



ELSEVIER

Contents lists available at ScienceDirect

Ocean Engineering

journal homepage: www.elsevier.com/locate/oceaneng

Research paper

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

Daniel Vergara ^a, Shanshan Fu ^b, Wengang Mao ^{a,*}, Jonas W. Ringsberg ^a, Tsoulakos Nikolaos ^c

^a Division of Marine Technology, Department of Mechanical Engineering, Chalmers University of Technology, Gothenburg, 41296, Sweden

^b College of Transport and Communications, Shanghai Maritime University, Shanghai, 201308, China

^c Laskaridis Shipping Company Co. Ltd., Kifissia, Athens, 13677, Greece

ARTICLE INFO

Keywords:

Bayesian optimisation
Dynamic programming
Power allocation
Energy efficiency
Short-sea shipping

ABSTRACT

Optimising propulsion power along a ship's voyage is critical for energy efficient voyage planning, as it enables the distribution of propulsion settings along a route to minimise fuel consumption while satisfying operational constraints such as target arrival time. This study proposes a structured discrete-to-continuous optimisation framework to enhance propulsion power allocation strategies. Feasible solutions generated by a modified parallel coupling dynamic programming approach serve as prior knowledge (PCDP), and are subsequently refined using constraint-aware Bayesian optimisation (BO). When combined with a smooth exponential penalty function (BO + EP), the Bayesian optimiser embeds convex constraints directly within the optimisation process, improving convergence behaviour. The framework is validated using full-scale data from four voyages of a chemical tanker. Results demonstrate that the proposed method achieves a tenfold reduction in computational cost compared with explicit constraint (EC), while providing up to around 3% additional fuel savings compared with dynamic programming. Overall, the proposed power allocation strategy reduces fuel consumption by up to 9% on average while satisfying arrival time requirements.

1. Introduction

The transport sector is the largest source of greenhouse gas emissions in the European Union (EU). Achieving the EU's 2050 decarbonisation strategy requires close cooperation across all transport sectors to improve overall efficiency. Promoting inland and short-sea shipping is a key strategy, as it provides a sustainable alternative to road transport by alleviating congestion and reducing infrastructure dependence in coastal regions (Fadda et al., 2020). In response, strict emissions regulations for these waters have been enacted by both the IMO and the EU (DNV, 2024a,b). Moreover, many short-sea routes in the EU fall within designated Emission Control Areas. The Emissions Trading System (ETS) is expected to further incentivise improvements in shipping energy efficiency by increasing the economic pressure to reduce fuel consumption (Hu et al., 2022).

Among available energy efficiency measures, voyage optimisation demonstrates significant potential to lower fuel consumption, operating costs, and air emissions (Chen and Mao, 2024). Ship voyage optimisation aims to determine the optimal route and/or operational settings,

such as vessel speed and engine power, along the voyage to minimise fuel consumption and emissions while ensuring timely arrival (Chen et al., 2025). By sailing an optimal route with more favourable environmental conditions, e.g., calm sea or current assistance, ships can sail faster on reduced energy.

Various optimisation algorithms have been investigated for voyage planning problems, which are often classified as continuous or discrete methods. Continuous algorithms include gradient-based methods or metaheuristic approaches. For example, evolutionary algorithms (EAs) have been widely employed, and genetic algorithms (GA) can handle complex, multi-objective optimisation problems effectively (Wang et al., 2021). Those methods have been applied to optimise routes, speeds, and other operational parameters such as trim or engine speed (Lee et al., 2018). Their capabilities have been further extended through integration with complementary techniques. For instance, Szlapeczynska and Szlapeczynski (2019) developed an evolutionary multi-objective framework to balance fuel efficiency, voyage duration, and safety. While Li et al. (2024) proposed a hybrid GA long short-term memory (GA-LSTM) model combined with the

* Corresponding author.

E-mail addresses: danver@chalmers.se (D. Vergara), ssfu@shmtu.edu.cn (S. Fu), wengang.mao@chalmers.se (W. Mao), jonas.ringsberg@chalmers.se (J.W. Ringsberg), tsoulakos@laskaridis.com (T. Nikolaos).

<https://doi.org/10.1016/j.oceaneng.2026.125082>

Received 30 September 2025; Received in revised form 2 March 2026; Accepted 12 March 2026

Available online 31 March 2026

0029-8018/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Nomenclature

ANN	Artificial Neural Network	N	Total number of voyage legs
BO	Bayesian Optimisation	MAE	Mean absolute error
DP	Dynamic Programming	MAPE	Mean absolute percentage error
EC	Explicit constraint	MARS	Meteorological Archival and Retrieval System
EP	Embedded penalty	MCR	Maximum continuous rating
ERA5	ECMWF climate reanalysis data	MS-PELT	MeteOcean Score-PELT
ETA	Estimated Time of Arrival	MSE	Mean squared error
ETS	Emissions Trading System	PCDP	Parallel Coupling Dynamic Programming
JIT	Just-In-Time	R^2	Coefficient of determination
α	BO LCB acquisition function	RMSE	Root mean squared error
α (XGB)	L1 regularisation parameter	κ	BO covariance kernel
α_C	Constraint acquisition function in BO + EC	λ	L2 regularisation parameter
$\alpha_{current}$	Sea current direction	μ	GP mean function
α_{wave}	Sea wave direction	ν	GP smoothness parameter
α_{wind}	Wind direction	$\xi_{i,j}$	Waypoint coordinates
γ	Minimum loss reduction for leaf partition	σ	GP variance function
Γ	Gamma function	Φ	Gaussian probability distribution
Δt	Temporal resolution / sailing time	\Pr	Probability
ϵ	GP white noise	\mathbb{Q}	Set of rational numbers
C_{ETA}	ETA Constraint	\mathbb{R}	Set of real numbers
$C_{ETA}^{up/low}$	ETA constraint bounds	P	Engine power
D_{tol}	Tolerable ETA difference	\mathbf{P}	Vector of propulsion powers
\hat{f}	Predicted performance value	$\hat{\mathbf{P}}$	Optimal propulsion power vector
f_{fuel}	Fuel consumption rate function	r	BO decay coefficient
H_s	Significant wave height	T_{ETA}	Expected sailing time
\mathbf{I}	Identity matrix	T_m	Mean wave period
i, j	Indices for legs and waypoints	T_{sail}	Total sailing time
ites	Optimisation iteration count	V	Ship speed
J	DP objective function	V_{CP}	Charter-party speed
k	BO exploration parameter	$V_{current}$	Sea current speed
k_1	Amplitude coefficient in EP	V_{wind}	Wind speed
k_2	Growth rate coefficient in EP	$W_{i,j}$	MetOcean conditions
k_i	Number of waypoints in leg i	w_{fuel}	Penalty weighting function
\mathbf{K}	Covariance matrix	x	Longitude
K_v	Modified Bessel function	\mathbf{X}	Set of power observations
l	GP length scale	y	Latitude
L	Sequence of voyage legs	\mathbf{Y}	Set of fuel cost evaluations
L_i	Voyage leg i	\mathbf{Z}	Set of sailing time observations
		M_{fuel}	Voyage fuel consumption

non-dominated sorting GA (NSGA-III) for multi-criteria optimisation. Particle swarm optimisation (PSO) has also been applied to route and speed optimisation under environmental constraints (Wang et al., 2020). More recent PSO formulations integrate higher-order update rules (Du et al., 2023) to address dynamic metocean conditions for reduced fuel consumption. Other recent PSO formulations integrate higher-order update rules (Du et al., 2023) to address dynamic metocean conditions for reduced fuel consumption. Additional advancements also include multi-objective PSO frameworks that optimise speed and trim simultaneously while adjusting optimisation frequency based on sea condition clustering (Zhang et al., 2025). Furthermore, the integration of reinforcement learning (RL), in the form of Proximal Policy Optimisation (PPO), has enabled adaptive waypoint density control, further to refine continuous route planning (Lee et al., 2025). Despite these advances, continuous optimisation methods often require a large number of function evaluations, resulting in substantial computational cost. Their performance can also be sensitive to hyperparameter settings and initialisation, potentially leading to unstable convergence.

Discrete algorithms include graph search techniques and integer programming methods, with discrete sets typically representing predefined ship speed increments, specific route waypoints, or operational modes. They are widely used in actual ship voyage planning systems due to

their computational efficiency. For example, the three-dimensional Dijkstra method has been used to jointly optimise ship speed and route for fuel efficiency (Wang et al., 2019). Other refinements to Dijkstra's classical algorithm have also been considered, wherein network weights are iteratively updated in response to metocean conditions to enhance energy efficiency (Bahrami and Siadatmousavi, 2024). The algorithm has been further extended using the Floyd-Warshall algorithm to optimise paths across varying sea states for minimum operational cost (Bahrami et al., 2025). Similarly, while the isochrone method is a staple of discrete search, recent Isochrone-A* strategies have improved convergence by utilising augmented cost functions and midpoint strategy switching (Chen et al., 2024). Dynamic programming (DP), based on Bellman's principle of optimality, has likewise been widely accepted. DP methods systematically explore discrete states and transitions to identify the minimum-cost optimisation variables. DP has been designed to address three-dimensional voyage optimisation with environmental constraints in Zaccone et al. (2018), and to tanker speed allocation under variable weather conditions (Du et al., 2019). In addition, DP-based frameworks have been enhanced by coupling them with artificial neural networks (ANNs) to optimise speed and trim jointly (Fan et al., 2022) or by using dynamic segmentation to cope with diminishing forecast accuracy (Tzortzis and Sakalis, 2021). Despite these advances, discrete




Fig. 1. Limitations of discrete optimisation algorithms.

optimisation methods inherently depend on predefined resolution levels. Consequently, the true optimum may remain unattainable when it lies between discrete decision points as illustrated in Fig. 1, leading to resolution-induced suboptimality. Increasing discretisation density can partially alleviate this issue, but often results in a rapid growth of the search space and computational burden.

In addition to optimisation algorithms, constraint-handling techniques are essential, as real-world optimisation problems must be solved within operational limits. Constraint-handling approaches can guide the optimisation process towards the feasible region and help to ensure that feasible solutions are found (Fiacco and McCormick, 1968). Common strategies include penalty functions, which prevent the optimiser from entering unfeasible regions by increasing the cost near constraint boundaries (Gelbart et al., 2014), and augmented Lagrangian methods, which incorporate penalties through the use of Lagrange multipliers (Nocedal and Wright, 2006). While these methods often rely on manually tuned penalty terms, solution may struggle with highly non-linear or unknown constraint landscapes, and can result in poor exploration near constraint boundaries (Wang et al., 2021). Increasing the resolution of the discretisation can help, but at increased computational cost.

For short sea shipping, ships' navigational conditions exhibit limited variation, so fuel consumption is influenced less by optimal route planning than by operational parameters such as speed or engine RPM (Jimenez et al., 2022; Vergara, 2025). As ship speed is controlled (indirectly) by the propulsion system, maintaining a constant or stepwise-constant speed can entail frequent engine power fluctuations under varying metocean conditions (Li et al., 2022). Additionally, the economic implications of such operational adjustments, such as those required by Just-In-Time (JIT) arrival systems, must be modelled to ensure the continued competitiveness of short-sea shipping routes (Tsioumas et al., 2025). In the current study, a discrete-to-continuous optimisation framework is proposed to refine ship propulsion power allocation. Within this framework, a modified parallel coupling dynamic programming (PCDP) method is used to generate feasible discrete solutions that serve as prior knowledge (Vergara et al., 2025). The proposed constraint-aware BO framework systematically refines these solutions toward improved feasible outcomes. Specifically, the proposed BO framework integrates an embedded penalty mechanism, that allows convex constraints to be incorporated directly into the optimisation process. In addition, sequential domain reduction is introduced to further enhance convergence by focusing the search on promising regions of the solution space. The optimisation framework is formulated as a pre-voyage planning tool, and the impact of weather forecast uncertainty on optimisation outcomes has also been investigated.

The proposed method is applied to a constrained voyage power allocation problem, where improved feasibility and optimality are demonstrated in comparison to four actual voyage cases. The remainder of this paper is organised as follows. Section 2 introduces the optimisation framework proposed in this study. Section 3 presents the case study ship and voyages, as well as the fuel consumption and speed models used in the optimisation process. The optimisation results and the correspond-




Fig. 2. Main components of a route, including the departure, destination, waypoints and legs.

ing fuel consumption and emissions savings are presented in Section 4. Finally, the conclusions are presented in Section 5.

2. Proposed methodology for propulsion power optimisation

2.1. Overall framework and problem formulation

Consider a discretised ship voyage represented as a sequence of legs, denoted by $L = (L_1, L_2, \dots, L_N)$. Each leg is associated with relatively stable metocean conditions (without large changes). This environmental consistency implies approximately constant propulsion requirements along each leg, allowing propulsion power to be treated as constant for stable engine operation. For the short-sea shipping in this study, the ship service route/path for each voyage is often fixed. Then, the voyage optimisation problem becomes to find the sequence of optimal propulsion powers $P = (P_1, P_2, \dots, P_N)$ to minimise the fuel consumption along the entire voyage $M_{\text{fuel}}(L : P)$ as

$$\hat{P} = (\hat{P}_1, \hat{P}_2, \dots, \hat{P}_N) = \arg \min_P M_{\text{fuel}}(L : P) = \arg \min_P \sum_{i=1}^N M_{\text{fuel}}(L_i : P_i), \quad (1)$$

where \hat{P}_i denotes the optimised propulsion power for the i th leg L_i . Each leg is characterised by a fixed propulsion power P_i and a sequence of waypoints $\xi_{i,j} = (x_{i,j}, y_{i,j})$, where x and y represent the longitude and latitude of each waypoint, as illustrated in Fig. 2.

In this optimisation problem, $M_{\text{fuel}}(L_i : P_i)$ represents the fuel consumption along the i th leg of the voyage, which can be estimated by

$$M_{\text{fuel}}(L_i : P_i) = \sum_{j=1}^{k_i} f_{\text{fuel}}(P_i | \mathbf{W}_{i,j}, V_{i,j}) \cdot \Delta t_{i,j}, \quad (2)$$

where f_{fuel} denotes the cost function used to estimate the fuel consumption rate, which depends on the metocean conditions $\mathbf{W}_{i,j}$ and the sailing speed $V_{i,j}$ at the j th waypoint of the i th leg $\xi_{i,j}$. Here, k_i represents the number of waypoints within the i th leg, and $\Delta t_{i,j}$ denotes the sailing time between consecutive waypoints $\xi_{i,j}$ and $\xi_{i,j+1}$. The location of each waypoint along the fixed route is estimated from the sailing speeds $V_{i,j}$ and arrival times $t_{i,j}$. The temporal spacing between adjacent waypoints, defined as $\Delta t = t_{i,j+1} - t_{i,j}$, is set to match the temporal resolution of the weather inputs (e.g., 1 h). Given $\xi_{i,j}$ and $t_{i,j}$, the corresponding metocean conditions $\mathbf{W}_{i,j}$ at waypoint (i, j) can be obtained. The metocean conditions comprise wave, wind, and current information, namely,

$$\mathbf{W}_{i,j} = \{ H_s^{i,j}, T_m^{i,j}, \alpha_{\text{wave}}^{i,j}, V_{\text{wind}}^{i,j}, \alpha_{\text{wind}}^{i,j}, V_{\text{current}}^{i,j}, \alpha_{\text{current}}^{i,j} \}, \quad (3)$$

where H_s , T_m and α_{wave} represent significant wave height, mean wave period, and wave direction. V_{wind} , V_{current} , α_{wind} , and α_{current} represent the speeds and directions of wind and current, respectively. All directional variables are defined as relative angles with respect to the ship heading: 0° corresponds to head sea, head wind, or head current (i.e., approaching from the bow), and the angle increases clockwise, with 180° indicating following sea, following wind, or following current. Those values are determined using trilinear interpolation of ocean weather data based on the location and time of the waypoints.

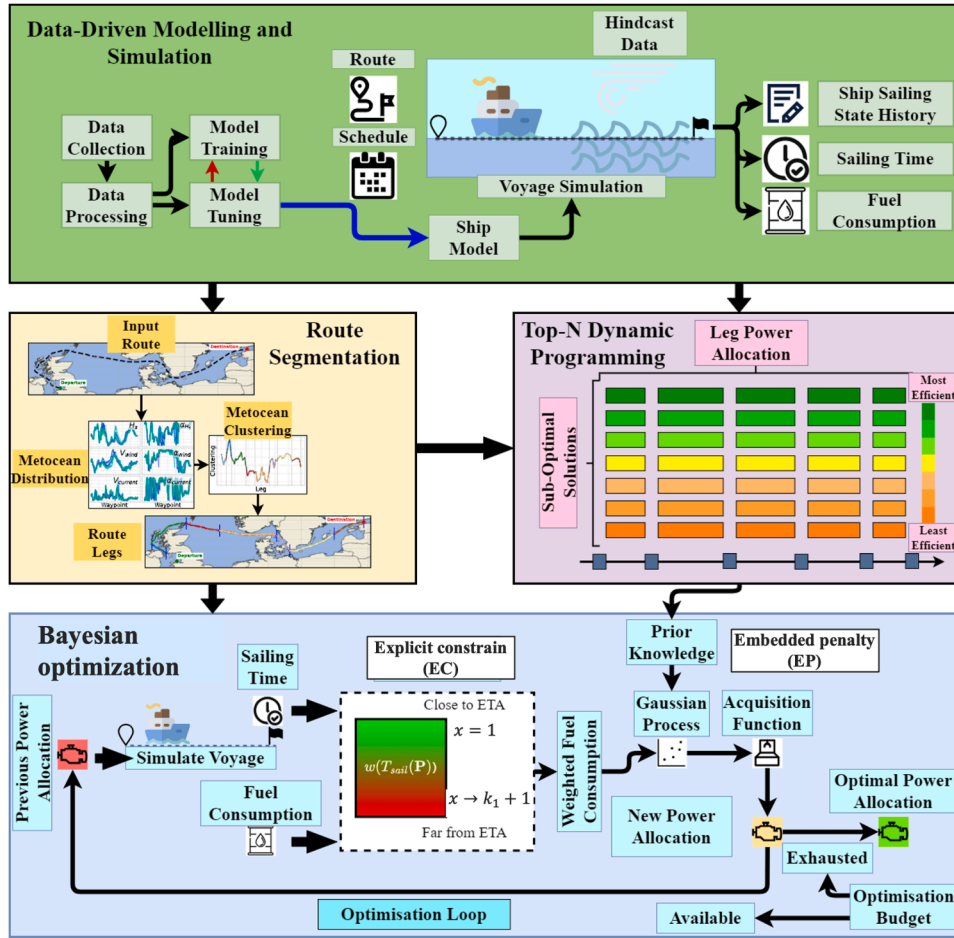


Fig. 3. General voyage optimisation framework with proposed strategies to refine optimisation results.

The candidate voyages, with total sailing time T_{sail} , selected during the voyage optimisation process, are constrained by

$$C_{ETA}(P) := \left| 1 - \frac{T_{sail}(P)}{T_{ETA}} \right| \cdot 100\% \leq D_{tol}, \quad (4)$$

where T_{ETA} is the expected sailing time for the specific voyage planning, and D_{tol} is the tolerable time difference between the planned voyage's sailing time and the required ETA. The tolerance can be taken so that the difference between the average sailing speed satisfies the actual contract speed V_{cp} in a charter party (Kasi, 2022) as

$$D_{tol} \leq \frac{0.5 \text{ kn}}{V_{CP}} \cdot 100\%. \quad (5)$$

In this study, the tolerance is fixed conservatively at 1% regardless of the contractual speed V_{cp} . In addition to the ETA requirement, several operational constraints are considered in the optimisation framework to ensure practical feasibility. Specifically, engine power is bounded by minimum and maximum limits, and vessel speed is restricted to operational ranges that reflect propulsion system capabilities. These constraints prevent infeasible operating points and enhance the realism of the optimisation outcomes.

The proposed optimisation framework for the short-sea shipping problem is illustrated in Fig. 3. The ship performance model used to evaluate the cost (objective) function in Eq. (2), together with the corresponding meteocean conditions defined in Eq. (3) (upper box), is assumed to be known.

The route is segmented into legs L_i , $i = 1, 2, \dots, N$, and the preliminary leg power allocation (middle boxes) is obtained using the methodology developed by Vergara et al. (2025), specifically the MS-PELT al-

gorithm. In that study, MS-PELT was benchmarked against state-of-the-art multivariate time-series segmentation approaches and was shown to produce fewer, longer, and more operationally meaningful segments characterised by homogeneous environmental conditions. Additionally, the algorithm demonstrated substantially lower computational cost, supporting real-time or near-real-time applications. Owing to these advantages, MS-PELT is adopted in the present study for voyage segmentation.

To mitigate the discretisation limitations inherent in traditional voyage optimisation, this study develops a structured discrete-to-continuous optimisation framework. The proposed methodology is organised into three steps:

- (1) generation of feasible discrete solutions using modified dynamic programming,
- (2) continuous refinement through Bayesian optimisation,
- (3) constraint-handling strategies to ensure operational feasibility.

This framework bridges discrete feasibility with continuous optimality while maintaining computational efficiency. The details of each step are presented in the following subsections.

2.2. Discrete solution generation by PCDP

The first step of the proposed discrete-to-continuous optimisation framework is to generate a set of feasible discrete solutions to initialise the Bayesian optimisation process. This step is particularly important because BO can become prohibitively expensive when a large number of iterations is required, primarily due to the covariance matrix inversion inherent in it. Random initialisation of BO is typically inefficient.

Instead, learning from already reasonable solutions can substantially improve convergence efficiency (Wang et al., 2022). Therefore, BO is initialised with so-called prior knowledge in this work, and a modified PCDP approach is employed to generate feasible discrete solutions that serve as this prior. The PCDP determines the sub-optimal propulsion power settings for each voyage leg while minimising the total fuel consumption M_{fuel} .

Unlike traditional sequential coupling, which requires exact timing dependencies between legs, the PCDP technique introduces flexible departure time intervals for each leg to allow for computation of parallel scenarios (Vergara, 2025). This approach provides greater flexibility in power allocation and improves computational efficiency by decoupling the interconnections among voyage segments across different optimisation variables. The optimisation problem is then solved using dynamic programming, where the Bellman equation is taken to be

$$J_i^* = \min_{P_i} \left\{ J_{i-1} + M_{\text{fuel}}(P_i, \mathbf{W}_{i,1 \rightarrow k_i}) \right\}, \quad (6)$$

where $J^*(i)$ is the optimal accumulated fuel consumption up to the i th leg, and $M_{\text{fuel}}(P_i)$ is the fuel consumption for the i th leg. For such DP-based optimisation, a specific number of the sub-optimal solutions that satisfy the ETA constraint from Eq. (4) are used for the next step and search for the optimum. These sub-optimal solutions are introduced into the distillation refinement and passed to the BO as the prior, thus initialising the algorithm more efficiently.

2.3. Continuous refinement using Bayesian optimisation

The feasible discrete solutions generated by the modified PCDP are subsequently used as prior knowledge to initialise the Bayesian optimisation, thereby enabling a more informed and sample-efficient search. In this study, the BO algorithm is implemented using a modified version of the Bayesian optimisation library for Python (Nogueira, 2014). Rather than directly optimising the original objective function, BO constructs a probabilistic surrogate model that approximates the objective landscape. This surrogate allows the algorithm to systematically identify promising regions of the input space while maintaining an effective balance between exploration of uncertain regions and exploitation of high-performing candidates.

The surrogate function is chosen as a Gaussian process (GP) – a smooth, non-parametric probabilistic model that defines a distribution over functions. A GP is fully specified by its mean function $m(\mathbf{P})$ and covariance function (kernel) $\mathbf{K} = \kappa(\mathbf{P}, \mathbf{P}')$. Therefore, BO uses a set of observations of the search space $\mathbf{X} = \{\mathbf{P}_1, \mathbf{P}_2, \dots, \mathbf{P}_N\}$ and the cost evaluations $\mathbf{Y} = \{M_{\text{fuel}}(\mathbf{L}_1 : \mathbf{P}_1), M_{\text{fuel}}(\mathbf{L}_1 : \mathbf{P}_2), \dots, M_{\text{fuel}}(\mathbf{L}_N : \mathbf{P}_N)\}$, to model the objective function $M_{\text{fuel}}(\mathbf{P})$ as a sample from a GP as

$$M_{\text{fuel}}(\mathbf{P}) \sim \text{GP}(m(\mathbf{P}), \kappa(\mathbf{P}, \mathbf{P}')). \quad (7)$$

The GP is initialised with $M_{\text{fuel}}(\mathbf{P}) = m(\mathbf{P}) + \epsilon$ where $\epsilon \sim \mathcal{N}(0, \sigma^2)$ and a kernel that encodes all the prior knowledge or assumptions about the function and correlations between different inputs. On each iteration of the optimisation loop, the GP is updated with every new observation \mathbf{P}_* of the objective function $M_{\text{fuel}}(\mathbf{P}_*)$. The GP is updated to produce a posterior distribution over the function. For each new optimisation variable \mathbf{P}_* , its mean and variance under conditions of previous optimisation samples \mathbf{P} and evaluations $M_{\text{fuel}}(\mathbf{P})$ can be estimated by

$$\begin{aligned} m(\mathbf{P}_* | \mathbf{P}) &= m(\mathbf{P}_*) + \kappa(\mathbf{P}_*, \mathbf{X}) \cdot (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \cdot (\mathbf{Y} - m(\mathbf{X})), \\ \kappa(\mathbf{P}_* | \mathbf{P}) &= \kappa(\mathbf{P}_*, \mathbf{P}) - \kappa(\mathbf{P}_*, \mathbf{X}) \cdot (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \cdot \kappa(\mathbf{X}, \mathbf{P}_*). \end{aligned} \quad (8)$$

For the optimisation process, the Matérn kernel is chosen as it provides a flexible way to control the smoothness of the GP. It is parametrised by a smoothness parameter ν , a length scale l , and an amplitude parameter. The general form of the Matérn kernel between two optimisation variables \mathbf{P} and \mathbf{P}' is given by

$$\kappa(\mathbf{P}, \mathbf{P}') = \frac{2^{1-\nu}}{\Gamma(\nu)} \cdot (2 \cdot \sqrt{\nu} \|\mathbf{P} - \mathbf{P}'\|/l)^\nu \cdot K_\nu \cdot (2 \cdot \sqrt{\nu} \|\mathbf{P} - \mathbf{P}'\|/l), \quad (9)$$

where $\Gamma(\cdot)$ is the gamma function and K_ν is the modified Bessel function of the second kind. Once the GP posterior is defined, BO uses an acquisition function to guide the search. The values of $\nu = \frac{5}{2}$ and $l = 1$ are used, reducing the kernel to

$$\kappa_{\nu=\frac{5}{2}}(\mathbf{P}, \mathbf{P}') = \left(1 + 5\|\mathbf{P} - \mathbf{P}'\| + \frac{5\|\mathbf{P} - \mathbf{P}'\|^2}{3} \right) \cdot \exp(-5\|\mathbf{P} - \mathbf{P}'\|), \quad (10)$$

which is twice differentiable. Once the posterior of GP is established, the BO algorithm determines the next sampling point using an acquisition function $\alpha(\mathbf{P})$ to guide the search for the next optimisation variable candidate. The decaying lower confidence bound (LCB) acquisition function is used as

$$\alpha(\mathbf{P}_*) = m(\mathbf{P}_* | \mathbf{P}) - k \cdot \kappa(\mathbf{P}_*), \quad k = k_0 \cdot r^{\text{ites}}, \quad (11)$$

where k is the exploration parameter that sets the trade-off between exploration and exploitation. It is initialised at $k_0 = 5$ and with decay coefficient $r = 0.95$; ites is the number of iterations in the optimisation process. The next optimisation variable candidate \mathbf{P}_* is selected as the point that minimises the acquisition function $\alpha(\mathbf{P}_*)$ by

$$\hat{\mathbf{P}}_* = \arg \min_{\mathbf{P}_*} \alpha(\mathbf{P}_*). \quad (12)$$

After selecting the next candidate $\hat{\mathbf{P}}_*$, the resulting sample is added to the existing optimisation dataset \mathbf{P} . The GP is then updated with new information, refining the model and improving the posterior distribution over the objective function. This process is repeated until the predefined number of iterations is reached. After every iteration, the optimum is determined by tracking the best-performing $\hat{\mathbf{P}}_*$.

2.4. Constraint-handling strategies

2.4.1. Explicit constraints modelling

In constrained optimisation, constraint functions are often expensive to evaluate, non-differentiable, or only observable through black-box simulations. Surrogate modelling therefore provides a principled approach for incorporating feasibility within the optimisation process. Within the BO framework, a probabilistic model is fitted for each constraint in parallel with the objective function (Gelbart et al., 2014). This enables estimation of the probability that a candidate voyage plan \mathbf{P} satisfies the bounded constraints. The acquisition function is accordingly modified as

$$\alpha_C(\mathbf{P}_*) = \alpha(\mathbf{P}_*) \cdot \mathbb{P}_\setminus (T_{\text{sail}}(\mathbf{P}_*) \in C_{\text{ETA}}(\mathbf{P})), \quad (13)$$

where $C_{\text{ETA}}(\mathbf{P})$ denotes the set of power allocations satisfying the ETA constraint defined in Eq. (4). This approach, referred to as explicit constraint handling (BO+EC), is implemented by fitting GP models to represent each independent constraint as

$$\begin{aligned} T_{\text{sail}}(\mathbf{P}) &\sim \text{GP}(m_T(\mathbf{P}), \kappa_T(\mathbf{P}, \mathbf{P}')), \\ \mu_T(\mathbf{P}_*) &= m_T(\mathbf{P}_*) + \kappa_T(\mathbf{P}_*, \mathbf{X}) \cdot (\mathbf{K}_T + \sigma^2 \mathbf{I})^{-1} \cdot (\mathbf{Z} - m_T(\mathbf{X})), \\ \sigma_T^2(\mathbf{P}_*) &= \kappa_T(\mathbf{P}_*, \mathbf{P}) - \kappa_T(\mathbf{P}_*, \mathbf{X}) \cdot (\mathbf{K}_T + \sigma^2 \mathbf{I})^{-1} \cdot \kappa_T(\mathbf{X}, \mathbf{P}_*), \end{aligned} \quad (14)$$

where m_T and κ_T are the mean and co-variance of sailing time for the choice of power allocation space \mathbf{P} , $\mu_T(\mathbf{P}_*) := m_T(\mathbf{P}_* | \mathbf{P})$ and $\sigma_T^2(\mathbf{P}_*) := \kappa_T(\mathbf{P}_* | \mathbf{P})$ represent the mean and variance of sailing time for the given power allocation candidate \mathbf{P}_* under the previous power samples \mathbf{P} , while $\mathbf{K}_T = \kappa_T(\mathbf{P}, \mathbf{P}')$, and the corresponding sailing time for the set of observations in the search space \mathbf{X} is $\mathbf{Z} = \{T_{\text{sail}}(\mathbf{L}_1 : \mathbf{P}_1), T_{\text{sail}}(\mathbf{L}_1 : \mathbf{P}_2), \dots, T_{\text{sail}}(\mathbf{L}_N : \mathbf{P}_N)\}$.

Once the GP is trained, the probability of satisfying the constraint at a candidate point \mathbf{P}_* is estimated as

$$\mathbb{P}_\setminus (T_{\text{sail}}(\mathbf{P}_*) \in C_{\text{ETA}}) = \Phi\left(\frac{C_{\text{ETA}}^{\text{up}} - \mu_T(\mathbf{P}_*)}{\sigma_T(\mathbf{P}_*)}\right) - \Phi\left(\frac{C_{\text{ETA}}^{\text{low}} - \mu_T(\mathbf{P}_*)}{\sigma_T(\mathbf{P}_*)}\right), \quad (15)$$

where $\Phi(\cdot)$ is the Gaussian probability distribution. This formulation guides the acquisition function toward regions that are both promising

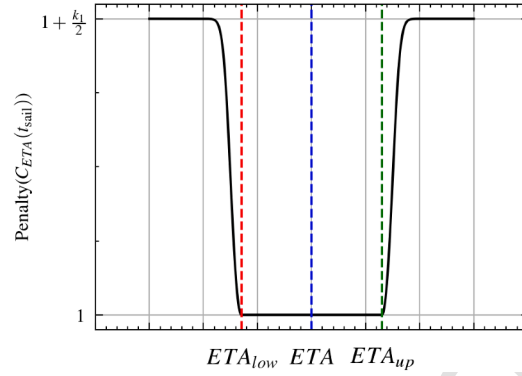
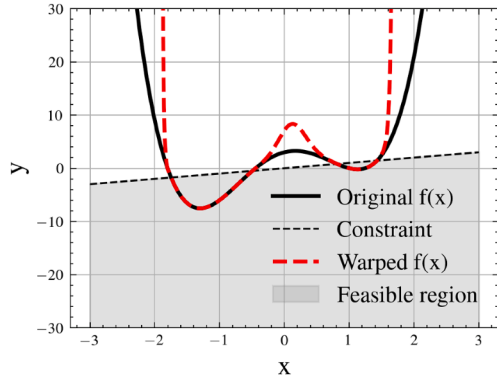


Fig. 4. Illustrative example of the impact of the multiplicative penalty on the objective function (left), and penalty function to embed the optimisation constraint (right).

for minimising M_{fuel} and likely to satisfy the ETA constraint, thereby improving feasibility during optimisation.

2.4.2. Embedded penalty mechanism

Unlike explicit constraint modelling, the embedded penalty approach incorporates feasibility directly into the objective function. By learning from the penalised objective, the GP surrogate naturally prioritises regions that satisfy the constraints while discouraging infeasible solutions. This strategy is referred to as Bayesian optimisation with an embedded penalty (BO + EP). In this formulation, the original objective function is modified as

$$f(\mathbf{P}) = M_{fuel}(\mathbf{P}) \cdot w_{fuel}(T_{sail}(\mathbf{P})), \quad (16)$$

where the penalty function w_{fuel} embeds the sailing-time constraint C_{ETA} into the objective, transforming the constrained problem into a smoothly penalised optimisation as

$$w_{fuel}(T_{sail}(\mathbf{P})) = \begin{cases} 1 & T_{sail} \in C_{ETA}, \\ 1 + k_1 \cdot \left(\frac{1}{1 + \exp(-k_2 \cdot C_{ETA}(\mathbf{P})^2)} - 0.5 \right) & T_{sail} > C_{ETA}^{up}, \end{cases} \quad (17)$$

where $0 < k_1, k_2 < \infty$. The parameters k_1 and k_2 control the magnitude and growth rate of the penalty near the feasibility boundary (Fig. 4). This smooth formulation preserves continuity in the search space while discouraging constraint violations, enabling the optimiser to efficiently identify feasible high-quality solutions.

3. Case study ship, cost models and test voyages

To demonstrate the applicability and effectiveness of the proposed optimisation method for reducing fuel consumption in short-sea shipping, a chemical tanker operating between ports in European waters is selected as the case study. The main particulars of the vessel are provided in Table 1, together with the measurement period and sampling frequency of the full-scale performance monitoring data used in this study.

Table 1
Chemical tanker characteristics.

Parameter	Value
Length between perpendiculars	138.22 m
Design draft	9.27 m
Maximum continuous rating (MCR)	7200 kW
Service speed	14 knots
Data collection period	No. 2020–Mar 2024
Data frequency	1 sample per minute

3.1. Performance model of the case study ship

A key component of the proposed voyage optimisation framework is a cost model for estimating the fuel consumption rate f_{fuel} under varying operational conditions (P, V) and metocean environments \mathbf{W} , as required by Eq. (2). In this work, the cost model is established using three years of full-scale performance monitoring data from the case-study tanker, whose sailing routes are shown in Fig. 5.

The full-scale performance monitoring data from the chemical tanker were preprocessed by downsampling to 10-minute intervals, reducing noise while preserving the underlying signal characteristics and ensuring high data quality for subsequent model development. Each waypoint $\xi_{i,j}$ is assumed to represent a ship's energy performance within those 10-minute intervals (same as the sampling interval from full-scale measurements), and the corresponding fuel consumption rate is modelled by

$$f_{fuel}(P|\mathbf{W}, V) = f_{fuel}(P, RPM, V, H_s, T_m, \alpha_{wave}, V_{wind}, \alpha_{wind}, V_{current}, \alpha_{current}, \dots), \quad (18)$$

where the sailing speed $V_{i,j}$ and metocean conditions $\mathbf{W}_{i,j}$ are assumed constant during the 10-minute time interval. To reconstruct the actual environmental conditions encountered during the voyages and enable fair validation against measured fuel consumption while minimising environmental uncertainty, the framework is evaluated using hindcast data. The metocean conditions for each waypoint $\mathbf{W}_{i,j}$ is determined from hindcast using the ECMWF Reanalysis v5 (ERA5) service (Hersbach et al., 2020) and the Global Ocean Physics Analysis and Forecast (International, 2024). For the voyage optimisation process, in addition to the cost model, i.e., the above Eq. (2) to estimate fuel consumption as the cost function, another model, which estimates the ship speed along the discretised legs for any candidate propulsive power P , is also needed to estimate the corresponding sailing times. Therefore, similar procedures

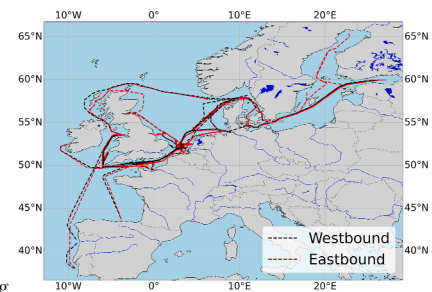


Fig. 5. All sailing routes for three years' sailing of the chemical tanker.

Table 2
Hyperparameters and search bounds for the ship models.

Hyperparameter	Bounds	Hyperparameter	Bounds
Number of estimators	500–5000	Min. child Weight	1–5
Learning Rate	0.01–0.3	γ	0–5
Max depth	5–15	Regularisation α	0–0.5
Subsample	0.8–1	Regularisation λ	0.1–5
Column sample by tree	0.7–1		

Table 3
Optimised hyperparameters for the speed and fuel models.

Parameter	Speed model f_V	Fuel model f_{fuel}
Learning rate	0.0235	0.1687
Max depth	11	6
Number of estimators	1077	480
γ	0.9489	0.8016
Regularisation α	0.4863	0.4545
Regularisation λ	0.8084	0.7834
Min. child weight	15.6505	25.0
Subsample	0.8823	0.9794
Column sample by tree	0.5853	0.8703

Table 4
Performance metrics for the fuel and speed models.

Model		MAE	MAPE	RMSE	R^2
Fuel	Training	0.2608 kg/h	1.77%	0.4394 kg/h	0.9836
	Validation	0.7209 kg/h	5.15%	1.0320 kg/h	0.8486
Speed	Training	0.2323 knots	2.12%	0.3438 knots	0.9362
	Validation	0.3808 knots	3.87%	0.6034 knots	0.7404

will be performed to establish the so-called ship speed model as

$$V = f_V(P, \text{RPM}, H_s, T_m, \alpha_{\text{wave}}, V_{\text{wind}}, \alpha_{\text{wind}}, V_{\text{current}}, \alpha_{\text{current}}, \dots). \quad (19)$$

The aforementioned two supervised machine learning models are developed using the XGBoost algorithm (Chen and Guestrin, 2016; Lang et al., 2022a; Vergara et al., 2023) to predict fuel consumption rate and ship speed from the input features defined in Eqs. (18) and (19), respectively. The input features listed in Eqs. (18) and (19) were selected based on well-established ship propulsion principles and the impact of metocean conditions on ship performance. Model accuracy is sensitive to the choice of XGBoost hyperparameters. The search bounds used for hyperparameter tuning are listed in Table 2. Bayesian optimisation is employed to tune hyperparameters that minimise the mean squared error (MSE) between predictions and measurements (Lang et al., 2022b) as

$$\text{MSE} = \sum_{i,j} (\hat{f}_{i,j} - f_{i,j})^2, \quad (20)$$

where \hat{f} and f represent the predicted and measured performance, respectively. For the performance modelling by XGBoost, the optimal hyperparameters are listed in Table 3.

Table 4 summarises the predictive performance of the final models. For fuel consumption prediction, the model achieves a MAPE of 1.77% on the training set and 5.15% on the validation set. In comparison, the speed prediction model attains a training MAPE of 2.12% and a validation MAPE of 3.87%. The fuel consumption model yields R^2 values of 0.9836 on the training set and 0.8486 on the validation set, confirming its strong predictive capability. For the speed model, the R^2 decreases from 0.9362 to 0.7404. This reduction is expected, given the higher variability and complexity of ship speed dynamics under varying metocean conditions, particularly in extrapolative operation conditions. The observed decline indicates a moderate generalisation gap when the model is applied to unseen operational conditions. Overall, the results demonstrate that both models provide sufficiently accurate and robust predictions to support the subsequent voyage optimisation framework.




Fig. 6. An illustration of how the waypoints along the route are obtained.

Table 5
Summary of sailing cases with distance, ETA, and Fuel consumption.

Case ID	Sailing Area	Distance [km]	ETA [hours]	Fuel cost [tons]
R1	Baltic and North Sea	3286.90	141.41	121.62
R2	Baltic and North Sea	2965.21	140.94	86.23
R3	English Channel	1958.70	104.00	70.94
R4	English Channel	1632.63	76.57	56.13

3.2. Estimation of optimisation cost function

To evaluate fuel consumption (or equivalently CO₂ emissions) along candidate voyages under different power allocations, the fixed service route is discretised into waypoints as illustrated in Fig. 2 and Fig. 6.

At each waypoint $\xi_{i,j}$, the fuel consumption rate and speed are estimated using Eqs. (18) and (19), respectively. The metocean conditions $\mathbf{W}_{i,j}$ are obtained by interpolating hindcast data (ERA5), which provides reliable environmental inputs for the relatively short planning horizons considered in short-sea shipping. Voyage-level performance is then evaluated following Algorithm 1 (Fig. 6).

Algorithm 1 Ship performance estimation over a voyage.

```

Start at Departure
Set Leg  $i \leftarrow 1$ 
while not(Destination) do
  Set leg power  $P_i$ 
  while not(End Leg) do
    Interpolate  $\mathbf{W}_{i,j}$  from ERA5 hindcast
    Find  $V_{i,j}$  using Eq. (19)
    Find  $m_{\text{fuel}}(P_i, \mathbf{W}_{i,j})$  using Eq. (18)
    Hold  $V_{i,j}$  constant for 1 h
    Determine next waypoint along the route as sketched in Fig. 6
  end while
  Set Leg  $i \leftarrow i + 1$ 
end while

```

3.3. Case study voyages

Four case study voyages from the tanker are selected for evaluation (Fig. 7). Key information for these voyages, including distance, ETA, and fuel consumption, is summarised in Table 5. These routes are used as reference cases in the proposed optimisation framework, as they are representative of the vessel's operational patterns shown in Fig. 5.

In this section, the proposed BO + EC and BO + EP methods are applied to the four case study voyages and evaluated against the measured full-scale monitoring data. Since the Bayesian optimisation is initialised using feasible discrete solutions generated by PCDP, the corresponding DP results are also included as a benchmark to quantify the benefit of continuous refinement. First, BO + EC and BO + EP are benchmarked in terms of computational efficiency and constraint satisfaction.




Fig. 7. The four case study voyages across the Baltic Sea, North Sea and English Channel.

Table 6

Computational time comparison between BO + EP and BO + EC methods.

Voyage	BO + EP [s]	BO + EC [s]	Speedup
R1	~ 90	~ 1100	~ 12×
R2	~ 80	~ 900	~ 11×
R3	~ 40	~ 800	~ 20×
R4	~ 50	~ 900	~ 18×

Second, based on this comparison, the selected method is compared with the baseline DP solutions and the measured voyages for routes in the Baltic and North Seas (R1–R2) and the English Channel (R3–R4). Finally, recognising that practical pre-voyage applications rely on weather forecasts rather than hindcast data, the impact of metocean forecast uncertainty on the optimisation outcomes is also investigated.

3.4. Computational performance and constraint-handling comparison

First, the computational performance of the two constraint handling strategies is evaluated for the four case study voyages (R1–R4), for which no ice conditions were observed. Figs. 8–11 show the evolution of the ETA constraint metric, objective value, and cumulative runtime over the Bayesian optimisation iterations. For both BO + EP (blue) and BO + EC (red), the initial samples are predetermined based on the prior knowledge described in Section 2.2, generated using the PCDP approach. Each method is subsequently executed for 100 optimisation iterations. Across all case studies, BO + EP achieves objective values and power allocation profiles comparable to those obtained with BO + EC, while requiring approximately one-tenth of the computational time.

Table 6 summarises the computational time across all case studies. All computations were performed on a standard laptop equipped with an Intel Core i5 CPU. These runtimes correspond solely to the Bayesian optimisation procedure and exclude the subsequent PCDP-based discrete optimisation step. For voyage R1, BO + EP completes the Bayesian optimisation in approximately 90 s, whereas BO + EC requires about 1100 s. For voyage R2, BO + EP requires around 80 s compared to 900 s for BO + EC (Figs. 8 and 9). Similar efficiency gains are observed for R3 and R4, where BO + EP completes the optimisation in approximately 40 s and 50 s, respectively.

Regarding the effectiveness of the optimised solutions, the optimal fuel consumption (objective values) for R1 are 112.74 t (BO + EP) and 112.72 t (BO + EC), while for R2 the corresponding values are 79.26 t and 79.16 t, respectively. For R3 and R4, the fuel consumption obtained using the two methods is similarly consistent, with only negligible differences observed. Moreover, as illustrated in the right plots of Figs. 8–11, the power allocation profiles generated by BO + EP are similar to those produced by BO + EC. Given its substantially lower computational cost without any noticeable degradation in optimisation performance, BO + EP emerges as the preferred method. In the subsequent

Table 7

Performance comparison for Routes R1 and R2.

	Source	R1	R2	R1	R2
ETA	Actual	141.41	140.94	Δ_{ETA} w.r.t. Actual	
	DP	142.54	140.80	0.80%	-0.10%
	BO + EP	143.24	141.12	1.29%	0.13%
	BO + EC	142.80	141.92	0.98%	0.70%
M_{fuel}	Actual	121.62	86.23	$\Delta_{M_{fuel}}$ w.r.t. Actual	
	DP	113.99	80.26	-6.27%	-6.92%
	BO + EP	112.74	79.26	-7.30%	-8.08%
	BO + EC	112.72	79.16	-7.32%	-8.20%

analysis, BO + EP is therefore adopted for the detailed voyage optimisation study across all four case voyages, and its performance is compared with that of the DP method.

3.5. Comparative evaluation against DP in the Baltic and North Seas

For short-sea passages crossing the Baltic and North Seas, two case study voyages, R1 and R2 in Fig. 5, are optimised using the BO + EP method. Both voyages are long-distance routes of approximately 3000 km, each requiring more than 140 hours of sailing time.

As illustrated in Fig. 12, voyage R1 is divided into nine legs. The optimal solution obtained using the conventional DP method is shown in red, while the refined power allocation derived from BO + EP is shown in blue. The proposed BO + EP method achieves a potential fuel consumption reduction of 8.88 t, corresponding to a 7.30% decrease relative to the measured fuel consumption, as summarised in Table 7. The optimised power allocation strategy reveals clear adaptation to evolving metocean conditions. Higher engine power is applied during the early legs of the voyage under relatively calm wave conditions. A notable reduction in power occurs in Leg 7, where head waves exceeding 2 m are encountered, increasing resistance and making speed reduction more fuel-efficient. In Leg 8, the wave direction gradually transitions from head seas to beam seas, which impose lower resistance compared with bow or head waves; consequently, engine power is increased. In the final leg (Leg 9), decreasing wave height allows the power setting to return to levels comparable to those of the early voyage stages.

BO + EP further refines the initial DP solution by slightly reducing power in Legs 1, 7, and 8, while marginally increasing power in Legs 2 and 5. The optimised voyage duration is approximately 143.24 h, representing a modest increase of 0.7 h compared with the DP solution (142.54 h), but yielding an additional fuel saving of 1.03%.

The optimal solution for voyage R2 is presented in Fig. 13, where the route is partitioned into seven legs. The optimal power allocation obtained using BO + EP is largely consistent with the DP solution for Legs 1–3. In Leg 4, the proposed BO + EP method further reduces the power setting compared with the DP solution due to the presence of bow waves of approximately 2 m height, which increase added resistance in waves. In Leg 5, engine power is slightly increased since the

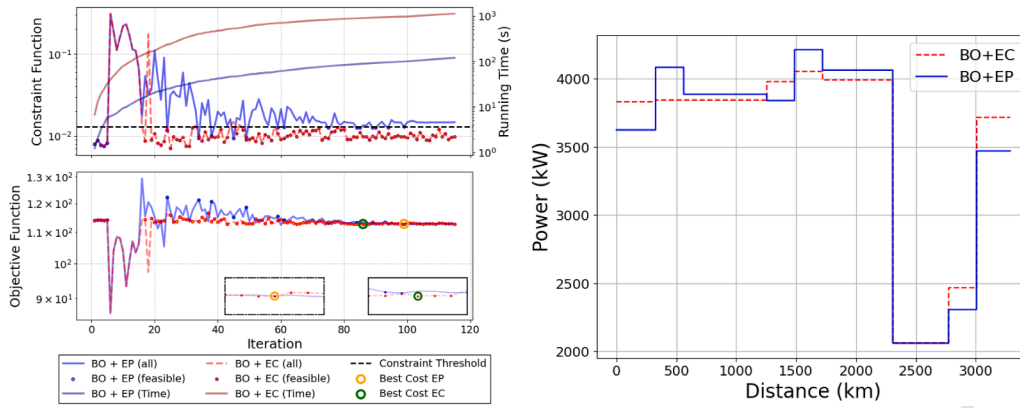


Fig. 8. Comparison of the efficiency and effectiveness of the BO + EC and BO + EP methods for voyage R1. Left: constraint and runtime evolution. Right: resulting optimal power profiles.

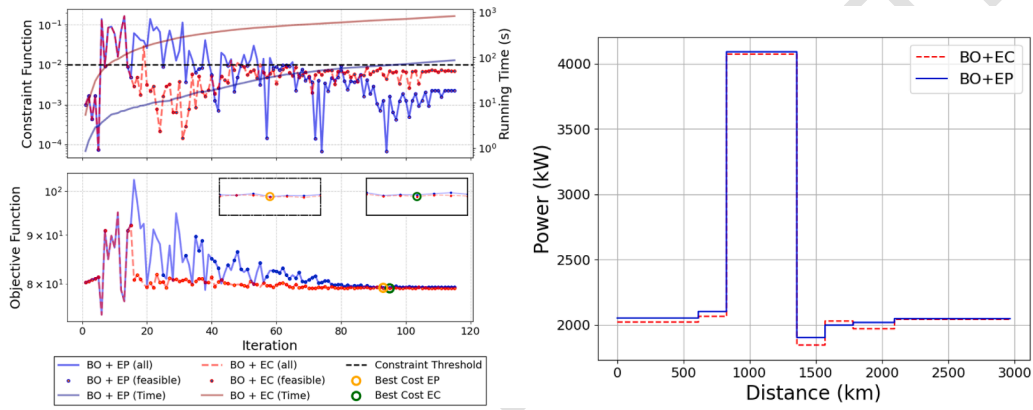


Fig. 9. Comparison of the efficiency and effectiveness of the BO + EC and BO + EP methods for voyage R2. Left: constraint and runtime evolution. Right: resulting optimal power profiles.

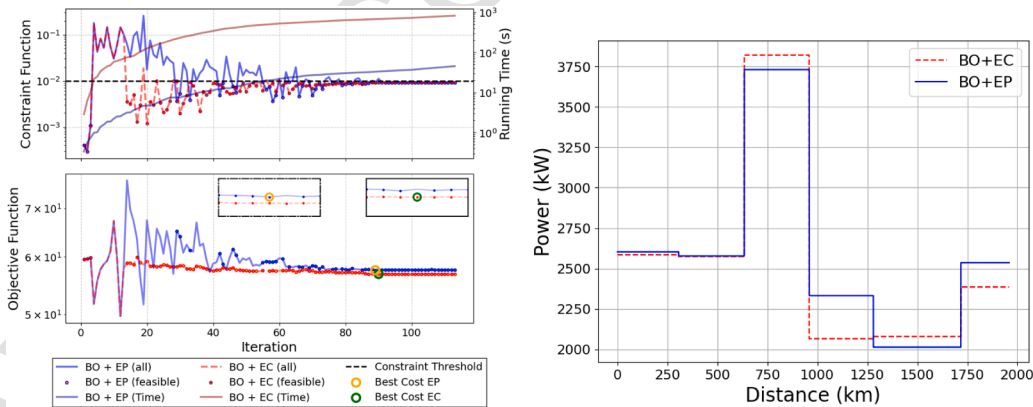


Fig. 10. Comparison of the efficiency and effectiveness of the BO + EC and BO + EP methods for voyage R3. Left: constraint and runtime evolution. Right: resulting optimal power profiles.

wave height reduced. During the final leg, neither method increases the power setting despite the presence of following seas with relatively low wave height. This behaviour reflects the requirement to satisfy the prescribed ETA constraint rather than pursuing additional fuel consumption savings. Overall, the BO + EP method reduces fuel consumption by 8.08% relative to the measured operation, decreasing fuel usage from 86.23 t to 79.26 t, and achieves an additional 1.16% reduction compared with the DP solution. The vessel still arrives well within the 1% ETA tolerance, as summarised in Table 7. The corresponding fuel consumption profiles for R1 and R2 are illustrated in Fig. 14 for reference.

3.6. Comparative evaluation against DP in the English channel

Voyages R3 and R4 pass through the English Channel, as illustrated in Fig. 5. Although the two voyages follow different routes, the sailing environment in this region is characterised by stronger and more variable current conditions compared with voyages R1 and R2. The optimised power allocation strategy for R3 is presented in Fig. 15.

Route R3 is partitioned into six legs, as illustrated in Fig. 15. Both the proposed BO + EP method and the DP approach increase the power setting in Leg 3, where wave conditions are relatively calm and both current speed and wind speed are low, creating favourable operating

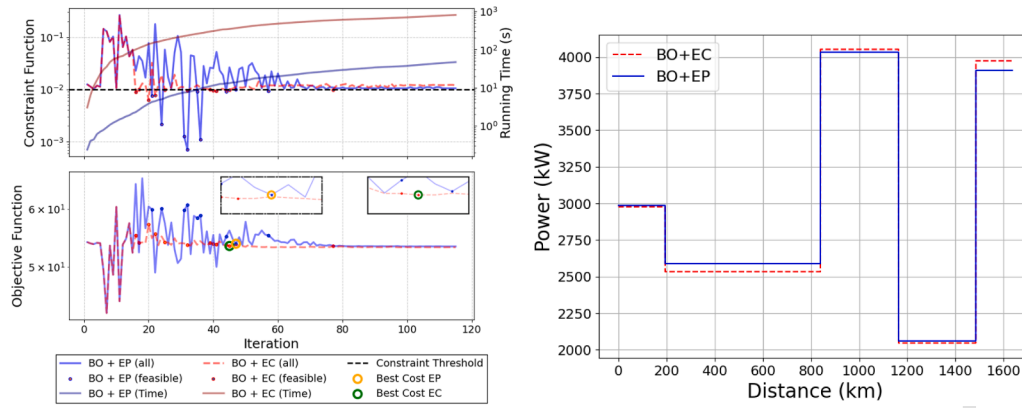


Fig. 11. Comparison of the efficiency and effectiveness of the BO+EC and BO+EP methods for voyage R4. Left: constraint and runtime evolution. Right: resulting optimal power profiles.

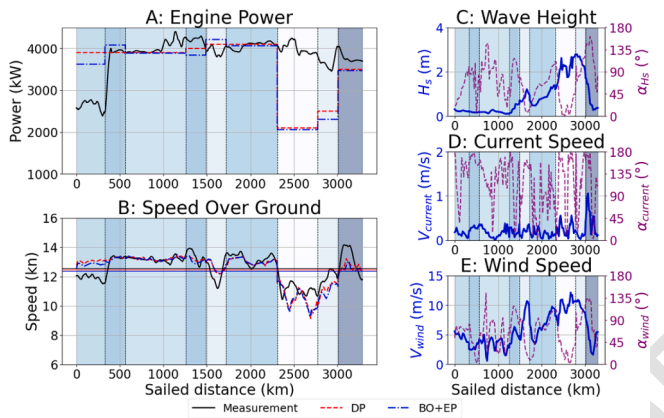


Fig. 12. Optimisation results and metocean conditions for voyage R1.

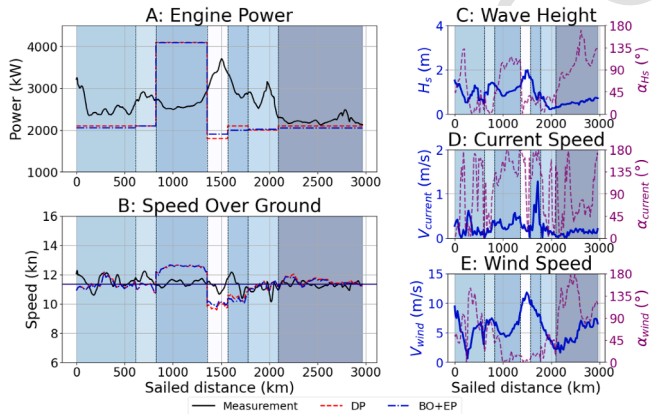


Fig. 13. Optimisation results and metocean conditions for voyage R2.

conditions. Compared with the DP solution, BO+EP selects a slightly lower power setting in Leg 3. In Leg 2, where two noticeable increases in current speed are observed, BO+EP again adopts marginally lower power levels than DP, indicating improved sensitivity to environmental assistance. For Leg 5 and the final leg, distinct adjustments are observed. In Leg 5, BO+EP slightly reduces engine power to enhance fuel savings. In contrast, during the final leg, BO+EP increases power due to the relatively calm wave conditions, allowing the vessel to compensate for earlier adjustments and ensure compliance with the ETA requirement. It reflects a more refined trade-off between fuel efficiency and schedule adherence. Overall, the BO+EP achieves fuel consumption saving from

Table 8

Performance comparison for voyages R3 and R4.

	Source	R3	R4	R3	R4
ETA	Actual	104.00	76.57	Δ_{ETA} w.r.t. Actual	
	DP	104.03	77.51	0.03%	1.23%
	BO+EP	104.96	77.37	0.92%	1.05%
	BO+EC	104.96	77.49	0.92%	1.21%
M_{fuel}	Actual	70.94	56.13	$\Delta_{M_{fuel}}$ w.r.t. Actual	
	DP	59.50	54.04	-16.12%	-3.72%
	BO+EP	57.50	53.32	-18.94%	-5.01%
	BO+EC	56.83	53.18	-19.88%	-5.25%

70.94 t to 57.50 t compared with the measured operation, as detailed in Table 8. Furthermore, it achieves an additional 2.82% fuel saving relative to the DP solution. The vessel still arrives well within the prescribed 1% ETA tolerance.

Finally, the optimised power allocation strategy for R4 is presented in Fig. 16, where the route is segmented into five legs. In Leg 2, the proposed BO+EP method selects a higher power setting than the DP solution. This adjustment is primarily driven by the decrease in wave height during the latter half of the leg, which reduces added resistance. In Leg 3, both methods substantially increase engine power since the reduced wind speed and current speed, creating favourable operating conditions. For Leg 4 and the final leg, the optimisation pattern is similar to that observed in R3. In Leg 4, where higher wave conditions are encountered, BO+EP noticeably reduces the power setting to improve fuel efficiency. In contrast, during the final leg, when wave conditions become milder, BO+EP increases engine power to ensure compliance with the ETA constraint.

The actual voyage duration is approximately 76.57 h. The BO+EP approach reduces fuel consumption by 5.01% compared with the measured operation, decreasing fuel consumption from 56.13 t to 53.32 t, as summarised in Table 8. In addition, it achieves a further fuel saving of approximately 1.3% relative to the DP solution. The corresponding fuel consumption profiles for R3 and R4 are illustrated in Fig. 17 for reference.

3.7. Sensitivity to metocean forecast uncertainty and operational implications

In the previous comparisons with measured fuel consumption and the DP, hindcast metocean data were adopted to demonstrate the effectiveness of the proposed optimisation framework. The use of hindcast data enables reconstruction of the actual environmental conditions encountered during the voyages, allowing a fair validation against measured fuel consumption while minimising environmental uncer-

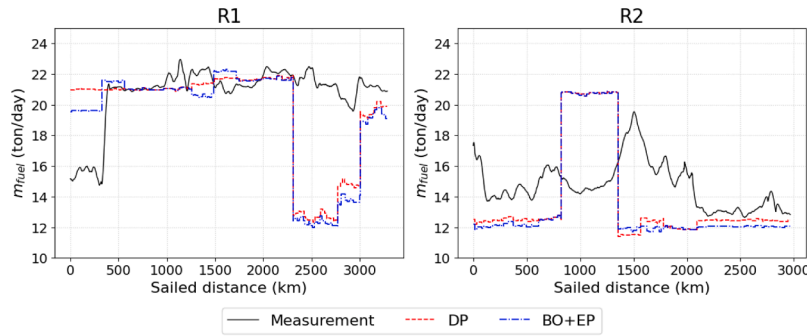


Fig. 14. Fuel consumption before and after optimisation for R1 and R2.

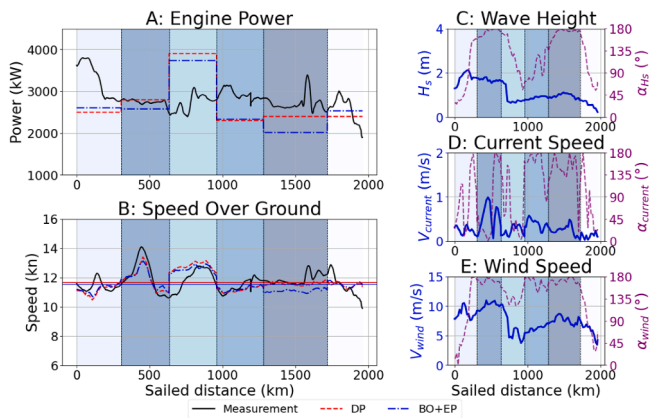


Fig. 15. Optimisation results and meteocean conditions for voyage R3.

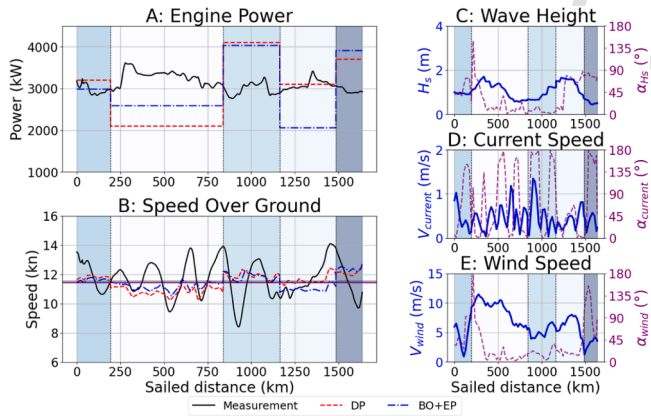


Fig. 16. Optimisation results and meteocean conditions for voyage R4.

tainty. In practical operations, however, voyage planning relies on weather forecasts rather than hindcast data, and forecast information is inherently subject to uncertainty. To evaluate the operational robustness of the proposed optimisation framework, this section investigates optimisation performance using forecast data retrieved from the Meteorological Archival and Retrieval System (MARS)(ECMWF, 2024), thereby representing a more realistic pre-voyage planning scenario. Specifically, voyage optimisation is conducted based on the weather forecast available on the day of departure. The MARS forecast data have a spatial resolution of 0.50 0.50, with a temporal resolution of 3 hours for the first 144 hours (6 days) and 6 hours thereafter up to day 10.

Voyage R1 in the Baltic and North Seas and voyage R3 in the English Channel are selected for the uncertainty investigation. Although

the forecast products are updated every 12 hours, in the following analysis, the forecast is updated once per day to compare the difference with the hindcast. Figs. 18 and 19 present the forecasted significant wave height at the start of the voyage in comparison with the corresponding hindcast data. The rolling forecast updates from Day 2, Day 3, and subsequent days until voyage completion are also illustrated.

Voyage R1 commenced on 6 June 2021 at 19:20 and concluded on 12 June 2021 at 16:10. For this analysis, the weather forecasts issued at 12:00 (UTC) each day were adopted, resulting in a total of seven consecutive daily forecasts over the voyage duration. As shown in Fig. 18, the Day 1 forecast (i.e., the forecast available at 12:00 on the departure day) matches the hindcast data well during the first three days of sailing. However, on the fourth day, a notable deviation emerges, with the forecasted significant wave height exceeding the hindcast by nearly 1 m. This discrepancy is further illustrated in Fig. 20, where the marker “” denotes the voyage starting point. It can be observed that in the latter half of the voyage, particularly in the North Sea region, the Day 1 forecast exhibits a clear divergence from the hindcast conditions. As the voyage progresses and updated forecasts become available, the deviation is significantly reduced. In particular, the forecasts issued on Day 4 and Day 5 substantially decrease the prediction error, as illustrated in Fig. 18. This demonstrates the benefit of rolling forecast updates in reducing meteocean environmental uncertainty.

For voyage R3, as shown in Fig. 19, the Day 1 forecast already exhibits deviations of up to 0.5 m from the hindcast conditions at the beginning of the voyage. The Day 2 forecast reduces this discrepancy. Overall, however, the general wave pattern along the entire voyage is well captured, particularly during the latter stages. This observation is further supported by Fig. 21.

The optimisation is then conducted using the proposed BO+EP framework based on the weather forecast available at the departure time. The resulting power allocation strategies are presented in Fig. 22 for voyage R1 and Fig. 23 for voyage R3. It can be observed that, when forecast data are used, voyage R1 is still segmented into nine legs. However, Legs 6 and 7, which appear as separate segments under forecast conditions, correspond to a single leg in the hindcast based optimisation. This difference arises because the forecast data predict higher wave heights and larger wind speed fluctuations compared with the hindcast conditions. Under these forecasted harsher head wave and head wind conditions, the optimisation prioritises maintaining the schedule connection between voyage legs. Consequently, in Leg 6 the engine power is maintained rather than reduced, unlike the hindcast based optimisation. In the subsequent leg, a power setting similar to that of the hindcast based optimisation is adopted to achieve fuel savings. As a result, optimisation using forecast data achieves a fuel saving of 2.53%, compared with the 7.30% potential saving obtained when hindcast data are used.

As shown in Fig. 23, voyage R3 is discretised into five legs when forecast data are used, whereas the hindcast based optimisation results in six legs (Fig. 15). The primary difference arises from the smoother

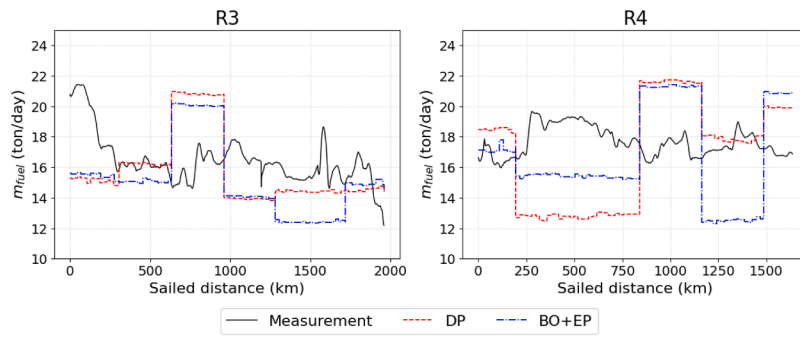


Fig. 17. Fuel consumption before and after optimisation for R3 and R4.




Fig. 18. Rolling comparison of forecasted and hindcast significant wave height for voyage R1, illustrating daily forecast updates.




Fig. 19. Rolling comparison of forecasted and hindcast significant wave height for voyage R3, including daily forecast updates.

environmental evolution predicted by the forecast data in Leg 3, which does not capture the pronounced decrease and subsequent recovery in wave height and wind speed observed in the hindcast conditions. Consequently, these environmental variations are grouped into a single clustered leg under forecast conditions. As a result, the mid-voyage power increase observed in the hindcast based optimisation does not occur

when using forecast data. Instead, engine power is increased in the final leg to satisfy the ETA constraint. Despite these discrepancies, the BO + EP method remains effective under forecast uncertainty, achieving a fuel saving of 15.83%. Although this is slightly lower than the 18.94% reduction obtained using hindcast data, it still represents a substantial improvement over the measured operation.

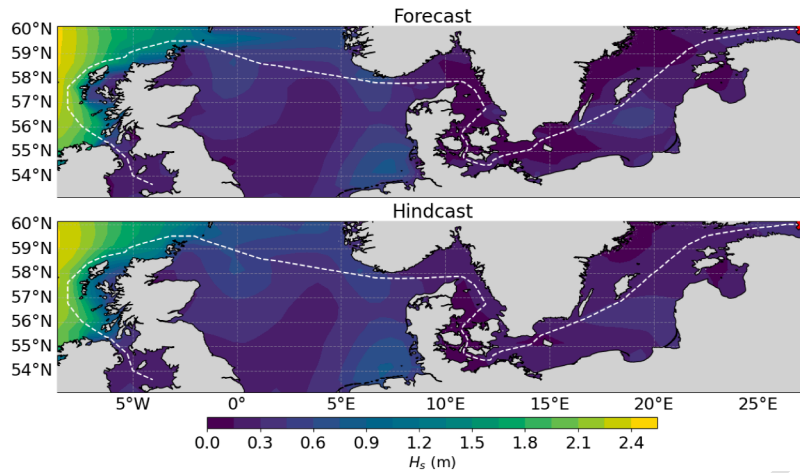


Fig. 20. Spatial comparison of forecasted and hindcast significant wave height for voyage R1 at departure time.

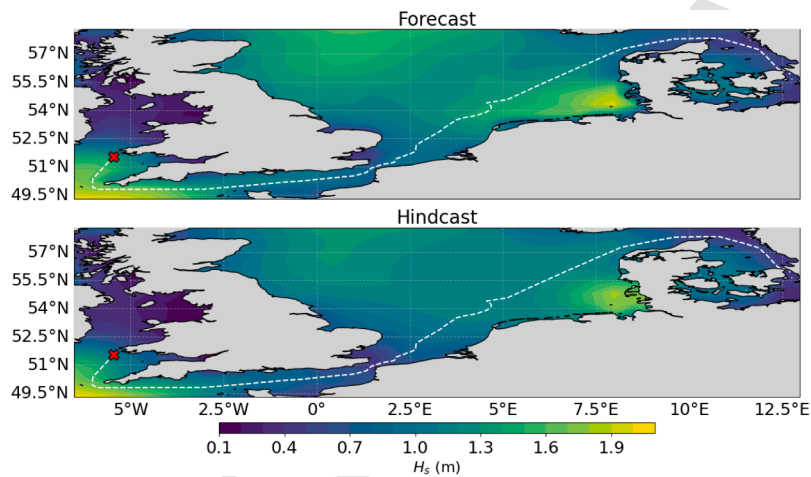


Fig. 21. Spatial comparison of forecasted and hindcast significant wave height for voyage R3 at departure time.

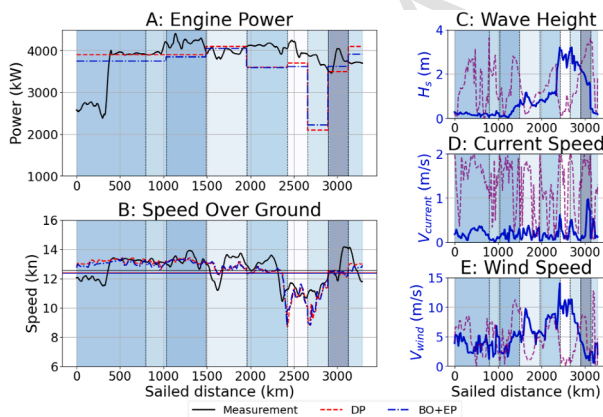


Fig. 22. Optimisation results and metocean conditions for voyage R1 when using forecast data.

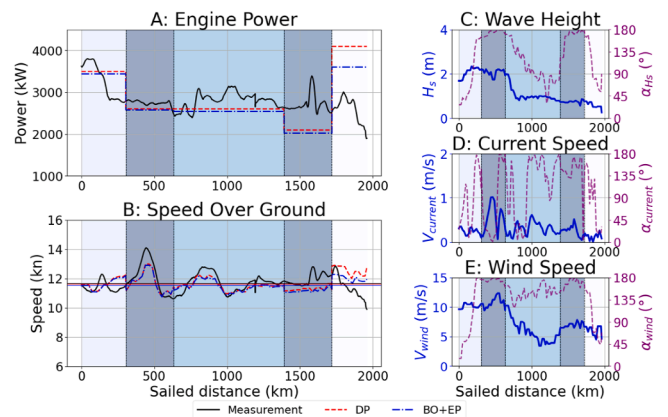


Fig. 23. Optimisation results and metocean conditions for voyage R3 when using forecast data.

The present investigation considers optimisation based solely on the weather forecast available at the departure time. Owing to the high computational efficiency of the proposed framework, this pre-voyage optimisation strategy can be naturally extended to real-time optimisation as updated weather forecasts become available during the voyage. Although such updates may lead to adjustments

in voyage segmentation, the availability of progressively more accurate metocean information enables improved power allocation decisions. Consequently, rolling re-optimisation has the potential to achieve fuel saving performance comparable to that obtained using hindcast data while maintaining operational practicality.

4. Conclusions

This study developed a novel structured discrete-to-continuous power allocation optimisation framework for short-sea shipping. Feasible solutions generated by a modified parallel coupling dynamic programming approach are used as prior knowledge, and subsequently refined using constraint-aware Bayesian optimisation. The proposed framework employs an embedded penalty mechanism that enables convex operational constraints to be handled explicitly within the optimisation process. As a result, the method achieves a good balance between computational efficiency and fuel saving performance.

The framework was validated using full-scale operational data from four voyages of a chemical tanker. For the case study voyages in the Baltic and North Seas, the optimised power allocation strategy achieves fuel consumption reductions of approximately 7–8% compared with measured operations. For voyages crossing the English Channel, fuel savings range from approximately 5% to 18%, depending on the prevailing environmental conditions. Compared with the conventional dynamic programming method, the proposed framework provides up to an additional 3% reduction in fuel consumption. The optimisation framework is formulated as a pre-voyage planning tool and was primarily validated using hindcast metocean data to ensure fair comparison with measured fuel consumption. To assess practical applicability, this study further investigated the impact of weather forecast uncertainty on optimisation outcomes. Differences between forecast and hindcast conditions were shown to influence voyage segmentation and corresponding power allocation strategies, leading to reduced achievable fuel savings when only departure-time forecasts are used. Owing to the high computational efficiency of the proposed framework, the method can be naturally extended to rolling or real-time optimisation as updated weather forecasts become available during the voyage. Such rolling re-optimisation enables progressively improved decision-making based on increasingly accurate metocean information, potentially achieving fuel saving performance comparable to hindcast based optimisation while maintaining operational practicality.

One limitation of the present study lies in the machine learning-based speed prediction model. The model may exhibit reduced generalisation capability under previously unseen environmental or operational conditions, which can introduce inaccuracies in voyage time estimation and ETA constraint evaluation. However, this limitation does not affect the validity of the proposed optimisation framework itself, as the framework remains applicable to improved or physics-informed performance models. Future work will focus on enhancing model generalisation through hybrid physics-informed learning approaches, and implementing real-time rolling optimisation under continuously updated forecast conditions to further support practical onboard deployment.

CRedit authorship contribution statement

Daniel Vergara: Writing – original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Shanshan Fu:** Writing – review & editing, Resources, Investigation, Conceptualization; **Wengang Mao:** Writing – review & editing, Supervision, Project administration, Funding acquisition; **Jonas W. Ringsberg:** Writing – review & editing, Validation, Investigation, Conceptualization; **Tsoulakos Nikolaos:** Writing – review & editing, Data curation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. Given his role as Editor-in-Chief, Jonas Ringsberg had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal

editor. Given their roles as Deputy Editor and Executive Guest Editor, Wengang Mao had no involvement in the peer-review of this article and has no access to information regarding its peer-review. Full responsibility for the editorial process for this article was delegated to another journal editor.

Acknowledgment

This work was supported by Trafikverket (Swedish Transport Administration) [grant No. TRV2023/98101]; the Vinnova (Swedish Governmental Agency for Innovation Systems) [grant No. 2021–02768]. The authors also acknowledge the ship owner for providing the full-scale measurement data.

References

- Bahrami, N., Siadatmousavi, S.M., 2024. Ship voyage optimisation considering environmental forces using the iterative Dijkstra's algorithm. *Ships Offshore Struct.* 19 (8), 1173–1180. <https://doi.org/10.1080/17445302.2023.2231200>
- Bahrami, N., Siadatmousavi, S.M., Kholgh, A.K., 2025. Algorithmic optimisation of ship routes for improved fuel economy and reduced carbon footprint across varying sea states. *Ships Offshore Struct.* <https://doi.org/10.1080/17445302.2025.2502694>
- Chen, T., Guestrin, C., 2016. XGBoost: a scalable tree boosting system. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, pp. 785–794. <https://doi.org/10.1145/2939672.2939785>
- Chen, Y., Mao, W., 2024. An isochrone-based predictive optimization for efficient ship voyage planning and execution. *IEEE Trans. Intell. Transp. Syst.* 25(11), <https://doi.org/10.1109/TITS.2024.3416349>
- Chen, Y., Tian, W., Mao, W., 2024. Strategies to improve the isochrone algorithm for ship voyage optimisation. *Ships Offshore Struct.* 19 (12), 2137–2149.
- Chen, Y., Zhang, C., Guo, Y., Wang, Y., Lang, X., Zhang, M., Mao, W., 2025. State-of-the-art optimization algorithms in weather routing - ship decision support systems: challenge, taxonomy, and review. *Ocean Eng.* 331, 121198.
- DNV, 2024a. FuelEU maritime - requirements, compliance strategies, and commercial impacts. DNV Report.
- DNV, 2024b. Maritime forecast to 2050 - energy transition outlook 2024. DNV Report.
- 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 (Part E), 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. *Transp. Res. Part B Methodol.* 122, 88–114.
- ECMWF, 2024. Meteorological Archival and Retrieval System. Accessed: 2026-02-22. <https://confluence.ecmwf.int/display/UDOC/MARS+user+documentation>.
- 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. *Comput. Indus. Eng.* 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. *Transp. Res. Part D Transp. Environ.* 113, 103527.
- Fiacco, A.V., McCormick, G.P., 1968. *Nonlinear Programming: Sequential Unconstrained Minimization Techniques*. John Wiley & Sons.
- Gelbart, M., Snoek, J., Adams, R.P., 2014. Bayesian optimization with unknown constraints. <https://arxiv.org/abs/1403.5607>.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., et al., 2020. The ERA5 global reanalysis. *Quarterly J. R. Meteorol. Soc.* 146 (730), 1999–2049. <https://doi.org/10.1002/qj.3803>
- Hu, Y., Li, R., Du, L., Ren, S., Chevallier, J., 2022. Could SO₂ and CO₂ emissions trading schemes achieve co-benefits of emissions reduction? *Energy policy* 170, 113252.
- International, M.O., 2024. Global ocean physics analysis and forecast. <https://www.mercator-ocean.eu>. Last metadata update: 18 June 2024.
- Jimenez, V.J., Kim, H., Munim, Z.H., 2022. A review of ship energy efficiency research and directions towards emission reduction in the maritime industry. *J. Cleaner Prod.* 366, 132888. <https://doi.org/10.1016/j.jclepro.2022.132888>
- Kasi, A., 2022. *Speed-Consumption Claims: A Guidebook. The Marine Law Box by Dr. Arun Kasi*. <https://arunkasico.com>.
- Lang, X., Wu, D., Mao, W., 2022a. Comparison of supervised machine learning methods to predict ship propulsion power at sea. *Ocean Eng.* 245, 110387. <https://doi.org/10.1016/j.oceaneng.2021.110387>
- Lang, X., Wu, D., Mao, W., 2022b. A machine learning ship's speed over ground prediction model and sailing time control strategy. *Int. J. Offshore Polar Eng.* 32 (4), 386–393. <https://doi.org/10.17736/ijope.2022.jc876>
- Lee, J., Park, Y., Eom, J., Hwang, H., Kim, S., 2025. Ship voyage route waypoint optimization method using reinforcement learning considering topographical factors and fuel consumption. *J. Mar. Sci. Eng.* 13 (8), 1554.
- 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 Eng.* 157, 301–312. <https://doi.org/10.1016/j.oceaneng.2018.03.068>
- 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. *J. Marine Sci. Eng.* 10 (12), 1835. <https://doi.org/10.3390/jmse10121835>

- 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 Eng.* 312, 119190. <https://doi.org/10.1016/j.oceaneng.2024.119190>
- Nocedal, J., Wright, S.J., 2006. *Numerical Optimization*. Springer, 2nd edition. <https://doi.org/10.1007/978-0-387-40065-5>
- Nogueira, F., 2014. Bayesian optimization: Open source constrained global optimization tool for python. Available at: <https://github.com/fmfn/BayesianOptimization>.
- Szlapczynska, J., Szlapczynski, R., 2019. Preference-based evolutionary multi-objective optimization in ship weather routing. *Appl. Soft Comput.* 84, 105742. <https://doi.org/10.1016/j.asoc.2019.105742>
- Tsioumas, V., Lyridis, D.V., Stavroulakis, P.J., Papadimitriou, S., 2025. Modeling just-in-time arrival effects on short sea shipping. *Marit. Bus. Rev.* 10 (3), 282–295.
- Tzortzis, G., Sakalis, G., 2021. A dynamic ship speed optimization method with time horizon segmentation. *Ocean Eng.* 226, 108840. <https://doi.org/10.1016/j.oceaneng.2021.108840>
- Vergara, D., 2025. *Voyage Segmentation and Propulsive Power Allocation: A Data-Driven Approach for Short Sea Shipping*. Licentiate thesis. Chalmers University of Technology. Department of Mechanics and Maritime Sciences, Marine Technology.
- Vergara, D., Alexandersson, M., Lang, X., Mao, W., 2023. A machine learning-based bayesian decision support system for efficient navigation of double-ended ferries. *J. Ocean Eng. Sci.* In Press. <https://doi.org/10.1016/j.joes.2023.11.002>
- Vergara, D., Lang, X., Zhang, M., Alexandersson, M., Mao, W., 2025. Reduced environmental impact of short sea shipping through optimal propulsion power allocation. *J. Cleaner Prod.* 513, 145683.
- Wang, H., Lang, X., Mao, W., 2021. Voyage optimization combining genetic algorithm and dynamic programming for fuel/emissions reduction. *Transp. Res. Part D Transp. Environ.* 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 Eng.* 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 Eng.* 216, 107591. <https://doi.org/10.1016/j.oceaneng.2020.107591>
- Wang, Z., Dahl, G.E., Swersky, K., Lee, C., Mariet, Z., Nado, Z., Gilmer, J., Snoek, J., Ghahramani, Z., 2022. Pre-training helps bayesian optimization too. <https://doi.org/10.48550/arXiv.2207.03084>.
- Zaccone, R., Ottaviani, E., Figari, M., Altosole, M., 2018. Ship voyage optimization for safe and energy-efficient navigation: a dynamic programming approach. *Ocean Eng.* 153, 215–224. <https://doi.org/10.1016/j.oceaneng.2018.01.100>
- Zhang, X., Yang, P., Gao, D., Bao, C., Ma, L., 2025. Two-stage energy efficiency optimization of ship speed and trim considering varying sea conditions and optimization frequency. *Ocean Eng.* 338, 121956.