THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Ship Surrogate Modelling and Voyage Optimisation for Short-Sea Shipping Fuel Efficiency

Daniel Vergara

Department of Mechanical Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden, 2026

Ship Surrogate Modelling and Voyage Optimisation for Short-Sea Shipping Fuel Efficiency

Daniel Vergara
ISBN: 978-91-8103-399-1
DOI: https://doi.org/10.63959/chalmers.dt/5856
Series number: 5856

Chalmers University of Technology
Department of Mechanical Engineering
Division of Marine Technology
SE-412 96, Gothenburg
Sweden
Telephone: +46 (0)31-772 1000
www.chalmers.se

Printed by Chalmers Reproservice
Gothenburg, Sweden, 2026

To the young me, who made the mistakes that brought me to where I am today.

No te arrepientas de las decisiones tomadas,
pues en tu inexperiencia fueron exploradas.
No todo fue bueno, ni todo fue malo,
fueron solo matices los altibajos.
Y si bien habrá algún día donde te arrepientas,
recuerda los buenos momentos, tente paciencia...

Le storie continuano, la vita cammina
tu diventi roscio come la pietra.
Forse un giorno capirai le mie parole,
non è mai tardi, il perdono è un raggio di sole

— *Your more experienced self*

Ship Surrogate Modelling and Voyage Optimisation for Short-Sea Shipping Fuel Efficiency

Daniel Vergara
Chalmers University of Technology
Department of Mechanical Engineering
Division of Marine Technology

Abstract

Short-sea shipping (SSS), essential for European transport logistics, is increasingly challenged by strict environmental regulations such as the EU Emissions Trading System, fuel price volatility, and the need to maintain tight schedules on short voyages. This thesis reframes SSS operations as a data-driven voyage-optimisation problem, developing an integrated framework that combines operational ship data, metocean data, machine learning (ML) models, and advanced optimisation algorithms to minimise voyage fuel consumption and emissions while meeting estimated time of arrival targets.

The framework was developed through the two interconnected fields of modelling and optimisation. (1) Modelling established surrogate models for ship performance: multiple ML regression algorithms were benchmarked for fuel consumption prediction, with XGBoost identified as the most stable and reliable. Independently, Gaussian process regression was employed to estimate added resistance in head waves for model-scale ships. When integrated into a grey-box neural network fuel model, it reduced prediction errors by a factor of 3 relative to semi-empirical methods. (2) Optimisation involved introducing a methodology to optimise the total fuel consumption of a voyage, validated across two case-study ships.

For a double-ended ferry, a complete decision-support system using Bayesian optimisation (BO) was implemented to determine the optimal power profile across an entire voyage. This achieved simulated fuel savings of up to 43%, with an 18% reduction confirmed during full-scale sea trials. For longer SSS voyages, the framework was extended to a chemical tanker case study. A metocean-aware segmentation algorithm, MS-PELT, was developed to divide routes into operationally meaningful legs, outperforming state-of-the-art methods whilst enabling near-real-time application. Voyage optimisation was subsequently performed using parallel coupled dynamic programming (PCDP), achieving up to 14.8% fuel savings. A final refinement step combining PCDP and BO achieved a potential fuel saving of 9.3% relative to measured voyage fuel consumption.

Keywords: Bayesian optimisation, decision support system, machine learning, modelling, short sea shipping, voyage optimisation, voyage simulation.

Preface

This thesis presents research conducted at the Division of Marine Technology, Department of Mechanical Engineering, Chalmers University of Technology, from June 2022 to May 2026. The research was supported by the Vinnova project 2021–02768.

Many people start this section by thanking their supervisors and examiners. I am, unfortunately, not like most people, and I will start this section by dedicating this whole work to my parents, Gladys and Leonardo. They are my world, alongside my girlfriend Elvira, my brother Alexander, my sister Solciree, and my nieces Fernanda and Sofia. Additionally, my cousin Leonardo has been an amazing support with respect to shipping. A special mention goes to Esterina, my mother from Italy, and Ornela, my other Italian mother... I love my family dearly. I have given them too many heartaches during these few years that I have been so far from them—often with little contact. Even so, they never turned their backs on me and made sure that I was healthy and well. Honestly, nothing in the world is worth more than knowing your parents are well, so enjoy their company while they are around[1].

On a similar note, the most important people in the world are my close friends, those I call an extended family. One of them deserves to be mentioned above all because, without him, I would not have continued in this programme. Kostas, our personalities and our visions of the world could not be more different[2], and yet we are able to look past all that and just be brothers, thank you for all you have done for me brother. As for the rest of my extended family, I list them in chronological order because I am using my power as the author to say so: Daniel, Julián, Luis, and Aníbal—my extended family since 2004, and those with whom I have shared the most stories. They are always there to fight with or for me... a bond stronger than diamond. Edgar, José, Jesús, Gustavo, Juan, Dina, and Fernando—by my side since 2009. I trust them with everything I have. They are also part of my extended family and I love them dearly. Sira, Martin, Vanessa, Ricardo, Raúl, Rosalía, Luis, Gabriela, Ronald and Micol—they are a very special extended family from a more mature period of my life around 2014, and friends who made my time in Rome unique; I love you guys. Abubakar and Chacho, from my time in the United Kingdom back in 2017: they are brothers who helped me mature from an immature kid who wanted an ideal world into a realist who understands that you must adapt to survive in this world[3]. Argenis and Rosalin: now we arrive at the pandemic times, 2020. They are an amazing pair who became part of my family by always being there, either for me or for my parents. Rahim, Leo, Ilias, Ana, and Christina: still from 2020, but part of my first time at Chalmers. With some of them, the bond has been broken, but

[1] Call them often; they will appreciate it.
[2] But our end vision of justice is quite similar. We just can't agree in how to reach it.
[3] Nothing is true; everything is permitted. You must always survive, no matter what.

not entirely. Especially Rahim and Leo have become part of my family. Laura, Joe, Manuela, Michelle, Patxi, Matheus, Daniela, David, and Angelika: these are from 2022 onwards, and they made my time special. My psychologist Giselle has a special mention because without her I would not have survived the stress of doing a PhD. In this group, I would also like to mention in this section Qais, Rui, and Mohammad, as they are very important to me.

Among my co-workers, the three most important are Qais, Rui, and Mohammad, which makes the most sense as we have been together for almost four years already. They became my brothers... and meeting them was definitely a reason I looked forward to the days when it was hard to get up and fight due to stress. They truly made the workplace a better place; they even made extracurricular activities better. For the rest: Stephan, Christopher, Nils, Azim, Malik, Oweis, Negin, Seyoung, Chengqian, Chi, and even Pontus. I have not interacted with some of you, but regarding Stephan, Jesus, Oweis, and Nils in particular, I have very fun memories. I know they will do great and they will preach the hegemony of steaks around the world[4]. A special co-worker who left the division far earlier that she should have to was Lucile; I know that you are better now. From the alumni, my favourites: Ioli, Mohsen, Mohammed^{sr}, Mehmet, Yuhan, and Artjoms. I had so much fun teasing them... it was lovely. From the faculty: Arash, Per, Martin, Rickard, Li, and Huadong. Please share the tobacco paste I brought you[5]. Finally, my supervisors and examiners: Xiao, Wengang, and Jonas. Xiao is the one I interacted with the most. He became as close as a friend and he helped me finish so many deadlines... Wengang, likewise, my main supervisor was a mentor and I know we have learned a lot from our shared experiences—I surely did. I admire his prowess in these subjects and his steadfastness about doing things without shortcuts in a way that is truly scientific. He is a very kind person once you get to know him. And, of course, Jonas: he got involved for the last half of my project and did so with his usual smile. I have learned enough from all of you. Thank you for trying your best with me.

These words were written with the intent of being remembered as a person, not just as an academic. Many of you knew me as the unhinged guy who loved to preach about meat and who brought joy to the team. That is the person I want to be remembered as, for my biggest achievements are my truest connections. Everything else we do in academia is just to satisfy ambition and ego. Please enjoy my thesis, and, in case I don't see ya, good afternoon, good evening, and good night!

No matter how good you were, someone was better. Live by that knowledge, and you would never grow so confident that you became sloppy.[6]

[4] Always medium rare. Beef and/or veal. Vegans/Activits/LinkedIn are our sworn enemies.
[5] It wasn't easy to get my hands on it... so Per please... don't forget about it.
[6] *Brandon Sanderson—The Emperor's Soul*

Notes

[1] *There is always another secret.* — Kaldor, [Secret Fighting, A]

Contents

List of appended papers

The author contributed substantially to all the attached papers, providing the main ideas and refining them in collaboration with the co-authors. The author managed data processing, conducted formal analysis, and/or developed algorithms to solve each research question, including clear visualisations to present the findings. The author wrote each manuscript, with final reviews and editing completed in collaboration with the co-authors.

Paper I D. Vergara, M. Alexandersson, X. Lang, and W. Mao. "Power allocation influence on energy consumption of a double-ended ferry." *In Proceedings of the 33rd International Offshore and Polar Engineering Conference*, Ottawa, Canada, 2023. ISBN 978-1-880653-80-7. ISSN 1098-6189.

Paper II C. Zhang, D. Vergara, M. Zhang, T. Nikolaos, and W. Mao. "A machine learning method to evaluate head sea induced weather impact on ship fuel consumption." *Energy*, vol. 328, 136533, 2025. DOI: https://doi.org/10.1016/j.energy.2024.136533.

Paper III D. Vergara, X. Lang, M. Zhang, M. Alexandersson, and W. Mao. "Reduced environmental impact of short sea shipping through optimal engine power allocation." *Journal of Cleaner Production*, vol. 513, 145683, 2025. DOI: https://doi.org/10.1016/j.jclepro.2024.145683.

Paper IV D. Vergara, M. Alexandersson, X. Lang, and W. Mao. "A machine learning-based Bayesian decision support system for efficient navigation of double-ended ferries." *Journal of Ocean Engineering and Science*, vol. 9, no. 6, pp. 605-615, 2024. DOI: https://doi.org/10.1016/j.joes.2023.11.002

Paper V D. Vergara, S. Fu, W. Mao, J. W. Ringsberg, and T. Nikolaos. "Forward dynamic programming-informed Bayesian optimisation for energy efficient power allocation in short-sea shipping." *Ocean Engineering*, vol. 356, Part 1, 125082, 2026. DOI: https://doi.org/10.1016/j.oceaneng.2026.125082.

Contribution to co-authored papers

Paper I

D. Vergara, M. Alexandersson, X. Lang, and W. Mao. "Power allocation influence on energy consumption of a double-ended ferry." *In Proceedings of the 33rd International Offshore and Polar Engineering Conference*, Ottawa, Canada, 2023. ISBN 978-1-880653-80-7. ISSN 1098-6189.

The author of this thesis served as the primary researcher and lead author. Responsibility for the complete processing of operational sensor data rested with the author. The author was responsible for the conceptualisation and mathematical formulation of the study's core features. Furthermore, the author implemented the predictive models, conducted the optimisation analysis, created all visualisations, and wrote the manuscript.

Paper II

C. Zhang, D. Vergara, M. Zhang, T. Nikolaos, and W. Mao. "A machine learning method to evaluate head sea induced weather impact on ship fuel consumption." *Energy*, vol. 328, 136533, 2025. DOI: https://doi.org/10.1016/j.energy.2024.136533.

The author was responsible for the data engineering pipeline, including the extraction, cleaning, and database management of the samples. The author led the model implementation and contributed to the feature selection and performance benchmarking. Additionally, the author participated in the drafting of the methodology and results sections, specifically those concerning wave resistance evaluation.

Paper III

D. Vergara, X. Lang, M. Zhang, M. Alexandersson, and W. Mao. "Reduced environmental impact of short sea shipping through optimal engine power allocation." *Journal of Cleaner Production*, vol. 513, 145683, 2025. DOI: https://doi.org/10.1016/j.jclepro.2024.145683.

The author was responsible for the data engineering of several years of operational and environmental data. The author developed the segmentation algorithms and designed the logic for integrating metocean data. The author was also responsible for the conceptualisation and implementation of the optimisation algorithms and the simulation framework. Finally, the author handled the original draft, created all visualisations, and performed the benchmarking against existing methods.

Paper IV

D. Vergara, M. Alexandersson, X. Lang, and W. Mao. "A machine learning-based Bayesian decision support system for efficient navigation of double-ended ferries." *Journal of Ocean Engineering and Science*, vol. 9, no. 6, pp. 605-615, 2024. DOI: https://doi.org/10.1016/j.joes.2023.11.002.

The author was responsible for the research design and the development of the decision support system. The author handled the data engineering and the interpolation of environmental datasets. The logic for the system was proposed and implemented by the author, who also executed the optimisation steps. While the full-scale experiments were managed by a co-author, the author was involved in the technical discussions and was responsible for processing the experimental data and evaluating the system performance.

Paper V

D. Vergara, S. Fu, W. Mao, J. W. Ringsberg, and T. Nikolaos. "Forward dynamic programming-informed Bayesian optimisation for energy efficient power allocation in short-sea shipping." *Ocean Engineering*, vol. 356, Part 1, 125082, 2026. DOI: https://doi.org/10.1016/j.oceaneng.2026.125082.

The author was responsible for the development of the optimisation framework. The author proposed the core methodological strategy and was responsible for the mathematical derivation and implementation of the constraint-handling functions. Additionally, the author conducted the performance benchmarking, created the visualisations, and wrote the manuscript.

Nomenclature

Greek notations

α	Relative angle between environment and vessel headings [°].
α_{ship}	Ship's heading [°].
$\alpha_{current}$	Sea current direction [°].
α_{wave}	Wave direction [°].
α_{wind}	Wind direction [°].
β	Direction Score (MS-PELT) [-].
γ	Peak-enhancement factor [-] or Overfitting factor (PELT).
$\mathbf{\Gamma}$	Discrete set of engine power settings [kW].
$\Gamma(\nu)$	Gamma function [-].
Γ_j	Power allocated to leg j [kW].
δ	Deviation from ETA constraint.
ϵ_t	Time corridor threshold for voyage optimisation [-].
ζ_a	Regular wave amplitude [m].
η_h	Hull efficiency [-].
η_o	Open water efficiency [-].
η_r	Relative rotative efficiency [-].
η_s	Shaft efficiency [-].
ι	Intensity Score (MS-PELT) [-].
λ	Regularisation parameter [-] or Wavelength [m].
$\mu(\xi_{k+1})$	Mean function of the GP [-].
ν	Matern kernel smoothness parameter [-].
Ξ	Scenario index for parallel Dynamic Programming [-].
ξ_i	Sampled data points in BO [-].
ρ_a	Air density $[kg/m^3]$.
ρ_w	Water density $[kg/m^3]$.
σ	Spectral width parameter [-].
σ^2	Variance parameter of the GP [-].
$\Sigma_{i,j}$	Covariance matrix [-].

τ	Changepoints (PELT) [-].
ϕ	Segment cost function (PELT) [-].
χ	Variable or search space element.
Θ	Inverse covariance matrix [-].
ω	Wave frequency [rad/s].
ω_p	Peak frequency [rad/s].

Latin notations

Adam	Adaptive Moment Estimation.
AI	Artificial Intelligence.
AIS	Automatic Identification System.
ANNs	Artificial Neural Networks.
B	Ship beam [m].
BO	Bayesian Optimisation.
$\mathcal{C}$	Cost function [-].
CatBoost	Categorical Boosting.
CFD	Computational Fluid Dynamics.
CO_2	Carbon Dioxide.
C_A	Air resistance coefficient [-].
C_B	Block coefficient [-].
C_{AW}	Non-dimensional added resistance coefficient [-].
C_F	Frictional resistance coefficient [-].
C_f	Fuel cost.
C_o	Operation cost.
C_p	Delay penalties.
DNN	Deep Neural Network.
DP	Dynamic Programming.
DSS	Decision Support System.
DWT	Deadweight Tonnage.
$d_{1\rightarrow j}$	Accumulated sailing distance between departure and leg j [km].
d_{sail}	Total sailing distance [km].
$\Delta d_{j,k}$	Haversine/Great circle distance between two waypoints [km].

$\Delta t_{j,k}$	Sailing time between two waypoints [h].
ΔT_{j+1}	Allowable time of arrival window for leg $j+1$ [h].
E	Emissions [gCO_2].
EA	Evolutionary Algorithms.
ECA	Emissions Control Area.
EI	Expected Improvement.
EMO	Evolutionary Multi-Objective Optimisation.
ETA	Estimated Time of Arrival [h].
EU ETS	European Union Emissions Trading System.
EU MRV	EU Monitoring, Reporting and Verification.
E_{max}	Maximum allowed emissions [gCO_2].
FFNNs	Feedforward Neural Networks.
FOC	Fuel Oil Consumption.
f	Generic objective function [-].
F_n	Froude number [-].
F_j	Cumulative fuel consumption up to leg j [tonne].
f_j	Fuel consumption function for leg j [tonne].
f_i	Measured quantity.
$\hat{f}_i$	Prediction estimation of f_i from χ_i.
f_{fuel}	Fuel consumption prediction model [tonne].
f_V	Speed overground prediction model [kt].
$\mathcal{F}_j$	Feasible set of states (time window) at stage j [h].
GA	Genetic Algorithm.
GBM	Grey Box Model.
GHG	Greenhouse Gas.
GP	Gaussian Process.
GPR	Gaussian Process Regressor.
g	Gravitational acceleration [m/s^2].
H_s	Significant Wave Height [m].
i	Generic index.
IMO	International Maritime Organization.
J	Generic objective function.

j	Index j, often refers to a leg.
k	Form factor [-].
K	Total number of discrete scenarios [-].
KNN	K-nearest neighbours.
K_{yy}	Radius of gyration [m].
k	Index k, often refers to a waypoint inside a leg.
LCB	Lower Confidence Bound.
l_d	GP characteristic length scale for feature d.
LOESS	Locally Estimated Scatterplot Smoothing.
LSTM	Long Short-Term Memory.
$\mathcal{L}$	Loss function [-].
L	Length between perpendiculars [m].
L_E/L	Effective entrance length ratio [-].
MC	Monte Carlo simulations.
MCR	Maximum Continuous Rating [kW].
MDO	Marine Diesel Oil.
MILP	(Mixed Integer) Linear Programming.
MLR	Multi-Linear Regression.
MPC	Model Predictive Control.
MS-PELT	Metocean Score-based Pruned Exact Linear Time.
M_{fuel}	Total fuel consumption of the voyage [tonne].
M_{east}	Fuel consumption eastbound.
M_{west}	Fuel consumption westbound.
MSE	Mean squared error.
m_{fuel}	Fuel consumption rate [$\frac{ton}{h}$ or $\frac{l}{h}$].
m_j	Total number of waypoints in a leg.
n	Index n.
n_j	Total number of legs.
NLCB	Negative Lower Confidence Bound.
NM	Nautical Miles.
NNs	Neural Networks.
NSGA-III	Non-dominated Sorting Genetic Algorithm III.

OLS	Ordinary Least Squares.
P	Engine power [kW].
P_E	Effective power [kW].
P_D	Propulsion/Delivered power [kW].
P_B	Brake power [kW].
$P_{\mathbf{bow}}$	Power bow engine [kW].
P_j	Power at leg j [kW].
$P_{\mathbf{max}}$	Maximum allowed power [kW].
$P_{\mathbf{min}}$	Minimum allowed power [kW].
$P_{\mathbf{stern}}$	Power stern engine [kW].
P_T	Thrust power [kW].
PSO	Particle Swarm Optimisation.
Q	Total number of reference Monte Carlo voyages.
PELT	Pruned Exact Linear Time.
r	Residual error [-].
$\mathcal{R}$	Route set [-].
Ro-Pax	Roll-on/roll-off passenger vessel.
R_{Calm}	Calm water resistance [kN].
R_{AW}	Added resistance due to waves [kN].
R_{Wind}	Aerodynamic resistance [kN].
R_F	Frictional resistance [kN].
R_{Earth}	Earth mean radius ($\tilde{6}371$ km).
R_T	Total resistance [kN].
Re	Reynolds number [-].
RPM	Engine speed [rpm].
R_P	Double-ended ferry power ratio [-].
ReLU	Rectified Linear Unit.
S	Ship state vector or Wetted surface area [m^2].
SSE	Sum of Squared Errors.
SSS	Short-Sea Shipping.
SVM	Support Vector Machine.
$SFOC$	Specific fuel oil consumption [$\frac{ton}{kWh}$].

T	Thrust [kN] or Ship draught [m].
t	Time [h].
$t^0_{j+1}, 1$	Nominal departure time from leg $j+1$ [h].
$t_{earliest}$	Earliest time of arrival [h].
t_{latest}	Latest time of arrival [h].
t_{sail}	Target total sailing time for the voyage [h].
$t^{(0)}_{j,m}$	Nominal sailing time at the end of leg j [h].
THETIS-MRV	EU information system for monitoring CO2 emissions.
TICC	Toeplitz Inverse Covariance-Based Clustering.
$T_{\mathbf{arrival}}$	Predicted time of arrival [h].
$T_{\mathbf{opt}}$	Optimal sailing time for a specific voyage [h].
T_{mean}	Mean draught [m].
T_z	Wave period [s].
UNCTAD	United Nations Conference on Trade and Development.
$V_{\mathbf{average}}$	Average speed [kt].
$V_{current}$	Surface current speed [$\frac{m}{s}$].
$V_{j,k}$	Speed overground between two waypoints [kt].
$V_{\mathbf{max}}$	Maximum allowed speed [kt].
$V_{\mathbf{min}}$	Minimum allowed speed [kt].
V_G	Ship speed overground [kt].
V_{water}	Speed through water [kt].
V_{wind}	Wind speed [$\frac{m}{s}$].
W	Weather/Metocean vector [-].
$W_{j,1:m_j}$	Sequence of metocean conditions for leg j.
w_{fuel}	Penalty factor in optimisation [-].
w.r.t.	With respect to.
x	Longitude [°].
$\hat{x}$	Approximate longitude [°].
XGBoost	Extreme Gradient Boosting.
y	Latitude [°].
$\hat{y}$	Approximate latitude [°].

CHAPTER 1

Introduction

1.1 Background and motivation

International trade, enabled by global supply chains and freight networks, underpins modern prosperity. However, the requisite transport infrastructure generates 22% of global greenhouse gas (GHG) emissions (see Fig. 1.1). While maritime transport is structurally cost-efficient and responsible for >80% of world trade (UNCTAD, 2025), it accounts for 11% of transport emissions (3% of total global emissions).

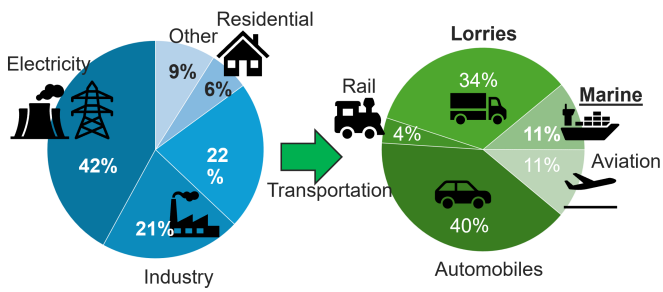

Figure 1.1: Distribution of GHG emissions from various industry and transport sectors (Jaramillo et al., 2022; Rodrigue, 2024).

The regulatory landscape governing maritime emissions is founded upon international frameworks established by the International Maritime Organization (IMO). Princi-

pally, MARPOL Annex VI mandates the prevention of air pollution from ships, prompting a fundamental shift in global fuel procurement strategies. The IMO 2020 regulation lowered the global upper limit on the sulphur content of ships' fuel oil from 3.50% to 0.50% m/m (IMO, 2019). This compelled operators to transition towards very low sulphur fuel oil (VLSFO) or invest in exhaust gas cleaning systems (EGCSs). Layered on top of these global standards are localised, highly stringent emission control areas (ECAs). Within sulphur emission control areas (SECAs), such as the Baltic and North Seas, fuel sulphur content is capped at 0.10% m/m. Concurrently, nitrogen emission control areas (NECAs) enforce Tier III NO_x standards, requiring an 80% reduction in NO_x emissions for newer ships relative to Tier I baselines (IMO, 2023).

The further tightening of international targets requires steep reductions in GHG emissions relative to 2008 levels (IMO, 2021). The IMO manages this by approaching efficiency from two angles: through the Energy Efficiency Existing Ship Index (EEXI), which governs "as-built" technical capabilities, and the Carbon Intensity Indicator (CII), which regulates "in-service" operational efficiency. Operating in tandem with the IMO, the European Union (EU) has also implemented rigorous regional measures through its Monitoring, Reporting, and Verification (MRV) regulation (DNV, 2025; Lloyd's Register, 2025), which represents the foundational pillar of the EU strategy.

Regulatory instruments have therefore fundamentally shifted decarbonisation costs onto day-to-day operations. Building upon the MRV regulation, the EU Emissions Trading System (ETS) integrated the maritime sector as of 1 January 2024 (European Parliament and Council, 2023a), operating on a direct "Tank-to-Wake" basis. Conversely, the FuelEU Maritime regulation enforces progressive GHG intensity limits starting in 2025, utilising a comprehensive lifecycle "Well-to-Wake" approach (European Parliament and Council, 2023b). While these regulations are motivated by the environmental objective of decarbonisation, the primary driver for shipowners to invest in operational efficiency remains strictly financial (DNV GL, 2015). The cost of non-compliance is deliberately engineered to exceed the cost of transition; failing to comply with EU Allowances incurs a penalty of €100 per tonne of CO_2, while FuelEU deficits are penalised at a severe rate of €2,400 per tonne of VLSFO equivalent, increasing by a 10% multiplier for each consecutive year of noncompliance. Within the EU ETS framework, the business case is often defined more by the immediate reduction of taxable fuel consumption—for example, saving 1.6 tonnes of MDO—than by the abstract target of mitigating the equivalent 5 tonnes of CO_2. Ultimately, the legal definition of the "Shipping Company" enables owners to push these substantial compliance burdens downstream to charterers and consumers via commercial contracts, making emissions management a central component of commercial maritime strategy.

Within the EU, short-sea shipping (SSS) handles 50–57% of all seaborne goods, shifting cargo from congested roads to a more energy-efficient mode (Eurostat, 2023; Styhre, 2010). SSS operates close to shore and within ECAs, where it faces stricter local environmental requirements and higher fuel and compliance costs, which can erode its competitiveness relative to trucking (Comi & Polimeni, 2020; Fadda et al., 2020). Reducing SSS emissions is therefore a priority, but services must maintain tight ETAs given frequent port calls and intermodal links. This requires voyage planning capable of responding to weather, market prices, and encountered conditions (Wang et al., 2020; Zaccone et al., 2018). Furthermore, the European short-sea fleet is ageing—ships have an average age of almost 25 years, and key coaster segments average approximately 27 years old—and capital-intensive retrofits and propulsion upgrades are often uneconomic for vessels nearing the end of life. Meanwhile, uncertainty over future fuels, high alternative-fuel costs, and shipyard constraints delay hardware decisions, making immediate technical pathways difficult to scale. In contrast, operational measures (e.g., ship performance monitoring and speed and voyage optimisation) provide low-cost, rapidly deployable savings that directly reduce fuel burn and EU MRV-related compliance costs, delivering immediate benefits for maritime transport (DNV GL, 2015).

SSS also exhibits operational characteristics that complicate efforts to improve efficiency. Most notably, SSS services must adhere to stringent, reliable ETA schedules due to frequent port calls, intermodal connections, and fixed timetables. These constraints reduce the solution space available to ship operators, thereby increasing the importance of intelligent, data-driven assistance. Voyage optimisation is therefore essential for ensuring operational efficiency and minimising fuel consumption by adjusting decisions (e.g., power allocation) in response to variations in weather, currents, and fuel costs (Wang et al., 2020; Zaccone et al., 2018).

However, many optimisation algorithms in the current literature rely on speed-based formulations that do not reflect how short-sea sailing ships are actually controlled. Rather than regulating speed directly, real-world propulsion systems manage power, torque, and propeller pitch, especially in slow-steaming scenarios aimed at reducing fuel costs (Wang et al., 2019). This motivates the use of power-based frameworks with accurate performance models to ensure ETA compliance while reducing fuel consumption. Short-sea voyages also benefit from high-quality short-horizon metocean forecasts and adaptive segmentation (instead of fixed waypoints), which produce route segments with more homogeneous conditions and stable power-allocation strategies. Finally, maritime digitalisation provides rich operational data, enabling high-fidelity ML/grey-box models that support reliable, real-time optimisation (Lang et al., 2022a).

1.2 Objectives and goals

To enhance shipping energy efficiency in SSS and facilitate the reduction of fuel consumption and emissions, this research develops an integrated, data-driven operational framework for optimal ship voyage planning. While ML techniques and optimisation methods have been independently developed for SSS, this research explores combining these approaches for direct engine-power allocation and voyage optimisation. In addition, the thesis develops a unified, data-driven approach that combines ML-based simulations with added-resistance modelling to provide a more reliable ship-performance model for optimising fuel consumption in SSS. To operationalise this framework, the present study is structured into the following objectives and research tasks:

1. **Develop and validate various principles and data-driven ship performance models.**
   This research first implements and validates high-accuracy black-box ML models to predict ship speed and fuel-consumption rates across various propulsion settings and metocean conditions using available sensor and hindcast data. To ensure robustness and generalisability, multiple ML methodologies are systematically compared to identify the most suitable algorithms for ship-performance prediction. Simultaneously, a general grey-box model is developed and calibrated for wave-added resistance coefficients, using empirical data to enhance prediction accuracy and improve overall model fidelity within the voyage optimisation framework.

2. **Develop refined voyage optimisation methods.**
   A set of hybrid optimisation methods is implemented and refined to determine the most energy-efficient engine power allocation for each voyage leg, achieving minimum fuel consumption whilst strictly meeting ETA constraints. Real-world case study voyages and, where available, full-scale trials are used to quantify the effectiveness of the SSS optimisation methods, demonstrating practical energy savings and emission reductions attributable to the optimal power-allocation strategies.

3. **Develop an integrated voyage optimisation framework.**
   The framework integrates the trained ML ship-performance models to accurately estimate total fuel consumption and sailing duration for complete short-sea voyages, enabling high-fidelity evaluation of operational scenarios. In conjunction, a novel clustering-based route-segmentation methodology is introduced to partition fixed routes into voyage legs defined by homogeneous metocean conditions. This ensures that the optimisation process systematically accounts for fluctuating environmental influences.

1.3 Assumptions and limitations

The development of a tractable, deployable, and scientifically rigorous optimisation framework necessitates several simplifying assumptions. These assumptions do not diminish the validity of the results; instead, they define the operational envelope within which the proposed methods are intended to function. This section details the principal assumptions and limitations that underline the modelling, simulation, and optimisation work conducted in this thesis.

- **Data availability and reliability**
  To enable large-scale model training and validation, this study assumes that, after appropriate statistical filtering, onboard propulsion and motion sensors provide sufficiently accurate ground-truth data. Whilst measurement noise is implicitly mitigated, long-term sensor drift and systematic bias are not explicitly modelled, reflecting common industrial practice in operational analytics. Environmental inputs are obtained from reanalysis hindcast databases, which are assumed to approximate the encountered wind, wave, and surface-current conditions with adequate fidelity for both performance modelling and voyage optimisation. Although such datasets contain inherent uncertainties and temporal–spatial smoothing, they remain the most widely used and validated resource for maritime operations and research. Because ML models are trained on historical operational and hindcast data, they inherit the uncertainty, sparsity, or imbalance of those data. These limitations are accepted as consistent with real-world short-sea operational data quality and availability.

- **Operational scope and dynamic constraints**
  The modelling framework is intentionally restricted to open-water, quasi-steady operations, representing the dominant portion of SSS activity. Transient phenomena, including manoeuvring, acceleration phases, berthing operations, shallow water hydrodynamics, and extreme sea states (e.g., storms or ice), are excluded due to their complexity and the limited availability of consistent sensor data. Furthermore, the ship is modelled as an isolated entity; interactions with other maritime traffic (e.g., collision avoidance manoeuvres and routing deviations due to congestion) are explicitly omitted, as incorporating multi-ship interactions falls outside the scope of this work. This assumption ensures that the ML models remain aligned with the physical regimes represented in the training data. Engine power optimisation is also constrained to feasible propulsion settings within manufacturer-specified limits on torque, shaft speed, and pitch (where applicable). Fast transient engine dynamics, governor behaviour, and control-system delays are excluded because the optimisation relies on steady-state (or quasi-steady) performance characteristics. Within each voyage segment, ship displacement, trim, and operational parameters are

assumed to be constant, reflecting the relatively small variations typically observed over short-sea legs.

- **Physical modelling assumptions**
  The performance models use wave, wind, and ocean current data as the primary environmental drivers of resistance and fuel consumption. Other environmental variables (e.g., water temperature, salinity variations) are considered second-order and are omitted from the modelling scope. The hull condition is assumed to remain consistent with the training data; modelling the dynamic progression of fouling would require specialised long-term datasets that were not available for this study. Fouling and degradation effects are thus not explicitly included. The wave-added resistance component is modelled using a generalised grey-box formulation calibrated to published model-scale experiments. Although care is taken to ensure physical generality, transferability is assumed only for ships with similar hull forms and within the same Froude-number range. Finally, voyage segmentation into legs is based on hindcast or forecast metocean fields. These conditions are assumed to remain statistically representative over the duration of each segment, consistent with short-sea voyage timescales, but may not hold for rapidly evolving weather systems.

1.4 Thesis outline

This thesis is structured around five research papers, each addressing components of the development, validation, and application of an integrated data-driven framework for performance modelling and voyage optimisation in SSS. Fig. 1.2 summarises these papers' relationships, methodological focus and case studies, organised along two dimensions: (i) the ship type and operational context and (ii) the technical objective of each study. This classification highlights how the research progresses from fundamental model development to advanced optimisation and system integration. The figure illustrates the interdependence of the papers, demonstrating a coherent progression from ML-based surrogate modelling through simulation to optimisation studies, culminating in full-scale validation and system-level demonstration.

Beyond the classification of the individual component papers, the broader structure and methodological logic of the thesis are captured in Fig. 1.3, which presents a complete overview of the framework developed throughout the doctoral project. The diagram illustrates the two central pillars of the research.

Data-driven performance modelling encompasses both black-box and grey-box approaches to predicting ship resistance, propulsion power, speed, and fuel consumption under varying environmental and operational conditions. Voyage optimisation,

which applies these predictive models to evaluate alternative operational strategies, performs power-allocation optimisation, ultimately recommending fuel-efficient voyages that satisfy strict ETA requirements.

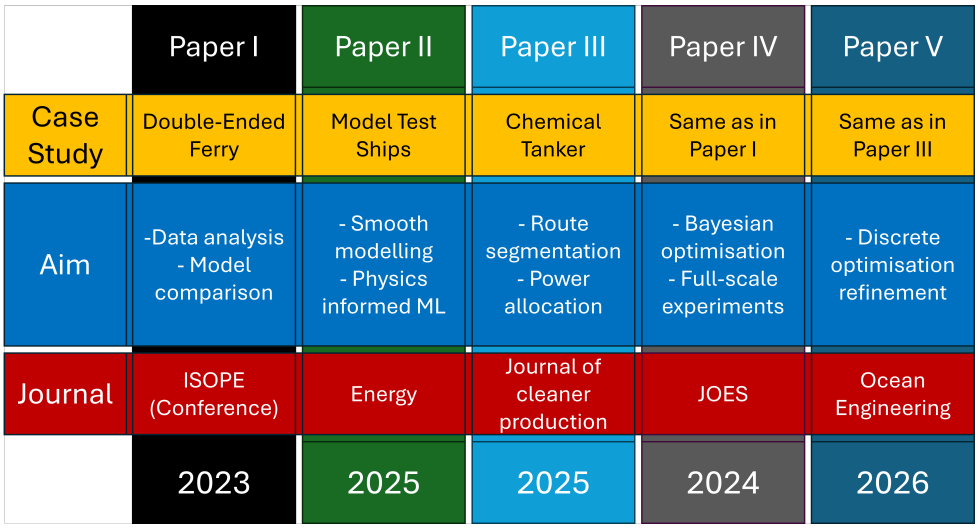

Figure 1.2: Breakdown of the research papers.

The figure shows how these components interact: raw sensor and metocean data feed into the modelling stage; the resulting models enable high-fidelity voyage simulations; and the simulation engine becomes the core of the advanced optimisation algorithms. These components are further synthesised into a unified decision-support system (DSS), thereby representing a complete pipeline from data acquisition to operational recommendations.

The thesis structure mirrors this methodological organisation. Chapter 2 establishes the theoretical and methodological foundations of data-driven modelling, discussing ML algorithms, black-box regression techniques, grey-box formulations for added resistance, and procedures for curating and pre-processing operational data. This chapter develops the predictive tools that underpin the remainder of the thesis.

Chapter 3 builds directly upon the modelling framework, introducing voyage-level simulation and optimisation methodologies. This includes formulating a constant-ETA simulation environment, developing clustering-based segmentation to define homogeneous voyage legs, and implementing hybrid optimisation methods for power allocation. Chapters 2 and 3 are intentionally interconnected, with the optimisation results depending fundamentally on the predictive accuracy of the previously con-

structed models.

Chapter 4 synthesises these methodologies, presenting the integrated data-driven modelling and optimisation framework developed across the five appended studies. Chapter 5 then shifts from methodology to application, providing a detailed overview of the data sources and case study ships used to validate the proposed framework. It describes the ships' operational characteristics, the available sensor systems, and the metocean datasets, thereby establishing the empirical basis for the subsequent research papers.

Chapter 6 synthesises and summarises the appended research papers, highlighting their individual contributions whilst situating them within the broader framework. Each paper is discussed in terms of its objectives, results, methodological innovations, and role in advancing the overall thesis. Finally, Chapters 7 and 8 present the conclusions and suggest avenues for future research. Chapter 7 summarises the key findings and evaluates the contributions of the thesis in relation to the research objectives, whilst Chapter 8 outlines potential extensions of the framework, including integration with real-time weather forecasts, adaptive control architectures, and applications to other ship types and operational contexts. This structure provides a coherent and progressive narrative that moves from foundational modelling through system-level integration to practical validation, reflecting both the scientific depth and applied relevance of the research.

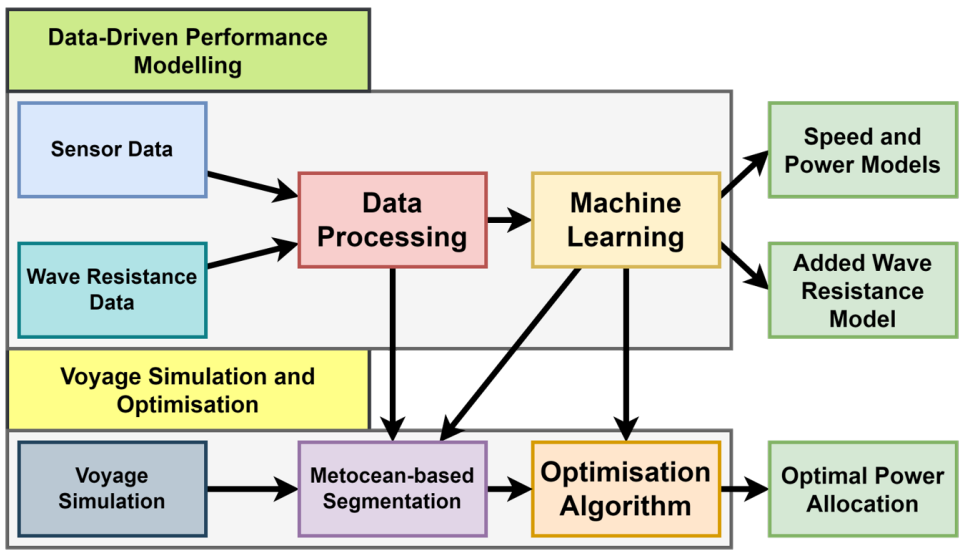

Figure 1.3: Complete thesis framework.

CHAPTER 2

Literature review

This chapter reviews the state-of-the-art literature relevant to the scope of this thesis, covering ship performance modelling, voyage optimisation, route segmentation, and ship added resistance in waves.

2.1 Ship performance modelling

Mathematical models of ship performance play an important role in the shipping industry by estimating voyage fuel consumption and thereby facilitating the determination of voyage viability and profitability. The simplest fuel consumption model, the Cubic Law, assumes that a ship's fuel consumption varies cubically relative to the design condition (Carlton, 2007). These models are classified as (1) white-box models, (2) black-box models, and (3) grey-box models. Physics-based or white-box models characterise ship behaviour using physical principles, matching empirical resistance correlations to energy consumption. A summary of relevant modelling research is presented in Table 2.1.

Holtrop and Mennen (1982) and Hollenbach (1998) conducted foundational work in ship performance modelling, establishing the standard approaches for estimating calm-water propulsion power in traditional merchant ships, including bulk carriers and tankers. However, these models are typically constrained by assumptions of steady-state operation, constant loading conditions, and simplified resistance components. Furthermore, these traditional formulations are inadequate for certain

ship types, such as double-ended ferries, which necessitate specialised modelling approaches. Subsequent refinements of these models, such as the inclusion of additional parameters and metocean conditions, have improved them over time (Mao et al., 2016).

Advances in onboard instrumentation and the increasing availability of high-frequency operational data have motivated a shift toward data-driven ship performance modelling. Purely data-driven approaches, commonly referred to as black-box models (Lang et al., 2024), adopt a machine learning (ML) regression framework based on processed operational data. Unlike white-box models, ML frameworks can capture the complex, nonlinear relationships among propulsion power, ship speed, and environmental data. Several comparative studies have evaluated various ML algorithms for ship performance modelling. For example, Lang et al. (2022a) assessed several algorithms for propulsion power prediction, finding that extreme gradient boosting (XGBoost) achieved the highest accuracy under varying sea states. Laurie et al. (2021) compared AdaBoost, multiple linear regression, K-nearest neighbours, and random forest models, identifying random forest as most effective for simulating the effects of biofouling. Similarly, Abebe et al. (2020) evaluated various ML methods for ship performance modelling using automatic identification system data, and Su et al. (2025) investigated the performance of categorical boosting in predicting the fuel consumption of Ro-Ro carriers. Chung et al. (2025) reported that XGBoost outperformed five other models in terms of both ship fuel consumption prediction accuracy and error reduction.

Artificial neural networks (ANNs) have also been extensively investigated for ship performance modelling. Beşikçi et al. (2016) developed an ANN-based decision-support system to estimate fuel consumption using noon-report data. Parkes et al. (2018) predicted shaft power using a three-month operational dataset, demonstrating that deeper network architectures improved predictive performance. Karagiannidis and Themelis (2021) employed feed-forward neural networks to estimate ship power and speed, and proposed a new method for handling missing data. However, the nonlinear complexity of ANNs may limit their generalisation capability. Kim et al. (2021c) highlighted this issue, comparing linear regression and ANN-based models for fuel consumption estimation and observing cases where simpler linear models achieved comparable or superior performance. To construct fuel consumption models for new vessels with limited operational data, Luo et al. (2025) leveraged knowledge from seven other container ships by applying an ANN-based transfer learning approach.

Data quality significantly impacts performance modelling. High-quality data are typically obtained by eliminating outliers and removing noise using methods such

Table 2.1: Summary of representative studies in ship performance modelling.

Reference	Methodology	Data Characteristics	Key Conclusions
Carlton (2007), Holtrop and Mennen (1982), Hollenbach (1998), Mao et al. (2016)	Empirical & white-box	Model-scale, design parameters, metocean	Established standard physics-based approaches for calm-water power, later improved with metocean data, but constrained by steady-state assumptions.
Cleveland (1979), Savitzky and Golay (1964), Bassam et al. (2022)	Data preprocessing	High-frequency operational data	Emphasises that model suitability strongly depends on data quality; filtering and outlier removal are critical for prediction accuracy.
Abebe et al. (2020), Laurie et al. (2021), Lang et al. (2022a), Su et al. (2025), Chung et al. (2025)	Black-box (tree-based & ML)	AIS and operational data	Tree-based ensembles (especially XGBoost and random forest) consistently achieve highly accurate fuel and power prediction under varying conditions.
Beşikçi et al. (2016), Parkes et al. (2018), Karagiannidis and Themelis (2021), Kim et al. (2021c), Luo et al. (2025)	Black-box (ANNs & transfer learning)	Noon-reports, operational datasets	Deeper networks improve accuracy and transfer learning mitigates data scarcity, although simpler linear models can sometimes match ANN performance due to generalisation limits.
Leifsson et al. (2008), Coraddu et al. (2017), Odendaal et al. (2023), Lang et al. (2024), Zhang et al. (2025b)	Grey-box (physics + ML)	Operational data / public wave data	Serial and parallel grey-box architectures provide superior extrapolation capabilities compared to pure black-box models.

as locally estimated scatter plot smoothing (LOESS) regression (Cleveland, 1979), Savitzky–Golay filters (Savitzky & Golay, 1964), or moving average filters (Karagiannidis & Themelis, 2021; Lang et al., 2022a). Feature selection and correlation analysis further improve prediction accuracy (Abebe et al., 2020; Kim et al., 2021c). Bassam et al. (2022) demonstrated that model suitability strongly depends on the characteristics of the available data and emphasised the importance of appropriate data preprocessing and hyperparameter selection.

Grey-box modelling has emerged as an effective alternative to address the limitations of purely data-driven approaches, particularly their poor extrapolation capability. Grey-box models integrate the physical principles from white-box models with data from black-box models. Two primary architectural configurations are commonly used: serial and parallel implementations (Leifsson et al., 2008). Although grey-box models typically outperform white-box models in interpolation and black-box models in extrapolation, they are often inferior to pure black-box models for interpolation tasks (Coraddu et al., 2017; Lang et al., 2024; Odendaal et al., 2023; Zhang et al., 2025a).

The choice of an appropriate modelling approach relies heavily on a ship's operational profile. Pure black-box modelling strategies are often sufficient for ships sailing along predetermined, relatively short routes, as such ships are likely to encounter similar operating and environmental conditions, resulting in dense, homogeneous datasets. In contrast, ocean-going ships are frequently exposed to novel metocean conditions, resulting in sparse data coverage and reduced prediction accuracy.

2.2 Voyage optimisation

Voyage optimisation aims to minimise fuel consumption and emissions whilst satisfying service constraints such as schedule adherence and safety (Chen et al., 2025). Studies have found that voyage optimisation can also help reduce the long-term accumulation of fatigue damage in ships (Lang et al., 2020). Traditional methods treat ship speed as the primary control variable and employ computational optimisation techniques to derive optimal speed or power profiles along a route. Commonly used approaches include evolutionary algorithms (EAs), dynamic programming (DP), particle swarm optimisation (PSO), and mixed-integer linear programming (MILP). Selected research in this area is comprehensively summarised in Table 2.2.

EAs, especially genetic algorithms (GAs), have been widely used for multi-objective optimisation in voyage planning. Lee et al. (2018) proposed a GA-based approach where path and speed optimisation are performed using rotational speed (RPM) and heading as decision variables. Szlapczynska and Szlapczynski (2019) employed

evolutionary multi-objective optimisation for weather routing, aiming to achieve balanced trade-offs among sailing time, fuel consumption, and navigational safety. Wang et al. (2021a) developed a multi-objective solution that combines GA with DP to reduce fuel consumption and greenhouse gas emissions. More recently, Li et al. (2024) introduced a GA-LSTM model coupled with the NSGA-III algorithm for the multi-objective optimisation of path, speed, and trim. Zhao et al. (2025) proposed a data-driven heuristic framework for emission-aware voyage optimisation, reporting an average ship emissions reduction of approximately 8%. Li et al. (2026) developed a multi-objective optimisation framework that balances sailing time and carbon emissions by employing an adaptive GA.

DP has also been extensively applied to voyage optimisation problems. Zaccone et al. (2018) proposed a three-fold DP model for the joint optimisation of ship routes and speeds. Du et al. (2019) applied DP for tanker speed optimisation under weather constraints, and Ma et al. (2020) extended this to account for emission control areas (ECAs) and sea-state variability. Tzortzis and Sakalis (2021) then further enhanced this approach by segmenting the planning horizon and explicitly addressing the uncertainty associated with long-term predictions. Additionally, Fan et al. (2022) combined ANNs with DP to jointly optimise ship speeds and trim for a predetermined route.

PSO has also proven effective for voyage optimisation. Wang et al. (2020) applied PSO to optimise ship speed and routing on a fixed spatial grid, accounting for dynamic environmental conditions. Dai et al. (2022) integrated PSO with a support vector machine (SVM)-based power prediction model for optimal power allocation. Du et al. (2023) enhanced the PSO framework by considering second-order oscillations, which improved convergence performance under different sea-state conditions and enabled the determination of the optimal route and speed for minimising fuel consumption and CO_2 emissions. Zhang et al. (2025a) proposed a PSO–based voyage planning method, harnessing depth-aware speed control to reduce the fuel consumption of inland ships in shallow waters.

Alternative optimisation approaches have also been explored, including Dijkstra-based route optimisation methods (Bahrami & Siadatmousavi, 2024; Wang et al., 2019), ant colony optimisation for Arctic navigation (Zhang et al., 2022), model predictive control (MPC) and swarm intelligence approaches for dynamic routing (Wang et al., 2021b), as well as reinforcement learning (RL) frameworks for electric propulsion scheduling (Shang et al., 2024). Bayesian optimisation (BO) has also been applied; for instance, Vergara et al. (2023) integrated BO with XGBoost to optimise power allocation for a double-ended ferry, demonstrating its effectiveness for operational decision-making, and Yu et al. (2024) employed spline-based surrogate models

Table 2.2: Summary of representative studies in voyage optimisation.

Reference	**Methodology**	**Control Variables & Objectives**	**Key Conclusions**
Lee et al. (2018), Szlapczynska and Szlapczynski (2019), Wang et al. (2021a), Li et al. (2024), Zhao et al. (2025), Li et al. (2026)	Evolutionary Algorithms (GA, NSGA-III)	Route, Speed, RPM, Trim (*Objectives: Fuel, Emissions, Safety, ETA*)	Effective for multi-objective trade-offs, balancing fuel savings, safety, and emissions.
Zaccone et al. (2018), Du et al. (2019), Ma et al. (2020), Tzortzis and Sakalis (2021), Fan et al. (2022)	Dynamic Programming (DP)	Route, Speed, Trim (*Objectives: Fuel, ETA*)	Robust deterministic framework; handles weather uncertainty well when combined with ML or segmented horizons.
Wang et al. (2020), Dai et al. (2022), Du et al. (2023), Zhang et al. (2025a)	Particle Swarm Optimisation (PSO)	Route, Speed, Power (*Objectives: Fuel, Emissions*)	Handles dynamic environments well; integrates easily with data-driven power models.
Wang et al. (2019), Bahrami and Siadatmousavi (2024), Zhang et al. (2022), Wang et al. (2021b)	Graph, Heuristics & MPC	Route (*Objectives: Route*)	Suitable for spatial pathfinding and dynamic routing in specific operational environments.
Shang et al. (2024)	Reinforcement Learning	Power, Propulsion (*Objectives: Fuel*)	Enables dynamic electric propulsion scheduling and continuous operational adaptation.
Vergara et al. (2023)	Bayesian Optimisation	Power (*Objectives: Fuel, ETA*)	Highly effective for operational decision-making by integrating with ML (XGBoost) for power allocation.
Yu et al. (2024)	Splines	Trim (*Objectives: Fuel*)	Successfully optimises specific ship parameters (trim) while highlighting the mechanical wear of speed-based control.

to optimise ship trim.

One aspect of voyage optimisation that is often neglected is the handling of strict operational constraints, such as the estimated time of arrival (ETA). Constraint handling approaches generally fall into two categories: penalty functions and explicit probabilistic modelling. Penalty methods transform constrained problems into unconstrained ones by adding a cost term for constraint violations (Fiacco & McCormick, 1968); however, this approach risks warping the optimisation landscape. Conversely, explicit constraint handling directly models the probability of feasibility for candidate solutions, guiding the optimiser towards valid regions of the search space (Gelbart et al., 2014).

Despite their established effectiveness, most existing voyage optimisation approaches treat ship speed as a directly controllable variable, overlooking that speed is ultimately governed by propulsion power and propeller dynamics. In real operations, the maintenance of a constant or piecewise target speed requires continuous engine load adjustments, which can increase energy losses and contribute to mechanical wear (Sørensen, 2013; Sørensen et al., 1997; Yu et al., 2024). Selecting a voyage optimisation algorithm therefore entails trading off the ability to handle the non-convexity of the objective function against the associated computational complexity. For long sailing routes, DP can be combined with suitable approximations that decompose the voyage into a sequence of state transitions, providing a deterministic and computationally efficient framework for selecting control variables. Conversely, when the solution space is discrete, graph-based or heuristic, optimisation methods may provide sufficiently accurate solutions at significantly lower computational cost.

2.3 Route segmentation

Segmenting a maritime route into sections with consistent environmental and operational characteristics can significantly enhance both performance analysis and optimisation by enabling adaptive speed or power allocation per segment to minimise fuel consumption (Vergara et al., 2025). In practice, such segmentation may be achieved heuristically by dividing the route into legs for which speed is held constant for a predefined duration. However, these heuristics become impractical when planning a voyage in advance, because the segments must be defined before the optimisation algorithm can determine the optimal speed or power profile (Hsieh et al., 2025). Table 2.3 compiles research implementing route segmentation.

Traditional segmentation schemes often divide routes uniformly by sailing distance (Wang et al., 2019; Zaccone et al., 2018) or by port-to-port legs for liner services (Guericke & Tierney, 2015; Qi & Song, 2012; Wang & Meng, 2012; Wu, 2020). However,

Table 2.3: Summary of representative studies on route segmentation.

Reference	**Methodology**	**Data Characteristics & Focus**	**Key Conclusions**
Hsieh et al. (2025)	Conceptual & Heuristics	Voyage planning	Highlights the necessity of defining segments before optimisation to enable adaptive power/speed allocation.
Zaccone et al. (2018), Wang et al. (2019), Wang and Meng (2012), Qi and Song (2012), Guericke and Tierney (2015), Wu (2020)	Uniform division	Distance-based or port-to-port legs	Simple to implement, but fundamentally ignores localised variations in the metocean conditions that dictate performance.
Zhang et al. (2018), Wen et al. (2019)	Spatial turning-point	Course changes and spatial waypoints	Effectively segments based on trajectory geometry, but overlooks the impact of dynamic weather on propulsion.
Wang et al. (2017), Yan et al. (2018), Zhang et al. (2024), Li et al. (2022), Wang et al. (2020), Li et al. (2023)	Data-driven clustering	Metocean parameters and sea states	Captures environmental trends, but risks over-segmentation and ignoring nonlinear dependencies between weather and fuel.
Killick et al. (2012), Hallac et al. (2017), Vergara et al. (2025)	Advanced time-series (PELT, TICC)	Multivariate temporal dependencies	Provides a highly scalable framework by efficiently detecting statistical change points in weather conditions.

these simple segmentation approaches ignore localised variations in wind, wave, and current conditions, which significantly affect ship performance. Methods for addressing these limitations include the turning-point algorithm (Wen et al., 2019; Zhang et al., 2018), which effectively segments routes based on course changes or spatial waypoints. However, this approach overlooks the impact of metocean conditions, which vary significantly even along predefined trajectories. For power allocation optimisation, route segments should ideally be defined by homogeneity in metocean conditions rather than by distance or heading, because fluctuating wave and wind forces require different propulsion responses to maintain operational efficiency (Li et al., 2022).

Data-driven segmentation methods utilise clustering and statistical change detection. K-means clustering has been widely adopted for grouping sequential waypoints based on the similarity of metocean parameters (Wang et al., 2017; Yan et al., 2018). Turning-point detection approaches have also been used to identify sudden sea-state transitions (Li et al., 2022; Zhang et al., 2024). While these methods capture general trends, they may over-segment routes or ignore nonlinear dependencies between environmental variables and fuel consumption, limiting operational practicality (Li et al., 2023; Wang et al., 2020).

More advanced time-series segmentation methods from other domains offer promising alternatives. The pruned exact linear time (PELT) algorithm (Killick et al., 2012) efficiently detects statistical change points by minimising a penalised cost function, maintaining scalability even with large datasets. Toeplitz inverse covariance clustering (TICC) (Hallac et al., 2017) extends this concept by modelling temporal dependencies and correlations within multivariate time series. Both approaches can identify transition zones between various metocean regimes, thereby providing a computationally efficient, performance-aware segmentation framework applicable to voyage optimisation.

It should be noted that using the turning point method bluntly or performing a fixed-interval segmentation can trigger higher fuel consumption because they ignore metocean conditions along the route. To address the inherent challenges of metocean environments, a segmentation approach should be used that adapts the length of the generated legs to the metocean conditions the ship may encounter.

2.4 Added resistance in waves

Because ships predominantly operate at sea, weather factors can account for up to 15% of total fuel consumption (Godet et al., 2024). Therefore, the accurate quantification of a ship's wave resistance is fundamental to ship performance assessment and simulation (Godet et al., 2024; Xiao et al., 2022). Existing methods for predicting wave-induced added resistance can be categorised into three groups: (1) numerical and experimental approaches, (2) semi-empirical formulations, and (3) data-driven models (Zhang et al., 2025b). Selected research in this area is summarised in Table 2.4.

Table 2.4: Summary of representative studies on predicting added resistance in waves.

Reference	Methodology	Data Focus	Key Conclusions
Godet et al. (2024), Xiao et al. (2022), Lee et al. (2019a), Kim and Kim (2011), Chen et al. (2021)	Numerical & experimental (CFD/EFD)	High-fidelity flow, wave diffraction, weather impact	Establishes that weather dictates up to 15% of fuel use; yields high accuracy but carries prohibitive computational costs.
Fujii and Takahashi (1975), Liu and Papanikolaou (2016), Liu and Papanikolaou (2020), Lang and Mao (2020), Lang and Mao (2021)	Semi-empirical formulations	Hull geometry, regression calibration	Enables rapid estimation for design phases, but struggles to generalise across novel ship types or rough sea states.
Cepowski (2020), Cepowski (2023), Duan et al. (2022), Sun et al. (2022), Zhang et al. (2025b)	Data-driven (ANNs, DFN, GPR)	Geometric parameters, Froude numbers	Offers excellent efficiency and handles nonlinearities, but often neglects correlations within experimental test conditions.

Experimental testing represents the most reliable method for determining added resistance, although it requires expensive facilities and scale models. Using computational fluid dynamics (CFD) to capture flow details is a common method, although it incurs substantial computational costs. For instance, Lee et al. (2019a) employed CFD simulations to investigate diffraction effects on added resistance and wake characteristics under regular head waves. Kim and Kim (2011) used a Rank-

ine panel method to model added resistance under irregular wave conditions. Kim et al. (2021b) investigated added resistance induced by oblique regular waves using combined experimental and numerical approaches. Chen et al. (2021) used a boundary element method to quantify wave-induced forces. Despite their high accuracy, the computational cost of these high-fidelity approaches limits their applicability to real-time operational scenarios.

To improve computational efficiency, empirical and semi-empirical formulations have been developed to approximate added resistance based on experimental regression data. Fujii and Takahashi (1975) introduced semi-empirical formulations to describe the increase in resistance caused by wave reflection. Subsequent studies refined these models by incorporating ship geometry and hydrodynamic correction factors. Liu and Papanikolaou (2016, 2020) enhanced existing formulations by incorporating terms related to the block coefficient and heading angle. Lang and Mao (2020, 2021) proposed semi-empirical models with correction factors for significant wave height and included frequency to account for nonlinear effects under rough sea conditions. While these methods enable rapid estimation that is suitable for design and operational analysis, their reliance on regression calibration often limits their generalisation across ship types and sea states.

Recent research increasingly employs data-driven models to overcome these limitations. ANNs have been predominant in this context, offering both computational efficiency and flexibility in handling nonlinear relationships. Cepowski (2020) applied a two-layer ANN using ship geometric parameters and Froude number inputs to predict added resistance in head waves. This approach was subsequently extended to improve prediction accuracy (Cepowski, 2023). Duan et al. (2022) introduced a deep feed-forward neural network (DFN) method for the rapid prediction of added resistance in heading waves, addressing both the complexity of semi-empirical formulations and the limitations of shallow neural architectures. Other ANN-based methods have been developed for specialised environments, such as ice-covered waters (Sun et al., 2022). Despite their advantages, most ANN models treat each frequency test as independent, neglecting correlations within experimental datasets that share operating conditions.

CHAPTER 3

Fundamental principles

This chapter presents the theoretical foundations of the surrogate-modelling and voyage-optimisation framework developed in this thesis. It first outlines the relationships governing ship fuel consumption and wave-induced added resistance, providing the physical basis for the subsequent modelling components. It then introduces the mathematical principles underlying the supervised learning techniques, optimisation algorithms, and clustering methods employed in the research. These principles inform the selection of the input variables and model structures used for training the ML regression models presented in later chapters.

3.1 Ship resistance and propulsion

The total resistance R_{Total} acting on a ship can be expressed as the sum of calm-water and environment-induced components (Birk, 2019):

$$R_{Total} = R_{Calm} + R_{Wind} + R_{AW} + R_O, \tag{3.1}$$

where R_{Calm} represents the calm-water resistance, R_{Wind} represents the aerodynamic resistance, R_{AW} represents the wave-induced added resistance, and R_O represents resistance caused by miscellaneous factors.

Calm-water resistance R_{Calm} denotes the steady-state resistance encountered by a ship operating in still water, unaffected by wind, waves, or current. This parameter can be estimated using model-scale experiments, CFD simulations, or empirical for-

mulations, including those proposed by Holtrop and Mennen (1982) and Hollenbach (1998). According to the ITTC (2017) procedure, R_{Calm} can be decomposed into viscous and wave-making components as follows:

$$R_{Calm} = (1 + k)R_F + R_W, \tag{3.2}$$

where R_F represents the frictional resistance, k the form factor, and R_W the wave-making resistance. The frictional resistance is given by:

$$R_F = \tfrac{1}{2} \cdot \rho_w \cdot V^2 \cdot S \cdot C_F, \tag{3.3}$$

where ρ_w represents the water density and S represents the wetted surface area of the hull. The frictional resistance coefficient can be determined using:

$$C_F = \frac{0.075}{(\log_{10}(Re - 2))^2}. \tag{3.4}$$

The aerodynamic resistance R_{Wind} can be estimated empirically as a quadratic function of the ship's speed:

$$R_{Wind} = \tfrac{1}{2} \cdot \rho_a \cdot C_{Wind} \cdot A_T \cdot V^2, \tag{3.5}$$

where ρ_a is the air density, A_T denotes the projected frontal area of the ship above the waterline, and C_{Wind} represents an empirical coefficient typically ranging from 0.8×10^{-3} to 2.0×10^{-3}, depending on the ship's superstructure and above-water geometry.

R_O represents the resistance caused by factors not considered above, including ice, hull and propeller fouling, shallow water effects, rudder resistance, and bank effects. R_O is included in the formulation for the sake of completeness, but is disregarded for the purposes of the analysis, as several of these factors are of limited relevance to the specific operational conditions of the ships considered. For instance, ice is not commonly encountered on the Northern European routes under study, and the ships navigate primarily in open, deep waters, which minimise shallow water and bank effects.

3.1.1 Added resistance in waves

When a ship is sailing in real-world sea environments, the corresponding added resistance in waves, referring to the increase in total resistance relative to calm water conditions, is computed statistically via the wave spectrum. This is because the incoming waves are not monochromatic but rather comprise a broad range of frequencies. The ocean environment encountered during a voyage is typically idealised

as a series of stationary sea states that last from tens of minutes to a few hours. Each stationary sea state is characterised by a directional (or unidirectional) wave spectrum parameterised by the wave height H_s and peak period T_p. This thesis employs the JONSWAP spectrum to describe the distribution of wave energy as a function of the wave frequency ω (Hasselmann et al., 1973):

$$S(\omega|H_s,T_p,\gamma) = \frac{320\,H_s^2}{T_p^4\,\omega^5}\exp\left(-\frac{1950}{T_p^4\,\omega^4}\right)\gamma^{\exp\left[-\frac{(\omega-\omega_p)^2}{2\sigma^2\,\omega_p^2}\right]}, \tag{3.6}$$

where ω_p is the peak frequency corresponding to T_p viz., $\omega_p = 2\pi/T_p$, γ is the peak-enhancement factor and σ is the spectral width parameter ($\sigma = 0.07$ for $\omega \leqslant \omega_p$ and $\sigma = 0.09$ when $\omega > \omega_p$). The peak enhancement factor γ represents the ratio of the maximum spectral density to that of the corresponding Pierson–Moskowitz spectrum for a fully developed sea. γ typically ranges from 1 to 7, and $\gamma = 1$ means the JONSWAP spectrum changes to the Pierson–Moskowitz spectrum. In this thesis, γ is set to the standard value of 3.3.

When encountering a specific sea state denoted by $S(\omega|H_s,T_p)$ with default $\gamma = 3.3$ at forward ship speed V and relative heading angle θ in a seaway, can be estimated by:

$$R_{AW}(\omega|H_s,T_p,V,\theta) = 2\int_0^\infty S\left(\omega|H_s,T_p,\right)\frac{R_{AW}\left(\omega|V,\theta\right)}{\zeta_a(\omega)^2}d\omega, \tag{3.7}$$

where $\zeta_a(\omega)$ is the amplitude of the regular wave of frequency ω that induces the added resistance $R_{AW}(\omega)$ under ship operational conditions at forward speed V and relative wave angle β ($\beta = 0^\circ$ indicates head sea). In this description, $\zeta_a(\omega)$ is often obtained from ship model tests in wave basins. R_{AW}/ζ_a^2 is also known as the added resistance transfer function (RAOs) for the given operational conditions.

To establish semi-empirical models or numerical methods for estimating a ship's added resistance in waves, as defined by Eq. 3.7, the associated transfer functions are often transformed into a non-dimensional format as:

$$C_{AW}(\omega \mid V,\theta) = \frac{R_{AW}(\omega \mid V,W)}{\rho_w \cdot g \cdot \zeta_a^2 \cdot B^2/L}, \tag{3.8}$$

where ρ_w is the water density, g is the gravitational acceleration, B is the ship's beam, and L is the length between perpendiculars. The transfer function C_{aw} primarily depends on encounter frequency, wave heading, Froude number, and hull geometry. Section 4.1 details the modelling of the non-dimensional transfer function developed in this thesis.

3.1.2 Ship speed-power modelling

When sailing at sea, a ship's total resistance must be compensated by the thrust force provided by the marine engine-powered propellers. The power required to move a ship sailing forward at a specific speed V through water, overcoming total resistance R_T, is called the effective power P_E (MAN, 2023), and is computed by:

$$P_E = R_T \cdot V, \tag{3.9}$$

where P_E is transmitted from propellers that generate the propulsion power P_D to rotate the propeller, denoted by:

$$P_B = \frac{P_D}{\eta_s} = \frac{P_E}{\eta_s \cdot \eta_h \cdot \eta_r \cdot \eta_o}, \tag{3.10}$$

where η_s represents shaft efficiency, P_B refers to the brake power the marine engine provides to rotate the ship's propellers, η_h represents the hull efficiency, η_o denotes the open water efficiency and η_r the relative rotative efficiency. Fuel consumption can then be derived from the engine brake power P_B, which represents the result of energy conversion within the engine, transforming chemical energy (stored in the fuel) into mechanical energy (engine power) through a combustion process. Thus, the consumed fuel mass flow rate of the main engine required to deliver P_B can be estimated as a function of the delivered power:

$$\dot{m}_f = P_B \cdot \mathrm{SFOC}(P_B, RPM), \tag{3.11}$$

where SFOC is the specific fuel oil consumption, which typically increases at very low or very high engine loads due to reduced propulsion efficiency. Here, the SFOC is a function of, for example, engine load (brake power) and engine speed (RPM) (Latarche, 2021). Depending on the engine type and propeller properties of a specific ship, propulsion efficiency is frequently provided by the manufacturer or ship owner and can be used to calculate overall fuel consumption.

3.1.3 Ship power performance modelling

In modelling a ship's speed and power relationship, full-scale measurements indicate that, for moderate Froude numbers and constant displacement, the total resistance varies approximately with the square of the velocity (Carlton, 2007). Consequently, the delivered power exhibits a cubic dependency on ship speed, which is known as the Admiralty formula (Carlton, 2007):

$$P_B \propto V_G^3, \tag{3.12}$$

which is frequently used to quickly estimate the required power or fuel savings when speed is reduced.

Under adverse environmental conditions, the required power increases due to additional resistance components (Eq. 3.1). The total brake power can therefore be expressed as:

$$P_B \propto V_G^3 + \Delta P_W(H_s, \alpha_{H_s}, T_z, V_{\text{wind}}, \alpha_{wind}, V_{\text{current}}, \alpha_{current}), \tag{3.13}$$

where ΔP_W represents the additional power demand due to incident wave height H_s, wave heading α_{H_s}, wave period T_z, and environmental velocities of wind and current.

The collection of metocean conditions is instead represented by the vector:

$$\mathbf{W} = [H_s, \alpha_{H_s}, T_z, V_{\text{wind}}, \alpha_{wind}, V_{\text{current}}, \alpha_{current}]. \tag{3.14}$$

3.2 Data-driven approaches

In ship performance modelling, data-driven approaches vary in the degree to which they incorporate physical knowledge into the model structure. The surrogate models developed in this thesis follow the general supervised-regression formulation that represents the relationship between a vector of input features $\mathbf{x}$ (such as the metocean conditions $\mathbf{W}$) and a target variable y (such as the brake power P_B) as:

$$y = f(\mathbf{x}) + \varepsilon, \tag{3.15}$$

where $f(\mathbf{x})$ denotes the unknown mapping to be learned from data and ε represents the model uncertainty. This research explored the following modelling approaches: linear/polynomial regression, ANNs, XGBoost, and Gaussian process regression (GPR). These modelling philosophies are commonly grouped into three categories:

- **White-box models** rely on physics-based or semi-empirical formulations. While their parameters have clear physical interpretations, their predictive accuracy is limited under complex operational conditions.
- **Black-box models** learn input–output mapping from data alone, without enforcing any physical structure. They typically achieve high predictive accuracy but offer limited interpretability.
- **Grey-box models** combine physics-based components with data-driven corrections. They retain some of the interpretability of white-box models whilst improving flexibility and accuracy by deploying data-driven elements.

Linear and polynomial regression

Linear and polynomial regression are fundamental supervised learning methods for quantifying simple relationships between dependent variables and one or more explanatory variables (Chatterjee & Simonoff, 2013). As a basis for more advanced algorithms, they merit first exploration in data-driven methodologies.

For a dataset with m samples, the linear regression model can be written as:

$$\hat{y}_i = \mathbf{x}_i^T \boldsymbol{\beta}, \quad i = 1, \ldots, m, \tag{3.16}$$

where $\hat{y}_{(i)}$ is the predicted target, and $\boldsymbol{\beta}$ represents the regression coefficients. The coefficients are obtained via the ordinary least squares (OLS) method, which minimises the sum of squared errors (SSE):

$$\text{SSE} = \sum_{i=1}^{M} (y_i - \hat{y}_i)^2. \tag{3.17}$$

It is widely acknowledged that linear models cannot adequately capture complex nonlinear physical relationships. An extension of this modelling approach is to augment the feature set with polynomial transformations of the original variables. The k-th order polynomial model may be expressed as:

$$\hat{y}_i = \beta_0 + \sum_{r=1}^{k} \sum_{1 \le j_1 \le \cdots \le j_r \le p} \beta_{j_1 \cdots j_r} \, x_{ij_1} \cdots x_{ij_r}, \quad i = 1, \ldots, m, \tag{3.18}$$

where p represents the number of original features. The higher-order terms enable the model to capture nonlinearities in the target variable whilst retaining a linear parametrisation that can be determined using the OLS method. Higher-order polynomials may be used when stronger nonlinearities are present; however, complexity increases rapidly with model order.

Artificial Neural Networks

ANNs constitute a flexible class of nonlinear regression models capable of learning complex relationships between input and output variables. Unlike polynomial models, which rely on predefined basis functions, ANNs learn these transformations directly from data, enabling them to represent highly nonlinear dependencies (Mehlig, 2021). A fully connected feedforward network comprises several layers that successively transform the input via weighted linear combinations, followed by nonlinear activation functions. The model can be written in a compact form as:

$$\hat{y}_i = F(\boldsymbol{\Theta}, \mathbf{x}_i) = f\Big(\boldsymbol{\Theta}_l, \, f_{l-1}\big(\boldsymbol{\Theta}_{l-1}, \, \ldots, f_1(\boldsymbol{\Theta}_1, \mathbf{x}_i)\big)\Big), \quad i = 1, \ldots, m. \tag{3.19}$$

As shown in Eq. (3.19), the neural network defines a feedforward mapping where data flow from the input feature vector $\mathbf{x}_i$ through multiple hidden layers to the predicted target $\hat{y}_i$. Each layer, indexed by l, contains a set of parameters $\mathbf{\Theta}_l$ and a nonlinear activation function f. This study uses the rectified linear unit (ReLU) activation function due to its ability to mitigate the vanishing-gradient problem.

eXtreme Gradient Boosting (XGBoost)

XGBoost (Chen & Guestrin, 2016) is an ensemble learning approach that combines multiple weak learners to produce a strong predictive model. It adds shallow regression trees sequentially, with each new tree trained to correct the residual errors of the previous ensemble. Assuming that the ensemble model comprises K decision trees, the prediction for sample i is given by:

$$\hat{y}_i = \sum_{k=1}^{K} f_k(\mathbf{x}_i), \tag{3.20}$$

where f_k denotes the k-th decision tree. For each tree, the prediction can be written as $f_k(\mathbf{x}_i) = w_{q(\mathbf{x}_i)}$, where $q(\mathbf{x}_i)$ is the index of the leaf node reached by sample $\mathbf{x}_i$, and $w_{q(\mathbf{x}_i)}$ represents the corresponding leaf weight.

Standard loss functions (e.g., mean squared error) measure prediction accuracy but fail to account for model complexity. XGBoost incorporates an explicit regularisation term to control tree complexity. Its objective function is defined as:

$$\text{Obj} = \sum_{i=1}^{m} \ell(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k), \tag{3.21}$$

where the first term represents the empirical loss measuring the discrepancy between the true value y_i and the prediction $\hat{y}_i$, and the second term penalises model complexity based on tree structure. XGBoost is trained in an additive manner. When a new tree is added at iteration n, the loss becomes $\ell\left(y_i^{(n)}, \hat{y}_i^{(n-1)} + f_n(\mathbf{x}_i)\right)$, where f_n is the tree introduced at iteration n. Using a second-order Taylor expansion, the objective at iteration n can be approximated as:

$$\text{Obj}^{(n)} = \sum_{i=1}^{m} \left[g_i f_n(\mathbf{x}_i) + \frac{1}{2} h_i f_n^2(\mathbf{x}_i) \right] + \Omega(f_n), \tag{3.22}$$

where g_i and h_i are the first- and second-order derivatives of the loss function with respect to the previous prediction $\hat{y}_i^{(n-1)}$. Assume a tree has T leaves indexed by j, each with weight w_j. Let $I_j = \{\, i \mid q(\mathbf{x}_i) = j \,\}$ denote the set of samples assigned to

leaf j. The optimal leaf weight that minimises $\text{Obj}^{(n)}$ is:

$$w_j^* = -\frac{\sum_{i \in I_j} g_i}{\sum_{i \in I_j} h_i + \gamma}, \tag{3.23}$$

where γ is the L^2 regularisation parameter used to control overfitting.

Gaussian Process Regression

GPR (Rasmussen & Williams, 2006) provides a probabilistic framework for non-linear regression, providing both predictions and uncertainty estimates. It can be interpreted as defining a prior distribution over functions, which is updated to a posterior distribution after observing the data. Formally, the unknown function $\mathbf{f}(\mathbf{x})$ is modelled as a random process:

$$\mathbf{f} = [f_1, f_2, ..., f_M] \sim \mathcal{GP}\big(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\big), \tag{3.24}$$

where $\mu(\mathbf{x})$ is the mean function and $k(\mathbf{x}, \mathbf{x}') = cov(\mathbf{f}(\mathbf{x}), \mathbf{f}(\mathbf{x}'))$ represents the covariance (kernel) function that encodes prior assumptions regarding smoothness and correlation structure.

For a training set with inputs $X = \{\mathbf{x}_i\}_{i=1}^M$ and outputs $\mathbf{y} = [y_1, \ldots, y_n]^\top$, a Gaussian likelihood is assumed:

$$y_i = f(\mathbf{x}_i) + \epsilon_i, \qquad \epsilon_i \sim \mathcal{N}(0, \sigma_n^2), \tag{3.25}$$

which leads to the joint distribution:

$$\mathbf{y} \sim \mathcal{N}\big(\mu(X), K(X, X) + \sigma_n^2 I\big), \tag{3.26}$$

where $K(X, X)$ is the covariance matrix with elements $[K(X, X)]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$.

For a new input $\mathbf{x}^*$, the joint prior of $\mathbf{y}$ and $f^* = f(\mathbf{x}^*)$ share a joint Gaussian distribution given by:

$$\begin{bmatrix} \mathbf{y} \\ f_* \end{bmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu(X) \\ \mu(\mathbf{x}_*) \end{pmatrix}, \begin{pmatrix} K(X, X) + \sigma_n^2 I & K(X, \mathbf{x}^*) \\ K(\mathbf{x}^*, X) & k(\mathbf{x}^*, \mathbf{x}^*) \end{pmatrix} \right). \tag{3.27}$$

Within the Bayesian formulation, the posterior distribution remains a Gaussian distribution with predictive mean and variance:

$$\overline{\mathbf{f}}^* = K(\mathbf{x}^*, X)\big[K(X, X) + \sigma_n^2 I\big]^{-1}\mathbf{y}, \tag{3.28}$$

$$\text{Var}(\mathbf{f}^*) = k(\mathbf{x}^*, \mathbf{x}^*) - K(\mathbf{x}^*, X)\big[K(X, X) + \sigma_n^2 I\big]^{-1}K(X, \mathbf{x}^*). \tag{3.29}$$

These expressions define both the model prediction and its associated uncertainty at unseen inputs.

The choice of covariance function $k(\mathbf{x}, \mathbf{x}')$ strongly influences model behaviour. This work employs the Matérn kernel with smoothness parameter $\nu > 0$ due to its flexibility in representing functions with varying degrees of roughness. The general Matérn form is given by:

$$k_{\mathrm{M}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \cdot \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu}\, r\right)^{\nu} K_\nu\left(\sqrt{2\nu}\, r\right), \tag{3.30}$$

where $K_\nu(\cdot)$ is the modified Bessel function of the second kind, and:

$$r = \sqrt{\sum_{d=1}^{D} \frac{(x_d - x'_d)^2}{\ell_d^2}}. \tag{3.31}$$

Here, D represents the number of input features, x_d and x'_d are their values for two samples, and ℓ_d represents the characteristic length scale associated with feature d. The parameters σ_f, σ_n, and ℓ_d are typically optimised by maximising the log marginal likelihood of the observed data.

Grey-box approaches

Grey-box modelling integrates prior physical knowledge (white-box) with pure data-driven techniques (black-box) to improve predictive performance, particularly in extrapolation regimes, where purely data-driven models often fail (Duarte et al., 2004; Lindskog & Ljung, 1994). Two primary architectures—serial and parallel structures—are commonly employed to fuse these components (Sævarsdóttir et al., 2005).

In a serial or cascaded architecture, the physical model functions as a feature extractor or pre-processor. The output of the physical model, $\hat{y}_{\mathrm{phys}}$, is fed as an input into the ML algorithm alongside the original inputs $\mathbf{x}$:

$$y = f_{\mathrm{ML}}(\mathbf{x}, \hat{y}_{\mathrm{phys}}) + \varepsilon. \tag{3.32}$$

Conversely, in a parallel architecture, the ML model is trained to compensate for the physical model's structural deficiencies. The physical model provides a baseline prediction $\hat{y}_{\mathrm{phys}}$, and the data-driven model learns the residual error (bias) $r(\mathbf{x})$ between the physical estimate and the observed data:

$$y = \hat{y}_{\mathrm{phys}}(\mathbf{x}) + \underbrace{f_{\mathrm{ML}}(\mathbf{x})}_{\text{Residual}} + \varepsilon. \tag{3.33}$$

This approach facilitates the model's retention of the generalisation capabilities of physical laws whilst correcting for the unmodelled dynamics or simplified assumptions inherent to the white-box formulation. This representation is further explained in Paper II.

3.3 Optimisation algorithms

Optimisation describes the mathematical process of identifying the best possible solution from a set of available alternatives, subject to given constraints. Each optimisation problem is defined by three elements: (1) an objective function that quantifies the quality of the solution, (2) decision variables that control the outcome of the objective function, and (3) constraint functions that limit the feasibility of the decision variables. An optimisation algorithm is necessary for efficiently navigating the set of feasible decision variables that satisfy both the objective and constraint functions. Two algorithms are used throughout the research underpinning this thesis: dynamic programming for discrete problems and Bayesian optimisation for continuous problems.

Dynamic Programming

Dynamic programming (DP) is an optimisation method for solving discrete sequential decision problems by decomposing them into a series of staged subproblems. This approach is founded upon Bellman's Principle of Optimality (Bellman, 1957), which states that an optimal policy has the property that, regardless of the initial state and decision, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision. Unlike heuristic methods such as GAs, which rely on stochastic search mechanisms and cannot guarantee optimality, DP guarantees a globally optimal solution within the discretised state-space by exhaustively evaluating all state-decision combinations at each stage. It should be noted that this optimality is relative to the discretised grid, and the fidelity of the solution to the true continuous optimum is limited by discretisation error (further elaborated in Paper V). This thesis employs a forward-propagating DP approach to determine the optimal sequence of decisions from a fixed initial state towards a fixed destination.

Let $j \in \{1, \dots, M\}$ index the discrete stages of the problem. The state of the system at the completion of stage j is denoted by $s_j \in \mathcal{S}$. The optimal accumulated-cost function, $J_j(s_j)$, represents the minimum cost required to reach state s_j from the initial condition. The value of this function is determined by the forward Bellman recursion:

$$J_j(s_j) = \min_{a \in \mathcal{A}} \big[J_{j-1}(s_{j-1}) + c_j(s_{j-1}, a) \big], \tag{3.34}$$

where a is a discrete control, $\mathcal{A}$ is the set of allowable control actions and $c_j(s_{j-1}, a)$ is the transition cost associated with moving from the preceding state s_{j-1} to the current state s_j. The system's evolution is governed by the state transition function given by the recursion:

$$s_j = h(s_{j-1}, a) = s_{j-1} + \delta(s_{j-1}, a), \tag{3.35}$$

where $\delta(s_{j-1}, a)$ represents the incremental change in the state (e.g., a distance covered or a time elapsed due to the control action a). To mitigate the computational complexity of exhaustive searches, two primary strategies are employed within the DP framework:

1. **State aggregation:** The state-space is discretised such that multiple trajectories arriving at the same state s_j are compared and only the trajectory with the lowest accumulated cost is retained. This ensures that the number of stored paths remains proportional to the density of the discrete state-space rather than growing exponentially with the number of stages.

2. **Search space pruning:** The feasibility of a state is evaluated at each stage against a set of constraints. A state is only considered for further propagation if it belongs to the feasible set $\mathcal{F}_j$:

$$s_j \in \mathcal{F}_j. \tag{3.36}$$

The feasibility set corresponds to each state s_j that is reachable from a previous state s_{j-1} through the control action a, that is:

$$\mathcal{F}_j = \{s_j \mid s_j = h(s_{j-1}, a), \quad s_{j-1} \in \mathcal{F}_{j-1}, \quad a \in \mathcal{A}\}. \tag{3.37}$$

Therefore, the algorithm focuses computational resources solely on regions of the state-space that can potentially satisfy the global boundary conditions (e.g., power limits, speed limits, time constraints).

The global optimal solution is identified by selecting the state that yields the minimum total cost whilst satisfying the terminal boundary conditions and tolerances:

$$J^* = \min\{J_M(s_M) \mid s_M \in \mathcal{T}\}, \tag{3.38}$$

where $\mathcal{T}$ represents the set of terminal states that satisfy the required convergence criteria. The DP algorithm converges exhaustively within the defined discretised state-space. As long as the terminal constraints are reachable, the algorithm is guaranteed to identify the global optimum relative to the state-space resolution. Convergence speed is primarily governed by the number of stages, the resolution of

the discrete control variable, and the computational cost of evaluating the transition cost function.

Bayesian Optimisation

Bayesian optimisation (BO) provides a sample-efficient framework for the global optimisation of computationally expensive black-box objective functions (Mockus et al., 1978; Snoek et al., 2012). BO is particularly effective when analytical gradient information is unavailable, as it leverages a probabilistic surrogate to guide the search process. The unknown objective function $f(\xi)$, where ξ denotes a vector of continuous decision variables, is modelled as a Gaussian process (GP):

$$f(\xi) \sim \mathcal{GP}\big(\mu(\xi), k(\xi, \xi')\big), \tag{3.39}$$

where $\mu(\xi)$ and $k(\xi, \xi')$ represent the mean and covariance functions, respectively. The GP posterior provides both a prediction and a measure of uncertainty, $\sigma(\xi)$, for any unobserved point, enabling a principled trade-off between exploration and exploitation. The objective of this work is to find the global minimum of the function. Consequently, a new evaluation point is selected at each iteration by minimising the lower confidence bound (LCB) acquisition function:

$$\alpha_{\mathrm{LCB}}(\xi) = \mu(\xi) - \beta \cdot \sigma(\xi), \tag{3.40}$$

where $\beta > 0$ represents a hyperparameter that controls the balance between exploitation (low mean) and exploration (high uncertainty). The next evaluation point, ξ_{k+1}, is selected by solving the sub-optimisation problem of finding the global minimum of the acquisition function:

$$\xi_{k+1} = \arg\min_{\xi} \alpha_{\mathrm{LCB}}(\xi). \tag{3.41}$$

This iterative process continues until either the convergence criteria are met or the evaluation budget is exhausted. Unlike DP, which operates over a discretised state-space, BO handles continuous variables directly. It achieves high sample efficiency by balancing exploration of uncertain regions with exploitation of known minima using the acquisition function. Whilst the number of required function evaluations is low, the computational cost per iteration increases as the surrogate model incorporates more data, leading to a trade-off between search precision and execution time.

CHAPTER 4

Data-driven modelling and optimisation framework

This chapter presents the methodological framework developed across the five studies that comprise this thesis. The research progresses from data-driven modelling of ship propulsion performance to the optimisation of power allocation and energy efficiency under realistic operational and environmental conditions. The methods used combine the fundamental principles of naval architecture with modern data-driven regression, probabilistic modelling, and numerical optimisation techniques, all of which are defined in Chapter 3.

Figure 4.1 summarises the relationships among the different research areas, datasets, specific algorithms, and their corresponding publications. Each publication contributes a distinct methodological element; collectively, these elements define a hierarchical framework for ship performance modelling and voyage optimisation:

- **Paper I** compares different regression models and establishes the foundations of the data-driven performance modelling used throughout the research.
- **Paper II** develops a GPR model for added resistance in waves and integrates it into a grey-box framework.
- **Paper III** solves for optimal power using DP, and proposes the MS-PELT algorithm for adaptive voyage segmentation.
- **Paper IV** integrates data-driven models into a BO framework for decision support.

- **Paper V** combines discrete DP optimisation with continuous BO for solution refinement, establishing a hybrid approach valuable for high-resolution voyage planning.

	Case study	Data volume	Modelling algorithm	Optimisation algorithm	Energy savings
Paper I	Doble ended ferry	1 Year (sensors)	OLS, ANN, XG Boost		Simulation: 35%
Paper II	Chemical tanker	2096 (model-tests) 3 Years (sensors)	GPR, Grey-Box models		Potential: 8%
Paper III	Chemical tanker	3 Years (sensors)	XG Boost	Forward DP	Simulation:<15%
Paper IV	Doble ended ferry	1 Year (sensors)	XG Boost	BO	Simulation: 40% Experiment: 15%
Paper V	Chemical tanker	3 Years (sensors)	XG Boost	Forward DP + BO	Simulation:<15%

Figure 4.1: Relationships between the different research subjects.

The data used in this research comprise full-scale measurements collected from the ships' on-board continuous monitoring systems (Papers I, III, IV and V) and public model-scale tests for added wave resistance (Paper II). These data are detailed further in Chapter 5. In addition, the energy savings from each paper are further elaborated in Chapter 6.

4.1 Data-driven modelling

This thesis employs a multi-tiered modelling approach tailored to the specific requirements of each case study. In general performance prediction, when large volumes of operational data are available, gradient boosting is used for its computational efficiency and its ability to capture non-linear interactions. The use of GPR and grey-box models addresses the specific problem of predicting added resistance in waves, where experimental data are sparse and uncertainty quantification is critical.

4.1.1 Data pre-processing pipeline

Before using the data for model training, a pre-processing pipeline was applied to ensure data quality and consistency. This process involved first identifying steady-state operating conditions using first-derivative thresholds for propulsion power and ship speed, with the transient and manoeuvring phases excluded from the analysis.

Following the work of Lang et al. (2022b), time-series data were downsampled from 15-second (ferry) or 1-minute (tanker) averages to 10-minute averages to suppress measurement noise and reduce temporal autocorrelation. Residual fluctuations in continuous variables were smoothed using a second-order Savitzky–Golay filter. To further ensure data quality, outliers were removed iteratively by applying a three-sigma criterion for each signal:

$$\chi_i \in \left[\bar{\chi}_i - 3\sigma_{\chi_i},\, \bar{\chi}_i + 3\sigma_{\chi_i}\right], \tag{4.1}$$

where $\bar{\chi}_i$ and σ_{χ_i} denote the mean and standard deviation of variable χ_i, respectively.

In parallel to the processing of ship sensor data, metocean conditions W, as defined in Eq. 3.14, were retrieved from hindcast databases. Wind and wave data were obtained from the ERA5 reanalysis, and surface currents were extracted from the Copernicus/Mercator Global Ocean Physics products. Tri-linear interpolation enabled spatial and temporal alignment of the environmental variables with ship positions and timestamps. Relative angles between environmental and ship headings were then computed as:

$$\alpha = \theta_{\text{env}} - \theta_{\text{ship}}, \tag{4.2}$$

and expressed in degrees within the range $[-180°, 180°]$. Finally, all variables were normalised to the range $[0, 1]$ using min–max scaling:

$$\hat{\chi}_i = \frac{\chi_i - \min(\chi_i)}{\max(\chi_i) - \min(\chi_i)}. \tag{4.3}$$

4.1.2 XGBoost for speed and fuel prediction

XGBoost was selected as the primary regression algorithm for the full-scale ship performance models in Papers I,III, IV and V because of its robust predictive performance and its fast model training (Lang et al., 2022b). As an ensemble learning method, XGBoost constructs a strong predictive model by iteratively combining multiple weak learners (decision trees). Based on the energy flow equations derived in Chapter 3 (see Eq. 3.11 for fuel mass flow and Eq. 3.13 for the power relationship), an input/output simplification for fuel consumption and speed was achieved by mapping the relevant features of the energy conversion process. The functional forms for fuel consumption (m_{fuel}) and speed over ground (V_G) are defined as:

$$\begin{aligned} m_{\text{fuel}} &= f(P_B, V_G, W), \\ V_G &= g(P_B, W), \end{aligned} \tag{4.4}$$

where P_B is the brake power and W represents the vector of environmental conditions (wind, wave, current) defined in Eq. 3.14. To ensure generalisability and prevent over-

fitting, the hyperparameters of the XGBoost models were not manually selected and were instead optimised using a BO strategy. The objective function of the tuning process was to minimise the mean squared error (MSE) on a held-out validation set. This approach allows the model to capture the complex non-linear relationships among propulsion power, speed, and weather conditions without memorising noise in the training data. The specific hyperparameters included in the optimisation search space are listed in Table 4.1.

Table 4.1: XGBoost hyperparameters tuned via Bayesian optimisation.

Hyperparameter	**Description**
Learning rate (`eta`)	Step size shrinkage used in updates to prevent overfitting.
Max depth	Maximum depth of a tree.
Number of estimators	Number of boosting rounds (trees).
Min child weight	Minimum sum of instance weight needed in a child.
Gamma	Minimum loss reduction to make a further partition.
Subsample	Subsample ratio of the training instances.
Colsample By Tree	Subsample ratio of columns when constructing each tree.
Alpha (L1)	L1 regularisation term on weights.
Lambda (L2)	L2 regularisation term on weights.

4.1.3 GPR for added wave resistance

In Paper II, the focus shifts to the specific hydrodynamic problem of predicting the added wave resistance coefficient, C_{AW}. Unlike the large operational datasets used for the ferry and tanker, experimental data for added resistance are relatively sparse and highly correlated within specific test series. To address this, a GPR model was developed. This probabilistic formulation models the unknown function that maps ship parameters to added resistance as a random process:

$$f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')), \tag{4.5}$$

where $\mu(\mathbf{x})$ represents the mean function (assumed to be zero prior to observation) and $k(\mathbf{x}, \mathbf{x}')$ represents the covariance kernel function. The GPR provides not only a point prediction for C_{AW} but also a variance estimate that quantifies the uncertainty of the prediction in regions of the input space where experimental data are lacking.

The choice of kernel function dictates the model's smoothness and periodicity assumptions. In Paper II, the Matérn kernel (with smoothness parameter $\nu = 2.5$) is selected, providing flexibility in modelling differentiable but not infinitely smooth physical processes, which is appropriate for hydrodynamic coefficients that may ex-

hibit sharp variations near resonance frequencies (see Eq. 3.8). Correspondingly, the inputs $\mathbf{x}$ for the GPR model comprise seven non-dimensional parameters derived from the ship's geometry and operational state:

$$C_{AW} = f\left(\frac{\lambda}{L}, F_n, C_B, \frac{B}{T}, \frac{L}{B}, \frac{K_{yy}}{L}, \frac{L_E}{L}\right), \tag{4.6}$$

where λ/L represents the wavelength-to-ship-length ratio, F_n is the Froude number, and L_E/L represents the effective entrance length ratio, a critical parameter that the research presented in Paper II uses to capture the bow's interaction with incident waves.

To assess the model's capacity to generalise beyond the training data, two distinct data partitioning strategies were employed, as Fig. 4.2 illustrates. Strategy 1: Random Split (Interpolation Test) pooled and randomly shuffled the entire dataset of 2096 samples, with 10% held out for testing. This strategy evaluates the model's ability to interpolate resistance values for known ship types at unobserved frequencies or speeds. Strategy 2: Ship-Type Split (Extrapolation Test) excluded from the training set entire experimental datasets corresponding to specific hull geometries, using them exclusively for testing. This strategy simulates a real-world design scenario, testing the model's ability to predict the resistance curve for a novel ship design that does not exist in the database.

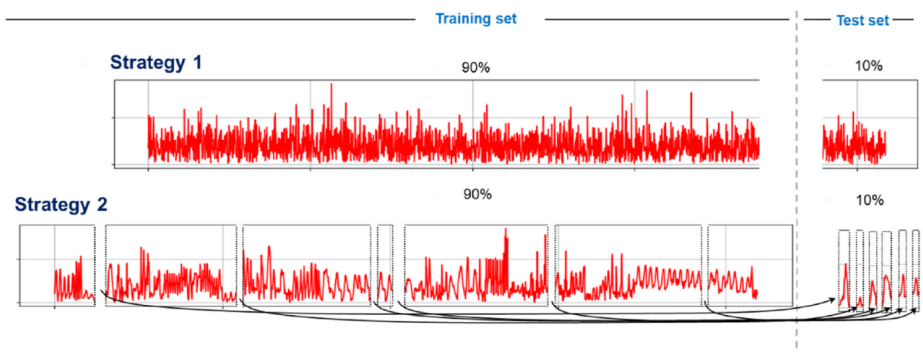

Figure 4.2: Data partitioning strategies used to validate the GPR model.

4.1.4 Grey-box modelling for fuel consumption

Paper II implements a grey-box model (GBM) framework to predict the ship's total fuel consumption. This approach addresses the limitations of pure black-box models, which lack physical constraints, and of pure white-box models, which often underestimate consumption due to unmodelled losses such as fouling. The GBM employs

a parallel architecture in which a baseline rooted in physics is corrected by a data-driven residual model following the pipeline visualised in Fig. 4.3.

The physical component of the model, denoted as the white box (WB), estimates the fuel consumption by first calculating the total resistance R_T acting on the hull. This is defined by Eq. 3.1 as the sum of three primary components:

$$R_T = R_{\text{Calm}} + R_{\text{Wind}} + R_{\text{Wave}}^{\text{GPR}}, \tag{4.7}$$

where $R_{\text{Wave}}^{\text{GPR}}$ is calculated by integrating the transfer function predicted by the GPR model (C_{AW}) over the wave energy spectrum $S(\omega)$:

$$R_{\text{Wave}}^{\text{GPR}} = 2 \int_0^{\infty} S(\omega \mid H_s, T_z) \frac{C_{aw}^{\text{GPR}}(\omega) \cdot \rho_w g \zeta_a^2 B^2 / L}{\zeta_a^2} \, d\omega. \tag{4.8}$$

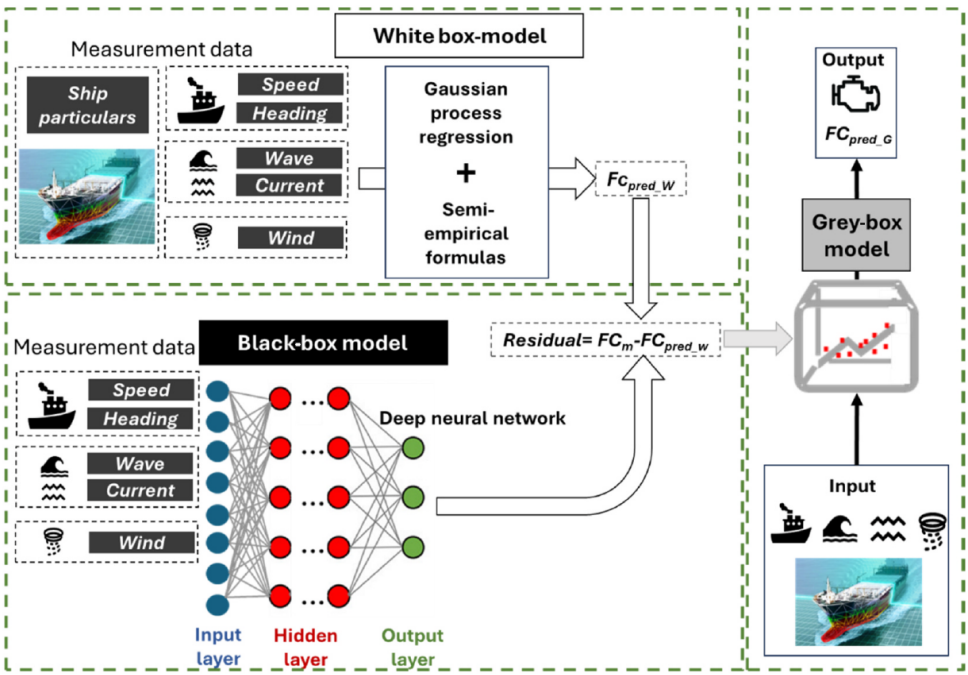

Figure 4.3: GBM pipeline.

The total resistance R_T is transformed into effective power P_E (Eq. 3.9), then into brake power P_B (Eq. 3.10), and finally into the white-box fuel consumption estimate M_{fuel}^{WBM} using the engine's SFOC curve.

However, Paper II reports that the white-box model (WBM) consistently underestimated actual fuel consumption, likely due to unmodelled factors such as hull and propeller fouling, and engine wear. To address this, a deep neural network (DNN) was employed to model the residual error r. The residual target for training is defined as:

$$r = M_{fuel}^{\text{Measured}} - M_{fuel}^{\text{WBM}}. \tag{4.9}$$

where r is the residual error, $M_{fuel}^{Measured}$ is the measured fuel consumption, and M_{fuel}^{WBM} is the fuel consumption predicted by the WBM. The DNN was then trained to map the operational state vector $\mathbf{x}$ (speed, draught, weather) to this residual:

$$\hat{r} = f_{DNN}(P_s, V_{water}, \alpha_{ship}, T, SFOC, H_s, T_z, V_{wind}, \alpha_{wind}), \tag{4.10}$$

where $\hat{r}$ is the predicted residual, P is the engine shaft power, V_{water} is the speed through water, α_{ship} is the heading, T_m is the mean draught, $SFOC$ is the specific fuel oil consumption, H_s is the significant wave height, T_m is the mean wave period, V_{wind} is the wind speed, and α_{wind} is the wind direction.

The network architecture comprises five hidden layers of 48 neurons each, using ReLUs to capture non-linear error dynamics. With both components defined, the final prediction of the GBM (M_{fuel}^{GBM}) is the sum of the physics-based estimate and the data-driven residual correction:

$$M_{fuel}^{\text{GBM}} = M_{fuel}^{\text{WBM}} + \hat{r}. \tag{4.11}$$

This additive structure ensures that the model retains the generalisation capabilities of physical laws (essential for extrapolation beyond the training set) whilst leveraging ML to capture the high-frequency non-linearities and system degradation that physical formulas alone cannot represent. The hyperparameters used for the DNN are summarised in Table 4.2.

Table 4.2: DNN hyperparameters.

Hyperparameter	**Value/Setting**
Number of hidden layers	5
Neurons per hidden layer	48
Activation function	ReLU
Optimiser	Adam
Learning rate	0.002
Loss function	MSE

4.2 Power allocation optimisation

Power allocation optimisation refers to the process of planning a ship's power usage before a journey to minimise operational costs (e.g., fuel consumption and crew salaries), maximise sailing efficiency, and meet operational and contractual requirements. It can be described mathematically as a solution to the optimisation problem from Eq. 4.12:

$$\begin{aligned}
&\min_{P(t)} J = \int_{t_0}^{t_f} \Big(C_f(V(t), P(t)) + C_o(t) + C_p(D(t)) \Big)\, dt \\
&\text{s.t.} \\
&V_{\min} \leq V(t) \leq V_{\max}, \\
&P_{\min} \leq P(t) \leq P_{\max}, \\
&(x(t), y(t)) \in \mathcal{R}, \\
&t_f \in [t_{\text{earliest}}, t_{\text{latest}}], \\
&E(t) \leq E_{\max},
\end{aligned} \tag{4.12}$$

where C_f is the fuel cost, C_o is the cost to run the operation, C_p is the cost of delays, t_0 is the departure time, t_f is the arrival time, V is the sailing speed, P is the engine power, $\mathcal{R}$ is the sailing route, and E represents the emissions. This thesis is concerned with minimising C_f under the ETA constraint whilst sailing on a fixed $\mathcal{R}$. As such, Eq. 4.12 simplifies to:

$$\begin{aligned}
&\min_{P(t)} J = \int_{t_0}^{t_f} \Big(C_f(V(t), P(t))) \Big)\, dt \\
&\text{s.t.} \\
&V_{\min} \leq V(t) \leq V_{\max}, \\
&P_{\min} \leq P(t) \leq P_{\max}, \\
&(x(t), y(t)) \in \mathcal{R}, \\
&t_f \in [t_{\text{earliest}}, t_{\text{latest}}].
\end{aligned} \tag{4.13}$$

Solving this problem requires the creation of a digital copy of the ship, simulating how it sails along $\mathcal{R}$. The digital copy, which corresponds to the data-driven surrogate models, serves as a virtual representation of the ship's energy system, allowing for the rapid evaluation of the ship's performance under different power settings.

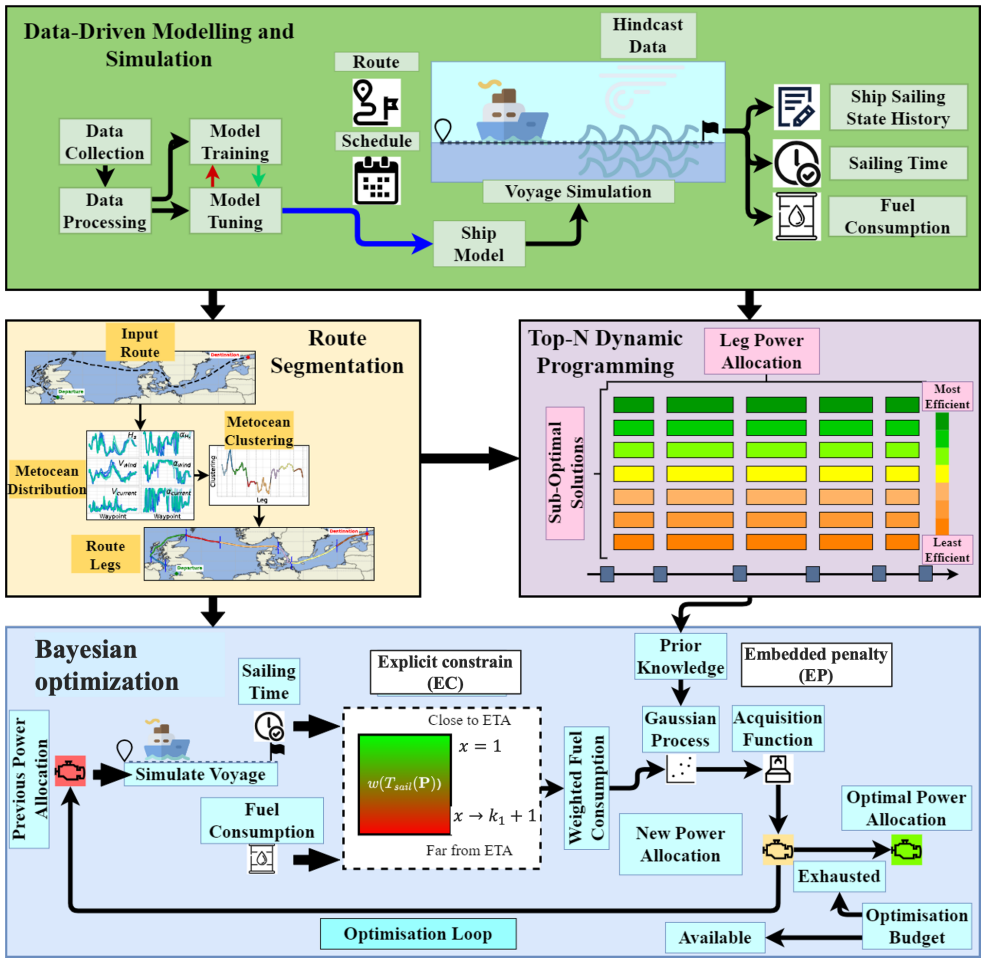

Figure 4.4: The proposed optimisation framework.

Figure 4.4 summarises the optimisation framework proposed in this thesis. It represents a hybrid discrete-to-continuous multi-stage algorithm for refining the optimal power allocation for minimising fuel consumption during a voyage. The algorithm stages correspond to:

1. **Data-driven modelling and performance prediction**: At the core of the framework are data-driven surrogate models (XGBoost-based), serving as the ship's digital copy. These models predict the fuel consumption rate and speed over ground based on control parameters, ship loading, and metocean conditions, enabling the determination of ship fuel consumption under various

operating conditions.

2. **Route segmentation**: The route is segmented using the MS-PELT algorithm to reduce the number of control inputs along the voyage.

3. **Parallel coupling dynamic programming**: The PCDP algorithm is used to determine the global discrete solution of the optimisation problem.

The previous stages correspond to the optimisation framework presented in Paper III. An additional step is introduced in Paper V with the discrete-to-continuous optimisation refinement:

4. **Bayesian optimisation**: Building upon the discrete solution, the framework employs BO to explore the continuous state-space. The results from the PCDP phase are used as prior knowledge to seed the BO search, significantly reducing the number of iterations required to converge on the global minimum. This discrete-to-continuous transition enables the system to refine power-allocation set points with high precision.

5. **Constraint handling**: A smooth exponential penalty function is integrated into the objective function to maintain operational feasibility. This approach enforces the ETA requirements by penalising candidate solutions that approach or violate these boundaries, thereby steering the optimiser towards a valid power configuration.

4.2.1 Voyage simulation

Due to the exclusive focus on minimising fuel consumption, the ship is simplified to a point mass where only two-degrees-of-freedom (2-DOF) are considered, i.e., surge and yaw. This approach allows the optimisation algorithm to focus on the total voyage fuel consumption calculation without the computational overhead of a full six-degree-of-freedom (6-DOF) state-space model. The simplified voyage is therefore described by a general state vector $\mathbf{S}$ that includes its coordinates (latitude and longitude) and timestamp. Furthermore, the voyage is divided into legs to simplify the optimisation problem. When the ship has reached the k-th waypoint within the j-th leg, it is described as:

$$\mathbf{S}_{j,k} = [x_{j,k}, y_{j,k}, t_{j,k}], \tag{4.14}$$

where $x_{j,k}$ and $y_{j,k}$ are the spatial coordinates and $t_{j,k} \geq t_0$ is the timestamp. An initial problem arises in that real sailing occurs in real time, whilst computers operate in discrete time. Therefore, the digital description of the voyage requires a zero-order

hold, meaning the ship's state is only updated at the next waypoint visited:

$$\mathbf{S}_{j,k}(t) = \mathbf{S}_{j,k}, \quad \text{for } t \in [t_{j,k}, t_{j,k+1}). \tag{4.15}$$

It would be convenient for state times $t_{j,k}$ to correspond to full hours to simplify the interpolation of the metocean conditions, as this would reduce the order of the interpolation from tri-linear to bi-linear due to the state time matching that of the hindcast data. Because the resolution of the hindcast database is 1h, all the waypoints are collocated as:

$$t_{j+1,k} = \lfloor t_{j,k} + 1h \rfloor. \tag{4.16}$$

This simplification comes at the cost of complicating the identification of the ship's position in $\mathcal{R}$. The problem of finding the sailing time between two adjacent waypoints $(x_{j,k}, y_{j,k})$ and $(x_{j,k+1}, y_{j,k+1})$ is straightforward because the Haversine distance provides the distance directly, as illustrated in Fig. 4.5:

$$\Delta d_{j,k} = 2 \cdot R_{Earth} \cdot \arcsin\left(\sqrt{\sin^2\left(\frac{y_{j,k+1} - y_{j,k}}{2}\right) + \cos(y_{j,k}) \cdot \cos(y_{j,k+1}) \cdot \sin^2\left(\frac{x_{j,k+1} - x_{j,k}}{2}\right)}\right), \tag{4.17}$$

where R_{Earth} is the average radius of the Earth, and the sailing time is:

$$\begin{aligned} \Delta t_{j,k} &= \frac{\Delta d_{j,k}}{V_{j,k}} \\ t_{j,k+1} &= t_{j,k} + \Delta t_{j,k}. \end{aligned} \tag{4.18}$$

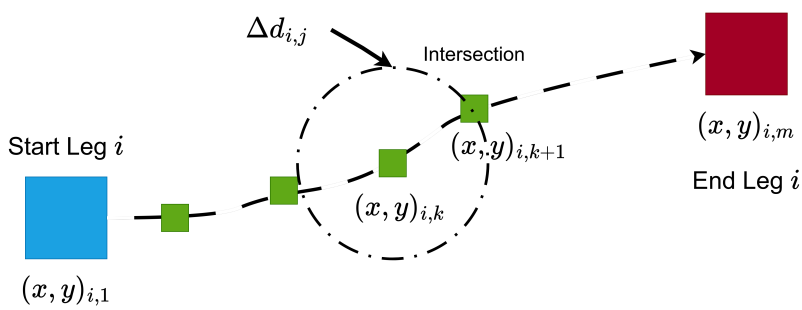

Figure 4.5: Waypoint generation.

The inverse problem of finding the position $(x_{j,k+1}, y_{j,k+1})$ after the ship has sailed at speed $V_{j,k}$ for duration $\Delta t_{j,k}$ is more difficult because it requires a function that follows the curvature of $\mathcal{R}$, takes inputs $(x_{j,k}, y_{j,k})$ and $\Delta d_{j,k}$, and returns the new position. No such function exists for a general $\mathcal{R}$ in a closed analytical form. To address this, the problem is approximated using a linear parametrisation of a finely discretised version of $\mathcal{R}$. This approach replaces the continuous curve with a series of discrete waypoints, approximating the distance between consecutive waypoints by the Euclidean distance. The solution involves approximating a starting point $(\hat{x}_{j,k}, \hat{y}_{j,k})$ within the discrete sequence and iteratively accumulating distances along the path until the ship has covered at least a total distance of $\Delta d_{j,k}$. A correction factor can be used to adjust the resulting position $(\hat{x}_{j,k+1}, \hat{y}_{j,k+1})$.

4.2.2 Objective and constraint functions

The objective function from Eq. 4.13 is also simplified by the zero-order hold of Eq. 4.15. The discretised objective function is given by:

$$M_{fuel} = \sum_{j=1}^{n} f_j\left(P_j, W_{j,1:m_j}\right) = \sum_{j=1}^{n_j} \sum_{k=1}^{m_j-1} m_{\text{fuel}}\left(P_j, W_{j,k}\right) \cdot \Delta t_{j,k}, \tag{4.19}$$

where n is the total number of legs the $\mathcal{R}$ was segmented into, m_j denotes the total number of waypoints within the j-th leg, and $W_{j,k}$ represents the metocean conditions at state $\mathbf{S}_{j,k}$, given by:

$$W_{j,k} = \big[H_{s(j,k)}, \alpha_{H_s(j,k)}, T_{z(j,k)}, V_{\text{wind}(j,k)}, \alpha_{\text{wind}(j,k)}, V_{\text{current}(j,k)}, \alpha_{\text{current}(j,k)}\big] \tag{4.20}$$

The optimisation constraints ensure that power setting P_j and speed V are within operational limits, and the total travel time aligns with the ETA:

$$C(P, W) = \begin{cases} P_{\min} \leq P_j \leq P_{\max}, \\ V_{\min} \leq V(P_j, W_{j,k}) \leq V_{\max}, \\ 0.99 \leq \dfrac{\sum_{j=1}^{n} \sum_{k=1}^{m_j-1} \Delta t_{j,k}}{\text{ETA}} \leq 1.01. \end{cases} \tag{4.21}$$

The power allocation $\mathbf{P}$ represents a set of valid engine power allocation values for the voyage:

$$\mathbf{P} = [P_1, P_2, \ldots, P_n]. \tag{4.22}$$

4.2.3 Parallel coupling dynamic programming

To solve the voyage power allocation problem, the general DP framework introduced in section 3.3 is adapted for the SSS minimal fuel consumption problem. The generic DP components are mapped directly to the physical voyage parameters defined in sections 4.2.1 and 4.2.2:

- **Stage** (j) $\rightarrow$ Voyage leg.
- **State** (s_j) $\rightarrow$ Cumulative sailing time t_j at the end of leg j.
- **Control action** (a) $\rightarrow$ Discrete engine power setting $P_j \in \mathbf{\Gamma}$.
- **Accumulated cost** (J_j) $\rightarrow$ Cumulative fuel consumption $F_j(t_j)$.

Since each leg has fixed departure and arrival coordinates—$(x_{1,j}, y_{1,j})$ and $(x_{m,j}, y_{m,j})$, respectively—the state vector from Eq. 4.14 reduces to only the cumulative sailing time t_j. The state transition function from Eq. 3.35 then becomes:

$$t_j = t_{j-1} + \Delta t_j(P_j), \tag{4.23}$$

where $\Delta t_j(P_j)$ is the time required to complete leg j under power setting P_j.

The cost function represents the accumulated fuel consumption over leg j as defined in Eq. 4.19. The optimisation objective is to determine the power setting P_j for each leg that minimises total fuel consumption whilst satisfying the ETA constraint:

$$M_j\left(P_j, W_{j,1:m_j}\right) = \sum_{k=1}^{m_j-1} m_{\text{fuel}}\left(P_j, W_{j,k}\right) \cdot \Delta t_{j,k}. \tag{4.24}$$

In a standard DP formulation, the legs are sequentially coupled (Fig. 4.6), meaning the departure time of leg j is strictly determined by the arrival time of leg $j-1$. Whilst exact, this coupling enforces sequential computation and limits parallelisation.

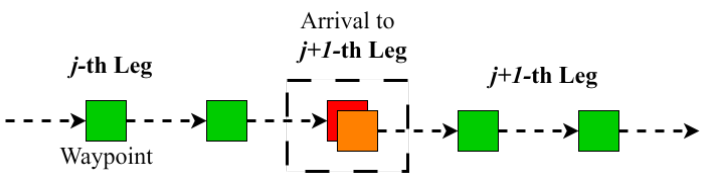

Figure 4.6: Exact coupling between consecutive legs.

To overcome this limitation, a parallel scenario-based DP approach is proposed, in which the legs are decoupled and each leg is simulated independently within a dynamically pruned time interval (Fig. 4.7). The pruning is based on a reference nominal

sailing time, computed as the cumulative time at the end of leg j, assuming a constant average speed V_{average} that satisfies the ETA:

$$t_{j,m}^{(0)} = \frac{d_{1\to j}}{V_{\text{average}}} = \frac{d_{1\to j}}{d_{sail}} \cdot t_{sail} = t_{j+1,0}^{(0)}, \tag{4.25}$$

where $t_{j,m}^{(0)}$ and $d_{1\to j}$ represent the nominal sailing time and sailing distance for leg j, d_{sail} is the total sailing distance, and t_{sail} is the target sailing time. Note that the nominal sailing time for leg j also corresponds to the departure time of leg $j+1$, which is given by $t_{j+1,0}^{(0)}$.

Around this nominal time, a set of discrete scenarios $\Xi \in \{1, 2, \ldots, K\}$ is generated for each leg, spanning a symmetric time window of width ΔT_{j+1}:

$$t_{j+1,0}^{(\Xi)} \in \left[t_{j+1,0}^{(0)} - \frac{\Delta T_{j+1}}{2},\ t_{j+1,0}^{(0)} + \frac{\Delta T_{j+1}}{2} \right]. \tag{4.26}$$

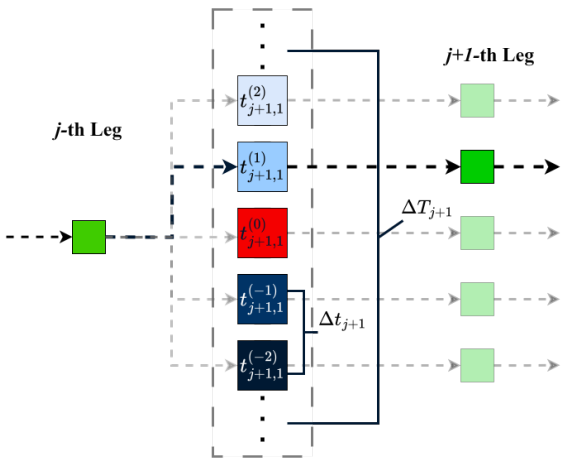

Figure 4.7: Parallel scenarios for optimising power allocation across voyage legs.

To ensure global feasibility across the full voyage, the time window is bounded by a percentage threshold ϵ_t applied symmetrically around the nominal time:

$$\begin{aligned} t_{\min,j} &= t_{j,m}^{(0)} \cdot (1 - \epsilon_t), \\ t_{\max,j} &= t_{j,m}^{(0)} \cdot (1 + \epsilon_t), \end{aligned} \tag{4.27}$$

where ϵ_t controls the width of the admissible time corridor (e.g., $\epsilon_t = 0.25$ corresponds to a $\pm 25\%$ deviation from the nominal time). The feasible set $\mathcal{F}_j$ from Eq. 3.37 is

therefore formally defined as:

$$\mathcal{F}_j = \{t_j \mid t_{\min,j} \leq t_j \leq t_{\max,j}\}. \tag{4.28}$$

At the final leg n, these bounds are tightened to a narrow tolerance (e.g., $\pm 1\%$) to strictly enforce the ETA constraint.

Within the admissible time window, each scenario Ξ is paired with all power settings from the discrete set $\mathbf{\Gamma} = [\Gamma_1, \Gamma_2, \ldots, \Gamma_r]$, spanning the range $P_{\min} \leq P \leq P_{\max}$. The transition function and fuel cost for each scenario–power combination are pre-computed and stored in a multi-dimensional graph, as illustrated in Fig. 4.8. Each node in this graph corresponds to the sailing time $t_j^{\Xi}(P_j)$ of a pre-simulated sub-voyage under power setting P_j and scenario Ξ, and each edge represents the time elapsed between consecutive legs.

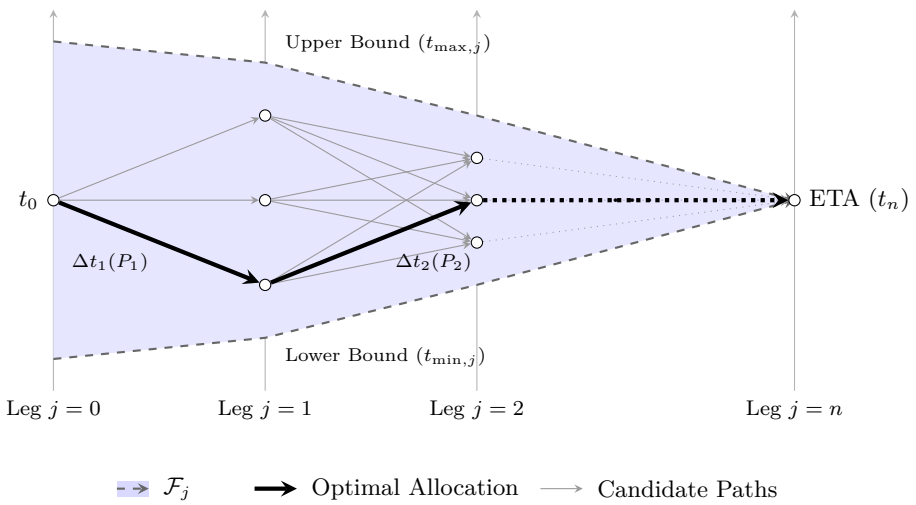

Figure 4.8: PCDP search space.

The DP optimisation then minimises total fuel consumption across all valid scenarios using the recursive Bellman equation:

$$F_j = \min_{P_j} \left(M_j(P_j, W_{j,1:m_j}) + F_{j-1} \right), \tag{4.29}$$

where F_j represents the cumulative fuel consumption up to leg j and $M_j(P_j, W_{j,1:m_j})$ represents the fuel consumption of leg j under power setting P_j and metocean conditions $W_{j,1:m_j}$. State aggregation is enforced throughout the forward search: when multiple candidate trajectories arrive at the same discretised time state t_j, only the trajectory with the minimum cumulative fuel consumption is retained as leg $j+1$.

4.3 Voyage segmentation

The previous sections assumed that $\mathcal{R}$ is divided into n legs. Segmenting the voyage is beneficial: although timely and a priori solutions for optimal instantaneous values of $P(t)$ are not possible, the optimal value for a finite number of segments can be determined. This section compares two data-driven segmentation methods for generating legs based on the environmental conditions expected along $\mathcal{R}$ during sailing between t_0 and t_f, whilst accounting for other operational requirements.

Metocean conditions significantly impact a ship's fuel consumption and sailing speed. Standard sailing strategies voluntarily reduce V in adverse metocean conditions only to later catch up with sailing. The objective is to initially determine the optimal segmentation based on metocean conditions before using the DP solver to guide navigation. This research therefore proposes a segmentation approach to determining the optimal number of legs, namely, the metocean score-based pruned exact linear time (MS-PELT) algorithm. For benchmarking purposes, the MS-PELT algorithm is compared with a multivariate time-series clustering method, Toeplitz inverse covariance-based clustering (TICC).

4.3.1 MS-PELT algorithm

The MS-PELT algorithm comprises four steps, as illustrated in Fig. 4.9. The first step involves calculating the target average speed V_{average} required to meet the ETA based on the voyage distance. Next, Monte Carlo simulations generate multiple reference voyages featuring different speed profiles that all meet the average speed requirement within a specified speed range $[V_{\min}, V_{\max}]$.

For each reference voyage, the metocean conditions, including wind, wave, and current characteristics, along $\mathcal{R}$ are obtained at each waypoint. A metocean score is computed for each waypoint, capturing the combined effects of environmental factors. The metocean score at the k-th waypoint for the q-th reference voyage is defined as:

$$\begin{aligned} \text{MS}_{q,k} = &\ \beta\left(\alpha_{H_s(q,k)}\right) \cdot \iota\left(H_{s(q,k)}\right) + \\ &\ \beta\left(\alpha_{\text{wind}(q,k)}\right) \cdot \iota\left(V_{\text{wind}(q,k)}\right) + \\ &\ \beta\left(\alpha_{\text{current}(q,k)}\right) \cdot \iota\left(V_{\text{current}(q,k)}\right), \end{aligned} \tag{4.30}$$

where H_s is the significant wave height, V_{wind} and V_{current} represent the wind and current speeds, respectively, and α_{H_s}, α_{wind}, and α_{current} each represent the respective direction factors. Each component reflects the relative impact of each metocean variable at the waypoint. The functions $\beta(\cdot)$ and $\iota(\cdot)$ correspond to the direction (Fig. 4.10) and intensity (Fig. 4.11) scores, respectively.

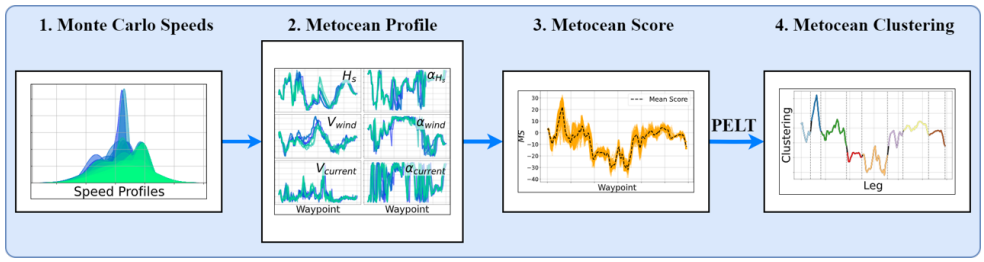

Figure 4.9: Workflow of the MS-PELT voyage segmentation method.

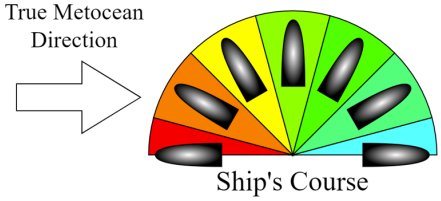

Metocean direction (°)	β_i
165-180	-5
135-165	-3
105-135	-1
75-105	0
45-75	1
15-45	3
0-15	5

Figure 4.10: Metocean direction score $\beta(\alpha)$.

V_{wind} (m/s)	0 - 0.2	0.2 - 1.5	1.5 - 3.3	3.3 - 5.4	5.4 - 7.9	7.9 - 10.7	10.7 - 13.8	13.8 - 17.1	17.1 - 20.7	20.7 - 24.4	24.4 - 28.4	28.4 - 32.6	32.6 +
Beaufort scale (ι_{wind})	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**
H_s (m)	0 - 0.1	0.1 - 0.5	0.5 - 1.25	1.25 - 2.5	2.5 - 4	4 - 6	6 - 9	9 - 14	14+				
Douglas scale (ι_{H_s})	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**				
$V_{current}$ (m/s)	0 - 0.1	0.1 - 0.2	0.2 - 0.5	0.5 - 1	1 - 1.5	1.5 +							
Current scale ($\iota_{current}$)	**0**	**2**	**4**	**6**	**8**	**10**							

Figure 4.11: Metocean intensity score $\iota(H_s)$.

The overall metocean score for waypoint k is obtained by averaging the metocean scores across all reference voyages:

$$\mathrm{MS}_k = \frac{1}{Q} \sum_{q=1}^{Q} \mathrm{MS}_{q,k}, \tag{4.31}$$

where Q is the total number of reference voyages.

The MS-PELT algorithm then segments $\mathcal{R}$ by identifying change points in the metocean scores, optimising the balance between fitting accuracy and complexity using the following objective function:

$$\arg\min_{\{\tau_j\}} \left(\sum_{j=1}^{b+1} \phi\left(\mathrm{MS}_{\tau_{j-1}+1:\tau_j}\right) + \gamma b \right), \tag{4.32}$$

where ϕ represents the cost function of each segment, γ is a penalty factor to prevent overfitting, b is the number of change points, and τ_j represents the change points detected by the algorithm.

4.3.2 TICC algorithm

The TICC algorithm (Hallac et al., 2017) is used as a comparison voyage segmentation method. Unlike MS-PELT, which scores metocean conditions based on ensemble values, TICC clusters multivariate time series subsequences of metocean conditions directly. For each waypoint, TICC generates an average matrix of metocean variables over all reference voyages:

$$\mathbf{W} = \begin{bmatrix} \overline{H_s}(1) & \overline{H_s}(2) & \dots & \overline{H_s}(a) \\ \overline{\alpha_{H_s}}(1) & \overline{\alpha_{H_s}}(2) & \dots & \overline{\alpha_{H_s}}(a) \\ \vdots & \vdots & \ddots & \vdots \\ \overline{\alpha_{\text{current}}}(1) & \overline{\alpha_{\text{current}}}(2) & \dots & \overline{\alpha_{\text{current}}}(a) \end{bmatrix}, \tag{4.33}$$

where a is the number of waypoints along $\mathcal{R}$. The time series data are divided into fixed-length subsequences, each represented by a matrix of size $s \times l$, where s is the number of metocean variables and l is the length of the subsequence.

Instead of clustering individual waypoints, TICC clusters subsequences to account for temporal dependencies. Each cluster is associated with an inverse covariance matrix forming a Markov random field that encodes structural patterns across segments. TICC aims to find optimal segment assignments by solving:

$$\arg\min_{\boldsymbol{\Theta},\sigma} \sum_{j=1}^{n} \left(\lambda \|\theta_j\|_1 + \sum_{\mathbf{W}_k \in \sigma_j} \left(-\log L(\mathbf{W}_k; \theta_j) + \Omega \mathbf{1}\left[\sigma_{k-1} \neq \sigma_j\right] \right) \right), \tag{4.34}$$

where $\boldsymbol{\Theta} = \{\theta_1, \dots, \theta_n\}$ represent the inverse covariance matrices for each cluster, $\sigma = \{\sigma_1, \dots, \sigma_n\}$ denotes the segment assignments, λ and Ω control sparsity and temporal consistency, $L(\mathbf{W}_k; \theta_j)$ is the likelihood function, and $\mathbf{1}[\cdot]$ is the indicator function.

4.4 Continuous optimisation refinement

Whilst DP effectively identifies the global strategy across the discretised search space, the resolution of the solution is inherently limited by the step size of the power settings Γ. As illustrated in Paper III, the true global optimum often lies between discrete states. Addressing this involves using BO to introduce a continuous refinement stage.

Standard BO is typically initialised with random samples, which can be computationally inefficient for high-dimensional problems. In this framework, the sub-optimal solution trajectory P_{DP} obtained from the parallel coupling DP (section 4.2.3) is used as *prior knowledge*. These discrete solutions are injected into the Gaussian Process (GP) surrogate model as the initial training set $\mathcal{D}_0$:

$$\mathcal{D}_0 = \{(P_{\text{DP}}^{(i)}, M_{\text{fuel}}^{(i)})\}_{i=1}^{k}, \tag{4.35}$$

where k represents the number of top-performing discrete paths retained. This "warm start" immediately focuses the BO search on the most promising regions of the continuous power landscape.

The primary challenge in refining the power allocation is satisfying the strict ETA constraint, defined as $|1 - T_{\text{sail}}(P)/T_{\text{ETA}}| \leq 1\%$. Two strategies for handling this constraint within the probabilistic framework were developed and compared.

The first method, explicit constraint modelling (BO+EC), treats the sailing time T_{sail} as an unknown black-box function, modelled by a separate, independent GP. The acquisition function $\alpha(P)$ is weighted by the probability that a candidate power setting satisfies the constraint C_{ETA}:

$$\alpha_C(P_*) = \alpha(P_*) \cdot \Pr(T_{\text{sail}}(P_*) \in C_{\text{ETA}}), \tag{4.36}$$

where the probability $\Pr(\cdot)$ is derived from the cumulative distribution function of the constraint GP posterior. This approach strictly separates the objective (fuel) from the constraint (time).

Conversely, the embedded penalty (BO+EP) strategy was developed to improve computational efficiency by integrating the constraint directly into the objective function. The fuel consumption objective is warped by a multiplicative penalty factor w_{fuel}:

$$f(P) = M_{\text{fuel}}(P) \cdot w_{\text{fuel}}(T_{\text{sail}}(P)). \tag{4.37}$$

The penalty function w_{fuel} is designed as a smooth, differentiable sigmoid-like function that escalates exponentially when the sailing time violates the allowable margin:

$$w_{\text{fuel}}(T) = \begin{cases} 1 & \text{if } T \in C_{\text{ETA}} \\ 1 + k_1 \cdot \left(\frac{1}{1+\exp(-k_2 \cdot \delta^2)} - 0.5 \right) & \text{if } T \notin C_{\text{ETA}}, \end{cases} \tag{4.38}$$

where δ represents the deviation from the ETA constraint, and k_1, k_2 are tuning parameters controlling the penalty stiffness. As demonstrated in Paper V, the BO+EP method achieves convergence speeds that are approximately 10 times faster than BO+EC by avoiding the computational overhead of maintaining a second GP model, simultaneously yielding practically identical optimal power allocations.

4.5 Optimisation of the double-ended ferry

Double-ended ferries pose unique optimisation challenges due to their dual-engine configuration, which requires the simultaneous adjustment of both engines' power levels. Each engine, typically labelled "bow" and "stern" depending on the ship's current heading, is controlled independently. Unlike conventional single-direction ships, these ferries change direction without turning around, complicating fuel management. Efficient operation requires that both engines be actively managed regardless of the ship's heading, as either may serve as the leading or trailing propulsion unit during a given voyage.

To address these challenges, the DSS framework depicted in Fig. 4.12 was developed. The DSS incorporates trip-specific historical and real-time data to inform optimal power allocation for each voyage utilising the following inputs:

- $\mathcal{R}$**:** Historical sailing waypoints, including longitude, latitude, heading angles, speed over ground V_G, and fuel consumption metrics for similar trips.

- **Metocean data:** Weather and sea state data interpolated along $\mathcal{R}$ and aligned with the ferry's schedule to represent operating conditions accurately.

- **Initial operational guess:** A pre-defined parameter set representing a likely optimal region in the search space, used to efficiently initialise the optimiser.

Given that the ferry operates on short, regular routes with minimal metocean variation during individual trips, no segmentation into multiple legs is applied. Each voyage is treated as a single optimisation unit, simplifying the objective of determining the optimal power settings for the bow and stern engines, as described in Eq. 4.19.

A key variable in the optimisation is the power distribution between the two engines. This distribution is quantified using the *power ratio*, R_p, defined as the proportion

of power distribution to the stern engine:

$$R_p = \frac{P_{\text{stern}}}{P_{\text{bow}} + P_{\text{stern}}}, \tag{4.39}$$

where P_{stern} and P_{bow} represent the respective power outputs of the stern and bow engines. The value of R_p lies between 0 and 1, with $R_p = 0$ indicating that all power is allocated to the bow engine, and $R_p = 1$ indicating full stern engine usage.

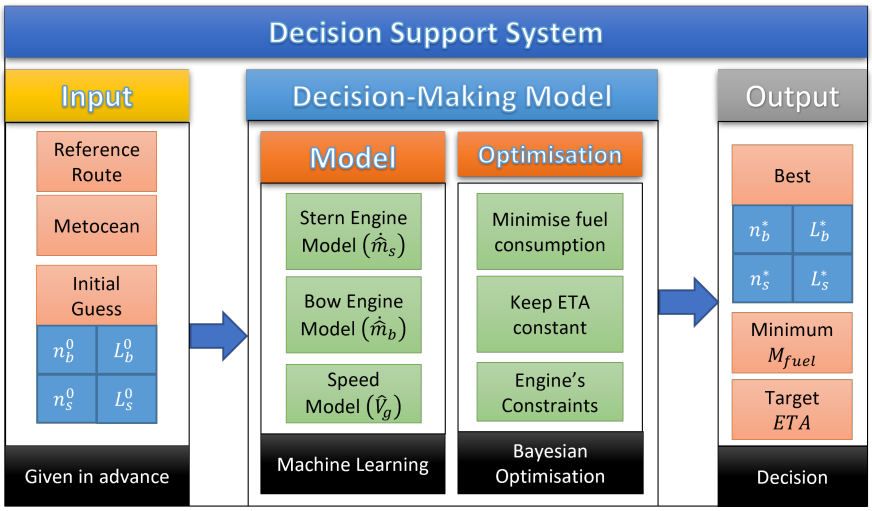

Figure 4.12: Decision-support system workflow for optimising power allocation.

This work explores the operational hypothesis that allocating more power to the stern engine reduces fuel consumption. This proposition is based on performance observations suggesting higher propulsion efficiency under stern-heavy power configurations. To assess this relationship, the time-averaged power ratio for each voyage is computed as:

$$\overline{R_p} = \frac{1}{n} \sum_{i=1}^{n} R_p^{(i)}, \tag{4.40}$$

where n is the total number of time samples during the trip and $R_p^{(i)}$ is the instantaneous power ratio at time i. The optimisation process aims to identify the optimal $\overline{R_p}$ value that minimises fuel use for each journey.

To evaluate the power ratio hypothesis and optimise fuel consumption, subject to voyage constraints (e.g., ETA), BO is used within the DSS framework. BO is well-suited to this task because it enables an efficient search of complex, expensive-to-evaluate

objective functions by building a probabilistic surrogate model. The optimisation proceeds iteratively as follows:

1. Initialise with a set of sample waypoints and evaluate fuel consumption at each waypoint.
2. Train a GPR surrogate model mapping power ratios to fuel consumption.
3. Use an acquisition function to select the next point to evaluate, balancing exploitation and exploration.
4. Evaluate fuel consumption at this new waypoint.
5. Update the surrogate model with the new data and repeat the process.

This framework enables the system to learn an optimal allocation strategy for bow- and stern-mounted engines for each trip.

CHAPTER 5

Data and case study ships

The methods used for this work are supported by multiple datasets collected from both full-scale ship operations and hydrodynamic experiments. This research considers three different cases: (1) A double-ended ferry operating on a very short, repeatable route, (2) a short-sea chemical tanker exposed to variable metocean conditions, and (3) model-scale data to account for added resistance in waves.

5.1 Operational data from full-scale measurements

The datasets used in Papers I, III, IV and V were retrieved from onboard monitoring systems installed on two commercial ships: the Ro–Pax ferry *M/S Uraniborg* and a mid-size chemical tanker. The onboard data acquisition systems continuously recorded operational parameters, including propulsion power, shaft speed, fuel flow, ship speed, and navigation information.

Ferry dataset

The first dataset covers one year of operational data collected from the double-ended Ro–Pax ferry *M/S Uraniborg*, operating between the island of Ven and the town of Landskrona in Sweden. The ship's energy management system (BlueFlow) recorded engine load, power output, propeller revolutions, ship speed, and fuel consumption at a 1-minute resolution. Table 5.1 summarises the main technical particulars of the ship.

The collected data reveal a distinct operational profile characteristic of short-sea commuter ferries. Each trip typically lasts between 25 and 30 minutes and comprises three distinct sailing phases (Fig 5.1): (1) an acceleration phase of approximately 5 minutes during which the ship ramps up to service speed, (2) a steady cruising phase where the ship maintains an average speed of 11.5 knots, and (3) a deceleration phase of roughly 4 minutes as the ship approaches the harbour. The energy modelling and optimisation in this thesis primarily focus on the cruising phase, when the ship is operating in a quasi-steady state.

Table 5.1: Principal particulars of the Ro–Pax ferry *M/S Uraniborg* (Papers I and IV).

Parameter	**Value**
Length overall	49.95 m
Beam	12.00 m
Draught	2.85 m
Installed Power	2 × 709 kW
Engine Type	Scania DI16 diesel engines
Engine Speed	1600 rpm
Service Speed	11.5 kn
Data Period	Jan 2021 – Jan 2022
Sampling Frequency	1 sample per minute

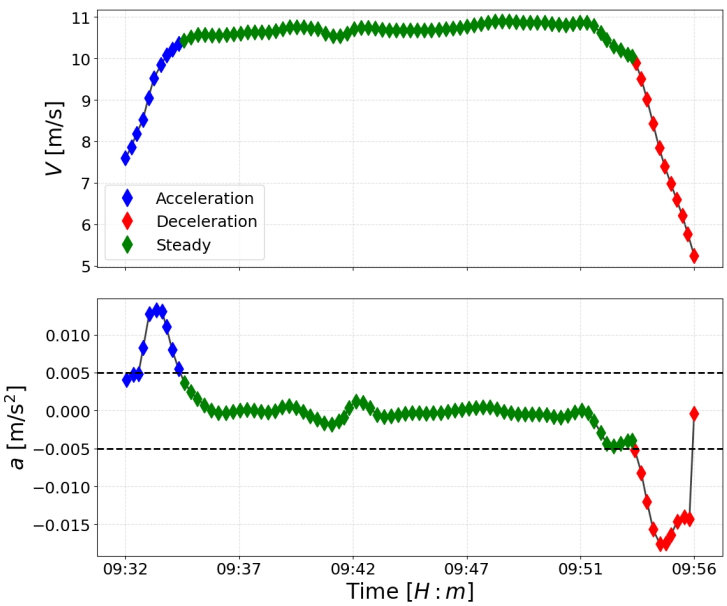

Figure 5.1: Phases of the voyage.

Tanker dataset

The second dataset comprises three years of measurements from a chemical tanker operating across Northern European waters, including the Baltic Sea, North Sea, and English Channel, as shown in Fig. 5.2. The ship is equipped with a single medium-speed diesel engine, and data were logged at a 1-minute resolution by the ship's monitoring and automation system. The principal particulars of the tanker are provided in Table 5.2.

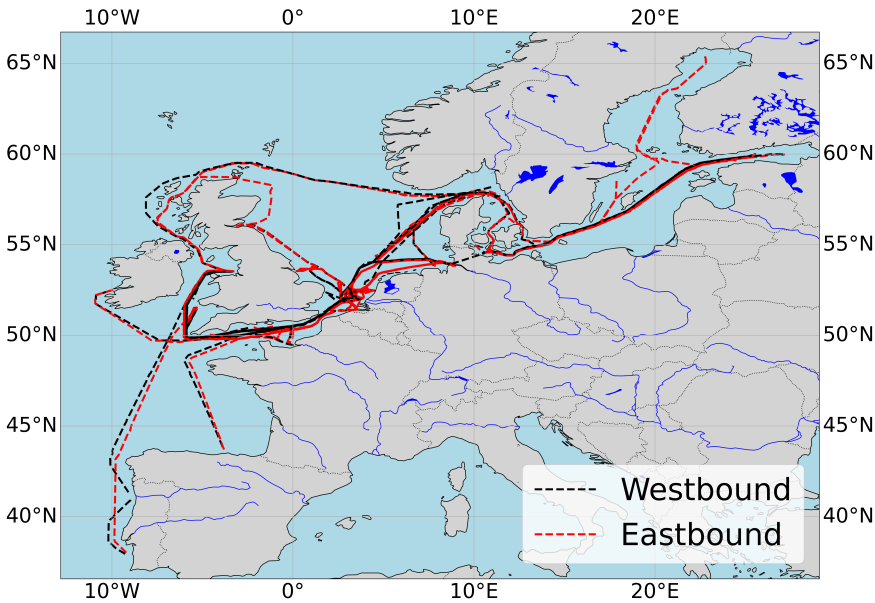

Figure 5.2: Sailing region of the chemical tanker.

Table 5.2: Principal particulars of the chemical tanker (Papers III and V).

Parameter	Value
Length between perpendiculars	138.22 m
Beam	23.76 m
Design draught	9.27 m
Displacement	25,174 m^3
Maximum continuous rating	7200 kW
Service speed	14 kn
Data period	Nov 2020 – Mar 2024
Sampling frequency	1 sample per minute

Unlike the ferry, the chemical tanker operates on longer, more irregular routes for durations ranging from 40 to over 140 hours. The dataset encompasses diverse metocean conditions, including significant wave heights exceeding 3 metres and various current interactions. The raw measurements data include propulsion system parameters such as shaft power (P), engine speed (RPM), fuel consumption rate (m_{fuel}) and navigational data speed over ground (V_g), heading (α_{ship}), and draught (T), as shown in Fig. 5.3. To ensure data quality for modelling, the raw 1-minute data were downsampled to 10-minute averages, and transient periods, such as manoeuvring in port or voluntary speed reductions, were filtered using a 3-sigma rolling standard deviation method.

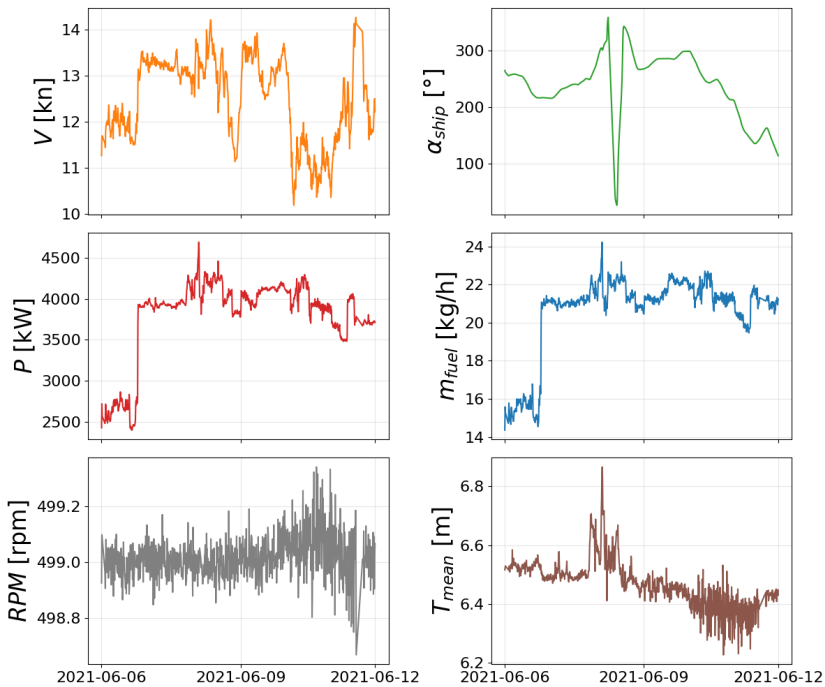

Figure 5.3: Example of the data used for regression modelling.

5.2 Experimental tests on added resistance data

Publicly available model-scale data, summarised in the appendix Table A.1, were used to develop GPR models for wave-induced added resistance, as described in section 4.1.3. The resulting database comprises 2,096 samples from 45 ships across 175 distinct experimental series. It includes the added resistance coefficient C_{AW}, defined in Eq. 3.8, for a range of different model-scale hull forms, tested under different

Froude numbers F_n and wavelengths λ. The distribution within this database is as follows:

- Froude numbers (F_n) ranging from 0 to 0.5.
- Wavelength-to-ship length ratios (λ/L) ranging from 0.2 to 3.0.
- Block coefficients (C_B) ranging from 0.55 to 0.87.
- Beam-to-draft ratios (B/T) ranging from 2.0 to 4.5.

Table 5.3 summarises the main particulars of the studied hulls, as well as the Froude numbers at which they were simulated.

Table 5.3: Principal particulars of the model-scale ships and Froude numbers.

Ship Type	L **[m]**	B **[m]**	T **[m]**	C_B	F_n
KVLCC2	320	58.0	20.8	0.81	0.100
S175	175	25.4	9.5	0.57	0.150
Hull 2020	187.3	32.28	12.0	0.82	0.120
JBC	280	45.0	16.5	0.86	0.156
SR108	3.5	0.508	0.19	0.57	0.250
Aframax	248	43.0	14.3	0.84	0.156
SCb87	178	32.26	14.46	0.87	0.166
Supramax	192	36.0	11.2	0.81	0.170

5.3 Metocean data along case study voyages

The metocean parameters required for the models described in section 4 were extracted from the ECMWF ERA-5 dataset for wind and waves, whilst the ocean current data were acquired from the Copernicus Marine Service. A sample of these data is presented in Fig. 5.4. The direction components are relative to the movement of the ship and are calculated as:

$$\alpha_i = \arccos\left(\frac{\mathbf{V}\cdot\mathbf{V}_i}{\|\mathbf{V}\|\cdot\|\mathbf{V}_i\|}\right), \tag{5.1}$$

with $\mathbf{V}$ the ship speed over ground, and $i \in \{\text{wind, current, wave}\}$.

The corresponding forecast data were retrieved from the ECMWF Meteorological Archival and Retrieval System (MARS). The utilisation of operational forecast data for SSS voyage optimisation is justified by the high fidelity of this weather forecast over a short time horizon, as visualised in Fig. 5.5 (wind) and Fig. 5.6. The forecast tracks the ground truth (hindcast) accurately for the first 4 days, exhibiting high reliability for optimisation over this period.

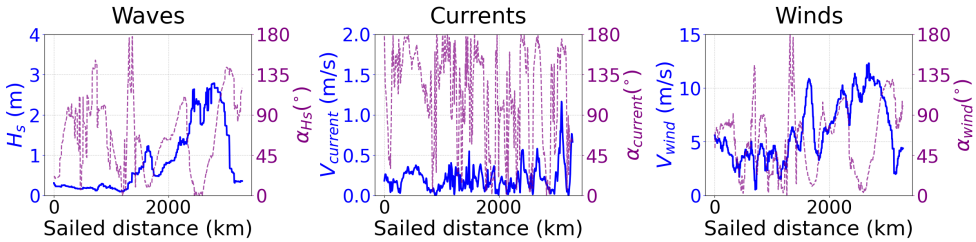

Figure 5.4: Metocean data.

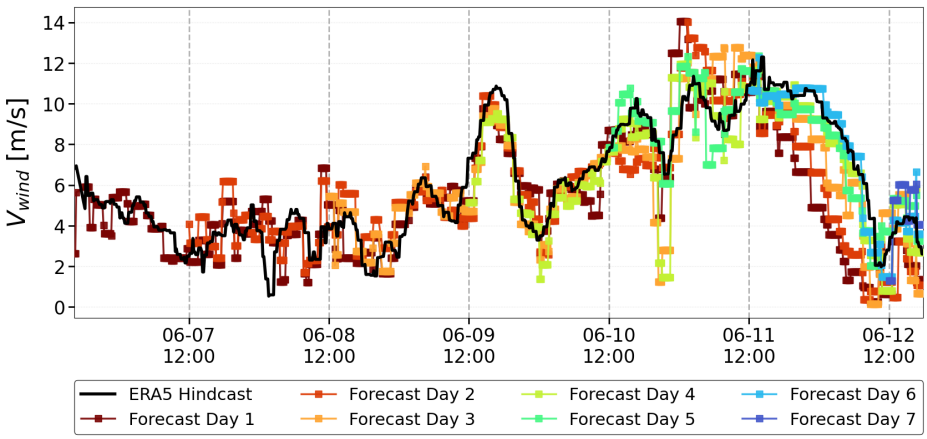

Figure 5.5: Reliability of the weather forecast data for short-sea voyages (winds).

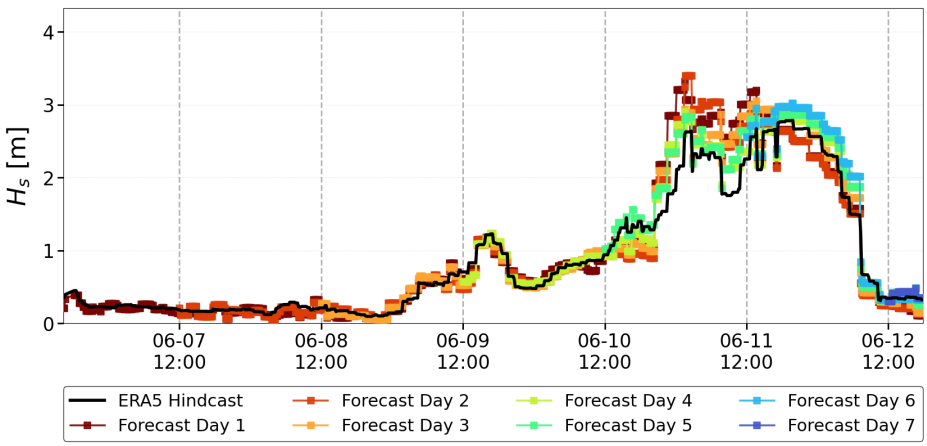

Figure 5.6: Reliability of the weather forecast data for short-sea voyages (waves).

Finally, it can be noted that even at the start of the voyage (Fig. 5.7), the forecast and hindcast data do not match perfectly. This is because the operational forecast is initialised using only the data available at the moment of execution. In contrast, the hindcast is a retrospective product generated subsequently, and it incorporates a significantly larger pool of observations to provide a more accurate representation of the actual conditions at all times.

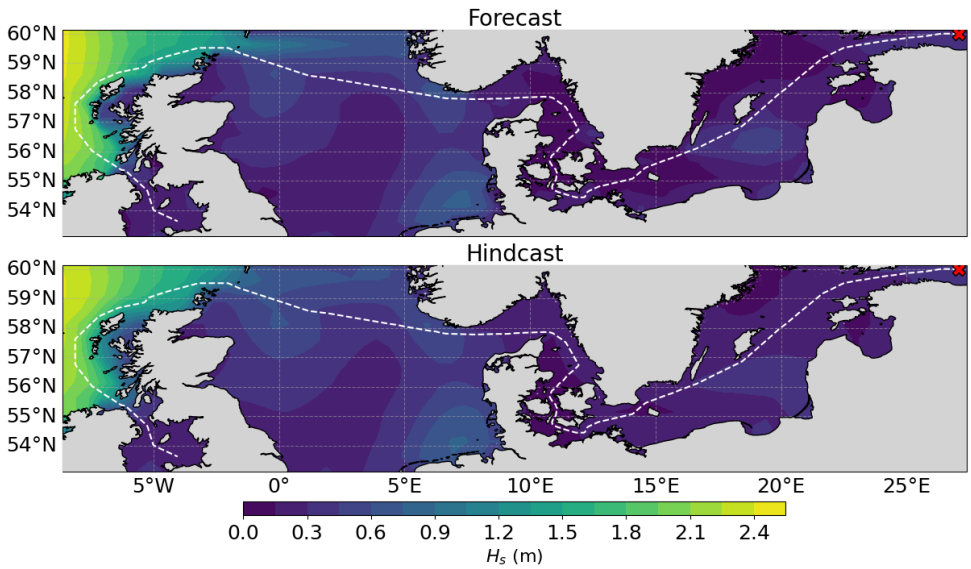

Figure 5.7: Reliability of the weather forecast data for short-sea voyages (start).

CHAPTER 6

Summary of appended papers

This chapter presents a summary of the appended papers, including the research conducted, a selection of the most important results, and highlights of the main achievements. The relationships between the five papers and the methodological progression they demonstrate are summarised in Fig. 6.1, which conceptually divides the research into a modelling phase and an optimisation phase, while also connecting them by case study.

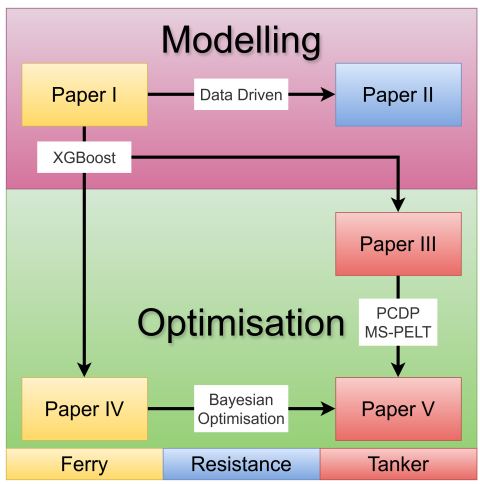

Figure 6.1: Paper hierarchy and research connections.

Highlight of Methodological Contributions

Paper I establishes a foundational framework for the exploratory data analysis and rigorous data processing utilised throughout the research. It introduces the power ratio metric (section 4.5) subsequently used in Paper IV, and the XGBoost modelling methodology used in Papers III to V, as described in section 4.1.2. Furthermore, the paper proposes an initial approach to operational fuel optimisation.

The data handling methodology described in Paper I is specifically applied in Paper II. Although Paper II primarily focuses on determining an ML approach to added wave resistance through the smooth regression described in section 4.1.3, it touches upon the use of these models for optimisation by predicting fuel consumption using a grey-box modelling approach. Whilst Paper II stands alone in terms of its specific subject, it is a direct and relevant branch of the foundational work in Paper I.

The XGBoost ship performance models established in Paper I are utilised in Paper III to develop a methodology for simulating a sailing ship for voyage planning. The main contribution of Paper III lies in its novel parallel coupling dynamic programming (PCDP) approach, described in section 4.2.3, alongside the route segmentation MS-PELT algorithm presented in section 4.3.1.

Paper IV utilises the XGBoost framework for voyage simulation but introduces a BO approach tailored for very short voyages. The significance of this paper lies in its validation of the power ratio hypothesis established in Paper I by providing a more detailed voyage simulation and a more precise optimisation algorithm. In addition, Paper IV studies where full-scale experiments are conducted to empirically verify theoretical methods.

Finally, Paper V serves as a methodological synthesis, integrating the XGBoost modelling from Paper I, the PCDP and route segmentation strategies from Paper III, and the BO framework from Paper IV. In addition, the study described in this paper implements advanced constraint handling to accelerate optimisation convergence, as described in section 4.4, delivering a highly efficient continuous power allocation framework that addresses the inherent limitations of DP. Furthermore, it analyses the reliability of hindcast data relative to forecast data, highlighting the necessity for refined optimisation methods to capitalise on the most accurate weather information available at any given moment.

Computational performance

The computational performance of each stage within the proposed framework is presented to illustrate the practical feasibility and relative costs of the data-driven and optimisation components (see Table 6.1). All models and optimisation algorithms were implemented in Python and executed on a local machine equipped with an Intel Core i9-10900 CPU @ 2.80GHz and 64 GB of RAM. The source code is available in this repository[1].

Table 6.1: Computational duration for the execution of various framework processes.

Process	Duration [s]
Generating interpolators	20
Generating weather ranges	15
MS-Pelt	2.00×10^{-2}
TICC	50
DP-Graph	60–80
DP-Solve	2–10
Bayes-Opt DP initialisation	2–10
Bayes-Opt EP	100–120
Bayes-Opt EC	600–1200
Training model XGBoost	600

The highest computational cost across the entire framework occurs during the initial model training (XGBoost). This cost is ideally paid only twice: once for the speed model and once for the fuel consumption model. The discrete components of the framework, such as the MS-PELT segmentation and the DP-Solve algorithm, operate at near-instantaneous speeds, facilitating rapid initial route generation. The 60–80 seconds required for DP-Graph construction represents a one-time cost per voyage. Finally, the explicit constraint method (BO+EC) constitutes a clear bottleneck, requiring almost 20 minutes of execution time (or more). However, adopting the BO+EP approach reduced this to 100–120 seconds, bringing the total refinement time into a practical range for operational use.

[1] https://github.com/Deftheros/PhD-Voyage-Optimisation.git

6.1 Summary of Paper I

"Power allocation influence on energy consumption of a double-ended ferry"

Paper I demonstrates the use of data analytics and ML regression to reduce fuel consumption in a double-ended ferry. The study investigates optimal engine loading conditions from a data-driven perspective and introduces a novel variable, the power ratio (R_P), to characterise the stern engine's loading relative to total engine utilisation. The results indicate that operating at higher R_P values offers substantial potential for fuel savings, with reductions of up to 35% achievable through appropriate operational adjustments. The research for Paper I was conducted as follows:

- **Exploratory data analysis** was used to analyse the sensor and environmental data to determine operational fuel consumption trends and patterns.

- **Regression analysis** was applied to model the relationships among power allocation, environmental data, and fuel consumption.

- **Machine learning** was employed to increase model precision.

The analysis was performed by plotting total fuel consumption data against R_P, resulting in a linear trend (see Fig. 6.2). The line fit is presented in Eqs. 6.1 and 6.2. The minimum of this regression occurs at the right extremum ($R_P = 1$), suggesting that allocating more power to the stern engine results in lower fuel consumption.

$$M_{\text{west}} = 83.87 - 53.834 \cdot R_P \tag{6.1}$$

$$M_{\text{east}} = 84.455 - 53.888 \cdot R_P \tag{6.2}$$

In addition to the data analysis, ML regression was employed to model ferry fuel consumption and to quantify the impact of the power ratio R_P with greater accuracy. Three regression models were evaluated (see Fig. 6.3), among which the XGBoost model achieved the highest predictive accuracy and was therefore selected for further analysis.

Various power allocation scenarios were simulated using the trained XGBoost model by manually prescribing different R_P values, as illustrated in Figs. 6.4a and 6.4b. The simulation results confirm that higher R_P values yield substantial fuel savings, aligning with the trends identified using linear regression analysis. In particular, a power ratio of $R_P = 1$ resulted in fuel savings of up to 18% for eastbound trips and 35% for westbound trips, as summarised in Table 6.2.

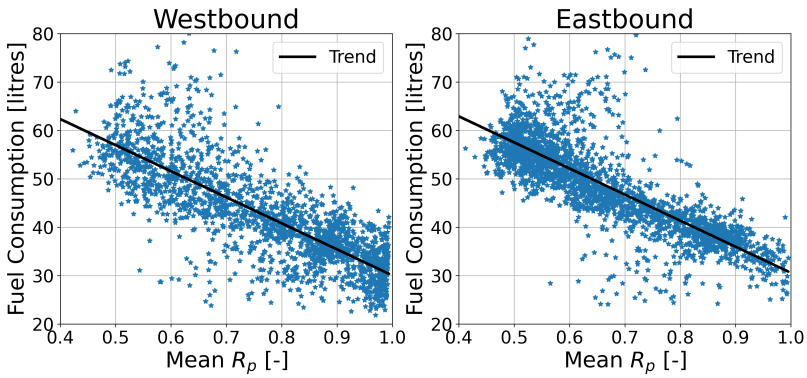

Figure 6.2: Fuel consumption vs power ratio for each trip and direction in 2020.

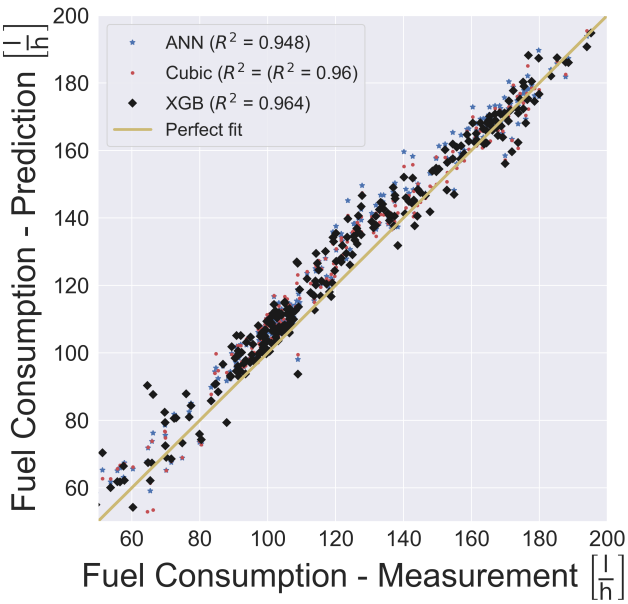

Figure 6.3: Performance of three machine learning algorithms.

Table 6.2: Summary of the selected simulation results.

Direction	Measured $\mathbf{M_{fuel}}$ [l]	Simulated $\mathbf{M_{fuel}}$ [l]	Savings
Westbound	54.90	34.72	-35.8%
Eastbound	41.36	35.5	-18.45%

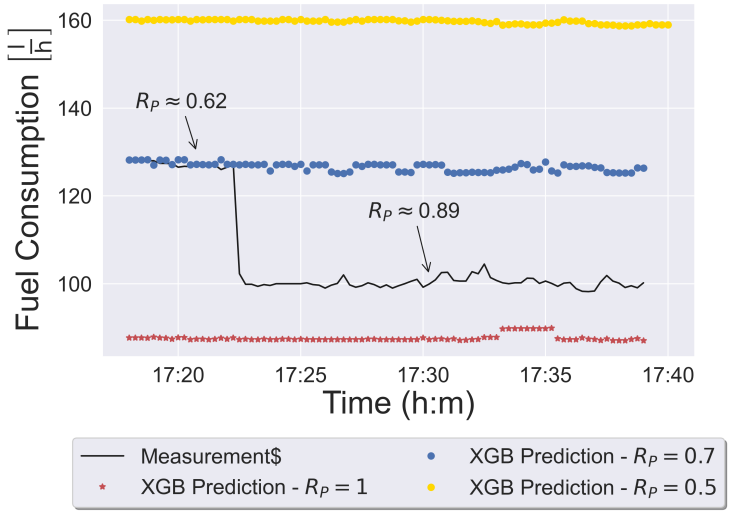

(a) R_P scenarios for an eastbound trip.

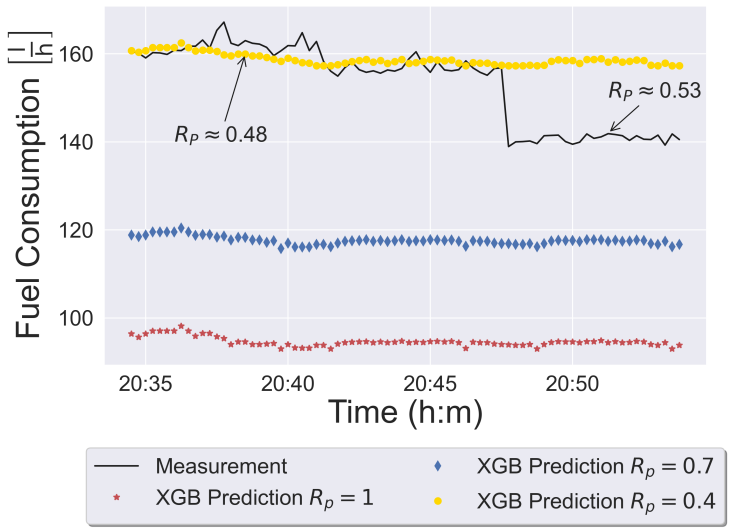

(b) R_P scenarios for a westbound trip.

Figure 6.4: Comparison of the power allocation impact on two trips.

6.2 Summary of Paper II

"A machine learning method to evaluate head sea induced weather impact on ship fuel consumption"

Paper II proposes a Gaussian process regression (GPR) model to describe the added resistance of a generic ship in head waves. Whilst vessel-specific ML models can achieve high accuracy under familiar operating conditions, they often exhibit poor generalisation to new hull forms or extreme sea states that are not represented in the training data. Given that adverse weather may account for up to 15% of a ship's total fuel consumption, accurate and computationally efficient prediction of wave-induced added resistance is essential for assessing weather-related fuel consumption. Traditional semi-empirical formulations often exhibit large uncertainties under real ocean conditions, and purely data-driven models frequently fail to capture the physical shape of the resistance curve, even during interpolation. The research for Paper II was conducted as follows:

- **GPR** was applied to develop a probabilistic model for predicting added resistance coefficients, capturing the hydrodynamic correlations inherent in experimental datasets.
- **Ship-type split validation** was implemented to rigorously evaluate the model's ability to extrapolate to previously unseen hull geometries.
- **Grey-box modelling** was employed to combine the validated GPR model with a deep neural network (DNN), forming a hybrid framework for full-scale vessel fuel consumption prediction.

The methodology involved processing a database of 2,096 experimental model-test points covering 45 ship types and 175 distinct experimental cases. The model uses dimensionless parameters such as λ/L, the Froude number (F_n), and the block coefficient (C_B) to account for variations in ship scale and outputs the added resistance coefficient (C_{aw}). Unlike standard ANNs, which treat these data points as independent samples, the GPR model explicitly models the covariance between inputs using a kernel function. This enables it to smoothly learn the underlying physical behaviour of the wave–hull interaction rather than simply memorising data points.

The results demonstrate that the GPR model exhibits more reliable extrapolation behaviour than baseline alternatives. Whilst ANN-based models may produce erratic or overly flat resistance curves for unseen ships (e.g., the S175 and JBC cases), the GPR captures resonance features characteristic of wave–hull interactions. Quantitatively, the GPR achieves an RMSE of 0.95, compared with 2.45 for the semi-empirical CTH method, as summarised in Table 6.3. Even when the RMSE is not the lowest,

the GPR yields smoother curve shapes (Fig. 6.5), which is desirable for operational use and subsequent optimisation.

Table 6.3: RMSEs of different methods.

Model	$\mathbf{F_n}$	GPR	ANN	CTH
S175	0.150	0.53	0.39	0.55
JBC	0.156	0.34	1.84	5.75
KVLCC	0.100	0.69	0.87	1.25
Supramax	0.170	0.71	0.61	1.25
HSVA	0.233	0.68	0.86	0.64
Aframax	0.156	0.39	0.69	2.07
2020 Hull	0.120	0.20	0.22	1.85
Scb87	0.166	0.45	0.54	1.48
SR108	0.250	0.88	0.82	2.21

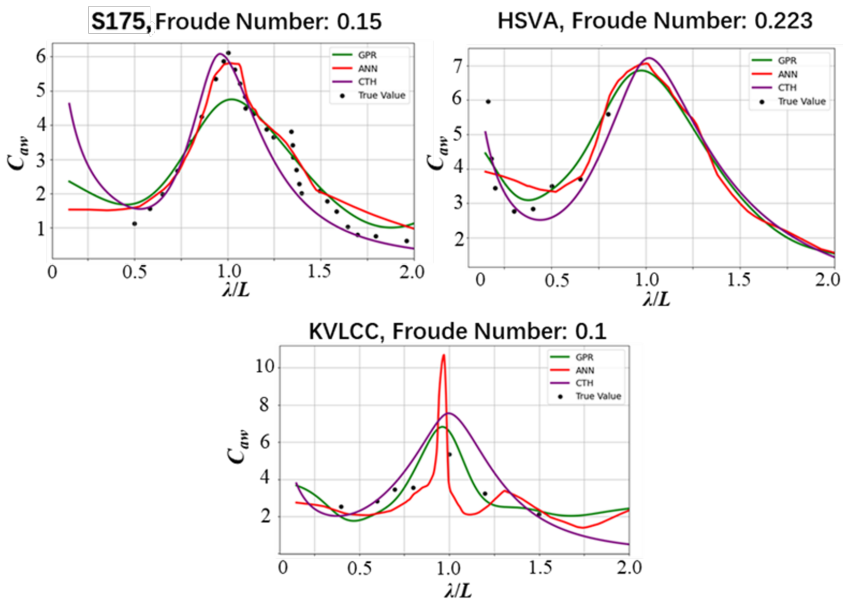

Figure 6.5: Comparison of predicted wave resistance coefficient (C_{aw}) curves.

Finally, the validated added-resistance model was integrated into an additive grey-box fuel-prediction framework for full-scale application. A physics-based baseline model was used to estimate fuel consumption (RMSE = 129.58 kg/h), and the GPR-predicted added resistance under head-sea conditions was incorporated into a DNN that learned the residual between the baseline estimate and measured fuel consump-

tion. When evaluated on a 45,000 DWT chemical tanker, the resulting grey-box model reduced the fuel-prediction RMSE by 65% to 45.65 kg/h and increased R^2 from 0.69 to 0.95 (Fig. 6.6). Furthermore, the framework enables a quantitative assessment of weather impacts, indicating that head-sea waves account for approximately 8% of total fuel consumption on average, reaching up to 39% under severe conditions.

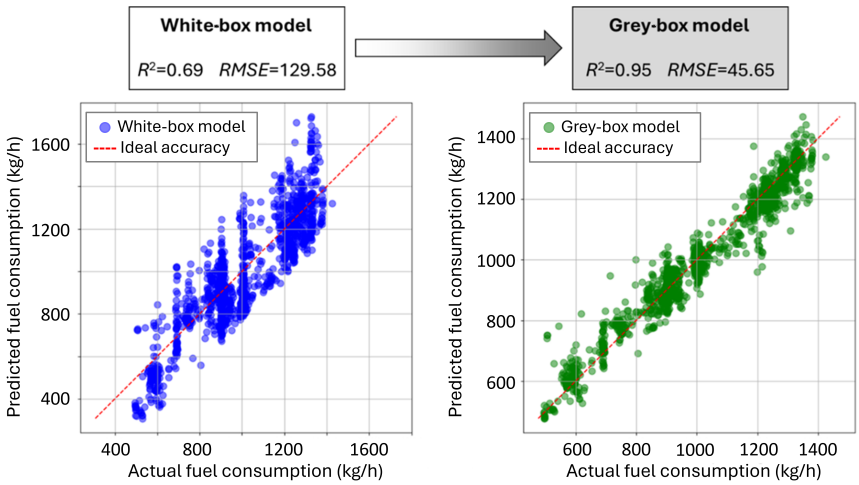

Figure 6.6: Prediction accuracy of white-box and grey-box models.

6.3 Summary of Paper III

"Reduced environmental impact of short sea shipping through optimal engine power allocation"

Paper III presents a two-stage framework for optimising voyages along fixed routes. In contrast to Papers I and IV, it focuses on a conventional merchant vessel, a chemical tanker operating across European waters. The operational dataset includes three years of historical sailing routes, as illustrated in Fig. 5.2. Although the literature predominantly focuses on optimising speed for voyage optimisation, this often requires continuous engine-load adjustments to maintain prescribed speeds under varying metocean conditions, potentially increasing inefficiencies. Paper III aims to address this limitation, enhancing ship energy efficiency by directly optimising engine power allocation whilst enforcing constraints on average sailing speed or total voyage duration. The proposed methodology comprises two sequential steps:

- **Route segmentation** was performed using the proposed MS-PELT algorithm based on hindcast metocean conditions to produce an appropriate number of partitions.
- **Parallel scenario dynamic programming** was applied to identify the optimal engine power allocation for each of the route segments.

The proposed MS-PELT segmentation algorithm is described in detail in section 4.3.1. The segmentation aims to identify sailing legs comprising waypoints with similar metocean conditions. In particular, the algorithm accounts for sea currents, wind speed, and wave height along the route, accounting for both their magnitudes and directions. MS-PELT and a state-of-the-art TICC method were applied to segment routes under varying metocean conditions. The segmentation results (see Fig. 6.7) demonstrate that MS-PELT generates larger and more consistent route segments, whereas TICC produces shorter and more frequent segments. Such oversegmentation, as exhibited by TICC, is undesirable in real operations because it may lead to frequent engine power adjustments, thereby increasing fuel consumption and operational complexity. MS-PELT also incurs a significantly lower computational cost, with an average runtime of approximately 10 ms, compared to around 50 s for TICC, making it more suitable for real-time and near-real-time applications.

Eight case studies were simulated to assess the effectiveness of the power allocation optimisation, with the results summarised in Tables 6.4 and 6.5. The scenario-based dynamic programming approach described in section 4.2.3 was employed to determine the optimal propulsion power for each voyage segment. Instead of adopting a sequential coupling strategy, the method precomputes multiple scenarios in parallel, thereby reducing inter-segment dependencies and enabling efficient parallelisation of the optimisation process.

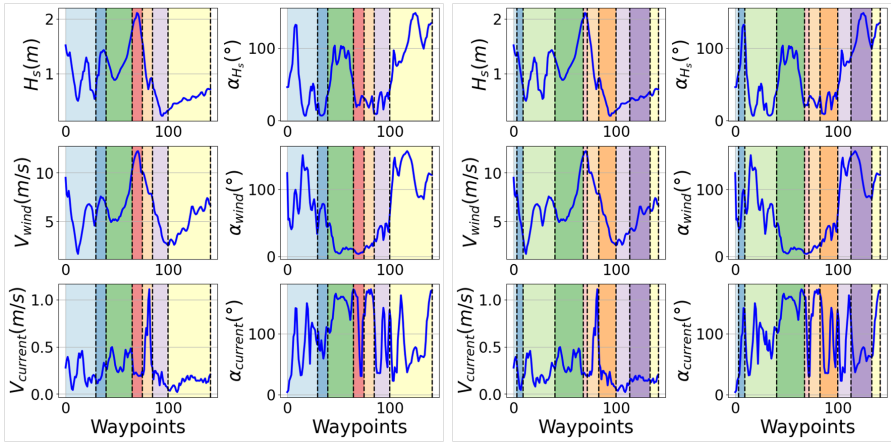

(a) PELT R1. **(b)** TICC R1.

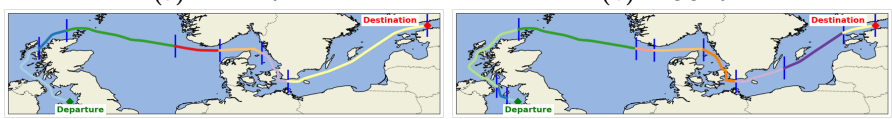

(c) PELT MAP R1. **(d)** TICC MAP R1.

Figure 6.7: Comparison between PELT and TICC for a sample route (Errata).

By integrating data-driven route segmentation with dynamic programming, the proposed framework achieves fuel savings and emission reductions of up to 16% on long voyages across European waters. At the same time, it incurs sailing time deviations of approximately 1–1.5% relative to the ETA. The framework's high computational efficiency supports real-time or near-real-time applications.

Table 6.4: Operational characteristics and baseline emissions for the selected voyages.

Voyage ID	Sailing Region	Distance [km]	ETA [hours]	Emissions [tonnes]
1	Baltic and North Sea	3388.04	145.50	383.16
2		3089.35	143.83	271.68
3	Baltic	1118.20	60.50	83.06
4		1153.58	55.33	90.24
5	North Sea and English Channel	2043.59	108.50	223.41
6		1718.97	80.33	176.90
7		1712.24	78.00	147.47
8	English Channel	980.23	47.00	69.10

A representative case was selected to illustrate the integrated segmentation optimisation framework. Case 1 covers a voyage across the Baltic and North Seas (see Fig. 6.8a). Along this route, the vessel encountered relatively harsh metocean conditions, with significant wave heights reaching up to 3 m. Despite these conditions, the optimised power allocation strategy achieved fuel savings of 6.27% whilst incurring an ETA delay of only 1 hour 10 minutes.

Table 6.5: Optimisation results (Errata).

Voyage ID	New Sailing Time [hours]	Optimised Emissions [tonnes]
1	146.66 (+0.8%)	359.14 (-6.3%)
2	143.69 (-0.1%)	252.88 (-6.9%)
3	61.14 (+1.1%)	70.71 (-14.9%)
4	55.92 (+1.1%)	84.09 (-6.8%)
5	108.53 (+0%)	187.40 (-16.1%)
6	81.32 (+1.2%)	170.32 (-3.7%)
7	79.00 (+1.3%)	130.04 (-11.8%)
8	47.47 (+1.0%)	64.37 (-6.9%)

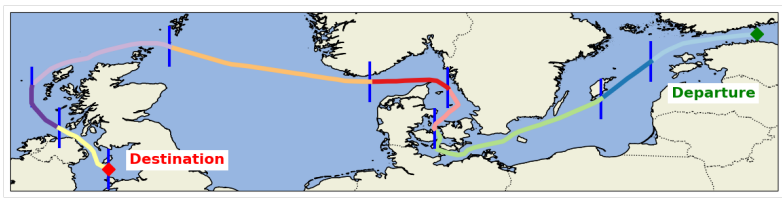

(a) Route Case 1.

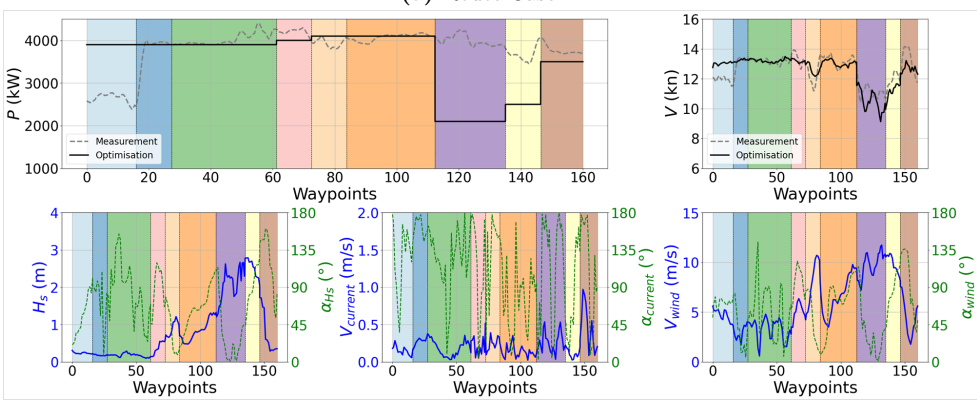

(b) Optimisation R1.

Figure 6.8: Optimisation results for Case 1 (Errata).

6.4 Summary of Paper IV

"A machine learning-based Bayesian decision-support system for efficient navigation of double-ended ferries"

Paper IV is a direct continuation of Paper I, proposing a real-time DSS for determining optimal propulsion settings for a double-ended ferry. The DSS identifies appropriate engine operating set-points, expressed in terms of the power ratio R_P and engine speed, whilst simultaneously satisfying an ETA constraint for the voyage. The research for Paper IV was conducted in three parts:

- **Machine learning modelling** was used to establish a performance model for the double-ended ferry.
- **BO** was applied to determine the optimal engine operating set-points in terms of the power ratio R_P and engine speed.
- **Full-scale experiments** were conducted using a high R_P by the shipping company to validate the proposed DSS.

With the objective of minimising overall fuel consumption, the proposed DSS identified optimal engine operating conditions by jointly optimising the power ratio R_P and the rotational speeds of the bow and stern engines, n_b and n_s. Fuel consumption over an entire voyage, denoted as M_{fuel}, was selected as the objective function and evaluated using voyage simulations. To enable this optimisation, XGBoost models were employed to estimate individual engines' fuel consumption rates as well as the vessel's speed over ground (Fig. 6.9).
BO was subsequently applied to optimise power allocation for six case-study voyages, including three westbound and three eastbound. The optimisation reduced fuel consumption by up to 43% without violating the ferry's ETA constraint. Table 6.6 provides a comprehensive summary of the simulated voyages. Notably, the optimal power ratio R_P did not always equate to unity for the models considered. Comparing with measured data for one representative westbound voyage (W1), the optimised results (see Fig. 6.10) showed a sailing time deviation of less than 2%.

To evaluate the performance of the proposed DSS, full-scale trials were conducted on board the ferry between 19 and 22 August 2022, with the objective of assessing the effect of various power allocation strategies on fuel consumption. The captain and the first mate alternated operating the vessel on successive round-trips to ensure comparable operational conditions. During the trials, the captain was instructed to prioritise power to the stern thruster, whereas the first mate, who was unaware of the experimental setup, operated the vessel as normal.

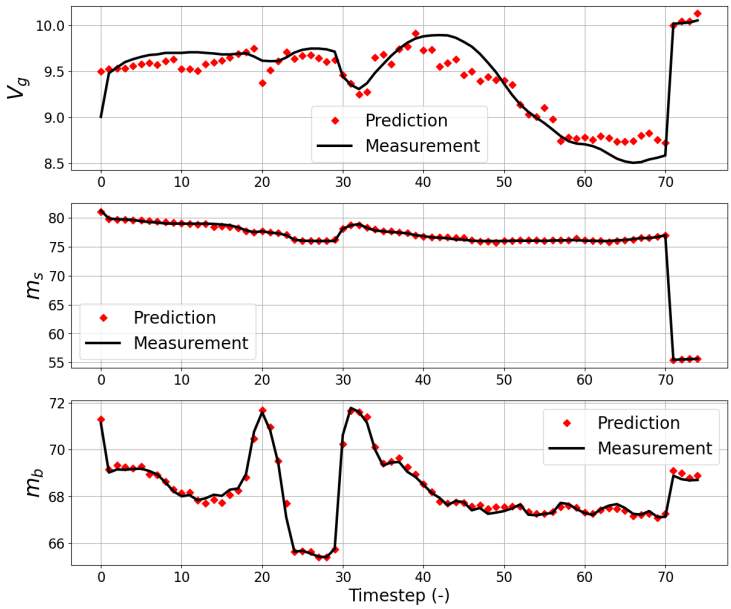

Figure 6.9: XGBoost model forecast for Trip W1.

Table 6.6: Verification of DSS for historical trips based on performance simulations.

Direction	Case	M_{fuel} [l]	R_p	M^*_{fuel} [l]	R'_p	V_g [%]	M^*_{fuel} [%]
Westbound	W1	45.12	0.472	25.45	0.898	+1.44 %	-43.59 %
Westbound	W2	40.08	0.549	25.21	0.947	+0.40 %	-37.10 %
Westbound	W3	39.78	0.480	35.62	0.523	+0.07 %	-10.46 %
Eastbound	E1	33.48	0.565	28.47	0.684	+1.39 %	-14.96 %
Eastbound	E2	31.93	0.649	23.11	0.958	+1.8 %	-27.62 %
Eastbound	E3	29.11	0.535	18.73	0.624	+0.44 %	-35.66 %

Round trips were scheduled at intervals of approximately 1–2 hours to limit variations arising from changing weather and traffic conditions. The subsequent analysis excluded voyages exhibiting irregular speed profiles, retaining ten valid trips per operator for comparison (Fig. 6.11). In addition, reference data from routine ferry operations conducted by the same crew prior to the experiments were used as a baseline. The comparison indicates that the improved power allocation strategy reduced fuel consumption.

The results of the full-scale experiments are summarised in Table 6.7. The analysis

indicates that voyages operated by the captain achieved an average fuel consumption reduction of approximately 18% compared with those operated by the first mate, which are considered representative of standard operational practice. Relative to the reference baseline, fuel consumption decreased by approximately 15% for the captain and 3.5% for the first mate, respectively. These results further highlight the potential efficiency gains associated with higher levels of stern power allocation.

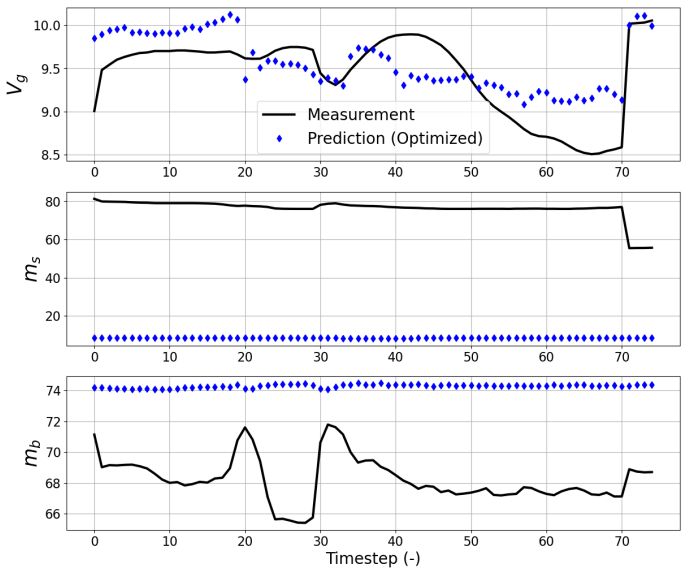

Figure 6.10: Optimised operation for trip W1.

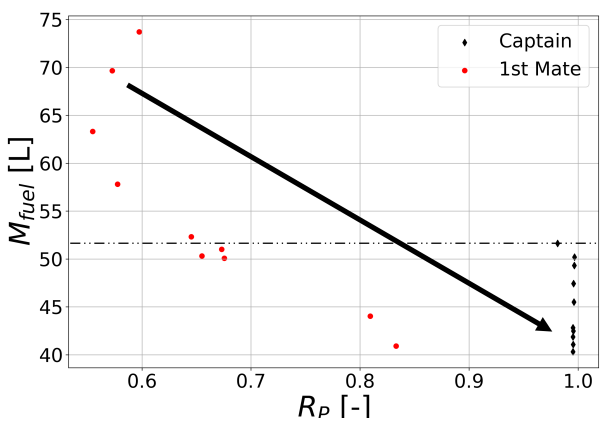

Figure 6.11: Observed trend in the full-scale experiments.

Table 6.7: Summary of full-scale test results.

Operator	Period	N-Trips	V_g $[kn]$	R_p	M_{fuel} $[l]$	% Change w.r.t. Operator	% Change w.r.t. Reference
1st Mate	Reference	10	9.3	0.7	57	-	-
Captain	Reference	10	9.3	0.8	53	-7%	-
1st Mate	Experiment	10	9.4	0.7	55	-	-3.5%
Captain	Experiment	10	9.5	1.0	45	-18%	-15%

6.5 Summary of Paper V

"Forward dynamic programming-informed Bayesian method for optimal power allocation in short-sea shipping"

Paper V continues the research presented in Paper III, addressing a fundamental limitation of discrete optimisation methods in maritime applications. Although DP effectively identifies a global strategy, its precision is inherently constrained by the discretisation step size of the engine power settings. The true global optimum often lies between these discrete states, and increasing the grid resolution to identify it incurs prohibitive computational costs. To address this, the study proposes a structured discrete-to-continuous framework that refines the coarse DP solution into a precise continuous optimum without sacrificing computational speed. The research methodology is divided into two primary steps:

- **Continuous refinement** was achieved by utilising the discrete DP solution as a "warm start" prior to a BO search.
- **Constraint handling strategies** were developed to compare two distinct methods, explicit constraint modelling (BO+EC) and embedded penalty (BO+EP), to ensure strict adherence to the ETA.

The refinement process begins by using the optimal power sequence identified by the parallel coupling DP (from Paper III) to inject prior knowledge into the GP surrogate model. Rather than beginning the search with random samples, the BO starts with a high-confidence belief about the optimal region. This allows the algorithm to immediately focus its search on fine-tuning power allocation in the continuous domain, rather than exploring the entire solution space from scratch.

To manage the strict ETA constraints required for commercial shipping, the BO+EC method treats the sailing time as a black-box function, modelling it with a separate, independent GP to predict the probability of feasibility. Conversely, the BO+EP method integrates the constraint directly into the objective function. It employs a smooth, differentiable penalty function that deforms the fuel-consumption objective,

exponentially increasing when the sailing time exceeds the allowable margin. This effectively transforms the constrained problem into an unconstrained one, simplifying the search landscape.

The framework was validated on four short-sea voyages across the North Sea and English Channel. Figure 6.12 illustrates the progression of the constraint function, computational time, and objective function across each iteration of BO for one case study voyage. The initial iterations were predetermined for BO+EP (blue) and BO+EC (red), followed by ten randomised solutions. After this initialisation, each method was run for 100 optimisation iterations. BO+EP achieved results comparable to BO+EC in all cases, at only one-tenth of the computational cost. Specifically, BO+EP required 100 seconds compared to the 1,100 seconds needed by BO+EC. Across all cases, BO+EP demonstrated speed-ups ranging from 11× to 20×.

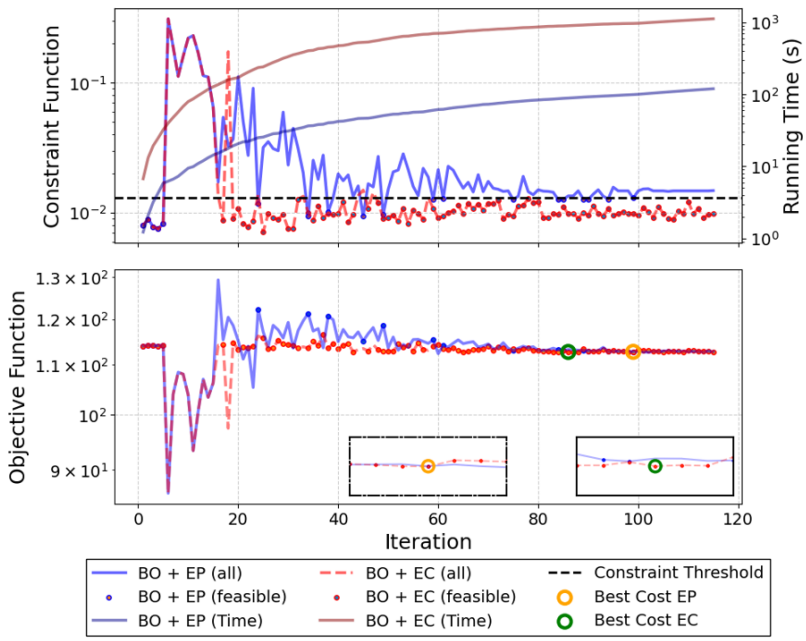

Figure 6.12: Optimisation speed comparison for BO+EC and BO+EP methods.

The power optimisation results for the case study voyage are illustrated in Fig. 6.13. The voyage was divided into nine legs. The optimal power allocation obtained from DP is shown in red, and the refined power allocation is shown in blue. The optimised strategy increases engine power during the initial legs of the voyage, when wave conditions are relatively mild, and applies a notable power reduction during legs 7 and 8 to reduce fuel consumption.

The BO+EP approach further refines the DP solution by slightly reducing power levels in legs 1, 4 and 9 and marginally increasing power in legs 2, 5 and 7. The optimised voyage duration is approximately 143.24 hours, representing an increase of about 0.7 hours compared to the DP solution; however, the vessel still arrives within approximately 1.5% of the ETA constraint. Under these conditions, the BO+EP strategy achieves a potential reduction of 8.88 t of fuel consumption, corresponding to a 7.30% decrease relative to the measured voyage fuel consumption. Additionally, the DP solution yields a potential reduction of 6.27%.

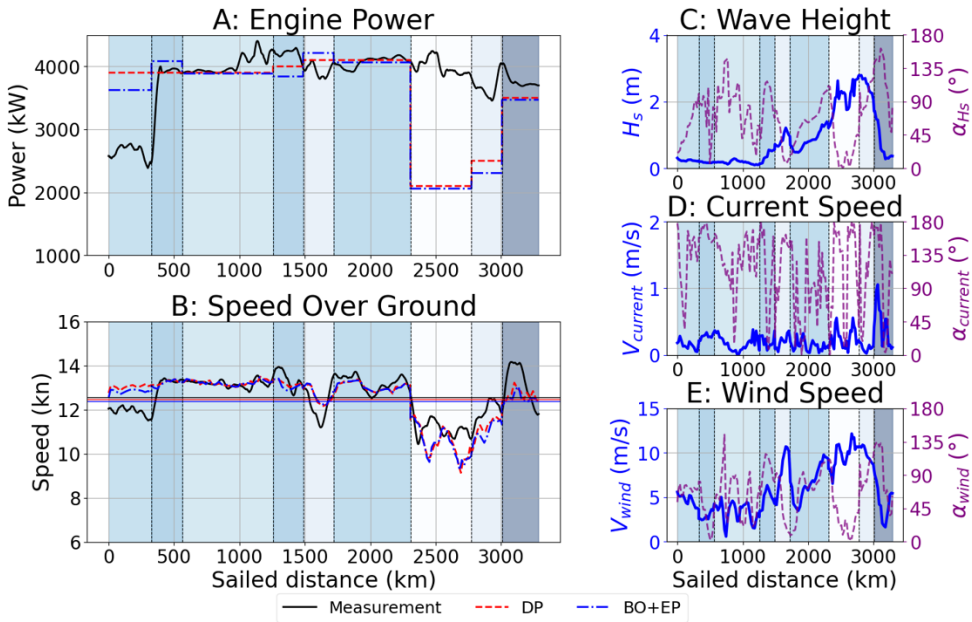

Figure 6.13: Optimisation results and metocean conditions for one case study voyage.

CHAPTER 7

Conclusions

This thesis has developed a comprehensive, data-driven framework for surrogate modelling and optimising ship fuel consumption in SSS. The work integrates ML, voyage simulation, route segmentation, and mathematical optimisation into a unified decision support framework. The main objective of the research was to improve energy efficiency whilst respecting operational constraints such as ETA. The key findings and contributions of the thesis are summarised below.

Machine learning-based ship performance surrogate modelling: ML models form the foundation of the proposed framework, enabling accurate prediction of ship fuel consumption and operational performance.

- A systematic comparison of supervised learning algorithms identified XGBoost as the most stable and reliable model for ship fuel consumption prediction.
- Ship speed prediction remains challenging during novel scenarios, highlighting the intrinsic difficulty of modelling speed using purely data-driven approaches.

Wave-induced resistance modelling based on Gaussian process regression (GPR): To address the limitations of purely data-driven and semi-empirical models under adverse weather conditions, the thesis proposed a method for predicting added resistance coefficients in head wave using GPR.

- The GPR model preserved the smoothness and physical consistency of resistance curves, outperforming conventional ANN-based models, particularly in terms of extrapolation to unseen hull forms.

- Compared with state-of-the-art semi-empirical methods, the GPR-based approach reduced prediction errors approximately threefold.
- When integrated into a grey-box fuel prediction model, the improved resistance estimation enabled more realistic quantification, indicating that head-sea waves account for approximately 8% of total fuel consumption on average.

Voyage segmentation based on metocean conditions: Accurate voyage segmentation is essential for linking environmental variability with power allocation and optimisation.

- A metocean-aware segmentation method, MS-PELT, was proposed to divide routes into segments with homogeneous environmental conditions.
- MS-PELT outperformed the state-of-the-art TICC method by producing fewer, longer, and more operationally meaningful segments.
- MS-PELT incurred considerably lower computational costs, with runtimes of approximately 10 ms compared to 50 s for TICC, enabling real-time or near-real-time application.

Power allocation optimisation: Building on the proposed performance models and segmentation framework, voyage optimisation strategies were investigated to improve operational energy efficiency.

- BO was applied within the DSS to optimise ferry power allocation, achieving fuel consumption reductions of up to 43% whilst keeping the sailing-time deviation within the imposed ETA margin.
- Full-scale trials validated the practical feasibility of the proposed operating strategy for the case study ferry, achieving an average fuel consumption reduction of approximately 18%.
- For the chemical tanker case studies, the two-stage framework combining MS-PELT and scenario-based dynamic programming enabled emission reductions of up to 14.8%.
- A discrete-to-continuous refinement strategy was also proposed, involving warm-starting BO from the DP solution. The BO+EP refinement achieved a potential reduction of 35.31 t of CO_2 (9.3% relative to the measured voyage emissions), compared with a 27.31 t reduction with the DP solution.

Overall, this thesis demonstrates that energy efficiency in maritime operations can be improved by combining machine learning ship performance models with physically consistent resistance prediction, metocean-aware route segmentation and optimisation-based power allocation. In practice, the realised savings will depend on operational constraints and human factors. The proposed framework establishes a solid foundation for future DSSs that reduce fuel consumption and emissions in SSS.

CHAPTER 8

Future research directions

To strengthen the scientific foundations and practical deployability of the proposed framework, future studies should investigate physically consistent, uncertainty-aware, and transferable optimisation systems. Although this thesis establishes a validated pipeline from data-driven modelling to voyage-level optimisation, its real-world effectiveness remains sensitive to model fidelity, environmental uncertainty, and operational execution. This warrants continued research to ensure both methodological robustness and industry relevance.

From the modelling perspective, future developments should focus on incorporating stronger physical structure into learning algorithms. Moving beyond black-box regression toward physics-informed or "glass-box" formulations would ensure that predictions remain consistent with hydrodynamic principles, especially in sparsely observed operational regimes. Furthermore, relaxing the assumption of constant displacement by incorporating dynamic cargo payload into the performance models could significantly improve fuel-prediction accuracy across varying capacity utilisation profiles, from fully laden to ballast conditions. Embedding uncertainty quantification, either via Bayesian neural networks or probabilistic surrogates, would allow optimisation algorithms to reason about prediction confidence rather than relying solely on point estimates. Enhancing model transferability is equally important; for example, techniques such as transfer learning, domain adaptation, and multi-task learning could enable models trained on high-quality datasets to be effectively deployed on vessels with limited or lower-resolution data. Data fusion approaches that integrate onboard sensors, noon-report data, loading computer outputs, AIS-derived

currents, and hindcast or forecast data will also improve robustness to distributional shifts.

Advancing toward closed-loop decision-making represents a promising direction for optimisation. Stochastic model-predictive control and safe reinforcement learning could support adaptive power-allocation strategies that respond to evolving metocean conditions, ETA constraints, and machinery limits whilst maintaining safety margins. Furthermore, expanding the optimisation boundary beyond the single vessel to incorporate multi-agent dynamics (e.g. traffic interactions, collision avoidance, and route congestion) would significantly enhance the framework's real-world applicability to dense short-sea corridors. Complementary improvements in route segmentation via adaptive, metocean-aware methods that detect regime changes would allow the optimiser to operate over segments with more homogeneous dynamics, improving solution stability and interpretability.

From an operational and economic perspective, future research should integrate voyage-optimisation objectives with the broader logistics and cost structures that govern short-sea transport. Incorporating fuel-price volatility, bunkering choices, port-slot availability, delays, and regulatory cost mechanisms (e.g., ETS allowances and GHG-intensity penalties) would enable multi-objective formulations aligned with real operator incentives. Critical to adoption will be the development of decision-support tools that interface seamlessly with bridge systems and voyage-planning workflows, along with the provision of interpretable recommendations rather than opaque black-box outputs. Finally, coupling voyage optimisation with long-horizon degradation models for hull fouling, propeller condition, and engine wear would extend the framework toward predictive maintenance and asset-management applications, enabling life-cycle optimisation rather than isolated voyage-level gains.

Collectively, these future directions aim to evolve the framework into a scientifically rigorous, operationally reliable, and economically compelling optimisation ecosystem capable of delivering sustained efficiency improvements at the fleet scale.

Appendix A - Data sources for added wave resistance

Table A.1: Ship hull models analysed per reference.

Reference	S175	JBC	KVLCC	Supramax	HSVA	Aframax	2020 Hull	Scb87	SR108
Ichinose et al. (2012)				✓					
Tsujimoto et al. (2012)				✓					
Sadat-Hosseini et al. (2013)	✓		✓						✓
Simonsen et al. (2013)							✓		
Söding et al. (2014)									
Oh et al. (2015)						✓			
Valanto and Hong (2015)					✓				
Park et al. (2016)	✓		✓						
Yu et al. (2017)				✓					
Iwashita and Kashiwagi (2018)				✓					
Chen et al. (2019)									
Diao et al. (2019)						✓			
Lee et al. (2019b)				✓					
Liu and Papanikolaou (2019)				✓				✓	
Liu et al. (2019)			✓		✓				
Park et al. (2019)	✓		✓	✓	✓				
Wicaksono (2019)				✓					✓
Kim et al. (2021a)									
Kobayashi et al. (2021)		✓							
Kim et al. (2022)	✓			✓	✓				

References

Abebe, M., Shin, Y., Noh, Y., Lee, S., & Lee, I. (2020). Machine learning approaches for ship speed prediction towards energy efficient shipping. *Applied Sciences*, *10*(7), 2325. https://doi.org/10.3390/app10072325

Bahrami, N., & Siadatmousavi, S. M. (2024). Ship voyage optimisation considering environmental forces using the iterative dijkstra's algorithm. *Ships and Offshore Structures*, *19*(8), 1173–1180. https://doi.org/10.1080/17445302.2023.2231200

Bassam, A. M., Phillips, A. B., Turnock, S. R., & Wilson, P. A. (2022). Ship speed prediction based on machine learning for efficient shipping operation. *Ocean Engineering*, *245*, 110449. https://doi.org/10.1016/j.oceaneng.2021.110449

Bellman, R. (1957). *Dynamic programming*. Princeton University Press.

Beşikçi, E. B., Arslan, O., Turan, O., & Ölçer, A. I. (2016). An artificial neural network based decision support system for energy efficient ship operations. *Computers & Operations Research*, *66*, 393–401. https://doi.org/10.1016/j.cor.2015.04.004

Birk, L. (2019). *Fundamentals of ship hydrodynamics: Fluid mechanics, ship resistance and propulsion*. John Wiley & Sons.

Carlton, J. S. (2007). *Marine Propellers and Propulsion* (Second). Butterworth-Heinemann. https://doi.org/10.1016/B978-075068150-6/50014-0.

Cepowski, T. (2020). The prediction of ship added resistance at the preliminary design stage by the use of an artificial neural network. *Ocean Engineering*, *195*, 106657.

Cepowski, T. (2023). The use of a set of artificial neural networks to predict added resistance in head waves at the parametric ship design stage. *Ocean Engineering*, *281*, 114744.

Chatterjee, S., & Simonoff, J. S. (2013). *Handbook of regression analysis*. John Wiley & Sons.

Chen, J., Duan, W., Ma, S., & Liao, K. (2021). Time-domain tebem method for mean drift force and moment of ships with forward speed under oblique seas. *Journal of Marine Science and Technology*, 1–13.

Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.

Chen, X., Ren, Y., Xiao, H.-s., Cai, X.-g., & Zhu, R.-c. (2019). A MDHOBEM Energy Radiated Method to Evaluate Added Wave Resistance of Ship.

Chen, Y. H., Zhang, C., Guo, Y. H., Wang, Y. Y., Lang, X., Zhang, M. Y., & Mao, W. G. (2025). State-of-the-art optimization algorithms in weather routing - ship decision support systems: Challenge, taxonomy, and review. *Ocean Engineering*, *331*. https://doi.org/10.1016/j.oceaneng.2025.121198

Chung, N., Balaji, G., Rudzki, K., Hoang, A. T., et al. (2025). Internet of things-driven approach integrated with explainable machine learning models for ship fuel consumption prediction. *Alexandria Engineering Journal*, *118*, 664–680.

Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, *74*(368), 829–836.

Comi, A., & Polimeni, A. (2020). Assessing the Potential of Short Sea Shipping and the Benefits in Terms of External Costs: Application to the Mediterranean Basin. *Sustainability*, *12*(13), 5383. https://doi.org/10.3390/su12135383

Coraddu, A., Oneto, L., Baldi, F., & Anguita, D. (2017). Vessels fuel consumption forecast and trim optimisation: A data analytics perspective. *Ocean Engineering*, *130*, 351–370.

Dai, X., Sheng, K., & Shu, F. (2022). Ship power load forecasting based on PSO-SVM. *Mathematical Biosciences and Engineering*, *19*(5), 4547–4567. https://doi.org/10.3934/mbe.2022210

Diao, F., Chen, J.-k., Duan, W.-y., Zhou, W.-x., Chen, J., & Wei, J.-f. (2019). Prediction of added resistance of a ship in waves at low speed. *Journal of Hydrodynamics*, *31*(6), 1231–1239. https://doi.org/10.1007/s42241-019-0062-9

DNV. (2025). *Mrv – monitoring, reporting and verification (eu and uk)* [Accessed: 2025-10-29]. https://www.dnv.com/maritime/insights/topics/mrv/

DNV GL. (2015). *Energy management study 2015* (Report). DNV GL. Høvik, Norway.

Du, W., Li, Y., Shi, J., Sun, B., Wang, C., & Zhu, B. (2023). Applying an improved particle swarm optimization algorithm to ship energy saving. *Energy*, *263*, 126080. https://doi.org/10.1016/j.energy.2022.126080

Du, Y., Meng, Q., Wang, S., & 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.

Duan, W., Yang, K., Huang, L., Jing, Y., & Ma, S. (2022). A dfn-based method for fast prediction of ships' added resistance in heading waves. *Ocean Engineering*, *245*, 110484.

Duarte, B., Saraiva, P. M., & Pantelides, C. C. (2004). Combined mechanistic and empirical modelling. *International Journal of Chemical Reactor Engineering*, *2*(1), Article A3.

European Parliament and Council. (2023a, May). Directive (eu) 2023/959 amending directive 2003/87/ec establishing a system for greenhouse gas emission allowance trading within the union [Official Journal of the European Union, L 130/134].

European Parliament and Council. (2023b, September). Regulation (eu) 2023/1805 on the use of renewable and low-carbon fuels in maritime transport, and amending directive 2009/16/ec [Official Journal of the European Union, L 234/48].

Eurostat. (2023). Maritime transport statistics - short sea shipping of goods [Accessed: 2025-10-28]. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Maritime_transport_statistics_-_short_sea_shipping_of_goods

Fadda, P., Fancello, G., Mancini, S., Pani, C., & Serra, P. (2020). Design and optimisation of an innovative two-hub-and-spoke network for the Mediterranean short-sea-shipping market. *Computers & Industrial Engineering*, *149*, 106847. https://doi.org/10.1016/j.cie.2020.106847

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

Fiacco, A. V., & McCormick, G. P. (1968). *Nonlinear programming: Sequential unconstrained minimization techniques*. John Wiley & Sons.

Fujii, H., & Takahashi, T. (1975). Experimental study on the resistance increase of a large full ship in regular oblique waves. *Journal of the Society of Naval Architects of Japan*, (137), 132–137. https://doi.org/10.2534/jjasnaoe1952.1975.137_132

Gelbart, M. A., Snoek, J., & Adams, R. P. (2014). Bayesian optimization with unknown constraints. *arXiv preprint arXiv:1403.5607*.

Godet, A., Wallner, L., Panagakos, G., & Barfod, M. (2024). Developing correction factors for weather's influence on the energy efficiency indicators of container ships using model-based machine learning. *Ocean Coastal Management*, *258*, 107390.

Guericke, S., & Tierney, K. (2015). Liner shipping cargo allocation with service levels and speed optimization. *Transportation Research Part E: Logistics and Transportation Review*, *84*, 40–60. https://doi.org/10.1016/j.tre.2015.10.002

Hallac, D., Leskovec, J., & Boyd, S. (2017). Toeplitz inverse covariance-based clustering of multivariate time series data. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 215–223. https://doi.org/10.1145/3097983.3098066

Hasselmann, K., Barnett, T. P., Bouws, E., Carlson, H., Cartwright, D. E., Enke, K., Ewing, J. A., Gienapp, H., Hasselmann, D. E., Kruseman, P., Meerburg, A., Muller, P., Olbers, D. J., Richter, K., Sell, W., & Walden, H. (1973). Measurements of wind-wave growth and swell decay during the joint North sea wave project (JONSWAP). *Ergnzungsheft zur Deutschen Hydrographischen Zeitschrift Reihe*.

Hollenbach, K. (1998). Estimating resistance and propulsion for single-screw and twin-screw ships. *Schiffstechnik/Ship Technology Research*, *45*(2), 72–76.

Holtrop, J., & Mennen, G. (1982). An approximate power prediction method. *International Shipbuilding Progress*, *29*, 166–170.

Hsieh, T. H., Li, Y. Z., Wang, S. Z., Liu, W., & Sun, Z. (2025). Ship ocean voyage weather routing optimization method based on weather clustering. *Applied Ocean Research*, *165*. https://doi.org/10.1016/j.apor.2025.104799

Ichinose, Y., Tsujimoto, M., Shiraishi, K., & Sogihara, N. (2012). Decrease of Ship Speed in Actual Seas of a Bulk Carrier in Full Load and Ballast Conditions. *Journal of the Japan Society of Naval Architects and Ocean Engineers*, *15*(0), 37–45. https://doi.org/10.2534/jjasnaoe.15.37

IMO. (2019). Frequently asked questions: The 2020 global sulphur limit. https://www.imo.org

IMO. (2021). *Fourth imo greenhouse gas study 2020* (Report No. MEPC 75/7/15). IMO. London, UK.

IMO. (2023). Mepc.1/circ.778/rev.4: List of special areas, emission control areas and particularly sensitive sea areas. *Marine Environment Protection Committee*.

ITTC. (2017). *Recommended procedures and guidelines – resistance tests and correlation of model test results* (tech. rep. No. 7.5-02-02-02 / 7.5-02-01-01) (ITTC Recommended Procedures and Guidelines, 28th ITTC, Wuxi, China). International Towing Tank Conference (ITTC). https://ittc.info/media/1219/75-02-02-02.pdf

Iwashita, H., & Kashiwagi, M. (2018). An Innovative EFD for Studying Ship Seakeeping.

Jaramillo, P., Kahn Ribeiro, S., Newman, P., Dhar, S., Diemuodeke, O. E., Kajino, T., Lee, D. S., Nugroho, S. B., Ou, X., Hammer Strømman, A., & Whitehead, J. (2022). Transport. In P. R. Shukla, J. Skea, R. Slade, A. Al Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas, R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz, & J. Malley (Eds.), *Climate change 2022: Mitigation of climate change. contribution of working group iii to the sixth assessment report of the intergovernmental panel on climate*

change (p. 1055). Cambridge University Press. https://doi.org/10.1017/9781009157926.012

Karagiannidis, P., & Themelis, N. (2021). Data-driven modelling of ship propulsion and the effect of data pre-processing on the prediction of ship fuel consumption and speed loss. *Ocean Engineering*, *222*, 108616. https://doi.org/10.1016/j.oceaneng.2021.108616

Killick, R., Fearnhead, P., & Eckley, I. A. (2012). Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, *107*(500), 1590–1598. https://doi.org/10.1080/01621459.2012.737745

Kim, B.-S., Oh, M.-J., Lee, J.-H., Kim, Y.-H., & Roh, M.-I. (2021a). Study on Hull Optimization Process Considering Operational Efficiency in Waves. *Processes*, *9*(5), 898. https://doi.org/10.3390/pr9050898

Kim, K., & Kim, Y. (2011). Numerical study on added resistance of ships by using a time-domain rankine panel method. *Ocean Engineering*, *38*(13), 1357–1367. https://doi.org/10.1016/j.oceaneng.2011.07.010

Kim, T., Yoo, S., & Kim, H. (2021b). Estimation of added resistance of an lng carrier in oblique waves. *Ocean Engineering*, *231*, 109068.

Kim, Y.-R., Esmailian, E., & Steen, S. (2022). A meta-model for added resistance in waves. *Ocean Engineering*, *266*, 112749. https://doi.org/10.1016/j.oceaneng.2022.112749

Kim, Y.-R., Jung, M., & Park, J.-B. (2021c). 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. https://doi.org/10.3390/jmse9020137

Kobayashi, H., Kume, K., Orihara, H., Ikebuchi, T., Aoki, I., Yoshida, R., Yoshida, H., Ryu, T., Arai, Y., Katagiri, K., Ikeda, S., Yamanaka, S., Akibayashi, H., & Mizokami, S. (2021). Parametric study of added resistance and ship motion in head waves through RANS : Calculation guideline. *Applied Ocean Research*, *110*, 102573. https://doi.org/10.1016/j.apor.2021.102573

Lang, X., & Mao, W. (2020). A semi-empirical model for ship speed loss prediction at head sea and its validation by full-scale measurements. *Ocean Engineering*, *209*, 107494. https://doi.org/10.1016/j.oceaneng.2020.107494

Lang, X., & Mao, W. (2021). A practical speed loss prediction model at arbitrary wave heading for ship voyage optimization. *Journal of Marine Science and Application*, *20*(3), 410–425. https://doi.org/10.1007/s11804-021-00224-z

Lang, X., Wang, H., Mao, W., & Osawa, N. (2020). Impact of ship operations aided by voyage optimization on a ship's fatigue assessment. *Journal of Marine Science and Technology*. https://doi.org/10.1007/s00773-020-00769-8

Lang, X., Wu, D., & Mao, W. (2022a). Comparison of supervised machine learning methods to predict ship propulsion power at sea. *Ocean Engineering*, *245*, 110387. https://doi.org/10.1016/j.oceaneng.2021.110387

Lang, X., Wu, D., & Mao, W. (2022b). Comparison of supervised machine learning methods to predict ship propulsion power at sea. *Ocean Engineering*, *245*, 110387. https://doi.org/10.1016/j.oceaneng.2021.110387

Lang, X., Wu, D., & Mao, W. (2024). Physics-informed machine learning models for ship speed prediction. *Expert Systems with Applications*, *238*, 121877. https://doi.org/10.1016/j.eswa.2023.121877

Latarche, M. (2021). Pounder's Marine Diesel Engines and Gas Turbines. Butterworth-Heinemann.

Laurie, A., Anderlini, E., Dietz, J., & Thomas, G. (2021). Machine learning for shaft power prediction and analysis of fouling related performance deterioration. *Ocean Engineering*, *234*, 108886. https://doi.org/10.1016/j.oceaneng.2021.108886

Lee, C., Park, S., Yu, J., Choi, J., & Lee, I. (2019a). Effects of diffraction in regular head waves on added resistance and wake using cfd. *International Journal of Naval Architecture and Ocean Engineering*, *11*(2), 736–749.

Lee, C.-M., Yu, J.-W., Choi, J.-E., & Lee, I. (2019b). Effect of bow hull forms on the resistance performance in calm water and waves for 66k DWT bulk carrier. *International Journal of Naval Architecture and Ocean Engineering*, *11*(2), 723–735. https://doi.org/10.1016/j.ijnaoe.2019.02.007

Lee, S.-M., Roh, M.-I., Kim, K.-S., Jung, H., & Park, J. J. (2018). Method for a simultaneous determination of the path and the speed for ship route planning problems. *Ocean Engineering*, *157*, 301–312. https://doi.org/https://doi.org/10.1016/j.oceaneng.2018.03.068

Leifsson, L. T., Sævarsdóttir, H., Sigurdsson, S. T., & Vésteinsson, A. (2008). Grey-box modeling of an ocean vessel for operational optimization. *Simulation Modelling Practice and Theory*, *16*(8), 923–932.

Li, H. J., Ma, Q. Q., Si, R. B., Wang, X. J., & Mao, W. G. (2026). Adaptive path planning for dry bulk carriers under sudden maritime disturbances: A network-based approach. *Ocean Engineering*, *343*. https://doi.org/10.1016/j.oceaneng.2025.123474

Li, X., Sun, B., Jin, J., & Ding, J. (2022). Speed optimization of container ship considering route segmentation and weather data loading: Turning point-time segmentation method. *Journal of Marine Science and Engineering*, *10*(12), 1835. https://doi.org/10.3390/jmse10121835

Li, X., Sun, B., Jin, J., & Ding, J. (2023). Ship speed optimization method combining fisher optimal segmentation principle. *Applied Ocean Research*, *140*, 103743. https://doi.org/10.1016/j.apor.2023.103743

Li, Z., Wang, K., Hua, Y., Liu, X., Ma, R., Wang, Z., & Huang, L. (2024). Ga-lstm and nsga-iii based collaborative optimization of ship energy efficiency for low-carbon shipping. *Ocean Engineering*, *312*, 119190. https://doi.org/10.1016/j.oceaneng.2024.119190

Lindskog, P., & Ljung, L. (1994). Tools for semi-physical modeling [Linköping University, Sweden]. *Proceedings of the 10th IFAC Symposium on System Identification (SYSID'94)*.

Liu, S., & Papanikolaou, A. (2016). Fast approach to the estimation of the added resistance of ships in head waves. *Ocean Engineering*, *112*, 211–225. https://doi.org/10.1016/j.oceaneng.2015.11.050

Liu, S., & Papanikolaou, A. (2020). Regression analysis of experimental data for added resistance in waves of arbitrary heading and development of a semi-empirical formula. *Ocean Engineering*, *206*, 107357. https://doi.org/10.1016/j.oceaneng.2020.107357

Liu, S., & Papanikolaou, A. (2019). Approximation of the added resistance of ships with small draft or in ballast condition by empirical formula. *Proceedings of the Institution of Mechanical Engineers, Part M: Journal of Engineering for the Maritime Environment*, *233*(1), 27–40. https://doi.org/10.1177/1475090217710099

Liu, S., Papanikolaou, A., Feng, P., & Fan, S. (2019). A Multi-Level Approach to the Prediction of the Added Resistance and Powering of Ships in Waves. *Volume 7B: Ocean Engineering*, V07BT06A037. https://doi.org/10.1115/OMAE2019-95113

Lloyd's Register. (2025). *What is the eu mrv & ets* [Accessed: 2025-10-29]. https://www.lr.org/en/services/statutory-compliance/eu-ets-and-eu-mrv/

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.

Ma, D., Ma, W., Jin, S., & Ma, X. (2020). Method for simultaneously optimizing ship route and speed with emission control areas. *Ocean Engineering*, *202*, 107170. https://doi.org/https://doi.org/10.1016/j.oceaneng.2020.107170

MAN. (2023, April). *Basic principles of ship propulsion: Optimisation of hull, propeller, and engine interactions for maximum efficiency* (tech. rep.). MAN Energy Solutions. Copenhagen, Denmark. https://www.man-es.com

Mao, W., Rychlik, I., Wallin, J., & Storhaug, G. (2016). Statistical models for the speed prediction of a container ship. *Ocean Engineering*, *126*, 152–162. https://doi.org/10.1016/j.oceaneng.2016.08.033

Mehlig, B. (2021). *Machine learning with neural networks: An introduction for scientists and engineers*. Cambridge University Press.

Mockus, J., Tiešis, V., & Žilinskas, A. (1978). The application of bayesian methods for seeking the extremum. In L. C. W. Dixon & G. P. Szegö (Eds.), *Towards global optimization 2* (pp. 117–129). North-Holland.

Odendaal, K., Alkemade, A., & Kana, A. (2023). Enhancing early-stage energy consumption predictions using dynamic operational voyage data: A grey-

box modelling investigation. *International Journal of Naval Architecture and Ocean Engineering*, *15*, 100484.

Oh, S., Yang, J., & Park, S.-H. (2015). Computational and Experimental Studies on Added Resistance of AFRAMAX-Class Tankers in Head Seas. *Journal of the Society of Naval Architects of Korea*, *52*(6), 471–477. https://doi.org/10.3744/SNAK.2015.52.6.471

Park, D.-M., Kim, Y., Seo, M.-G., & Lee, J. (2016). Study on added resistance of a tanker in head waves at different drafts. *Ocean Engineering*, *111*, 569–581. https://doi.org/10.1016/j.oceaneng.2015.11.026

Park, D.-M., Lee, J.-H., Jung, Y.-W., Lee, J., Kim, Y., & Gerhardt, F. (2019). Experimental and numerical studies on added resistance of ship in oblique sea conditions. *Ocean Engineering*, *186*, 106070. https://doi.org/10.1016/j.oceaneng.2019.05.052

Parkes, A., Sobey, A., & Hudson, D. (2018). Physics-based shaft power prediction for large merchant ships using neural networks. *Ocean Engineering*, *166*, 92–104. https://doi.org/10.1016/j.oceaneng.2018.07.060

Qi, X., & Song, D.-P. (2012). Minimizing fuel emissions by optimizing vessel schedules in liner shipping with uncertain port times. *Transportation Research Part E: Logistics and Transportation Review*, *48*(4), 863–880. https://doi.org/10.1016/j.tre.2012.02.001

Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. The MIT Press. http://www.gaussianprocess.org/gpml/

Rodrigue, J.-P. (2024). *The geography of transport systems* (6th). Routledge. https://doi.org/10.4324/9781003343196

Sadat-Hosseini, H., Wu, P.-C., Carrica, P. M., Kim, H., Toda, Y., & Stern, F. (2013). CFD verification and validation of added resistance and motions of KVLCC2 with fixed and free surge in short and long head waves. *Ocean Engineering*, *59*, 240–273. https://doi.org/10.1016/j.oceaneng.2012.12.016

Sævarsdóttir, H., Sigurðsson, S. Þ., Leifsson, L. Þ., & Vésteinsson, A. (2005). Grey-box modeling of an ocean vessel for operational optimization. *Proceedings of the 46th Conference on Simulation and Modelling (SIMS 2005)*.

Savitzky, A., & Golay, M. J. E. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, *36*(8), 1627–1639.

Shang, C., Fu, L., Bao, X., Xiao, H., Xu, X., & Hu, Q. (2024). Dynamic joint optimization of power generation and voyage scheduling in ship power system based on deep reinforcement learning. *Electric Power Systems Research*, *229*, 110165.

Simonsen, C. D., Otzen, J. F., Joncquez, S., & Stern, F. (2013). EFD and CFD for KCS heaving and pitching in regular head waves. *Journal of Marine Science and Technology*, *18*(4), 435–459. https://doi.org/10.1007/s00773-013-0219-0

Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems*, *25*.

Söding, H., Shigunov, V., Schellin, T. E., & Moctar, O. E. (2014). A Rankine Panel Method for Added Resistance of Ships in Waves. *Journal of Offshore Mechanics and Arctic Engineering*, *136*(3), 031601. https://doi.org/10.1115/1.4026847

Sørensen, A. J. (2013). *Marine cybernetics: Modelling and control of marine systems*. NTNU.

Sørensen, A. J., Fossen, T. I., & Berge, S. P. (1997). Marine dynamic positioning systems. *Control Engineering Practice*, *5*(4), 481–486.

Styhre, L. (2010). *Capacity utilisation in short sea shipping* [Doctoral dissertation, Chalmers University of Technology] [Department of Technology Management and Economics]. https://publications.lib.chalmers.se/records/fulltext/124588.pdf

Su, M., Lee, H. J., Wang, X., & Bae, S.-H. (2025). Fuel consumption cost prediction model for ro-ro carriers: A machine learning-based application. *Maritime Policy & Management*, *52*(2), 229–249.

Sun, Q., Zhang, M., Zhou, L., Garme, K., & Burman, M. (2022). A machine learning-based method for prediction of ship performance in ice: Part i. ice resistance. *Marine Structures*, *83*, 103181.

Szlapczynska, J., & Szlapczynski, R. (2019). Preference-based evolutionary multi-objective optimization in ship weather routing. *Applied Soft Computing*, *84*, 105742. https://doi.org/https://doi.org/10.1016/j.asoc.2019.105742

Tsujimoto, M., Kuroda, M., Shiraishi, K., Ichinose, Y., & Sogihara, N. (2012). Verification on the Resistance Test in Waves Using the Actual Sea Model Basin. *Journal of the Japan Society of Naval Architects and Ocean Engineers*, *16*(0), 33–39. https://doi.org/10.2534/jjasnaoe.16.33

Tzortzis, G., & Sakalis, G. (2021). A dynamic ship speed optimization method with time horizon segmentation. *Ocean Engineering*, *226*, 108840. https://doi.org/https://doi.org/10.1016/j.oceaneng.2021.108840

UNCTAD. (2025). *Review of maritime transport 2025: Staying the course in turbulent waters* (tech. rep. No. UNCTAD/RMT/2025) (Published September 24, 2025). United Nations Conference on Trade and Development. Geneva, Switzerland. https://unctad.org/publication/review-maritime-transport-2025

Valanto, P., & Hong, Y. (2015). Experimental Investigation on Ship Wave Added Resistance in Regular Head, Oblique, Beam, and Following Waves.

Vergara, D., Alexandersson, M., Lang, X., & Mao, W. (2023). A machine learning-based bayesian decision support system for efficient navigation of double-

ended ferries [In Press]. *Journal of Ocean Engineering and Science*. https://doi.org/10.1016/j.joes.2023.11.002

Vergara, D., Lang, X., Zhang, M. Y., Alexandersson, M., & Mao, W. G. (2025). Reduced environmental impact of short sea shipping through optimal propulsion power allocation. *Journal of Cleaner Production*, *513*. https://doi.org/10.1016/j.jclepro.2025.145683

Wang, H., Lang, X., & Mao, W. (2021a). Voyage optimization combining genetic algorithm and dynamic programming for fuel/emissions reduction. *Transportation Research Part D: Transport and Environment*, *90*, 102670. https://doi.org/10.1016/j.trd.2020.102670

Wang, H., Mao, W., & Eriksson, L. (2019). A three-dimensional dijkstra's algorithm for multi-objective ship voyage optimization. *Ocean Engineering*, *186*, 106131. https://doi.org/10.1016/j.oceaneng.2019.106131

Wang, K., Li, J., Huang, L., Ma, R., Jiang, X., Yuan, Y., Mwero, N. A., Negenborn, R. R., Sun, P., & Yan, X. (2020). A novel method for joint optimization of the sailing route and speed considering multiple environmental factors for more energy efficient shipping. *Ocean Engineering*, *216*, 107591. https://doi.org/10.1016/j.oceaneng.2020.107591

Wang, K., Xu, H., Li, J., Huang, L., Ma, R., Jiang, X., Yuan, Y., Mwero, N. A., Sun, P., Negenborn, R. R., & Yan, X. (2021b). A novel dynamical collaborative optimization method of ship energy consumption based on a spatial and temporal distribution analysis of voyage data. *Applied Ocean Research*, *112*, 102657. https://doi.org/https://doi.org/10.1016/j.apor.2021.102657

Wang, K., Yan, X., Yuan, Y., Jiang, X., Lodewijks, G., & Negenborn, R. R. (2017). Study on route division for ship energy efficiency optimization based on big environment data. *Proceedings of the 4th International Conference on Transportation Information and Safety (ICTIS 2017)*, 111–116. https://doi.org/10.1109/ICTIS.2017.8047752

Wang, S., & 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. https://doi.org/10.1016/j.tre.2011.12.003

Wen, Y., Sui, Z., Zhou, C., Xiao, C., Chen, Q., Han, D., & Zhang, Y. (2019). Automatic ship route design between two ports: A data-driven method. *Applied Ocean Research*, *88*, 102049.

Wicaksono, A. (2019). A Unified Computation Method for Seakeeping-Maneuvering of a Ship in Waves Using Slender-Ship Theory and MMG Model.

Wu, W.-M. (2020). The optimal speed in container shipping: Theory and empirical evidence. *Transportation Research Part E: Logistics and Transportation Review*, *136*, 101903. https://doi.org/10.1016/j.tre.2020.101903

Xiao, H., Ren, Y., Cai, X., & Zhu, R. (2022). Development, installation, and evaluation of energy efficiency measures for ships. *Ocean Engineering*.

Yan, X. P., Wang, K., Yuan, Y. P., Jiang, X. L., & Negenborn, R. R. (2018). Energy-efficient shipping: An application of big data analysis for optimizing engine speed of inland ships considering multiple environmental factors. *Ocean Engineering*, *169*, 457–468. https://doi.org/10.1016/j.oceaneng.2018.08.050

Yu, J.-W., Lee, C.-M., Lee, I., & Choi, J.-E. (2017). Bow hull-form optimization in waves of a 66,000 DWT bulk carrier. *International Journal of Naval Architecture and Ocean Engineering*, *9*(5), 499–508. https://doi.org/10.1016/j.ijnaoe.2017.01.006

Yu, Y., Zhang, H., Mu, Z., Li, Y., Sun, Y., & Liu, J. (2024). Trim and engine power joint optimization of a ship based on minimum energy consumption over a whole voyage. *Journal of Marine Science and Engineering*, *12*(3), 475. https://doi.org/10.3390/jmse12030475

Zaccone, R., Ottaviani, E., Figari, M., & Altosole, M. (2018). Ship voyage optimization for safe and energy-efficient navigation: A dynamic programming approach. *Ocean Engineering*, *153*, 215–224. https://doi.org/10.1016/j.oceaneng.2018.01.100

Zhang, C. Q., Zhang, C., Thies, F., Mao, W. A., & Ringsberg, J. W. (2025a). A voyage planning framework for energy performance analysis of autonomous inland waterway vessels. *Energy*, *335*. https://doi.org/10.1016/j.energy.2025.137906

Zhang, C., Vergara, D., Zhang, M., Tsoulakos, N., & Mao, W. (2025b). A machine learning method to evaluate head sea induced weather impact on ship fuel consumption. *Energy*, *328*, 132717. https://doi.org/https://doi.org/10.1016/j.energy.2025.132717

Zhang, C., Zhang, D., Zhang, M., Zhang, J., & Mao, W. (2022). A three-dimensional ant colony algorithm for multi-objective ice routing of a ship in the arctic area. *Ocean Engineering*, *266*, 113241. https://doi.org/https://doi.org/10.1016/j.oceaneng.2022.113241

Zhang, M., Tsoulakos, N., Kujala, P., & Hirdaris, S. (2024). A deep learning method for the prediction of ship fuel consumption in real operational conditions. *Engineering Applications of Artificial Intelligence*, *130*, 107425. https://doi.org/10.1016/j.engappai.2023.107425

Zhang, S., Shi, G., Liu, Z., Zhao, Z., & Wu, Z. (2018). Data-driven based automatic maritime routing from massive ais trajectories in the face of disparity. *Ocean Engineering*, *159*, 56–65.

Zhao, X., Guo, Y. H., & Wang, Y. Y. (2025). Green maritime navigation: A multi-objective voyage optimization approach based on data-driven heuristics and emission awareness. *Ocean Engineering*, *318*. https://doi.org/10.1016/j.oceaneng.2024.120138

Postscript

My PhD started in January 2022. Even though my contract says that it was in June, the foundations for this PhD started in January when, under the supervision of Wengang, I worked on a Master's thesis in the subject of modelling for double-ended ferries. During that moment, I knew I wanted to work for Wengang because I truly admired (and still do) his knowledge in the subjects of Machine Learning and mathematics in general. I was excited to start this chapter and, honestly, I underestimated how difficult doing a PhD was.

The first two years, thanks to Wengang's supervision, were extremely productive. We have so many fun memories, among them the ISOPE 2023 conference in Canada. The years passed and during 2024 Xiao joined the team at Wengang's suggestion and, under his instruction, Xiao began to help in my research project. Xiao indeed helped a lot and he became a friend through our interactions. I admire his hard work; I know he will get far in academia.

It was in 2024 when we became worried about my stunted progress; I went through so many problems that year. By 2025 I was a different person already. Yet somehow, most of the research was completed. It did require a lot of work on our part, but we did it and by March I had already defended my licentiate. That was a huge milestone, and I think that was the moment when we started to believe that I could make it to the end—at least I did. It was in 2025 when Jonas joined us to guide us through the last step.

If I were to start this PhD again I would likely have done many things very differently. I would have spent more time to improve my communication skills much earlier than I did. I now think that being able to communicate assertively is the most important skill an academic could develop as it saves you from so many misunderstandings. Learning to understand others and connect in a meaningful manner could have brought me further, earlier. I would have had less ambition so that I could have enjoyed the PhD process for longer. I would have cast aside my pride and embraced something that I do now more often, admitting defeat, admitting when I am wrong. There is nothing wrong to answer a question with "I honestly don't know" as there is no greater joy than to learn something.

My interactions with Wengang taught me that communication is the key to success. That you must always do good science. That you must be independent in your work. That you won't likely be able to replicate the method from other researchers because we don't share the data often (mainly because we aren't allowed by the data providers), but that you can learn from all of them ideas that you can then use to improve your own research. That there is value in most things we work on. Most importantly, that we should never give up on something we have started. I will carry on with these lessons for the rest of my life.

The PhD was a rollercoaster. It was full of ups and downs. But here I am, at the end. In the future I will see these days with nostalgia, and remember the times when the guys and me would sit in the lunch table and just discuss nonsense about anything. The division of MT became a family to me here in Sweden. My brothers Qais and Rui, Mohammad and Stephan, Christopher and Nils... I hope that you have learned something from me, and if you didn't, that at least you had fun in or shared moments. It was fun to interact with all of you.

This is my closing note. Thank you to every one of you. But especially to you, Wengang. We had good and bad memories, we had good and bad moments. We argued and laughed. Honestly, it saddens me that my PhD time has come to an end. But this thesis, this work, and this PhD wouldn't have been possible without your advice and guidance, Wengang. I couldn't just say it in the preface like a normal person just to play you one last prank before we part ways[1]. Thank you for not giving up on me when I wanted give up. Thank you for believing in my work, even through the uncertainty. Thank you for being strict when I needed it the most. Thank you for helping me grow as a person. I wouldn't want to go... but soon I will have to. Till the next time everyone.

Daniel Vergara
Gothenburg,
May 2026

"Move well, study well, play well, eat well, rest well -
That is the turtle master way!"

— Akira Toriyama, Dragon Ball, Vol. 1

[1] I hid this text at the very end of my thesis to make you worry one last time before I leave—sorry.