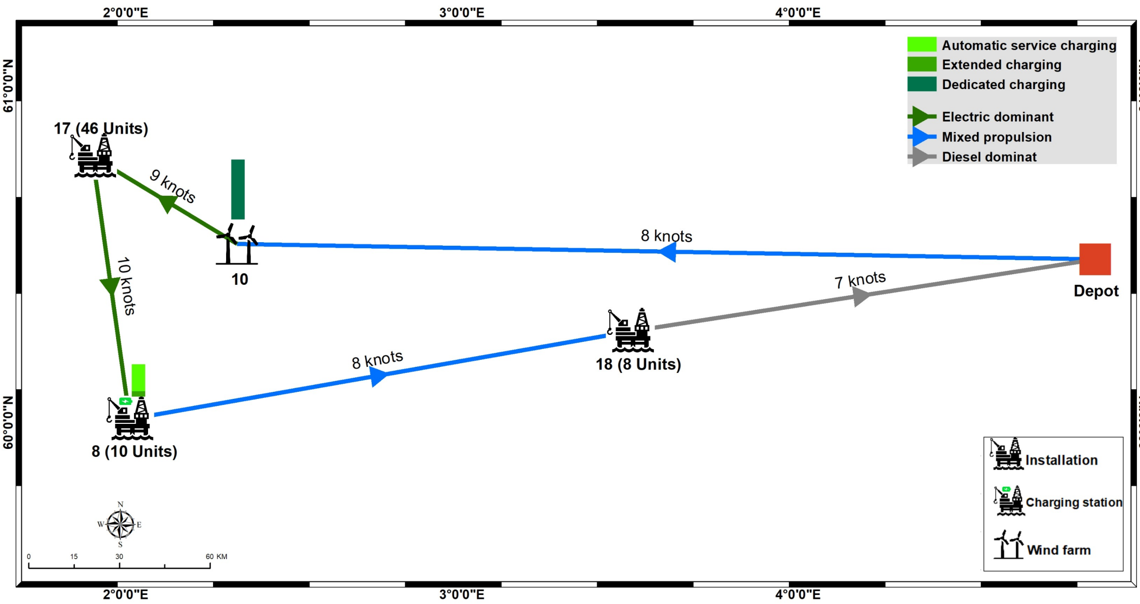

Figure 6.1: Illustrative PSV route demonstrating hybrid propulsion operations, charging strategies, and speed optimization. The route shows dedicated charging (green bar) at an offshore wind farm (installation 10) and servicing at regular installation 17. Furthermore, automatic charging (light green bar) takes place while servicing at installation 8, in addition to extended charging (medium green bar), before regular servicing at installation 18.

6.2 Modeling the HEPSVRP-SO

This section presents the mathematical formulation of the HEPSVRP-SO as an MIP model based on a time-space network representation. The formulation captures the complex interactions between routing decisions, speed optimization, energy management, and charging strategies while maintaining computational tractability. Section 6.2.1 describes the time-space network structure and arc generation procedure, Section 6.2.2 describes the mathematical notation, and Section 6.2.3 presents the complete optimization model with its objective function and constraints.

6.2.1 Time-Space Network

The time-space network construction begins by discretizing the planning horizon into uniform time periods, creating a temporal grid on which all operational events must occur. The time-space network consists of nodes and arcs, where each node is characterized by a pair (i, t) representing the completion of operations at installation or charging facility i at time t.[4] The onshore depot is represented through two special nodes: an origin node (o, t^*) representing the voyage start after vessel preparation is complete at time t^*, and a destination node (d, t), where t represents potential voyage completion times within the planning horizon.

The network contains two fundamental arc types that capture different operational decisions, i.e., service-transit arcs and battery charging arcs. Service-transit arcs represent combined sailing and service operations between different installations. A service-transit arc $((i, t), (j, t'))$ indicates that the PSV completes all operations (service and/or charging) at installation i at time t, sails to installation j, and completes its operations at installation j at time t'. For each pair of installations, multiple

[4]More precisely, (i, t) represents completion of service and automatic service charging at installations with both cargo and charging capability; charging only at charging facilities without cargo; or service only at installations without charging.

service-transit arcs may exist with different sailing speeds, enabling speed optimization within the time-space framework. Battery charging arcs are generated for installations equipped with charging facilities, connecting consecutive time points at the same installation. For installations with cargo demand, these arcs capture extended charging beyond the automatic service charging. For charging facilities without cargo demand, these arcs represent the duration of dedicated charging. A charging arc $((i,t),(i,t+1))$ indicates that the ship remains at installation or charging facility i for another time period for battery replenishment.

Figure 6.2 illustrates the concept of service-transit arcs and charging arcs in the time-space network, showing how charging decisions create temporal connections between consecutive time points at the same installation, enabling flexible energy management strategies. In this example, solid symbols represent nodes where the PSV completes operations, while dashed symbols indicate potential time-location combinations that are not utilized in this particular route. The PSV departs from the depot at $t = 0$, visits installation 1 at $t = 4$, and then proceeds to installation 2 where it remains for additional charging during $t = 5$ to $t = 7$ (shown by the red horizontal arcs). The black arrows represent service-transit arcs combining sailing and service operations, with segments from installation 3 to installation 4 demonstrating multiple speed options (4, 6, and 12 knots) over the 12 nautical mile distance, illustrating how the time-space network captures both routing decisions and speed optimization opportunities. Note that this is a simplified example, only for illustrational purposes, with a too coarse time discretization resulting in only three available speeds for this sailing leg.

For each potential service-transit arc, we evaluate feasibility based on minimum and maximum sailing speeds and service duration requirements. The energy consumption for each service-transit arc is calculated using a vessel-specific speed-power consumption function derived from the real power consumption data. The energy consumption calculation accounts for both sailing and service energy requirements. The sailing en-

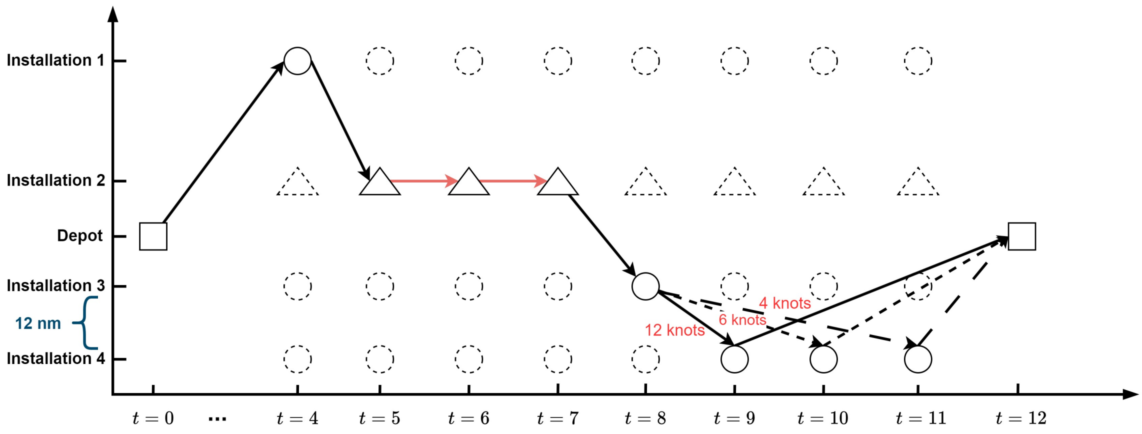

Figure 6.2: Time-space network structure with service-transit and battery charging arcs. Squares represent the depot, circles represent regular installations, and triangles represent installations with charging facilities.

ergy depends on the chosen speed, which is optimized within the feasible range to balance fuel consumption with schedule requirements. Service energy consumption remains constant for each installation, determined by the vessel's auxiliary power requirements during cargo operations. For charging arcs, the charging amount is calculated as the product of the charging rate and duration, adjusted for the energy required to keep vessel stationary and run on board systems while charging. This net charging amount determines the effective battery replenishment available from each charging decision. The detailed arc generation algorithm is presented in Appendix A.3.

6.2.2 Notation

This section presents the mathematical notation required to formulate the HEPSVRP-SO optimization model.

6.2.2.1 Sets and Indices

The mathematical model employs several sets to define the spatial and temporal structure of the problem. The basic sets include $\mathcal{N}$ representing all offshore installations, $\mathcal{N}^C$ for installations equipped with charging stations, and $\mathcal{N}^M$ for mandatory installations. The temporal structure is captured by $\mathcal{T}$, which comprises all discretized temporal points throughout the planning horizon, with $\mathcal{T}_i$ denoting possible operation completion times at each installation i. The arc set $\mathcal{A}$ contains all feasible connections in the time-space network.

The time-space network employs several specialized sets to manage the temporal relationships between installations and service operations. The set $\mathcal{T}_{ij}^S$ contains all feasible start times at installation i for arcs that connect installations i and j, enabling enumeration of all possible temporal connections between specific installation pairs. The set $\mathcal{T}_{ijt}^{SS}$ refines this further by containing feasible start times at installation i for arcs terminating at installation j at a specific time t. Similarly, $\mathcal{T}_{itj}^{SE}$ contains feasible end times at installation j for arcs originating from installation i at time t. Figure 6.3 illustrates these set relationships through a simplified example with two installations. In this example, the gray dashed lines represent all connections between i and j. Therefore, $\mathcal{T}_{ij}^S = \{5, 10, 15\}$. Furthermore, $\mathcal{T}_{ij15}^{SS} = \{5, 10\}$ is the set of feasible depature time at node i for arcs that terminate at node j at time $t = 15$, as displayed by the orange lines. Finally, $\mathcal{T}_{i15j}^{SE} = \{20, 25\}$ represents the set of feasible end time at node j for arcs that depart from node i at time $t = 15$, as shown by the blue lines. A summary of sets and indices is shown in Table 6.1.

6.2.2.2 Parameters

The parameters used in the HEPSVRP-SO model are organized into four categories as detailed in Table 6.2.

Table 6.1: Sets and indices

Symbol	Description
Indices	
o	Supply depot serving as the voyage starting point, represented as an origin node
d	Supply depot serving as the voyage endpoint, represented as a destination node
t, t'	Time point
i, j	Node indices
Sets	
$\mathcal{A}$	Set of arcs in the time-space network
$\mathcal{N}$	Set of all offshore installations
$\mathcal{N}^C$	Set of installations equipped with charging facilities
$\mathcal{N}^M$	Set of installations that must be visited
$\mathcal{T}$	Set comprising all discretized temporal points throughout the planning horizon, which is derived from the maximum allowable voyage duration (time budget)
$\mathcal{T}_i$	Set of possible operation completion time points at node i
$\mathcal{T}_{ij}^S$	Set of feasible departure time for arcs $((i,\cdot),(j,\cdot))$
$\mathcal{T}_{ijt}^{SS}$	Set of feasible departure time for arcs $((i,\cdot),(j,t))$
$\mathcal{T}_{itj}^{SE}$	Set of feasible arrival time for arcs $((i,t),(j,\cdot))$

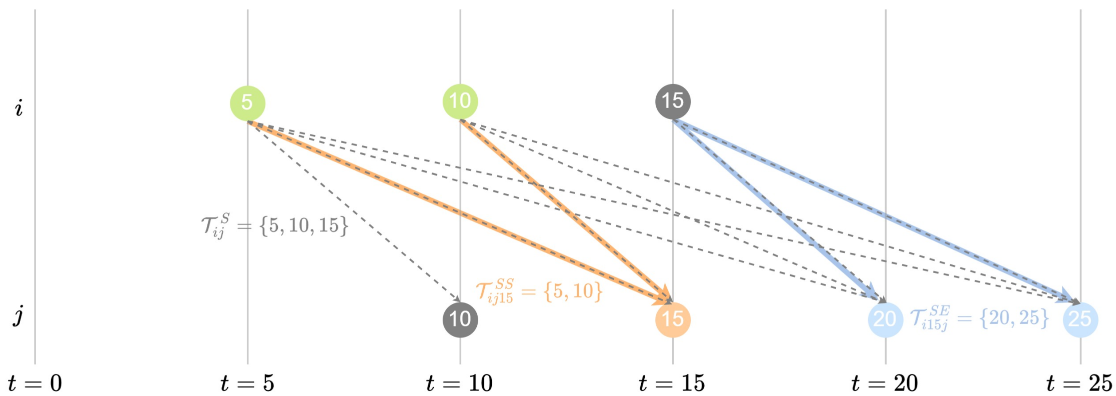

Figure 6.3: An example for temporal sets $\mathcal{T}_{ij}^S$, $\mathcal{T}_{ij15}^{SS}$, and $\mathcal{T}_{i15j}^{SE}$.

Cost parameters. The model incorporates three primary economic inputs. The costs of diesel fuel and purchased electricity are denoted by C^F and C^E, respectively. To account for environmental externalities, the cost of carbon emissions is included as C^{GHG}.

Energy consumption parameters. The model accounts for energy consumption during both sailing and charging operations. For each service-transit arc $((i,t),(j,t'))$, the energy consumption parameters $E_{itjt'}$ and $E^s_{itjt'}$ distinguish between total energy requirements and sailing-only energy requirements. The total energy $E_{itjt'}$ includes both sailing and service energy consumption, while the sailing energy $E^s_{itjt'}$ represents only the propulsion energy needed for vessel movement. This distinction is crucial for battery management at charging facilities, where the PSV only needs sufficient battery capacity to reach the next charging point rather than complete all operations. For each charging arc $((i,t),(i,t+1))$, the PSV consumes a constant auxiliary energy E_{aux} per time period to maintain its onboard system operations.

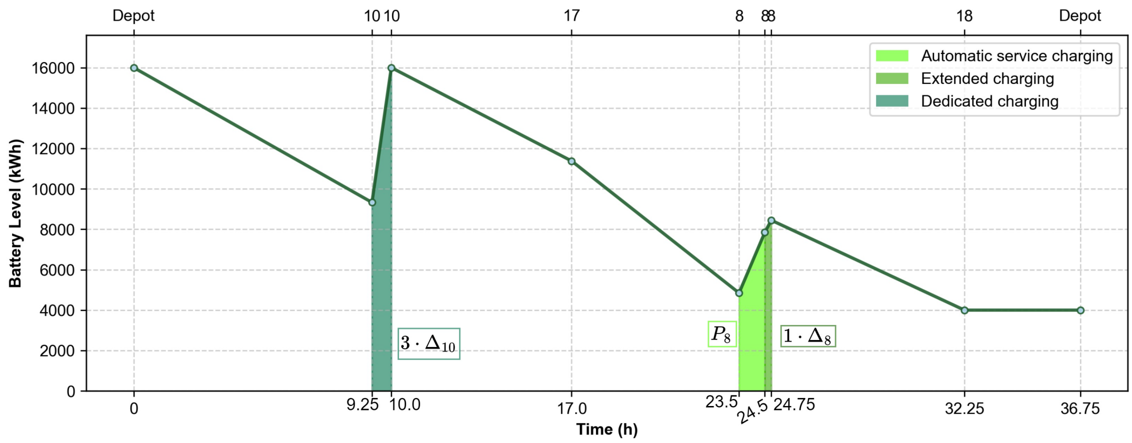

Figure 6.4: The battery level dynamics corresponding to the illustrative PSV route in Figure 6.1. The top x-axis indicates the installations visited by the PSV, while the bottom x-axis shows the corresponding time in hours. This illustration assumes a time period of 15 minutes. The figure shows the application of automatic service charging (P_8) and net charging per time period (Δ_{10}, Δ_8) at different installations.

Battery and charging parameters. The PSV's hybrid propulsion system is defined by its battery and charging characteristics. The battery's state of charge is constrained by an upper bound $\overline{Q}$ and a lower bound $\underline{Q}$, given as the shares of maximum battery capacity. At charging-equipped installations with cargo demand, two battery replenishment schemes are available. First, the PSV receives an automatic service charging at installation i, denoted by P_i, which is calculated as the product of the given charging rate and service duration (based on the requested cargo demand). Second, the PSV may opt for extended charging by staying longer after service, during which the net energy gained per time period is defined by Δ_i, after accounting for energy consumption during charging operations. In contrast, at installations without cargo demand, such as an offshore wind farm, the vessel makes dedicated charging stops. During these stops, the net energy gained per time period is also governed by Δ_i. To visually clarify the charging mechanisms, Figure 6.4 illustrates the battery level dynamics for the example route shown in Figure 6.1. As depicted, at installation 10, the PSV performs dedicated charging for three time periods, gaining a net energy of Δ_{10} in each time period. At installation 8, the PSV first receives an instantaneous automatic service charging of P_8, followed by an extended stay for one time period, gaining a net energy of Δ_8.

Table 6.2: Parameters

Symbol	Description
Cost parameters	
C^F	Cost of diesel [NOK/kWh]
C^E	Cost of electricity [NOK/kWh]
C^{GHG}	Cost of carbon emission [NOK/kg]
Energy consumption parameters	
$E_{itjt'}$	Total energy for the PSV to sail on arc $((i,t),(j,t'))$ and operate at j [kWh]
$E^s_{itjt'}$	Sailing energy for the PSV on arc $((i,t),(j,t'))$ [kWh]
E_{aux}	Auxiliary energy consumption per time period to maintain the PSV's stability during an extended charging stop [kWh]
Battery and charging parameters	
$\overline{Q}$	Upper bound of the battery capacity [kWh]
$\underline{Q}$	Lower bound of the battery capacity [kWh]
P_i	Automatic service charging amount at installation i [kWh]
Δ_i	Net charging amount per time period during an extended stay at installation i [kWh]
Model-related parameters	
E_{base}	Baseline fuel energy consumption [kWh]
α	Fuel quota factor
σ	Carbon emission factor
t^*	Time at which vessel preparation ends

Model-related Parameters. The operational boundaries of the problem are set by several key parameters. A baseline fuel energy consumption E_{base} is established from a diesel-only reference voyage, and the model's total fuel usage is then constrained to a fraction of this baseline, governed by the fuel quota factor α. In this chapter, we assume emissions are proportional to diesel energy via a fixed carbon emission factor σ, and therefore this fuel quota factor α effectively functions as an emissions quota factor.

6.2.2.3 Decision Variables

The mathematical model employs an arc-flow formulation with decision variables capturing routing decisions, energy management, and charging strategies. The binary variable $x_{itjt'}$ indicates whether arc $((i,t),(j,t'))$ is used, while u_i denotes whether installation (or charging facility) i is visited. The battery state of charge upon departure from installation i is represented by the continuous variable y_i.

Table 6.3: Decision variables

Symbol	**Description**
$x_{itjt'}$	Binary variable: 1 if arc $((i,t),(j,t'))$ is used by the ship, 0 otherwise
$x^f_{itjt'}$	Proportion of fuel usage on arc $((i,t),(j,t'))$
$x^e_{itjt'}$	Proportion of electricity usage on arc $((i,t),(j,t'))$
δ_{it}	Binary variable: 1 if the PSV charges at node i at time point t, 0 otherwise
u_i	Binary variable: 1 if node i is visited by the ship, 0 otherwise
y_i	Battery state of charge on departure from i [kWh]

The energy allocation is managed through continuous variables $x^f_{itjt'}$ and $x^e_{itjt'}$, representing the proportion of fuel and electric energy used for each arc, respectively. These variables enable the model to capture the hybrid propulsion system's flexibility, allowing optimal distribution between diesel and electric power based on operational constraints and economic objectives. The charging decision variable δ_{it} represents a key element of the optimization, enabling the PSV to remain at installation i for additional charging beyond automatic service charging during time period t. When $\delta_{it} = 1$, the PSV extends its stay at the installation i during time period t to $t+1$ for strategic battery replenishment. A summary of decision variables is presented in Table 6.3.

6.2.3 Mathematical Formulation

Based on the notation defined in Section 6.2.2, we formulate the HEPSVRP-SO as follows:

$$\min \sum_{((i,t),(j,t'))\in\mathcal{A}} \Big(E_{itjt'} \cdot C^F \cdot x^f_{itjt'} + \sigma \cdot E_{itjt'} \cdot C^{GHG} \cdot x^f_{itjt'} + E_{itjt'} \cdot C^E \cdot x^e_{itjt'} \Big) + \sum_{i\in\mathcal{N}^C} \sum_{t\in\mathcal{T}_i} \delta_{it} \cdot E_{\text{aux}} \cdot C^E \tag{6.1}$$

subject to

$$\left(\delta_{i(t-1)} + \sum_{j\in\mathcal{N}} \sum_{t'\in\mathcal{T}^{SS}_{jit}} x_{jt'it} \right) - \left(\delta_{it} + \sum_{j\in\mathcal{N}} \sum_{t'\in\mathcal{T}^{SE}_{itj}} x_{itjt'} \right) = 0, \quad i\in\mathcal{N}, t\in\mathcal{T}_i \tag{6.2}$$

$$\sum_{i\in\mathcal{N}\cup\{o\}} \sum_{t\in\mathcal{T}^S_{ij}} \sum_{t'\in\mathcal{T}^{SE}_{itj}} x_{itjt'} = u_j, \quad j\in\mathcal{N} \tag{6.3}$$

$$\sum_{j\in\mathcal{N}\cup\{d\}} \sum_{t'\in\mathcal{T}^{SE}_{ot^*j}} x_{ot^*jt'} = 1 \tag{6.4}$$

$$\sum_{i\in\mathcal{N}} \sum_{t\in\mathcal{T}^S_{id}} \sum_{t'\in\mathcal{T}^{SE}_{itd}} x_{itdt'} = 1 \tag{6.5}$$

$$x^f_{itjt'} + x^e_{itjt'} = x_{itjt'}, \quad ((i,t),(j,t')) \in \mathcal{A} \tag{6.6}$$

$$\sum_{((i,t),(j,t')) \in \mathcal{A}} E_{itjt'} \cdot x^f_{itjt'} \leq \alpha \cdot E_{\text{base}} \tag{6.7}$$

$$y_o = \overline{Q} \tag{6.8}$$

$$\begin{aligned} y_j \leq y_i - \sum_{t \in \mathcal{T}^S_{ij}} \sum_{t' \in \mathcal{T}^{SE}_{itj}} \left(x^e_{itjt'} \cdot E_{itjt'} \right) + P_j u_j + \sum_{\tau \in \mathcal{T}_j} \Delta_j \delta_{j\tau} \\ + \overline{Q} \left(1 - \sum_{t \in \mathcal{T}^S_{ij}} \sum_{t' \in \mathcal{T}^{SE}_{itj}} x_{itjt'} \right), \quad i \in \mathcal{N} \cup \{o\}, j \in \mathcal{N} \end{aligned} \tag{6.9}$$

$$\begin{aligned} y_i \geq \sum_{j \in \mathcal{N} \cup \{d\} \setminus \mathcal{N}^C} \sum_{t \in \mathcal{T}^S_{ij}} \sum_{t' \in \mathcal{T}^{SE}_{itj}} E_{itjt'} \cdot x^e_{itjt'} \\ + \sum_{j \in \mathcal{N}^C} \sum_{t \in \mathcal{T}^S_{ij}} \sum_{t' \in \mathcal{T}^{SE}_{itj}} E^s_{itjt'} \cdot x^e_{itjt'}, \quad i \in \mathcal{N} \end{aligned} \tag{6.10}$$

$$x_{itjt'} \in \{0,1\}, \quad ((i,t),(j,t')) \in \mathcal{A} \tag{6.11}$$

$$x^e_{itjt'}, x^f_{itjt'} \in [0,1], \quad ((i,t),(j,t')) \in \mathcal{A} \tag{6.12}$$

$$\delta_{it} \in \{0,1\}, \quad i \in \mathcal{N}^C, t \in \mathcal{T}_i \tag{6.13}$$

$$\underline{Q} \leq y_i \leq \overline{Q}, \quad i \in \mathcal{N} \cup \{d\} \tag{6.14}$$

$$u_i = 1, \quad i \in \mathcal{N}^M \tag{6.15}$$

$$u_i \in \{0,1\}, \quad i \in \mathcal{N} \setminus \mathcal{N}^M. \tag{6.16}$$

The objective function (6.1) minimizes total costs. The first term represents fuel costs, the second term represents carbon emission costs, the third term represents electricity costs for propulsion, and the fourth term represents electricity costs for maintaining vessel stability during charging operations. The inclusion of carbon emission costs in the objective function enables the model to capture the full economic impact of emission reduction strategies under carbon pricing policies, making the results directly applicable to policy evaluation and operator decision-making in carbon-constrained environments.

It is important to clarify the cost accounting mechanism for electricity. The cost of electricity transferred to the battery (either as automatic service charging P_i or during extended stays as Δ_i) is not accounted for at the moment of charging. Instead, this cost is realized when the stored energy is later consumed for electric propulsion on a sailing leg, which is captured by the term $E_{itjt'} \cdot C^E \cdot x^e_{itjt'}$. The final term in the objective function represents the cost of the auxiliary power required to remain stationary during these optional charging periods.

Constraints (6.2) preserve flow balance at each installation and time point, accounting for charging operations that allow the PSV to remain at the same location between consecutive time periods. Constraints (6.3) synchronize the number of visits to each installation with the service decision variables. Constraints (6.4) and (6.5) ensure that the PSV starts from the depot and returns to the depot exactly once.

Constraints (6.6) ensure that the energy supplied by diesel and electric propulsion together equals the total energy requirement for each arc when the arc is selected. Constraint (6.7) limits the total fuel energy consumption to at most α times the baseline fuel energy consumption, effectively forcing fuel reduction and increased electrification.

Constraints (6.8) ensure that the PSV departs from the depot with a fully charged battery. Constraints (6.9) update the battery state of charge when the PSV transitions between installations, accounting for energy consumption, service charging, additional charging operations, and the logical constraint that battery levels are only meaningful for visited installations.

Constraints (6.10) ensure that the PSV has sufficient battery charge to complete its next sailing leg, with different requirements for destinations with and without charging capabilities. For destinations without charging facilities, the PSV must have sufficient battery to complete the entire arc including service operations. For destinations with charging facilities, the PSV only needs sufficient battery to complete the sailing portion of the arc.

Constraints (6.11)–(6.16) specify the domains of the decision variables, including binary constraints for arc usage and charging decisions, continuous proportions for energy usage, and battery capacity bounds. Note that constraints (6.15) ensure that all mandatory installations are serviced. Therefore, as shown in Table 6.4, all instances are designed with total order sizes not exceeding the vessel's cargo capacity of 100 container units.

6.3 Experimental Design

This section presents the experimental framework developed to evaluate the performance of the proposed hybrid vessel routing model. We describe the test instances derived from real-world offshore operations, the scenario design incorporating multiple operational factors, and the implementation details of the computational study.

6.3.1 Test Instances

Our experiments draw on real-world operations from Norway's North Sea oil and gas sector. The test instances are derived from the offshore logistics network operating out of Mongstad, a major west coast supply base. This supply base ranks among Europe's most significant oil-and-gas harbors based on cargo throughput, serving as the operational hub for PSV operations to 21 offshore installations. The geographical distribution of these installations is shown in Figure 6.5. Three of the 21 installations, namely installations 3, 8, and 20, are equipped with conventional charging infrastructure (with green electricity supplied from onshore), providing charging rates of 5,000 kW, 3,000 kW and 7,000 kW, respectively. These facilities enable the PSV to recharge their batteries during regular operations. Additionally, installation 10 functions as a specialized wind farm charging station with an enhanced charging rate of 10,000 kW, representing the integration of renewable energy sources into the maritime logistics network. This installation operates exclusively as a charging facility without any cargo demand.

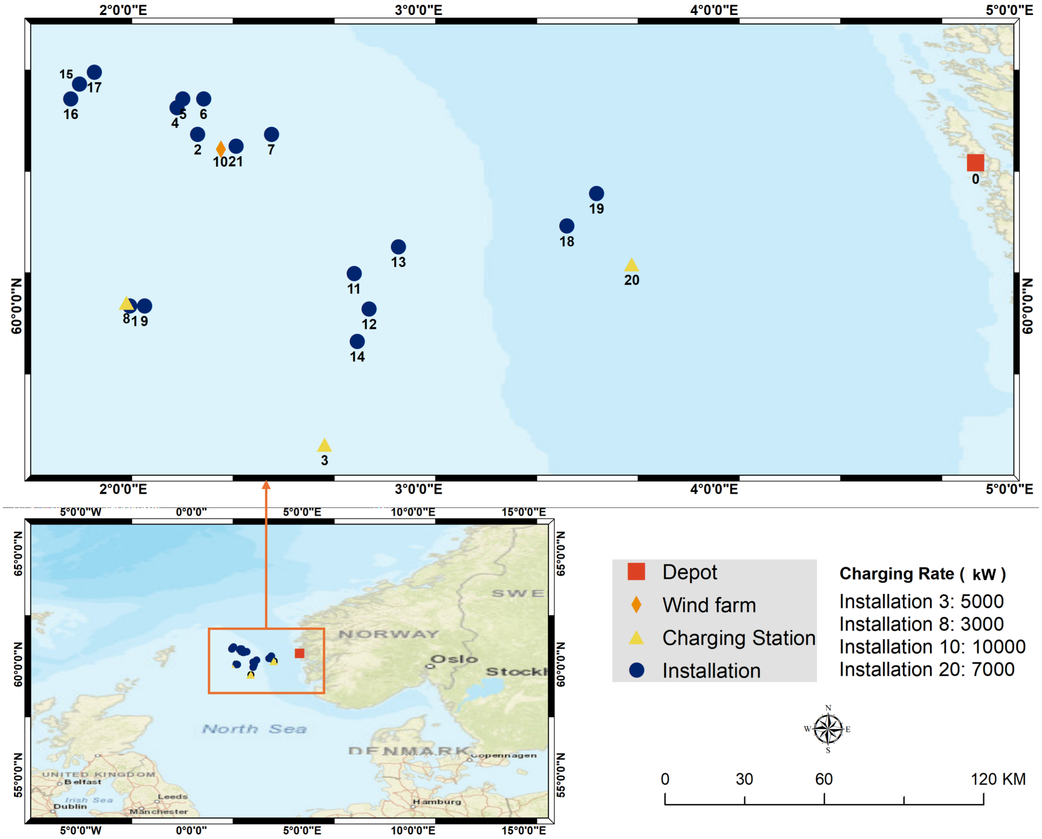

Figure 6.5: Geographic distribution of installations in the North Sea region. Different symbols represent installation types: circles for regular installations, triangles for charging stations, and diamonds for offshore wind farms. The square represents the onshore depot.

The instances are generated based on realistic offshore operations. The instances include installations with varying service requirements and charging capabilities. Notably, the installations equipped with conventional charging infrastructure can provide charging service for the PSV even if they are not mandatory service points (i.e., have no cargo demand). As shown in Table 6.4, we generate 15 representative instances covering different operational requirements and geographical configurations in the North Sea. In the second column of the table, installations that are equipped with charging capabilities (i.e., installations 3, 8, and 20) are marked in bold. Some instances include no charging facilities, such as instance 2, requiring the PSV to meet emission reduction requirements by using the initial battery charge and by making detours to offshore charging facilities for recharging. Other instances include charging facilities at selected installations, enabling charging during service operations. The "duration" column specifies the default planning horizon in days for each instance based on industry planning practices, as further explained in Section 6.3.2.

Table 6.4: Test instances

No.	Installations	Total order size	Duration (day)
1	1(25), **8**(12), 9(10), 11(16)	63	2
2	11(16), 12(18), 13(8), 14(18)	60	2
3	2(12), **3**(20), 6(21), 7(17), 21(4)	74	2
4	4(22), 5(23), 6(22)	67	2
5	11(16), 13(8), **20**(10)	34	3
6	**3**(20), 11(16), 12(18), 14(18)	72	3
7	1(25)	25	1
8	16(21), 17(42)	63	3
9	2(12), 6(22), 7(17), 18(8), 21(4)	63	3
10	18(8), 19(16)	24	2
11	2(12), **3**(20), 4(22), **20**(10), 21(4)	68	2
12	15(39), 17(46)	85	2
13	11(16), 12(18), 14(18)	52	2
14	1(25), 12(18)	43	2
15	1(24), 4(21), 5(22), 6(21), **8**(11)	99	3

Note: Numbers in parentheses indicate cargo demand in container units at each installation.

6.3.2 Scenario Design

The experimental framework employs a comprehensive scenario design to systematically evaluate the impact of various operational factors on hybrid routing solutions. Constraints (6.7) in the mathematical model limit the total fuel energy available for vessel operations, and therefore the right-hand side term $\alpha \cdot E_{\text{base}}$ can be regarded as the fuel energy (or emission) quota for the PSV. In current industry practice, where most vessels are diesel-driven, planners normally do not optimize sailing speeds but follow the vessel's design speed. Therefore, we establish a diesel-only reference where the PSV must follow its design speeds. However, we also add a second diesel-only reference that includes speed optimization, which corresponds to the current best practice. These two diesel-only references serve as benchmarks for determining fuel energy quotas in the hybrid routing model. In both references, the PSV operates only on diesel fuel without battery assistance. Under the *design-speed reference*, the PSV sails at its design speed throughout the voyage, which is implemented by filtering the time-space network to retain only arcs operating closest to the vessel's design speed. In contrast, the *speed-optimization reference* allows the PSV to adjust its speeds within the allowable speed range to minimize fuel consumption.

The temporal dimension of the hybrid routing problem is addressed through two distinct time budget settings, i.e., the default and custom time horizons. These time budgets define the planning horizon and consequently the temporal set $\mathcal{T}$ used in the time-space network construction (see Algorithm 4 in A.3 for how this affects arc generation). The *default time horizon* is derived from the current planning practice, which typically specifies voyage durations as an integer number of days, as shown in Table 6.4. Accounting for operational constraints at the onshore depot, particularly the 8-hour vessel preparation time required before departure, the actual available sailing time is calculated as 16, 40, or 64 hours for 1-, 2-, and 3-day voyages, respectively. This standardized approach reflects current industry practice, where voyage schedules are planned according to fixed daily cycles. To enable more flexible evaluation of the hybrid routing model's performance under varying operational conditions, we additionally implement a *custom time horizon* that adopts the actual voyage duration derived from the diesel-only reference. This adaptive time horizon, calculated as the ceiling of the diesel-only reference's operational duration, allows us to assess whether a hybrid vessel with a battery can achieve fuel reduction targets within stricter time constraints.

The wind farm availability represents a strategic consideration in the scenario design, examining the operational and economic implications of renewable energy integration. We investigate two configurations: 1) scenarios with wind farm availability include installation 10 as a charging option with 10,000 kW charging rate, while 2) scenarios without the wind farm restrict charging to conventional charging stations only. This comparison quantifies the value of investing in offshore renewable energy infrastructure for maritime operations.

Overall, the experimental framework systematically varies three key operational factors to comprehensively evaluate the hybrid vessel routing model's performance, including the type of fuel energy quota, the time budget setting, and wind farm availability. These factors combine to generate eight distinct scenarios, as shown in Table 6.5. $E_{\text{base}}^{\text{DS}}$ denotes the fuel energy quota obtained from the design-speed reference, while $E_{\text{base}}^{\text{SO}}$ represents the fuel energy quota obtained from the speed-optimization reference. Similarly, T^{DS} and T^{SO} denote the custom time horizons provided by the design-speed and speed-optimization references, respectively. As detailed in Table 6.5, scenarios 1–4 utilize the fuel energy quota $E_{\text{base}}^{\text{DS}}$, while scenarios 5–8 employ the fuel energy quota $E_{\text{base}}^{\text{SO}}$. Within each fuel quota group, scenarios are further distinguished by the time budget, where the custom time horizon (Default versus Custom) and wind farm availability (Yes versus No). The 15 test instances, shown in Table 6.4, are solved for each of these eight scenarios.

Table 6.5: Description of scenario design

.

Scenario	Fuel energy quota	Time budget	Wind farm availability
Scenario 1	$E^{\mathrm{DS}}_{\mathrm{base}}$	Default	Yes
Scenario 2	$E^{\mathrm{DS}}_{\mathrm{base}}$	Default	No
Scenario 3	$E^{\mathrm{DS}}_{\mathrm{base}}$	Custom (T^{DS})	Yes
Scenario 4	$E^{\mathrm{DS}}_{\mathrm{base}}$	Custom (T^{DS})	No
Scenario 5	$E^{\mathrm{SO}}_{\mathrm{base}}$	Default	Yes
Scenario 6	$E^{\mathrm{SO}}_{\mathrm{base}}$	Default	No
Scenario 7	$E^{\mathrm{SO}}_{\mathrm{base}}$	Custom (T^{SO})	Yes
Scenario 8	$E^{\mathrm{SO}}_{\mathrm{base}}$	Custom (T^{SO})	No

6.3.3 Parameter Settings

Having established the scenario framework, we now present the specific parameter values used in our computational experiments. The computational study employs realistic parameter values based on current offshore operations and emerging hybrid propulsion technologies. Cost parameters reflect current fuel and electricity costs in Norwegian offshore operations, while technical parameters are derived from representative vessel specifications and charging infrastructure capabilities. The planning horizon is discretized into 15-minute intervals, which, following the study by Ulsrud et al. (2022), ensures that departure and arrival occur at distinct time points for the vast majority of sailing legs at design speed. The fuel energy consumption limit is controlled through the parameter α, which varies from 0.1 to 0.9 with a step size of 0.1. This parameter defines the maximum allowable fuel energy consumption as a fraction of the diesel-only reference, where $r = 1 - \alpha$ represents the required fuel energy reduction rate relative to the diesel-only reference. Thus, $\alpha = 0.1$ implies a 90% fuel energy reduction, whereas $\alpha = 0.9$ corresponds to a 10% decrease. Given that emissions are proportional to diesel energy via the carbon emission factor σ, the parameter r represents both the fuel energy reduction rate and the targeted CO_2 reduction rate relative to the diesel-only baseline.

We consider a type of PSV that is representative for the operations in the North Sea, with a cargo capacity (for deck cargo) of 100 container units and an unloading rate of 10 units per hour, in accordance with the previous studies by Ulsrud et al. (2022) and Ormevik et al. (2023). Other ship-related parameters are shown in Table 6.6.

Table 6.6: Parameters for the computational experiment

Parameter	**Value**
Cost parameters	
Electricity price (C^E)	1.75 NOK/kWh
Fuel price (C^F)	0.62 NOK/kWh
Carbon emission price (C^{GHG})	1.2 NOK/kg CO_2
Model-related parameters	
Time granularity	15 minutes
Fuel quota factor	0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9
Carbon emission factor	3.2 kg CO_2/kg fuel
Ship-related parameters	
Design speed	10 knots
Minimum allowable sailing speed	7 knots
Maximum allowable sailing speed	13 knots
Unloading rate	10 container units/hour
Ship capacity	100 container units
Ship battery capacity	20,000 kWh

To establish the distance-specific energy consumption functions for the case study PSV, we fit an empirical model based on realistic power consumption data. The fitted quadratic function achieves a high coefficient of determination ($R^2 = 0.98$). The energy consumption per nautical mile (kWh/nm) as a function of vessel speed v is expressed as

$$P(v) = 3.36v^2 - 45.30v + 279.20, \tag{6.17}$$

which is subsequently used to calculate the energy consumption for each arc in the time-space network, accounting for different speed options between installations. Following Ormevik et al. (2023), we assume a constant auxiliary power requirement of 500 kW during vessel stays at installations, which contributes to the total energy consumption at each installation visit.

6.4 Results and Analysis

This section presents a comprehensive analysis of the solutions to the HEPSVRP-SO model across different instances, operational scenarios, and configuration parameters to evaluate the operational and economic performance of hybrid PSV operations. All optimization problems are solved using Gurobi 12.0.2 on a workstation with an Intel Core processor and 40 GB RAM. Solution times vary by instance complexity and emission reduction targets, ranging from 11 seconds on average for small instances (three installations or fewer plus potential charging installations) at low abatement rates ($r \leq 0.3$) to 29 minutes on average for larger instances (more than three installations) at high abatement rates ($r \geq 0.6$).

The analysis is organized into two main parts: Section 6.4.1 establishes baseline performance benchmarks and analyzes overall hybrid system performance across scenarios; and Section 6.4.2 examines key operational factors including wind farm availability and voyage time budget, which significantly influence the cost-effectiveness of hybrid routing solutions.

To facilitate systematic analysis across multiple scenarios, we establish the following notation. The set $\mathcal{Q} = \{\text{Default}, \text{Custom}\}$ denotes the set of time budgets, $\mathcal{W} = \{\text{Yes}, \text{No}\}$ denotes wind farm availability options, and $\mathcal{E} = \{E^{\text{DS}}_{\text{base}}, E^{\text{SO}}_{\text{base}}\}$ denotes the set of fuel energy quota benchmarks derived from design-speed (DS) and speed-optimization (SO) references, respectively. For instance k at abatement rate r, we define $C_k(r; e, q, w)$ as the total cost of the solution to the HEPSVRP-SO model under scenario (e, q, w), and $G_k(r; e, q, w)$ denotes the total emitted CO_2 under scenario (e, q, w), where $e \in \mathcal{E}$, $q \in \mathcal{Q}$, and $w \in \mathcal{W}$. For instance k, C_k^{DS} and C_k^{SO} denote the total cost of the design-speed and the speed-optimization references, respectively, while G_k^{DS} and G_k^{SO} represent the corresponding total emitted CO_2.

6.4.1 Comparative Performance Analysis Across Scenarios

This section compares the performance of hybrid PSV operations against conventional diesel-only benchmarks across different scenarios. First, in Section 6.4.1.1, we establish two diesel-only baselines that provide the fuel energy quotas for the HEPSVRP-SO model; Section 6.4.1.2 evaluates the cost characteristics of hybrid PSV operations relative to these baselines across all scenario configurations; and Section 6.4.1.3 analyzes the economic efficiency of emission reduction through abatement cost distributions. The analysis quantifies the emission reduction potential and cost implications of hybrid propulsion systems compared to conventional operations.

6.4.1.1 Diesel-only Baseline Comparison

To establish meaningful benchmarks for evaluating hybrid PSV operations, we first analyze the performance of conventional diesel-only operations under the two distinct speed strategies, i.e., the fixed design-speed operation and speed-optimization operation. These strategies represent the current range of industry practices and provide the foundation for defining fuel energy quotas in the hybrid PSV routing optimization.

Figure 6.6 summarizes the total cost under the two diesel-only references across all instances. For instance k, we compute the percentage saving of the total cost from the design-speed reference to the speed-optimization reference as $\frac{C_k^{\mathrm{DS}} - C_k^{\mathrm{SO}}}{C_k^{\mathrm{DS}}} \times 100\%$, which is displayed by the light-gray bar. Results show that the speed-optimization reference achieves a lower total cost than the design-speed reference across all instances, and the saving exceeds 20% for most instances. This pattern is consistent with well-known "slow-steaming" effects: by optimizing speed choices along each leg, the PSV operates more frequently at fuel-efficient speeds and eliminates unnecessary high-speed operations, thereby reducing fuel consumption and total costs. Consequently, the speed-optimization reference requires less diesel energy than the design-speed reference ($E^{\mathrm{SO}}_{\mathrm{base}} \leq E^{\mathrm{DS}}_{\mathrm{base}}$). Since these values define the fuel energy quotas for the hybrid PSV, adopting the fuel energy quota $E^{\mathrm{SO}}_{\mathrm{base}}$ creates more stringent fuel constraints than adopting $E^{\mathrm{DS}}_{\mathrm{base}}$ for the same abatement rate r.

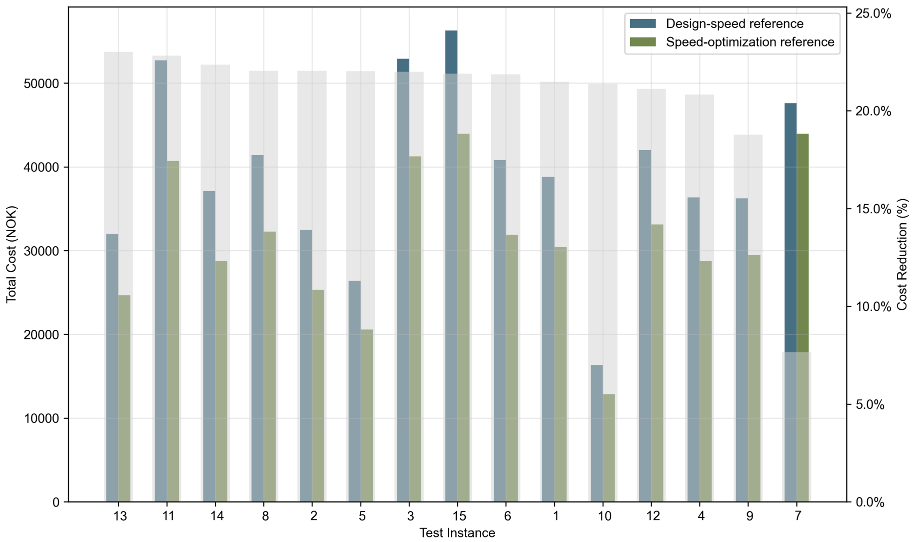

Figure 6.6: Total cost of all instances under the design-speed and speed-optimization references. The instances are sorted by the cost reduction achieved through speed optimization, as indicated by the light-gray bars.

6.4.1.2 Cost Analysis Relative to Diesel-only Baselines

To evaluate the economic performance of hybrid PSV operations across different scenarios, we analyze relative cost normalized by each instance's diesel reference. Under scenario (e, q, w), for instance k at abatement rate r, the relative cost is defined as $\frac{C_k(r;e,q,w)-C^b_k}{C^b_k} \times 100\%$, where $b \in \{\mathrm{DS}, \mathrm{SO}\}$. Figure 6.7 plots the average relative cost over the instances that are feasible at each abatement rate r, where each panel corresponds to one diesel reference and contains four curves formed by the time budget (default vs. custom) and the availability of wind farm. Across all scenarios, the average relative cost increases with abatement rate r, indicating that tighter fuel caps require more electric energy and more charging operations, raising the total cost relative to diesel references. In addition, whether the hybrid PSV operation results in a net cost saving or a penalty depends on the configuration of the fuel energy quota and time budget. The analysis first considers the scenarios without wind farm availability as a representative baseline.

When adopting the fuel energy quota $E^{\mathrm{DS}}_{\mathrm{base}}$, the use of a default time budget leads to cost savings for abatement rates $r \leq 0.7$, as the HEPSVRP-SO model enables the PSV to leverage the speed optimization strategy. These savings peak at approximately 20% for the low abatement targets ($r \leq 0.3$) and then gradually diminish as the reduction target becomes more stringent. Conversely, operating under the tighter custom time budget imposes a cost penalty, which starts at approximately 5% for the low abatement targets ($r \leq 0.3$) and escalates to over 30% for higher abatement targets. This occurs because the extra time under default time budgets enables slow-steaming and avoids costly charging detours, while tighter custom time budgets leave less room to slow down and trigger more expensive electric propulsion. When adopting fuel energy quota $E^{\mathrm{SO}}_{\mathrm{base}}$, average relative costs remain positive for both time budgets across the entire abatement range, because the speed-optimization reference already captures savings from optimized speeds, requiring any further abatement to come from electricity and charging. The magnitude of this cost penalty grows with the abatement rate: it remains within a range of 3–10% for $r \leq 0.3$; increases steadily to 15–25% for moderate abatement rates ($0.3 < r \leq 0.6$); and escalates to 35–45% for the most aggressive targets ($r > 0.6$).

The availability of the offshore wind farm offers a clear cost-saving advantage over the baseline scenarios described above, particularly for higher abatement rates ($r \geq 0.5$). Within the same time budget, curves with wind farm access lie below their no-wind-farm counterparts. The cost-saving effect induced by wind farm availability grows with the abatement rate: it remains approximately 5% for $r \leq 0.7$, and becomes around 10% for higher abatement rates. This trend reflects the increasing value of avoiding long charging detours as the vessel's fuel quota becomes more restrictive. A notable exception arises under $E^{\mathrm{DS}}_{\mathrm{base}}$ with custom time budget T^{DS}: for $r \geq 0.4$, the average relative cost with wind farm access exceeds that without wind farm. This does not mean that routes become more expensive when wind farms are available, but rather that wind farm availability makes some previously infeasible routes feasible at those high abatement targets and these newly feasible cases are costly, raising the average relative cost.

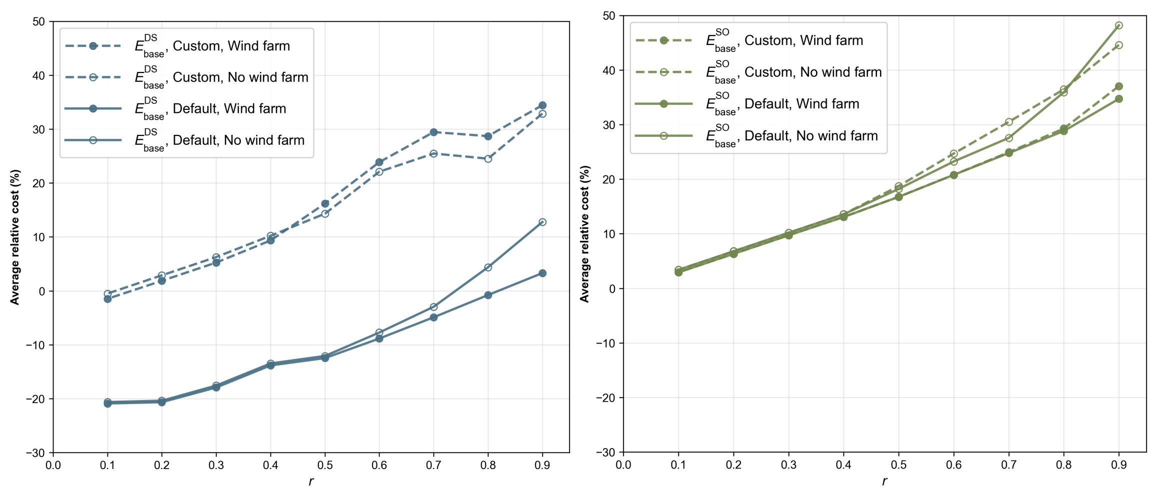

(a) (b)

Figure 6.7: Average relative cost vs abatement rate under the scenario of adopting the fuel energy quota (a)$E^{\mathrm{DS}}_{\mathrm{base}}$ and (b)$E^{\mathrm{SO}}_{\mathrm{base}}$.

6.4.1.3 Emission Abatement Cost Analysis

A critical metric for evaluating emission reduction strategies is the abatement cost, which measures the economic efficiency of achieving CO_2 reduction targets. This analysis examines the cost per unit of emission reduction across different scenario configurations, providing insights into the economic viability of various hybrid propulsion strategies and helping identify the most cost-effective approaches to meet environmental objectives. Figure 6.8 presents abatement cost distributions for each scenario configuration. For instance k at abatement rate r, the abatement cost under scenario (e, q, w) is calculated as $\frac{C_k(r;e,q,w)-C_k^b}{G_k(r;e,q,w)-G_k^b}$, where b refers to the relevant diesel-only baseline (either DS or SO). Note that the total costs include carbon emission costs, making this a policy-inclusive abatement cost that reflects the net economic impact under the assumed carbon pricing of 1.2 NOK/kg CO_2.

When using fuel energy quota $E_{\text{base}}^{\text{DS}}$ with the default time budgets, median abatement costs remain below zero up to $r \leq 0.8$ across wind farm configurations. This means that for operators currently using fixed design-speed operations, hybrid vessel operations with speed optimization can achieve substantial emission reduction targets (up to 80%) while actually reducing net operational costs due to the combined effects of fuel efficiency improvements and carbon cost savings. Under custom time budget settings, median abatement costs become positive from $r \geq 0.2$, reflecting that tighter time budgets limit slow steaming opportunities and increase reliance on purchased electricity and charging, thereby requiring additional expenditure.

When adopting fuel energy quota $E_{\text{base}}^{\text{SO}}$, median abatement costs are positive under all configurations, since speed optimization-related savings are already embedded in the diesel reference. This means that operators who have already implemented speed optimization strategies would need to invest additional capital to achieve emission reduction targets through hybrid propulsion systems and charging infrastructure. At low to moderate emission reduction rates ($r \leq 0.6$), median abatement costs remain stable at approximately 0.77 NOK/kg CO_2 across all scenario configurations. However, at higher reduction rates ($r \geq 0.7$), scenario-specific differences emerge, with wind farm availability providing marginal cost advantages. The median abatement cost ranges from 0.80 NOK/kg CO_2 (wind farm scenarios) to 0.91 NOK/kg CO_2 (no wind farm scenarios) at the highest feasible abatement rates, providing a benchmark for evaluating the economic viability of hybrid propulsion systems under current carbon pricing policies.

It should be noted that the abatement costs discussed here only include the energy costs and do not include the additional investment cost needed for the battery and charging infrastructure.

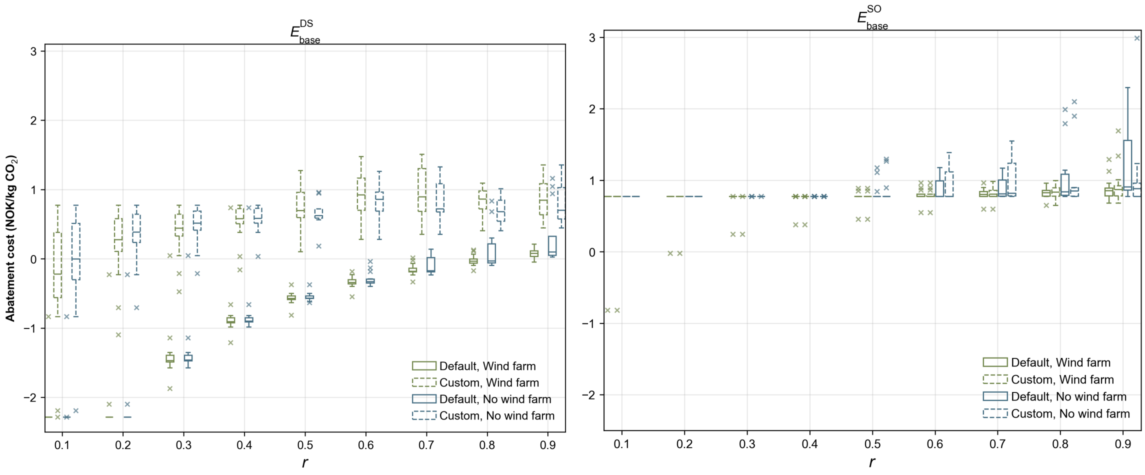

(a) (b)

Figure 6.8: Abatement cost vs abatement rate under the scenario of adopting the fuel energy quota (a) $E^{\mathrm{DS}}_{\mathrm{base}}$ and (b) $E^{\mathrm{SO}}_{\mathrm{base}}$.

6.4.2 Operational Factors Analysis

In this section, we examine how critical operational factors affect hybrid routing performance, focusing on wind farm availability and time budget configuration that significantly influence both the feasibility and cost-effectiveness of emission reduction targets. Section 6.4.2.1 quantifies the impacts of wind farm availability across different routes and operational conditions. Section 6.4.2.2 examines the cost implications of adopting different time budget configurations for emission reduction targets.

6.4.2.1 Impact of Wind Farm Availability

To explore the impacts of the wind farm on each instance, we calculate the percentage change in total cost between solutions with and without wind farm at abatement rate

r: $\frac{C_k(r;b,q,w=\text{No})-C_k(r;b,q,w=\text{Yes})}{C_(r;b,q,w=\text{No})} \times 100\%$. The wind farm has similar impacts on the instances with different fuel energy quotas and time budgets, but to different degrees. In this section, we only display the results derived from the $E^{\text{SO}}_{\text{base}}$ under default time budgets to ensure that we only analyze the effect of the use of hybrid propulsion.

As shown in Figure 6.9(a), cost reductions from wind farm availability vary significantly across instances. Some instances, such as instance 13, are geographically decoupled from the wind farm and show virtually no benefit at any abatement rate. Other instances include installations close to the wind farm and display increasing savings with higher abatement rates, consistent with the increasing value of intermediate charging as the diesel fuel constraints become more binding.

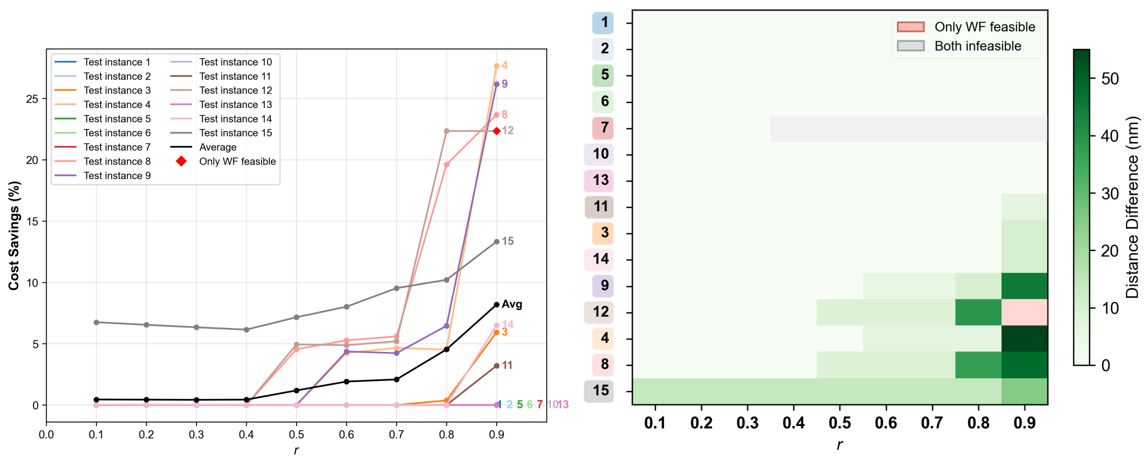

(a) (b)

Figure 6.9: Effect of wind farm under the default time budget with the fuel cap parameterized by the speed-optimization reference $E^{\text{SO}}_{\text{base}}$: (a) Relative cost difference when the wind farm is available and not available. Each curve is labeled with its instance number at the right end. Instances 1, 2, 5, 6, and 10 exhibit minimal cost savings (near 0%) across all emission reduction levels and their labels are horizontally offset at the bottom for clarity. Red diamond markers indicate scenarios where introducing the wind farm made previously infeasible cases feasible. (b) Sailing distance difference when the wind farm is available and not available. Instances are ordered by ascending sensitivity to wind farm availability. Salmon-colored cells indicate scenarios where only the wind farm configuration is feasible, while gray cells indicate scenarios where both configurations are infeasible. The background colors of the y-axis labels correspond to the instance color coding used in panel (a).

Figure 6.9(b) reports the difference in sailing distance under the same settings. The instances are ordered by their sensitivity to wind farm availability, measured as the average sailing distance difference between scenarios with and without wind farm access across all feasible emission reduction levels. Instances with minimal impact appear at the top, while those achieving substantial distance savings through wind farm charging appear at the bottom. Distance reductions concentrate in higher abatement rates ($r \geq 0.5$) and align with the instances that realize the largest cost savings in Figure 6.9(a). This evidence supports the mechanism: wind farm availability works primarily through detour avoidance. The same proximity effect expands the feasibility ratio by removing distance-induced detours that would otherwise violate time or energy limits; this feasibility gain is particularly pronounced under the custom time budget, where time slack is limited.

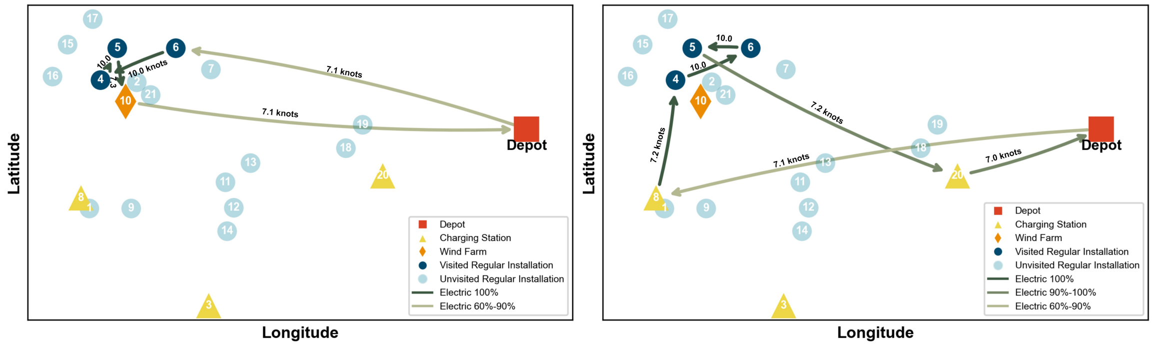

(a) (b)

Figure 6.10: Planned route of instance 4 at abatement rate $r = 0.9$ under (a) scenario 5 ($E^{\mathrm{DS}}_{\mathrm{base}}$, Default, With wind farm) and (b) scenario 6 ($E^{\mathrm{DS}}_{\mathrm{base}}$, Default, Without wind farm).

To illustrate this detour avoidance mechanism, we examine instance 4 at abatement rate $r = 0.9$ comparing scenarios with and without wind farm access. As shown in Figure 6.10(a), the service route includes mandatory installations 4, 6, and 5,[5]

[5]The routing sequence 6→4→5, rather than the geographically shorter path 6→5→4, results from time-space network discretization constraints. The arc from installation 5 to 4 requires a combined sailing time (approximately 0.15–0.29 hours at feasible speeds) and service duration (2.2 hours), totaling 2.35–2.49 hours. This duration does not align with any multiple of the 15-minute time interval, making the direct connection from installation 5 to 4 infeasible in the time-space network.

with additional visits to installations 8 and 20 solely for battery recharging, creating substantial detours. With access to the wind farm, however, the PSV uses the strategically positioned installation 10, reducing voyage time from 40.0 to 28.8 hours. Figure 6.11 shows the corresponding battery management: multiple charging sessions (4.5 hours total) without wind farm versus a single charging session (1 hour) with wind farm access.

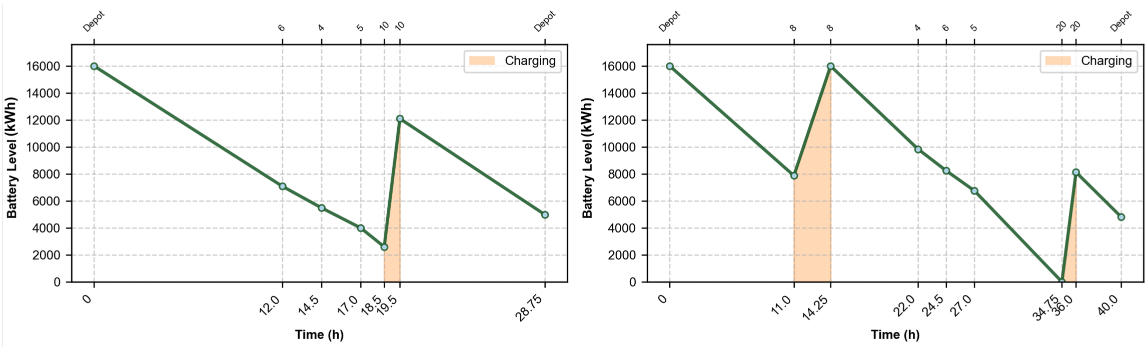

(a) (b)

Figure 6.11: Battery level of instance 4 at abatement rate $r = 0.9$ under (a) scenario 5 and (b) scenario 6.

Figure 6.12 demonstrates that both scenarios coordinate sailing speeds with energy sources: higher speeds (around 10 knots) for electric-dominant segments and lower speeds (around 7 knots) for fuel-intensive segments. The lower speeds correspond to the vessel's minimum allowable speed of 7 knots, which represents the most fuel-efficient operating point; sailing slower than this threshold would result in inefficient propulsion system operation. This explains why the hybrid PSV does not further exploit the available time budget to reduce speeds below 7 knots. Table 6.7 shows that while both scenarios achieve identical emission reductions, the routing inefficiencies without wind farm result in 29.2% higher electricity costs, accounting for the 27.7% total cost increase.

The magnitude of these savings depends on spatial proximity between the wind farm and service installations. Instances with installations located far from the wind farm show limited benefits, while those with nearby installations demonstrate substantial cost reductions, particularly at high abatement rates when charging needs intensify.

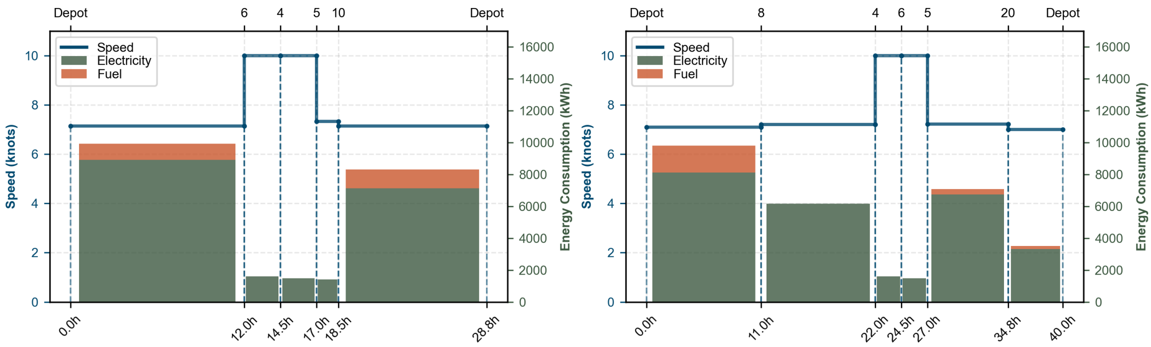

(a) (b)

Figure 6.12: Sailing speed and fuel-electric energy distribution for instance 4 (a) with wind farm access and (b) without wind farm access at $r = 0.9$. The upper x-axis shows visited installations, while the lower x-axis shows service completion times. Bars represent energy consumption for sailing segments between consecutive installations, with orange bars indicating fuel energy consumption and green bars indicating electric energy consumption.

Table 6.7: Cost breakdown for instance 4 under scenario 5 (with wind farm) and scenario 6 (without wind farm) at abatement rate $r = 0.9$

Cost breakdown (NOK)	Scenario 6	Scenario 5	Cost reduction
Fuel cost	1,365	1,365	0.0%
Electricity cost	51,949	36,785	29.2%
Carbon emission cost	1,512	1,512	0.0%
Total cost	54,826	39,662	27.7%

6.4.2.2 Impact of Time Budget Setting

Following the analysis of wind farm availability, this section examines the impact of a critical temporal factor: the voyage time budget. Industry practice often involves generous fixed schedules, similar to the *default time budget* used in this chapter, which provides significant operational slack. The extent of this slack is evident when comparing these default horizons to the tighter *custom time budgets* (T^{DS} and T^{SO}), which reflect the actual time required for diesel-only voyages, as illustrated in Figure 6.13. This analysis explores the economic implications when the hybrid PSV operates under these constrained voyage durations. To isolate the effect of time constraints, the following comparison utilizes the percentage cost increase between custom and default time budgets in scenarios without wind farm availability: $\frac{C_k(r;e,\text{Custom},w=\text{No})-C_k(r;e,\text{Default},w=\text{No})}{C_k(r;e,\text{Default},w=\text{No})}\times 100\%$, where $C_k(r;e,q,w)$ denotes the total cost for instance k at abatement rate r under scenario (e,q,w).

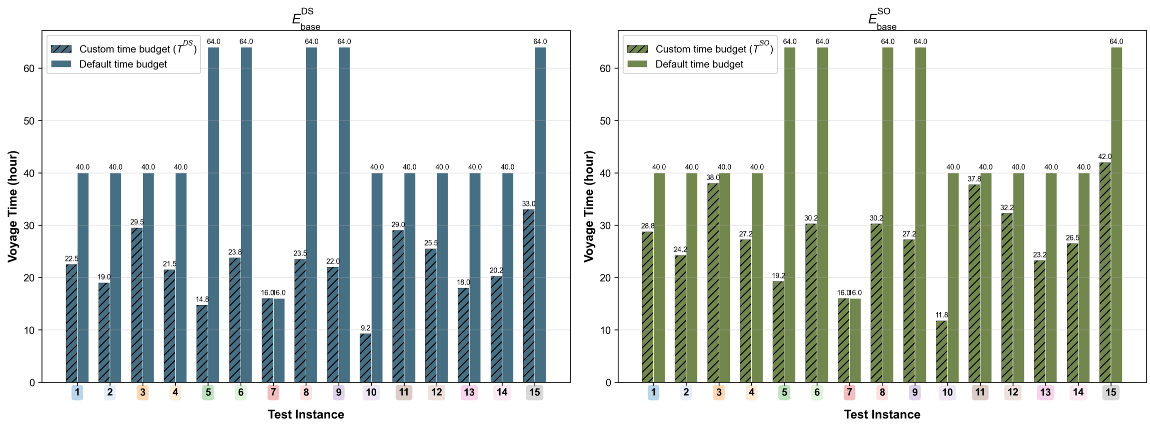

(a) (b)

Figure 6.13: Voyage durations under custom (hatched) and default time budgets when adopting the fuel energy quota (a) $E^{\text{DS}}_{\text{base}}$ and (b) $E^{\text{SO}}_{\text{base}}$.

When the fuel energy quota is derived from a design-speed reference ($E^{\text{DS}}_{\text{base}}$), operating hybrid PSVs under stricter custom time budgets (T^{DS}) leads to a significant cost increase compared to the default time budget across nearly all abatement rates. As shown in Figure 6.14(a), this cost penalty ranges from 25–40% for most instances at abatement rates $r \leq 0.5$ and can exceed 40 % at higher rates ($r \geq 0.6$). This cost penalty arises because the restrictive schedule eliminates the potential for slow-steaming. Consequently, emission reductions must be met through an increased reliance on electricity, which in turn necessitates charging. The magnitude of this penalty is directly linked to the route's dispersion and its proximity to charging facilities. For example, the mandatory stops (11, 12, 14) of instance 13 are geographically isolated from charging facilities, leading to a high initial cost penalty that rises sharply to over 45% at $r = 0.6$. In contrast, instance 10 (18, 19) exhibits a relatively low and stable cost penalty that remains under 20%. Its specific geographical configuration appears more resilient to the time pressure compared to more isolated routes. Beyond higher costs, however, a more severe consequence is operational infeasibility. Instances such as 4, 8, and 9 become infeasible beyond an abatement rate of $r = 0.3$ or $r = 0.4$. This premature infeasibility occurs because the custom time budgets for these instances are exceptionally short (as shown in Figure 6.13(a)), leaving insufficient time to complete the voyage while also performing the necessary charging to meet even moderate fuel reduction targets. Indeed, this issue becomes pervasive in the high-abatement zone, with the majority of instances failing to find a feasible solution beyond $r = 0.7$.

When the fuel energy quota is based on the speed-optimization reference ($E^{\mathrm{SO}}_{\mathrm{base}}$), the economic impact of a constrained time budget is markedly different. At abatement rates $r \leq 0.5$, operating a hybrid PSV in accordance with the custom time budget (T^{SO}) imposes no additional cost penalty compared to operating under a flexible default budget for most instances, as shown in Figure 6.14(b). This indicates that the hybrid PSV's initial battery charge is often sufficient to meet the targets for routes with moderate energy demand, resulting in no cost penalty for adhering to the efficient custom time budget (T^{SO}). Even for abatement rates approaching 0.7, the cost penalty for most instances stays below 5%, indicating that the combination of intrinsic energy demand and opportunistic charging is sufficient to meet even high abatement targets without requiring costly charging detours. However, for a few specific instances, namely instances 9 and 15, the cost penalty escalates sharply for $r \geq 0.7$, when external charging becomes a necessity. The magnitude of this penalty is dictated by the route's total energy demand and charging accessibility. For example, instance 15 involves a long voyage required to service installations 1, 4, 5, 6, and 8. Although the instance includes an opportunistic charging stop at installation 8, this single charge proves insufficient to satisfy the high energy requirements for stringent abatement targets ($r \geq 0.7$) under a tight schedule. The need for additional charging would require significant detours, explaining the sharp increase in its cost penalty before the problem becomes infeasible. In contrast, instance 5, a shorter and more compact route that also includes a charging-equipped installation 20, can fully accommodate its energy needs via this single opportunistic charge. This results in a much flatter cost increase curve, showcasing the immense value of route-integrated charging access.

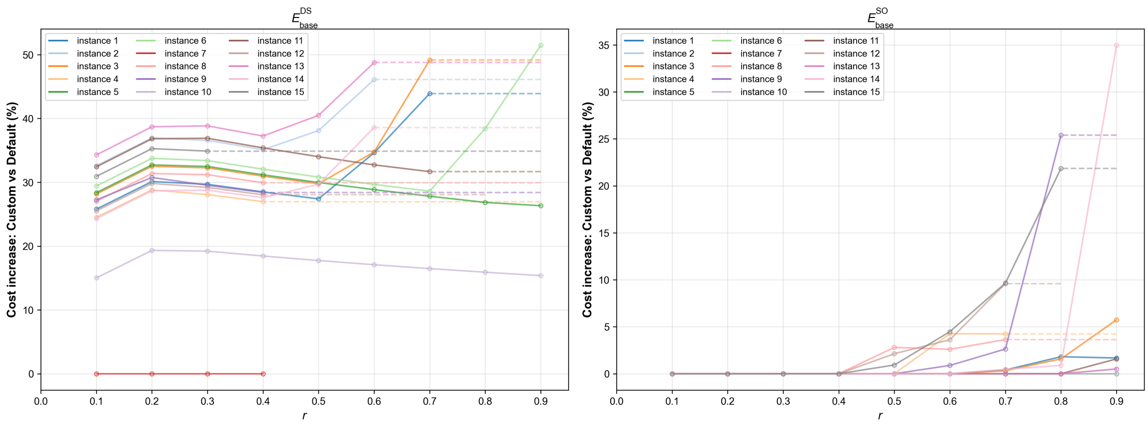

(a) (b)

Figure 6.14: Percentage cost increase of using a custom time budget versus a default time budget for (a) the design-speed reference ($E^{\text{DS}}_{\text{base}}$) and (b) the speed-optimization reference ($E^{\text{SO}}_{\text{base}}$). Solid lines with hollow circles indicate feasible solutions under both time budgets. Dashed lines represent instances where only the default time budget yields feasible solutions, and the absence of lines indicates that both time budget configurations are infeasible at those abatement rates.

These findings reveal that after operating a hybrid PSV, the choice of time budgets presents operators with distinct economic trade-offs. For operations benchmarked against a design-speed baseline, constraining the hybrid electric PSV's voyage duration to match its faster, diesel-only counterpart results in a significant cost penalty or, in many cases, makes the operation infeasible even at moderate abatement levels. Conversely, utilizing the operational flexibility within current scheduling practices to allow for longer, more efficient voyages leads to the more economically favorable outcomes. For operators already employing speed optimization, adhering to a time-constrained schedule for low to moderate emission reduction targets incurs no additional cost compared to operating under a more flexible schedule. However, for ambitious targets, a clear trade-off emerges: maintaining the same strict schedule for the hybrid PSV necessitates higher operational costs for electricity and charging.

6.5 Conclusion

This chapter develops an optimization framework for hybrid electric PSV operations under emission constraints, addressing the integrated decisions of route planning, speed selection, charging strategies, and energy management. The proposed HEPSVRP-SO model presents the unified approach to coordinate these tightly coupled decisions in offshore oil and gas logistics. Through computational experiments based on realistic North Sea operations, we demonstrate that hybrid PSV operations can achieve substantial emission reductions with economically viable abatement costs when compared against both design-speed and speed-optimization baselines. The analysis quantifies the critical role of operational factors in determining the cost-effectiveness of hybrid operations. Offshore wind farms are essential for achieving ambitious emission targets because they eliminate the need for costly detours to remote charging stations. The flexibility of voyage time budgets creates favorable trade-offs between operational efficiency and emission reduction costs. These findings demonstrate that the transition to hybrid propulsion systems requires not only technological investments but also strategic infrastructure development and operational policy adjustments to realize cost-effective decarbonization pathways in offshore supply operations.

Chapter 7

Conclusion

The global maritime industry faces mounting pressure to reduce GHG emissions while maintaining economic competitiveness under increasingly stringent environmental regulations. This thesis identifies three fundamental challenges in ship operational optimization: the uncertainty inherent in meteorological forecasts that inform voyage planning, the evolution of environmental conditions throughout voyages, and the operational complexity introduced by emerging hybrid propulsion technologies. This thesis develops methodologies that progressively advance from evaluating weather forecast uncertainty to enabling adaptive replanning and ultimately coordinating multi-dimensional decisions for hybrid propulsion systems. This chapter synthesizes the key accomplishments of this thesis and discusses directions for future investigation.

7.1 Summary of Accomplishments

(1) Innovative evaluation framework quantifying the operational value of ensemble weather forecast

This thesis develops a dual-model evaluation framework that quantifies the operational value of ensemble meteorological forecasts in ship speed optimization. The framework addresses the fundamental challenge that controlled experimental comparison between weather forecast types is impossible in real-world maritime operations, where vessels cannot simultaneously execute identical voyages under different forecasting regimes. The framework constructs two distinct fuel consumption prediction models to simulate the planning-execution gap inherent in maritime operations. An evaluator model is trained on ERA5 reanalysis data, serving as ground truth for simulating actual voyage conditions. A decision-making planner model is trained on simulated weather forecast data, mirroring the information available to operators during voyage planning.

Computational experiments on two nine-day voyages demonstrate that ensemble forecast-based optimization achieves approximately 1% additional fuel savings compared to deterministic forecast-based optimization when aggregated across both ballast and laden conditions. This framework provides shipping companies with quantitative evidence for evaluating the cost-effectiveness of different meteorological service products. It establishes a rigorous methodology for assessing forecast information value under operational uncertainty.

(2) Rolling horizon framework for ship speed optimization under evolving meteorological conditions

This thesis extends ship speed optimization from static pre-voyage planning to adaptive frameworks that exploit continuously updated meteorological forecasts throughout voyages. The rolling horizon optimization approach decomposes the multi-stage dynamic problem into a sequence of tractable deterministic subproblems solved at each forecast update point. This implements a "global planning, local execution" mechanism that adapts to evolving environmental conditions. To enable efficient solution of these subproblems, a specialized time-distance network formulation is developed. This formulation represents the speed optimization problem as a shortest path problem on a layered graph structure. Computational efficiency is ensured through systematic graph pruning techniques.

Computational experiments on two laden voyages demonstrate substantial fuel savings from the rolling horizon optimization framework. This framework achieves 10.86% fuel savings relative to fixed speed operations and 7.53% savings relative to static optimization using only initial forecasts. The method approaches within 2.85% to 4.10% of the theoretical optimum achievable with perfect foresight. These results quantify the substantial value of incorporating weather forecast evolution in ship speed optimization operations. They demonstrate that adaptive re-optimization frameworks can unlock significant efficiency gains beyond static planning approaches.

(3) Integrated optimization model for hybrid PSV operations under emission constraints

To address the operational complexity introduced by hybrid propulsion systems, this thesis develops an integrated optimization framework for PSVs operating under increasingly stringent emission reduction regulations. The framework coordinates routing, speed selection, charging strategies, and energy management decisions. The Hybrid Electric Platform Supply Vessel Routing Problem with Speed Optimization (HEPSVRP-SO) is formulated as an MIP model based on a time-space network. The model explicitly incorporates battery capacity constraints, voyage time budgets, and fuel energy quota restrictions that enforce predetermined emission reduction standards.

Computational experiments are performed based on realistic instances derived from offshore operations in the North Sea. Our analysis benchmarks hybrid PSV performance against two diesel-only references representing current industry practices. For operators already employing speed optimization, hybrid propulsion enables moderate emission reduction targets (up to 60%) with stable abatement costs of approximately 0.77 NOK/kg CO_2, rising to 0.80–0.91 NOK/kg CO_2 at the highest feasible abatement rates. Notably, our calculated abatement costs already incorporate carbon emission costs as part of total operational expenses. This reflects that operators who have already captured fuel efficiency gains through speed optimization must invest in electrification to achieve further emission reductions. For operators currently using fixed design speeds, the transition to hybrid operations presents a more compelling value proposition: the combined adoption of hybrid propulsion and speed optimization can achieve emission reductions up to 80% while simultaneously reducing total operational costs, provided that voyage schedules retain sufficient flexibility. However, it is important to recognize that these cost savings stem from various factors, such as speed optimization, electric propulsion, and operational flexibility, rather than hybrid technology alone. Under time-constrained schedules matching diesel baseline durations, the economic benefits diminish substantially or reverse into cost penalties.

The analysis further examines how key operational factors influence the performance of hybrid PSV operations. Offshore wind farm availability proves essential for achieving aggressive emission reduction targets ($r \geq 0.6$) as charging detours can be avoided. Time budget configuration significantly impacts the economics of hybrid PSV operations. When adopting tight schedules that match diesel-only voyage durations (versus flexible industry-standard schedules), operators face distinct trade-offs depending on their baseline: under design-speed fuel quotas, tight schedules incur 25–40% cost penalties at moderate abatement rates; under speed-optimization fuel quotas, tight schedules impose no additional costs for emission reduction targets $r \leq 0.4$ but require accepting higher costs for ambitious targets.

7.2 Recommended Topics for Further Research

While this thesis makes contributions to ship operational optimization, several limitations suggest promising directions for future research. This section identifies key constraints of the current methodologies and proposes extensions that could further advance the field toward more comprehensive optimization frameworks for maritime operations.

(1) Ship speed optimization under fuel consumption prediction uncertainty

The speed optimization frameworks developed in Chapters 4 and 5 follow a predict-then-optimize paradigm, where fuel consumption prediction models serve as the foundation for speed decisions. While the prediction models employed demonstrate satisfactory accuracy, prediction uncertainty can arise from multiple sources in practical applications. Limited training data availability, model complexity constraints, and inherent variability in operational conditions may all degrade prediction accuracy. More fundamentally, prediction models trained on historical data reflecting specific vessel configurations may become inadequate when vessels undergo technical upgrades such as hull modifications, propulsion system improvements, or the installation of energy-saving devices, as these modifications alter operational characteristics and introduce systematic prediction biases. Despite these potential sources of uncertainty, the current optimization frameworks do not explicitly account for fuel consumption prediction errors. Therefore, two critical aspects remain unexplored: the robustness of speed optimization outcomes to varying levels of prediction accuracy and the strategies for maintaining framework effectiveness when prediction quality degrades.

Future research should address prediction uncertainty through three complementary strategies. First, systematic sensitivity analysis should quantify optimization robustness by training prediction models with varying complexity levels or data quantities, producing a range of prediction accuracies, then evaluating how fuel savings and operational decisions degrade as prediction quality decreases. This would establish minimum accuracy thresholds required for reliable deployment. Second, building on these insights, robust optimization formulations could hedge against prediction uncertainty by defining uncertainty sets capturing plausible errors and optimizing solutions that remain feasible and near-optimal across all scenarios, or through chance-constrained programming ensuring constraint satisfaction with high probability. Third, to maintain prediction accuracy when vessel configurations evolve, adaptive modeling techniques such as transfer learning could efficiently update models using limited new operational data, online learning frameworks could continuously incorporate incoming observations, or physics-informed machine learning could embed hydrodynamic principles to preserve validity across technical modifications. These integrated developments would ensure that data-driven optimization frameworks remain both operationally reliable in the presence of prediction errors and sustainable over long-term deployment as vessel technologies evolve.

(2) Integration of ensemble weather forecasts into rolling horizon optimization frameworks

This thesis separately addresses two challenges: evaluating ensemble weather forecasts and optimizing ship's sailing speeds under changing meteorological conditions. Chapter 4 develops an evaluation framework for comparing different forecast types in ship speed optimization, while Chapter 5 develops a rolling horizon framework for ship speed optimization. These two approaches remain decoupled to isolate their respective contributions. Chapter 5 employs deterministic optimization at each decision point to quantify the value of adaptive speed planning without confounding effects from simultaneously changing both the optimization paradigm and the forecast input type. While this choice enables rigorous attribution of performance gains to the rolling horizon mechanism itself, it leaves unexploited the full probabilistic distribution information available in ensemble weather forecasts. Future research should develop stochastic rolling horizon optimization frameworks that explicitly incorporate ensemble weather forecasts at each decision point. Such frameworks could employ stochastic programming or robust optimization techniques to hedge against meteorological uncertainty dynamically throughout voyages. This integration would enable vessels to balance expected performance against worst-case scenarios at each re-optimization point, potentially achieving more robust operational performance under high forecast uncertainty conditions.

(3) Extension of adaptive optimization frameworks to hybrid electric vessel

operations

The hybrid electric vessel optimization model developed in Chapter 6 assumes deterministic environmental conditions and static voyage planning determined before departure. This limitation constrains the model's ability to respond to evolving conditions encountered during actual offshore operations. Future research should extend the rolling horizon optimization methodology to hybrid propulsion systems, creating adaptive optimization frameworks that update routing, speed, charging, and energy management decisions based on evolving weather forecasts and operational conditions. Such extensions must address substantial computational challenges, as the already complex mixed-integer programming formulation would need to be solved repeatedly throughout voyages. Developing efficient decomposition algorithms or heuristic algorithms could make real-time optimization computationally tractable.

(4) Joint optimization of investment and operational decisions

The optimization models developed in this thesis focus exclusively on operational level decisions, treating vessel specifications and infrastructure configurations as fixed parameters rather than decision variables. Battery capacity, charging station locations, and charging power capabilities are predetermined inputs rather than optimization outputs. Future research should develop integrated models that jointly optimize strategic investment decisions and tactical operational plans. Such models could determine optimal battery capacity configurations by trading off capital costs against operational savings across representative operational scenarios. Similarly, decisions regarding the deployment of charging infrastructure might be improved by assessing long-term operational benefits against installation and maintenance costs. These bilevel optimization frameworks would provide comprehensive decision support spanning both strategic investment planning and operational voyage execution, enabling shipping companies to make more informed technology adoption and infrastructure investment decisions.

(5) Adaptation of hybrid electric vessel optimization to ocean shipping contexts

The hybrid electric vessel optimization framework developed in Chapter 6 focuses on North Sea offshore supply operations, which represent a specific context with relatively short voyage distances, frequent stops, and accessible renewable energy infrastructure. Ocean shipping operations present different characteristics that limit the direct transferability of quantitative findings in this thesis. Ocean-going vessels operate on significantly longer routes with infrequent port calls, creating vastly different charging opportunity structures and battery capacity requirements. Despite these differences, the methodological contributions remain applicable, as the optimization framework addresses fundamental trade-offs inherent to hybrid propulsion systems, regardless of the operational domain. Future research should systematically adapt this framework to ocean shipping contexts. A particularly critical question is determining how many offshore charging stations are needed and where they should be strategically located along major trade routes to satisfy decarbonization demands while maintaining economic viability. Such research would require integrating infrastructure investment decisions with operational optimization across representative shipping networks, considering factors such as traffic density, voyage patterns, battery technology constraints, and carbon reduction targets. This bilevel optimization approach could provide actionable guidance for policymakers and industry stakeholders on the infrastructure requirements for enabling cost-effective maritime decarbonization at scale.

7.3 Closing

This thesis develops a suite of data-driven operational optimization frameworks that provide cost-effective pathways toward sustainable shipping. The research systematically addresses three challenges: (1) it quantifies the value of ensemble weather forecasts in ship speed optimization through a dual-model evaluation framework; (2) it enables adaptation to environmental dynamism via a rolling horizon optimization approach that leverages continuously updated weather forecasts; and (3) it manages hybrid propulsion system complexity with an integrated model that co-optimizes routing, speed, charging, and energy management. As the maritime industry continues its transition toward decarbonization, the integration of advanced forecasting capabilities, adaptive decision-making frameworks, and emerging propulsion technologies will be essential for achieving both regulatory compliance and competitive advantage in global maritime transportation.

Appendices

A.1 Algorithms for Rolling Horizon Speed Optimization

This appendix presents the detailed algorithmic for the rolling horizon optimization framework and the time-distance network construction described in Chapter 5.

A.1.1 Rolling Horizon Optimization Framework

Algorithm 2 describes the complete rolling horizon optimization process. The algorithm iteratively updates the speed profile as new meteorological forecasts become available, executing only the first segment speed at each iteration.

A.1.2 Time-Distance Network Construction

Algorithm 3 details the construction of the time-distance network at each optimization iteration. The algorithm builds the node set, establishes feasible arcs, computes arc weights using the ANN model, and prunes infeasible nodes.

Algorithm 2 Rolling Horizon Optimization Framework

Input: Total voyage distance L; Maximum sailing time $T_{\max}$; Departure time T_1^{start}; Speed limits $v_{\min}, v_{\max}$; Distance discretization Δd; Forecast update cycle $T_{\text{update}} = \Delta T$
Output: Executed speed profile $\mathbf{V}_{\text{exec}} = \{v_1^1, v_2^1, v_3^1, ..., v_{|\mathcal{H}|}^1\}$
Initialize: $h \leftarrow 1$, $L_1^{\text{rem}} \leftarrow L$, $T_1^{\text{rem}} \leftarrow T_{\max}$, $\mathbf{V}_{\text{exec}} \leftarrow \emptyset$
while $L_h^{\text{rem}} > 0$ **do**
  // Step 1: Forecast Acquisition
  Obtain forecast: $\mathbf{W}_h \leftarrow \text{GetForecast}(T_h^{\text{start}})$
  Determine T_h^{nearest} as most recent forecast issue time before T_h^{start}
  Compute time offset: $\Delta t_h \leftarrow T_h^{\text{start}} - T_h^{\text{nearest}}$
  // Step 2: Determine segment structure
  Calculate number of segments: $|\mathcal{S}_h| \leftarrow \left\lfloor \frac{T_h^{\text{rem}} - (\Delta T - \Delta t)}{\Delta T} \right\rfloor + 1$
  Compute segment durations $\{\Delta\tau_s^h\}_{s=1}^{|\mathcal{S}_h|}$ using Equation (5.1)
  // Step 3: Optimization
  Construct TDN (Algorithm 3): $\mathcal{G}_h \leftarrow \text{ConstructTDN}(L_h^{\text{rem}}, T_h^{\text{rem}}, \{\Delta\tau_s^h\}, \mathbf{W}_h)$
  Solve shortest path: $\pi_h \leftarrow \text{Dijkstra}(\mathcal{G}_h)$
  Extract speeds: $\{v_h^s\}_{s=1}^{|\mathcal{S}_h|} \leftarrow \text{ExtractSpeeds}(\pi_h, \{\Delta\tau_s^h\})$
  // Step 4: Execution and Rolling Forward
  $\mathbf{V}_{\text{exec}} \leftarrow \mathbf{V}_{\text{exec}} \cup \{v_h^1\}$
  Update: $L_{h+1}^{\text{rem}} \leftarrow L_h^{\text{rem}} - v_h^1 \times \Delta\tau_1^h$, $T_{h+1}^{\text{rem}} \leftarrow T_h^{\text{rem}} - \Delta\tau_1^h$
  $T_{h+1}^{\text{start}} \leftarrow T_h^{\text{start}} + \Delta\tau_1^h$, $h \leftarrow h + 1$
end while
return $\mathbf{V}_{\text{exec}}$

Algorithm 3 Time-Distance Network Construction

Input: Remaining distance L_h^{rem}; Remaining time T_h^{rem}; Segment durations $\{\Delta\tau_h^s\}_{s=1}^{|\mathcal{S}_h|}$; Forecast $\mathbf{W}_h$; Speed limits $v_{\min}$, $v_{\max}$; Discretization Δd; ANN model $\mathcal{F}_{\text{ANN}}(\cdot;\boldsymbol{\theta}_{\text{ANN}})$; Vessel state $\mathbf{u}$
Output: Network $\mathcal{G}_h = (\mathcal{V}_h, \mathcal{A}_h)$
// Phase 1: Node Construction
Initialize: $\mathcal{V}_h \leftarrow \emptyset$, $\tilde{t}_0 \leftarrow T_h^{\text{start}}$
for $s = 1$ to $|\mathcal{S}_h|$ **do**
$\tilde{t}_s \leftarrow \tilde{t}_{s-1} + \Delta\tau_h^s$
end for
$\mathcal{V}_h \leftarrow \mathcal{V}_h \cup \{(t_0, L_h^{\text{rem}})\}$
for $s = 1$ to $|\mathcal{S}_h| - 1$ **do**
$\text{lb}_s \leftarrow \max(0, L_h^{\text{rem}} - \sum_{i=1}^{s} v_{\max} \cdot \Delta\tau_h^i)$
$\text{ub}_s \leftarrow \max(0, L_h^{\text{rem}} - \sum_{i=1}^{s} v_{\min} \cdot \Delta\tau_h^i)$
if $\text{ub}_s - \text{lb}_s > 0$ **then**
$\mathcal{L}_h^s \leftarrow \{\text{lb}_s, \text{lb}_s + \Delta d, ..., \text{ub}_s\}$
for each $\ell \in \mathcal{L}_h^s$ **do**
$\mathcal{V}_h \leftarrow \mathcal{V}_h \cup \{(\tilde{t}_s, \ell)\}$
end for
else
$\mathcal{V}_h \leftarrow \mathcal{V}_h \cup \{(\tilde{t}_s, 0)\}$
end if
end for
$\mathcal{V}_h \leftarrow \mathcal{V}_h \cup \{(\tilde{t}_{|\mathcal{S}_h|}, 0)\}$
// Phase 2: Arc Construction
Initialize: $\mathcal{A}_h \leftarrow \emptyset$
for $s = 1$ to $|\mathcal{S}_h|$ **do**
for each $(\tilde{t}_{s-1}, \ell) \in \mathcal{V}_h$ at time $\tilde{t}_{s-1}$ **do**
for each $(\tilde{t}_s, \ell') \in \mathcal{V}_h$ at time $\tilde{t}_s$ **do**
$v \leftarrow (\ell - \ell')/\Delta\tau_s$
if $\ell' < \ell$ **and** $v_{\min} \leq v \leq v_{\max}$ **then**
Compute coordinates for ℓ using Equations (5.10)
Extract: $\mathbf{w} \leftarrow \mathbf{W}_h(s)$
if $\ell' = 0$ **and** $\ell < v_{\min} \times \Delta\tau_s$ **then**
$F \leftarrow \min_{v \in \{v_{\min}, v_{\min}+0.1, ..., v_{\max}\}} \mathcal{F}_{\text{ANN}}(v, \mathbf{w}, \mathbf{u}) \times \Delta\tau_h^s$
else
$F \leftarrow \mathcal{F}_{\text{ANN}}(v, \mathbf{w}, \mathbf{u}) \times \Delta\tau_h^s$
end if
$\mathcal{A}_h \leftarrow \mathcal{A}_h \cup \{((\tilde{t}_{s-1}, \ell), (\tilde{t}_s, \ell'), F)\}$
end if
end for
end for
end for
// Phase 3: Pruning
for $s = |\mathcal{S}_h| - 1$ down to 1 **do**
for each $(\tilde{t}_s, \ell) \in \mathcal{V}_h$ with no outgoing arcs **do**
Remove $(\tilde{t}_s, \ell)$ from $\mathcal{V}_h$ and its incoming arcs from $\mathcal{A}_h$
end for
end for
return $\mathcal{G}_h$

A.2 Detailed Voyage Information

Table A.1 and A.2 provide comprehensive day-by-day operational and meteorological data for the two validation voyages examined in Chapter 5. All meteorological data are derived from ERA5 reanalysis data and represent actual conditions experienced during the voyages. Directional measurements (wind direction, wave directions) are expressed relative to the vessel's true heading.

A.3 Time-space Network Generation Algorithm

Algorithm 4 presents the procedure for constructing the time-space network employed in the HEPSVRP-SO formulation of Chapter 6. The algorithm generates two types of arcs: transit arcs representing vessel movements between installations with speed-dependent energy consumption, and charging arcs modeling battery replenishment at charging-equipped installations. The discrete time representation with granularity Δt transforms the complex hybrid vessel routing problem into a tractable shortest path problem on a time-expanded graph.

Table A.1: Detailed operational and meteorological information for Voyage I

Date	FCR (mt/h)	Speed (kn)	Draft (m)	Sailing condition	WD (deg)	WS (m/s)	MWD (deg)	MPWW (s)	SHCWWS (m)	FSR (m)	2mT (°C)
2020/7/14	1.3478	10.13	17.5	Laden	164.2271	8.9022	7.1505	5.6554	1.5409	2.99×10^{-4}	20.3210
2020/7/15	1.3542	9.54	17.5	Laden	156.5666	7.6821	158.3993	6.3514	1.6400	3.08×10^{-4}	21.3732
2020/7/16	1.3542	9.33	17.5	Laden	138.2798	3.9703	43.9713	7.9327	1.6830	1.26×10^{-3}	20.8831
2020/7/17	1.3542	10.54	17.5	Laden	35.2410	3.0905	67.9185	6.6059	0.9274	1.74×10^{-1}	22.8000
2020/7/18	1.3478	10.96	17.5	Laden	169.7410	4.5898	3.7152	4.3399	0.7082	5.94×10^{-2}	23.4940
2020/7/19	1.3542	10.58	17.5	Laden	175.2500	5.6923	14.5434	3.9854	0.6232	5.99×10^{-5}	23.5432
2020/7/20	1.3542	10.71	17.5	Laden	178.1069	5.9768	21.2735	3.4014	0.5661	2.39×10^{-2}	24.8661
2020/7/21	1.3542	10.83	17.5	Laden	3.3383	1.1273	63.0173	3.0068	0.2739	4.32×10^{-2}	24.6442
2020/7/22	1.3478	10.87	17.5	Laden	59.4343	5.3178	169.4696	3.5963	0.6087	5.65×10^{-5}	25.9084
2020/7/23	1.3542	10.71	17.5	Laden	65.8526	6.0317	146.0129	4.0811	0.7314	5.25×10^{-3}	25.6781
2020/7/24	1.3542	10.67	17.5	Laden	52.1370	7.4369	131.3939	4.4977	1.0341	4.86×10^{-2}	24.7836
2020/7/25	1.3478	10.74	17.5	Laden	6.0543	5.9413	169.8358	4.4851	0.8542	6.11×10^{-5}	27.1862
Total sailing time: 261.0 h	Distance: 2701.56 nm										

Note: FCR = Fuel Consumption Rate; WD = Wind Direction (relative); WS = Wind Speed; MWD = Mean Wave Direction; MPWW = Mean Period of Wind Waves; SHCWWS = Significant Height of Combined Wind Waves and Swell; FSR = Forecast Surface Roughness; 2mT = 2-meter Temperature.

Table A.2: Detailed operational and meteorological information for Voyage II

Date	FCR (mt/h)	Speed (kn)	Draft (m)	Sailing condition	WD (deg)	WS (m/s)	MWD (deg)	MPWW (s)	SHCWWS (m)	FSR (m)	2mT (°C)
2019/2/14	1.3542	10.92	17	Laden	71.6087	7.5666	125.4316	7.1884	1.9317	2.13×10^{-2}	28.1699
2019/2/15	1.3542	10.92	17	Laden	67.4272	9.4594	114.3748	6.6539	1.9712	1.00×10^{-3}	27.9721
2019/2/16	1.3542	11.54	17	Laden	175.5482	6.9029	1.1794	5.4694	1.2431	4.02×10^{-3}	28.1760
2019/2/17	1.3542	10.04	17	Laden	169.3271	8.1297	0.1800	5.3415	1.1895	4.41×10^{-2}	27.6263
2019/2/18	1.3750	9.83	17	Laden	175.8656	7.1494	166.0797	6.5738	1.5390	1.14×10^{-4}	27.3241
2019/2/19	1.3540	10.80	17	Laden	135.4831	5.3609	95.2685	8.0560	1.7688	4.54×10^{-5}	27.6568
2019/2/20	1.3542	11.75	17	Laden	133.2608	4.0265	105.3969	8.4975	1.7343	2.46×10^{-5}	27.2466
2019/2/21	1.3542	12.08	17	Laden	103.2293	3.9204	109.9028	8.4649	1.7177	3.02×10^{-5}	27.0003
2019/2/22	1.3750	11.04	17	Laden	87.2013	6.3779	119.9324	7.6561	1.7526	8.04×10^{-5}	26.4363
2019/2/23	1.3540	10.80	17	Laden	110.2064	1.5615	149.4959	4.8875	0.6133	7.49×10^{-3}	27.0450
2019/2/24	1.3667	11.50	17	Laden	67.9050	4.6753	109.4991	3.9811	0.6849	1.26×10^{-2}	26.3614

Total sailing time: 242.0 h Distance: 2645.18 nm

Note: FCR = Fuel Consumption Rate; WD = Wind Direction (relative); WS = Wind Speed; MWD = Mean Wave Direction; MPWW = Mean Period of Wind Waves; SHCWWS = Significant Height of Combined Wind Waves and Swell; FSR = Forecast Surface Roughness; 2mT = 2-meter Temperature.

Algorithm 4 Time-Space Network Generation

Input: Installations $\mathcal{N}$, charging stations $\mathcal{N}^C$, origin o, destination d, set of time points $\mathcal{T}$, time granularity Δt, vessel preparation end time t^*, vessel return time T_{return}, distance matrix D, speed bounds $[v_{\min}, v_{\max}]$, charging rates r_i^C for $i \in \mathcal{N}^C$, service energy E_i^{service} for $i \in \mathcal{N}$, service duration s_i for $i \in \mathcal{N}$, service power P^{aux}
Output: Arc set $\mathcal{A}$ with transit and charging arcs
Initialize empty arc list $\mathcal{A} \leftarrow \emptyset$
$\mathcal{N} \leftarrow \mathcal{N} \cup \{o, d\}$
Generate Transit Arcs:
for each node $i \in \mathcal{N}$ where $i \neq d$ **do**
  $\mathcal{T}_i \leftarrow \{t^*\}$ if $i = o$, else $\mathcal{T}_i \leftarrow \mathcal{T}$
  for each node $j \in \mathcal{N}$ where $j \neq o$ and $i \neq j$ **do**
    Get distance d_{ij} from matrix D
    for each departure time $t \in \mathcal{T}_i$ **do**
      if $i \neq o$ and $t <$ earliest arrival time at i from depot **then**
        Skip to next t
      else if $t + d_{ij}/v_{\max} > T_{\text{return}}$ **then**
        Skip to next t
      end if
      $\tau_{\min} \leftarrow d_{ij}/v_{\max}$, $\tau_{\max} \leftarrow d_{ij}/v_{\min}$ ▷ Min/max sailing durations
      $t_{\min} \leftarrow t + \tau_{\min} + s_j$, $t_{\max} \leftarrow t + \tau_{\max} + s_j$ ▷ Min/max operation completion time
      for each $t' \in \mathcal{T}$ where $t_{\min} \leq t' \leq t_{\max}$ **do**
        $\tau_{sail} \leftarrow t' - t - s_j$, $v \leftarrow d_{ij}/\tau_{sail}$ ▷ Required speed calculation
        if $v_{\min} \leq v \leq v_{\max}$ **then**
          Calculate sailing energy: $E^s_{itjt'} \leftarrow P(v) \times d_{ij}$ using Eq.(6.17)
          Calculate total energy: $E_{itjt'} \leftarrow E^s_{itjt'} + E_j^{\text{service}}$
          Add transit arc $((i, t), (j, t'))$ to $\mathcal{A}$
        end if
      end for
    end for
  end for
end for
Generate Charging Arcs:
for each charging station $i \in \mathcal{N}^C$ **do**
  Extract operation completion times at i from existing arcs in $\mathcal{A}$
  Form set of consecutive time pairs $\mathcal{P}_i$
  for each consecutive time pair $(t, t + \Delta t) \in \mathcal{P}_i$ **do**
    $\Delta_i \leftarrow (r_i^C - P^{\text{aux}}) \times \Delta t$ ▷ Net charging amount
    Add charging arc $((i, t), (i, t + \Delta t))$ to $\mathcal{A}$
  end for
end for
return Arc set $\mathcal{A}$

Bibliography

Aas, B., Gribkovskaia, I., Halskau Sr, Ø., and Shlopak, A. (2007). Routing of supply vessels to petroleum installations. *International Journal of Physical Distribution & Logistics Management*, 37(2):164–179.

Beşikçi, E. B., Arslan, O., Turan, O., and Ölçer, A. I. (2016a). An artificial neural network based decision support system for energy efficient ship operations. *Computers & Operations Research*, 66:393–401.

Beşikçi, E. B., Kececi, T., Arslan, O., and Turan, O. (2016b). An application of fuzzy-ahp to ship operational energy efficiency measures. *Ocean Engineering*, 121:392–402.

Borthen, T., Loennechen, H., Wang, X., Fagerholt, K., and Vidal, T. (2018). A genetic search-based heuristic for a fleet size and periodic routing problem with application to offshore supply planning. *EURO Journal on Transportation and Logistics*, 7(2):121–150.

Chen, F.-W. and Liu, C.-W. (2012). Estimation of the spatial rainfall distribution using inverse distance weighting (IDW) in the middle of Taiwan. *Paddy and Water Environment*, 10(3):209–222.

Chin, C. S., Tan, Y.-J., and Kumar, M. V. (2022). Study of hybrid propulsion systems for lower emissions and fuel saving on merchant ship during voyage. *Journal of Marine Science and Engineering*, 10(3):393.

Cruz, R., Mendes, A. B., and Bahiense, L. (2023). Schedule robustness in the periodic supply vessels planning problem with stochastic demand and travel time. *International Transactions in Operational Research*.

Cuesta, E. F., Andersson, H., Fagerholt, K., and Laporte, G. (2017). Vessel routing with pickups and deliveries: an application to the supply of offshore oil platforms. *Computers & Operations Research*, 79:140–147.

de Bittencourt, G. C., Chagas, R. D. S., Silva, V. A., Vianna, I. G. P., Longhi, R. P., Ribas, P. C., and Ferreira Filho, V. J. M. (2021). A solution framework for the integrated problem of cargo assignment, fleet sizing, and delivery planning in offshore logistics. *Computers & Industrial Engineering*, 161:107653.

DNV GL (2016). Life cycle analysis of batteries in maritime sector. `https://www.nho.no/siteassets/nox-fondet/rapporter/2018/life-cycle-analysis-for-batteries-in-maritime-sector_final_v_0.1.pdf`. [Accessed 20 October 2025].

Du, Y., Chen, Y., Li, X., Schönborn, A., and Sun, Z. (2022a). Data fusion and machine learning for ship fuel efficiency modeling: Part II–voyage report data, ais data and meteorological data. *Communications in Transportation Research*, 2:100073.

Du, Y., Chen, Y., Li, X., Schönborn, A., and Sun, Z. (2022b). Data fusion and machine learning for ship fuel efficiency modeling: Part III–sensor data and meteorological data. *Communications in Transportation Research*, 2:100072.

Du, Y., Meng, Q., Wang, S., and Kuang, H. (2019). Two-phase optimal solutions for ship speed and trim optimization over a voyage using voyage report data. *Transportation Research Part B: Methodological*, 122:88–114.

ECMWF (2018). ERA5 hourly data on single levels from 1979 to present. https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels?tab=overview. [Accessed 20 October 2025].

European Commission (2025). Reducing emissions from the shipping sector. https://climate.ec.europa.eu/eu-action/transport-decarbonisation/reducing-emissions-shipping-sector_en. [Accessed 20 October 2025].

Fagerholt, K., Laporte, G., and Norstad, I. (2010). Reducing fuel emissions by optimizing speed on shipping routes. *Journal of the Operational Research Society*, 61(3):523–529.

Fagerholt, K. and Lindstad, H. (2000). Optimal policies for maintaining a supply service in the norwegian sea. *Omega*, 28(3):269–275.

Fan, A., Yang, J., Yang, L., Liu, W., and Vladimir, N. (2022). Joint optimisation for improving ship energy efficiency considering speed and trim control. *Transportation Research Part D: Transport and Environment*, 113:103527.

Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, 29(5):1189–1232.

Froger, A., Jabali, O., Mendoza, J. E., and Laporte, G. (2022). The electric vehicle routing problem with capacitated charging stations. *Transportation Science*, 56(2):460–482.

Gkerekos, C. and Lazakis, I. (2020). A novel, data-driven heuristic framework for vessel weather routing. *Ocean Engineering*, 197:106887.

Gkerekos, C., Lazakis, I., and Papageorgiou, S. (2018). Leveraging big data for fuel oil consumption modelling. In *17th Conference on Computer and IT Applications in the Maritime Industries*, pages 144–152.

Gribkovskaia, I., Laporte, G., and Shlopak, A. (2008). A tabu search heuristic for a routing problem arising in servicing of offshore oil and gas platforms. *Journal of the Operational Research Society*, 59(11):1449–1459.

Hafnia (2021). Hafnia Annual Report 2021. `https://annualreport2021.hafniabw.com/wp-content/uploads/sites/4/2022/04/Hafnia_AR2021_FINAL_WEB_compressed_290322.pdf`. [Accessed 20 October 2025].

Halvorsen-Weare, E. E., Fagerholt, K., Nonås, L. M., and Asbjørnslett, B. E. (2012). Optimal fleet composition and periodic routing of offshore supply vessels. *European Journal of Operational Research*, 223(2):508–517.

Hassoun, M. H. (1995). *Fundamentals of Artificial Neural Networks*. MIT press.

Hastie, T. (2009). *The elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer.

Hinnenthal, J. and Clauss, G. (2010). Robust pareto-optimum routing of ships utilising deterministic and ensemble weather forecasts. *Ships and Offshore Structures*, 5(2):105–114.

IMO (2011). Resolution MEPC.203(62): Inclusion of regulations on energy efficiency for ships In MARPOL Annex VI. `https://wwwcdn.imo.org/localresources/en/KnowledgeCentre/IndexofIMOResolutions/MEPCDocuments/MEPC.203(62).pdf`. [Accessed 20 October 2025].

IMO (2020). Fourth IMO GHG Study 2020. `https://www.imo.org/en/OurWork/Environment/Pages/Fourth-IMO-Greenhouse-Gas-Study-2020.aspx`. [Accessed 20 October 2025].

IMO (2021). Resolution MEPC.336(76): 2021 Guidelines on Operational Carbon Intensity Indicators and the Calculation Methods (CII Guidelines, G1). `https://wwwcdn.imo.org/localresources/en/KnowledgeCentre/`

IndexofIMOResolutions/MEPCDocuments/MEPC.336(76).pdf. [Accessed 20 October 2025].

IMO (2023). 2023 IMO Strategy on Reduction of GHG Emissions from Ships. https://wwwcdn.imo.org/localresources/en/OurWork/Environment/Documents/annex/MEPC%2080/Annex%2015.pdf. [Accessed 20 October 2025].

Inal, O. B., Charpentier, J.-F., and Deniz, C. (2022). Hybrid power and propulsion systems for ships: Current status and future challenges. *Renewable and Sustainable Energy Reviews*, 156:111965.

International Energy Agency (2023). Emissions from oil and gas operations in net zero transitions: A world energy outlook special report on the oil and gas industry and cop28. https://www.iea.org/reports/emissions-from-oil-and-gas-operations-in-net-zero-transitions. [Accessed 20 October 2025].

Jeon, M., Noh, Y., Shin, Y., Lim, O.-K., Lee, I., and Cho, D. (2018). Prediction of ship fuel consumption by using an artificial neural network. *Journal of Mechanical Science and Technology*, 32(12):5785–5796.

Jia, H., Adland, R., Prakash, V., and Smith, T. (2017). Energy efficiency with the application of virtual arrival policy. *Transportation Research Part D: Transport and Environment*, 54:50–60.

Jiang, Y., Hu, W., Gu, W., Yu, Y., and Xu, M. (2025). A multi-mode hybrid electric vehicle routing problem with time windows. *Transportation Research Part E: Logistics and Transportation Review*, 195:103976.

Kanellos, F. (2013). Optimal power management with ghg emissions limitation in all-electric ship power systems comprising energy storage systems. *IEEE Transactions on Power Systems*, 29(1):330–339.

Keskin, M. and Çatay, B. (2016). Partial recharge strategies for the electric vehicle routing problem with time windows. *Transportation Research Part C: Emerging Technologies*, 65:111–127.

Kim, J.-G., Kim, H.-J., Jun, H. B., and Kim, C.-M. (2016). Optimizing ship speed to minimize total fuel consumption with multiple time windows. *Mathematical Problems in Engineering*, 2016(1):3130291.

Kim, Y.-R., Jung, M., and Park, J.-B. (2021). Development of a fuel consumption prediction model based on machine learning using ship in-service data. *Journal of Marine Science and Engineering*, 9(2):137.

Kisialiou, Y., Gribkovskaia, I., and Laporte, G. (2018). The periodic supply vessel planning problem with flexible departure times and coupled vessels. *Computers & Operations Research*, 94:52–64.

Kisialiou, Y., Gribkovskaia, I., and Laporte, G. (2019). Supply vessel routing and scheduling under uncertain demand. *Transportation Research Part C: Emerging Technologies*, 104:305–316.

Kotzampasakis, M. (2025). Maritime emissions trading in the EU: Systematic literature review and policy assessment. *Transport Policy*, 165:28–41.

Lai, X., Wu, L., Wang, K., and Wang, F. (2022). Robust ship fleet deployment with shipping revenue management. *Transportation Research Part B: Methodological*, 161:169–196.

Lee, H., Aydin, N., Choi, Y., Lekhavat, S., and Irani, Z. (2018). A decision support system for vessel speed decision in maritime logistics using weather archive big data. *Computers & Operations Research*, 98:330–342.

Li, M., Xie, C., Li, X., Karoonsoontawong, A., and Ge, Y.-E. (2022a). Robust liner ship routing and scheduling schemes under uncertain weather and ocean conditions. *Transportation Research Part C: Emerging Technologies*, 137:103593.

Li, X., Du, Y., Chen, Y., Nguyen, S., Zhang, W., Schönborn, A., and Sun, Z. (2022b). Data fusion and machine learning for ship fuel efficiency modeling: Part I–voyage report data and meteorological data. *Communications in Transportation Research*, 2:100074.

Li, X., Sun, B., Guo, C., Du, W., and Li, Y. (2020). Speed optimization of a container ship on a given route considering voluntary speed loss and emissions. *Applied Ocean Research*, 94:101995.

Li, X., Sun, B., Zhao, Q., Li, Y., Shen, Z., Du, W., and Xu, N. (2018). Model of speed optimization of oil tanker with irregular winds and waves for given route. *Ocean Engineering*, 164:628–639.

Lin, Y.-H., Fang, M.-C., and Yeung, R. W. (2013). The optimization of ship weather-routing algorithm based on the composite influence of multi-dynamic elements. *Applied Ocean Research*, 43:184–194.

Lorenz, E. N. (1965). A study of the predictability of a 28-variable atmospheric model. *Tellus*, 17(3):321–333.

Luo, X., Wang, J., Wang, X., and Guan, X. (2025). Two-stage optimization of route, speed, and energy management for hybrid energy ship under sea conditions. *iEnergy*, 4(3):174–192.

Maersk (2021). Annual Report 2021. `https://ml-eu.globenewswire.com/Resource/Download/9135269a-6909-4fac-a06f-11fc0b222a97`. [Accessed 20 October 2025].

Mancini, S. (2017). The hybrid vehicle routing problem. *Transportation Research Part C: Emerging Technologies*, 78:1–12.

Medina, J. R., Molines, J., González-Escrivá, J. A., and Aguilar, J. (2020). Bunker consumption of containerships considering sailing speed and wind conditions. *Transportation Research Part D: Transport and Environment*, 87:102494.

Montoya, A., Guéret, C., Mendoza, J. E., and Villegas, J. G. (2017). The electric vehicle routing problem with nonlinear charging function. *Transportation Research Part B: Methodological*, 103:87–110.

NOAA (2017). Global ensemble weather forecast. `http://www.ncei.noaa.gov/products/weather-climate-models/global-ensemble-forecast`. [Accessed 20 October 2025].

Norwegian Ministry (2019). The Government's action plan for green shipping. `www.regjeringen.no`. [Accessed 20 October 2025].

Ormevik, A. B., Fagerholt, K., Meisel, F., and Sandvik, E. (2023). A high-fidelity approach to modeling weather-dependent fuel consumption on ship routes with speed optimization. *Maritime Transport Research*, 5:100096.

Paipa-Sanabria, E., González-Montoya, D., and Coronado-Hernández, J. R. (2025). Green technologies for environmental air and water impact reduction in ships: A systematic literature review. *Journal of Marine Science and Engineering*, 13(5):839.

Peng, Y., Liu, H., Li, X., Huang, J., and Wang, W. (2020). Machine learning method for energy consumption prediction of ships in port considering green ports. *Journal of Cleaner Production*, 264:121564.

Petersen, J. P., Jacobsen, D. J., and Winther, O. (2012a). Statistical modelling for ship propulsion efficiency. *Journal of Marine Science and Technology*, 17(1):30–39.

Petersen, J. P., Winther, O., and Jacobsen, D. J. (2012b). A machine-learning approach to predict main energy consumption under realistic operational conditions. *Ship Technology Research*, 59(1):64–72.

Psaraftis, H. N. and Kontovas, C. A. (2014). Ship speed optimization: Concepts, models and combined speed-routing scenarios. *Transportation Research Part C: Emerging Technologies*, 44:52–69.

Roh, M.-I. (2013). Determination of an economical shipping route considering the effects of sea state for lower fuel consumption. *International Journal of Naval Architecture and Ocean Engineering*, 5(2):246–262.

Ryder, S. and Chappell, D. (1980). Optimal speed and ship size for the liner trades. *Maritime Policy and Management*, 7(1):55–57.

Santos, A., Fagerholt, K., Laporte, G., and Soares, C. G. (2022). A stochastic optimization approach for the supply vessel planning problem under uncertain demand. *Transportation Research Part B: Methodological*, 162:209–228.

Sayarshad, H. R. (2024). Optimization of electric charging infrastructure: integrated model for routing and charging coordination with power-aware operations. *npj Sustainable Mobility and Transport*, 1(1):4.

Schiffer, M. and Walther, G. (2017). The electric location routing problem with time windows and partial recharging. *European Journal of Operational Research*, 260(3):995–1013.

Schneider, M., Stenger, A., and Goeke, D. (2014). The electric vehicle-routing problem with time windows and recharging stations. *Transportation Science*, 48(4):500–520.

Seyfi, M., Alinaghian, M., Ghorbani, E., Çatay, B., and Sabbagh, M. S. (2022). Multi-mode hybrid electric vehicle routing problem. *Transportation Research Part E: Logistics and Transportation Review*, 166:102882.

Shyshou, A., Gribkovskaia, I., Laporte, G., and Fagerholt, K. (2012). A large neighbourhood search heuristic for a periodic supply vessel planning problem arising in offshore oil and gas operations. *INFOR: Information Systems and Operational Research*, 50(4):195–204.

Silva, V. A., Wang, A., Ferreira Filho, V. J. M., and Gounaris, C. E. (2024). Routing and scheduling of platform supply vessels in offshore oil and gas logistics. *Computers & Operations Research*, 164:106556.

Skoglund, L., Kuttenkeuler, J., Rosén, A., and Ovegård, E. (2015). A comparative study of deterministic and ensemble weather forecasts for weather routing. *Journal of Marine Science and Technology*, 20(3):429–441.

Soner, O., Akyuz, E., and Celik, M. (2018). Use of tree based methods in ship performance monitoring under operating conditions. *Ocean Engineering*, 166:302–310.

Soner, O., Akyuz, E., and Celik, M. (2019). Statistical modelling of ship operational performance monitoring problem. *Journal of Marine Science and Technology*, 24(2):543–552.

Stålhane, M., Albjerk, N., Danielsen, T., Krey, S., and Fagerholt, K. (2019). A variable neighbourhood search heuristic for disruption management in offshore oil and gas logistics. *Journal of the Operational Research Society*, 70(4):588–600.

Tarelko, W. and Rudzki, K. (2020). Applying artificial neural networks for modelling ship speed and fuel consumption. *Neural Computing and Applications*, 32(23):17379–17395.

Thomas, C., Stillman, S., Davis, P., Mahon, I., Horstmann, B., Raphael, J., Moreau, K., Teller, E., Franklin, K., and Kroodsma, D. (2025). Exposing the environmental costs of offshore oil: Greenhouse gas emissions, oil slicks, and flaring. Technical report, Copernicus Meetings.

Tzortzis, G. and Sakalis, G. (2021). A dynamic ship speed optimization method with time horizon segmentation. *Ocean Engineering*, 226:108840.

Ulsrud, K. P., Vandvik, A. H., Ormevik, A. B., Fagerholt, K., and Meisel, F. (2022). A time-dependent vessel routing problem with speed optimization. *European Journal of Operational Research*, 303(2):891–907.

UNCTAD (2020). Review of maritime transportation 2020. `https://unctad.org/system/files/official-document/rmt2020_en.pdf`. [Accessed 20 October 2025].

Upasana, Khatri, A., Steer, N., and Reisch (2025). Offshore, off-limits: Making oceans off-limits to offshore oil and gas. `https://www.ciel.org/reports/oceans-off-limits-offshore-oil-and-gas-factsheet`. [Accessed 20 October 2025].

Uyanık, T., Karatuğ, Ç., and Arslanoğlu, Y. (2020). Machine learning approach to ship fuel consumption: A case of container vessel. *Transportation Research Part D: Transport and Environment*, 84:102389.

Wang, K., Yan, X., Yuan, Y., Jiang, X., Lin, X., and Negenborn, R. R. (2018a). Dynamic optimization of ship energy efficiency considering time-varying environmental factors. *Transportation Research Part D: Transport and Environment*, 62:685–698.

Wang, S., Ji, B., Zhao, J., Liu, W., and Xu, T. (2018b). Predicting ship fuel consumption based on LASSO regression. *Transportation Research Part D: Transport and Environment*, 65:817–824.

Wang, S. and Meng, Q. (2012). Sailing speed optimization for container ships in a liner shipping network. *Transportation Research Part E: Logistics and Transportation Review*, 48(3):701–714.

Wang, S., Meng, Q., and Liu, Z. (2013). Bunker consumption optimization methods in shipping: A critical review and extensions. *Transportation Research Part E: Logistics and Transportation Review*, 53:49–62.

Wang, Y., Liang, C., Aktas, T. U., Shi, J., Pan, Y., Fang, S., and Lim, G. (2023). Joint voyage planning and onboard energy management of hybrid propulsion ships. *Journal of Marine Science and Engineering*, 11(3):585.

Weintrit, A. and Kopacz, P. (2011). A novel approach to loxodrome (rhumb line), orthodrome (great circle) and geodesic line in ecdis and navigation in general. *TransNav-International Journal on Marine Navigation and Safety of Sea Transportation*, 5(4):507–517.

Williams, E. (2013). Aviation formulary V1.47. `https://edwilliams.org/avform147.htm`. [Accessed 20 October 2025].

World Meteorological Organization (2012). Guidelines on ensemble prediction systems and forecasting. `https://library.wmo.int/viewer/48473/download?file=wmo_1091_en.pdf&type=pdf&navigator=1`. [Accessed 20 October 2025].

Yan, R., Wang, S., and Du, Y. (2020). Development of a two-stage ship fuel consumption prediction and reduction model for a dry bulk ship. *Transportation Research Part E: Logistics and Transportation Review*, 138:101930.

Yan, R., Wang, S., and Psaraftis, H. N. (2021). Data analytics for fuel consumption management in maritime transportation: Status and perspectives. *Transportation Research Part E: Logistics and Transportation Review*, 155:102489.

Yang, D., Wu, L., Wang, S., Jia, H., and Li, K. X. (2019a). How big data enriches maritime research–a critical review of automatic identification system (ais) data applications. *Transport Reviews*, 39(6):755–773.

Yang, L., Chen, G., Rytter, N. G. M., Zhao, J., and Yang, D. (2019b). A genetic algorithm-based grey-box model for ship fuel consumption prediction towards sustainable shipping. *Annals of Operations Research*, 349:525–551.

Yang, L., Chen, G., Zhao, J., and Rytter, N. G. M. (2020). Ship speed optimization considering ocean currents to enhance environmental sustainability in maritime shipping. *Sustainability*, 12(9):3649.

Yao, Z., Ng, S. H., and Lee, L. H. (2012). A study on bunker fuel management for the shipping liner services. *Computers & Operations Research*, 39(5):1160–1172.

Yousuf, M. U., Al-Bahadly, I., and Avci, E. (2019). Current perspective on the accuracy of deterministic wind speed and power forecasting. *IEEE Access*, 7:159547–159564.

Yu, H., Fang, Z., Fu, X., Liu, J., and Chen, J. (2021). Literature review on emission control-based ship voyage optimization. *Transportation Research Part D: Transport and Environment*, 93:102768.

Yuan, J. and Nian, V. (2018). Ship energy consumption prediction with gaussian process metamodel. *Energy Procedia*, 152:655–660.

Zhu, Y., Zuo, Y., and Li, T. (2021). Modeling of ship fuel consumption based on multisource and heterogeneous data: Case study of passenger ship. *Journal of Marine Science and Engineering*, 9(3):273.

Zis, T. P., Psaraftis, H. N., and Ding, L. (2020). Ship weather routing: A taxonomy and survey. *Ocean Engineering*, 213:107697.