arXiv:2602.18403v1 [cs.LG] 20 Feb 2026

Scientific Knowledge–Guided Machine Learning for Vessel Power Prediction: A Comparative Study

Orfeas Bourchas[1], **George Papalambrou**[1]

[1] Laboratory of Marine Engineering N.T.U.A.
orfeasbourchas@mail.ntua.gr, george.papalambrou@lme.ntua.gr

Abstract

Accurate prediction of main engine power is essential for vessel performance optimization, fuel efficiency, and compliance with emission regulations. Conventional machine learning approaches, such as Support Vector Machines, variants of Artificial Neural Networks (ANNs), and tree-based methods like Random Forests, Extra Tree Regressors, and XGBoost, can capture nonlinearities but often struggle to respect the fundamental propeller law relationship between power and speed, resulting in poor extrapolation outside the training envelope. This study introduces a hybrid modeling framework that integrates physics-based knowledge from sea trials with data-driven residual learning. The baseline component, derived from calm-water power curves of the form $P = cV^n$, captures the dominant power–speed dependence, while another, nonlinear, regressor is then trained to predict the residual power, representing deviations caused by environmental and operational conditions. By constraining the machine learning task to residual corrections, the hybrid model simplifies learning, improves generalization, and ensures consistency with the underlying physics. In this study, an XGBoost, a simple Neural Network, and a Physics-Informed Neural Network (PINN) coupled with the baseline component were compared to identical models without the baseline component. Validation on in-service data demonstrates that the hybrid model consistently outperformed a pure data-driven baseline in sparse data regions while maintaining similar performance in populated ones. The proposed framework provides a practical and computationally efficient tool for vessel performance monitoring, with applications in weather routing, trim optimization, and energy efficiency planning.

1 Introduction

1.1 Classical power prediction modeling

The maritime transportation sector is the backbone of international trade, but its contribution to global carbon emissions has led to increasingly stringent regulatory frameworks aimed at reducing the industry's carbon footprint (IMO 2020). Addressing these environmental and economic pressures requires the optimization of vessel performance through accurate prediction of operational parameters, most critically, the main engine power consumption. Early approaches to power estimation relied on simple empirical regression of sea-trial curves or calm-water resistance models (Carlton 2018). These methods capture the dominant cubic relationship (propeller law) between speed and shaft power but fundamentally fail to account for complex environmental and operational deviations (e.g., wind, waves, hull fouling, or off-design drafts) encountered in real-world service. Classical machine learning models, such as Gaussian Processes and Support Vector Regression (Bassam et al. 2023; Gkerekos, Lazakis, and Theotokatos 2019) have been employed, demonstrating improved capability to model nonlinear relationships compared to simple regression. However, these methods still struggle with the high dimensionality and volatility of vessel data. More recently, Artificial Neural Networks, including Long Short-Term Memory (LSTM) variants, have been widely adopted due to their universal approximation capability (Hornik, Stinchcombe, and White 1989). Studies have shown that ANNs achieve high accuracy in predicting shaft power and fuel consumption (Guo et al. 2023; Cai et al. 2024; Chen et al. 2024; La Ferlita et al. 2024; Zhang et al. 2024; Mahmoodi, Böling, and Vettor 2025). Domain-augmented LSTMs have even shown more stable training characteristics (Cai et al. 2024). In the last few years, gradient-boosted decision trees, notably XGBoost (Chen and Guestrin 2016), have emerged as the state-of-the-art for tabular vessel performance datasets (Nguyen et al. 2023; Agand et al. 2023; Fan et al. 2024; Kiouvrekis et al. 2025; Wang et al. 2024). These models consistently show competitive accuracy, often outperforming ANNs and classical methods in the prediction of vessel performance parameters due to their robustness to nonlinearity and data noise. Despite their high in-sample accuracy, all purely data-driven models suffer from limited extrapolation ability. When applied outside the observed operational envelope (e.g., predicting power at speeds or drafts not seen during training), these models often produce predictions that violate the fundamental propeller-law scaling, thus becoming physically inconsistent and unreliable.

1.2 Physics-Informed and Residual-Based Modeling

The recognized shortcomings of purely data-driven models have spurred the development of methods that embed physi-

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved. This paper was accepted to the KGML Bridge at AAAI 2026 (non-archival)

cal knowledge into the learning process. Physics-Informed Neural Networks (PINNs), formally introduced by Raissi et al. (Raissi, Perdikaris, and Karniadakis 2017a,b), have been proposed for vessel performance modeling to enforce the underlying physical laws (Lang, Wu, and Mao 2024; Bourchas and Papalambrou 2025). While PINNs can enforce physical consistency, they are extremely hard to train and can struggle to capture steep gradients and discontinuities present in real-world operational data, filled with noisy measurements when applied in the forward problem. Hybrid Residual Strategies have emerged in the broader physics-informed learning literature as a promising alternative (Howard et al. 2023; Eshkofti and Barreau 2025). These methods decompose the complex prediction task: a simple physics-based component models the dominant relationship, and a data-driven component learns the residual or deviation caused by complex, unaccounted-for factors.

1.3 Our Contribution

Inspired by the residual-hybrid strategy, a novel hybrid modeling framework for main-engine power prediction is proposed. The new framework effectively combines the stability of physics-based models with the predictive power of modern ML. The vessel's sea-trial curves are utilized as a physically consistent baseline, which captures the dominant power-speed relationship, while a non-linear regressor is employed to learn the residual corrections—the difference between the observed power and the baseline—caused by environmental factors (e.g., waves, wind) and operational factors (e.g., trim, draft). By limiting the machine learning task to predicting only the residual corrections, our methodology significantly enhances extrapolation stability and physical consistency beyond conventional data-driven approaches. The main contributions of this work are:

- The development of a hybrid modeling framework that overcomes traditional extrapolation limitations in vessel power prediction by leveraging a physics-informed baseline coupled with a non-linear regressor.
- A comparative study applying the proposed hybrid framework to XGBoost, a simple Neural Network, and a Physics-Informed Neural Network (PINN) is conducted, demonstrating the framework's versatility and efficacy across diverse modeling paradigms.
- Validation using in-service data from real vessel operations provides a robust, computationally efficient tool for maritime optimization applications like weather routing and fuel efficiency planning.

2 Methods

2.1 Main Engine Power Modeling

Traditional machine learning approaches estimate the required main engine power by fitting a regression model directly to operational data. Formally, the predicted power can be expressed as:

$$\hat{P}(\bar{X}) = f(\bar{X}) \tag{1}$$

Here, the function $f(\cdot)$ is a complex machine learning model and $\bar{X}$ represents the input parameters to the model, which represent the operating condition. Usually, these parameters include the vessel's speed through water (S.T.W.), the vessel's mean draft (T), the vessel's trim, and the wind speed and direction, as well as several others. Even though this approach has yielded promising results in the literature, its effectiveness is tied to the quality and quantity of the available data. However, in most cases, the quantity of available clean data is scarce as the hull's and propeller's performance deteriorates over time, either by the accumulation of fouling or damage to the coating and paints.

2.2 Hybrid Main Engine Power Modeling

To address the limitations of purely data-driven models, particularly their inability to confidently extrapolate beyond the training data envelope, a hybrid modeling framework is proposed. This approach decomposes the complex power prediction problem into two distinct components: a physics-based baseline model and a data-driven correction model. This separation ensures that the fundamental physical relationships are captured by a robust, simple model, while the machine learning component is tasked with learning only the deviations from this baseline. This decomposition allows the model to confidently and accurately extrapolate because the core physics are already accounted for.

The baseline component is obtained from calm-water power curves measured at ballast and laden drafts during sea trials. These curves follow a power-law relation of the form

$$P = cV^n \tag{2}$$

where V is the vessel's speed through water, c and n are vessel-specific coefficients and are determined from the sea trial data, as presented in Appendix A. For intermediate drafts as seen in Figure 1, the baseline power is calculated by linear interpolation using equation 3:

$$P_{\text{sea trial}}(V,T) = \left(1 - \frac{T - T_b}{T_l - T_b}\right) \cdot P_b(V) + \left(\frac{T - T_b}{T_l - T_b}\right) \cdot P_l(V), \tag{3}$$

where T_b and T_l are the ballast and laden drafts, respectively, and $P_b(V)$ and $P_l(V)$ are the corresponding propeller power curves obtained from the sea trials. Even though the linear interpolation between drafts is a simplification of the underlying physics, it provides a robust baseline for intermediate drafts and satisfies the ballast and laden conditions simultaneously. The hybrid model's final prediction is the sum of this physics-based baseline and a learned correction term, $f(\bar{X})$, which accounts for all unaccounted factors such as weather, hull fouling, or engine degradation.

The total predicted power, as seen in Figure 2 can be expressed as in equation 4

$$\hat{P}(\bar{X}) = P_{\text{sea trial}}(V,T) + f(\bar{X}) \tag{4}$$

Here, $f(\cdot)$ is an XGBoost regressor, which is trained to predict the residual power. This residual is the difference between the measured main engine power and the sea trial baseline power. The input features to the machine learning model, $\bar{X}$, include parameters not captured by the sea trial data, such as wind speed and direction, wave height, vessel aging, and time since the last dry-docking. By constraining the machine learning model to learn only residuals, the

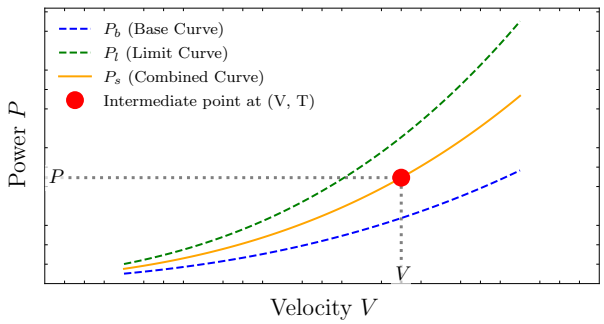

Figure 1: Calm-water baseline construction. The laden (P_l) and ballast (P_b) sea-trial curves define the upper and lower bounds, while the combined curve $P_{\text{sea trial}}(V,T)$ from Eq. (3) provides the interpolated baseline at intermediate drafts. The red point marks the baseline value at the operating condition (V,T).

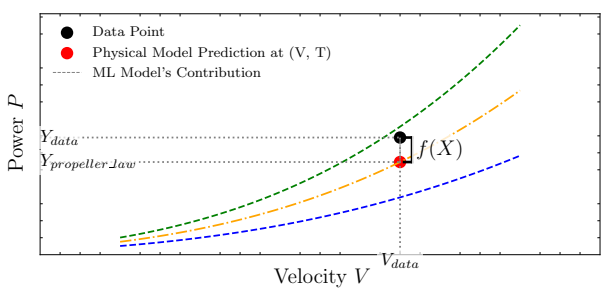

Figure 2: Hybrid model decomposition as expressed in Eq. (4). The laden and ballast sea-trial curves define the calm-water envelope, and the red point indicates the interpolated baseline at (V,T). The black point represents an actual measurement. The vertical bracket labeled $f(\mathbf{X})$ illustrates the learned residual correction added to the baseline to obtain the hybrid prediction.

prediction task becomes simpler and better regularized. This structure allows the model to generalize across a wider operational envelope, maintain consistency with physical laws, and extrapolate reliably beyond the training distribution.

2.3 Deriving PINN's Loss

The PINN is trained by minimizing a total loss that combines data fitting with a physics-based constraint.

The data loss is defined as

$$\mathcal{L}_{\text{data}} = \frac{1}{N}\sum_{i=1}^{N}\left(\hat{P}(\bar{X}_i) - P_i\right)^2, \tag{5}$$

and the physics loss $\mathcal{L}_{P.\text{law}}$ is given in Eq. (8). The total PINN loss is then

$$\mathcal{L}_{\text{PINN}} = \mathcal{L}_{\text{data}} + \lambda\,\mathcal{L}_{P.\text{law}}, \tag{6}$$

where λ is a weighting coefficient set to 100 in this work. Although fixed here for simplicity and stability, λ can be treated as a tunable hyperparameter in future studies.

As described in (Bourchas and Papalambrou 2025), the required power of a vessel's main engine P with respect to the vessel's speed through water V and the operating condition is given by

$$P = \frac{c_T \cdot W_{SA} \cdot V^3}{2 \cdot \eta_S \cdot \eta_B \cdot \eta_H} = c \cdot V^3 \tag{7}$$

and the proposed physics loss is defined in Eq. (8)

$$\mathcal{L}_{P.\text{law}} = \sum_{i=1}^{N}\left(\frac{\partial NN}{\partial V}\Big|_{x_i} - 3 \cdot c \cdot V_i^2\right)^2 \tag{8}$$

where $P \simeq NN$. In the hybrid framework, P is given by equation 4. Taking the partial derivative of P with respect to the vessel speed through water V:

$$\frac{\partial P}{\partial V} = \frac{\partial P_{\text{sea trial}}(V,T)}{\partial V} + \frac{\partial f(\bar{X})}{\partial V} \tag{9}$$

substituting equations 3, 7, and rearranging the terms:

$$\begin{aligned}\frac{\partial f(\bar{X})}{\partial V} &= 3 \cdot c \cdot V^2 \\ &- \left(1 - \frac{T - T_b}{T_l - T_b}\right) \cdot n_b \cdot c_b \cdot V^{n_b - 1} \\ &- \left(\frac{T - T_b}{T_l - T_b}\right) \cdot n_l \cdot c_l \cdot V^{n_l - 1}\end{aligned} \tag{10}$$

where,

$$\begin{aligned}\frac{\partial P_b(V)}{\partial V} &= n_b \cdot c_b \cdot V^{n_b - 1}, \\ \frac{\partial P_l(V)}{\partial V} &= n_l \cdot c_l \cdot V^{n_l - 1},\end{aligned} \tag{11}$$

All derivatives appearing in the physics loss are computed using `PyTorch`'s automatic differentiation engine (`autograd`).

3 Dataset and Model Training

3.1 Dataset

In the present work, the dataset from (Bourchas and Papalambrou 2025) was used. The dataset consists of approximately 40000 data points, representing five months of operational data and is composed of the parameters described in Table 1. The dataset was randomly split into 80%-10%-10% subsets representing training, validation, and testing sets. Drawing inspiration from the work of (Bourchas and Papalambrou 2025), the *Wind True Speed* and *Wind True Direction* were decomposed into the W_x and W_y components using the equation 12

$$\begin{aligned}W_x &= \text{WTS} \cdot \cos(\text{WTD}), \\ W_y &= \text{WTS} \cdot \sin(\text{WTD}),\end{aligned} \tag{12}$$

Table 1: Dataset variables used as model inputs.

Symbol	Description	Unit
t	Time of the recording (UTC date)	–
P	Main engine brake power	kW
V	Ship's speed through water (S.T.W.)	kn
T	Ship's mean draft	m
Trim	Longitudinal trim	m
WTS	Wind true speed	kn
WTD	Wind true direction	°

Table 2: Optimized hyperparameters for XGBoost models (Baseline vs. Hybrid).

Hyperparameter	Baseline	Hybrid
Learning rate (η)	0.25	0.1
Max depth	10	10
n-estimators	40	500
L1 regularization (α)	1	1
L2 regularization (λ)	0	100

3.2 Model Training and Hyperparameter Optimization

To ensure a robust and fair evaluation, both the baseline and hybrid configurations of all three model architectures—XGBoost, Neural Network (NN), and Physics-Informed Neural Network (PINN)—underwent a unified and rigorous hyperparameter optimization (HPO) procedure.

For the XGBoost implementation, the widely used `XGBoost` package was used, and hyperparameter tuning was carried out using `Scikit-learn`'s `RandomizedSearchCV`. The optimization explored parameter distributions including the learning rate (η) in $0.01, 0.05, \ldots, 0.30$, maximum tree depth in $3, \ldots, 10$, number of estimators (`n_estimators`) in $5, \ldots, 500$, and L1 (α) and L2 (λ) regularization coefficients in $0, 1, 100$. The negative mean squared error was used as the scoring metric. During training, both input and output variables were standardized using `Scikit-learn`'s `StandardScaler`, fitted on the training dataset.

The NN and PINN models were developed and trained within the `PyTorch` ecosystem, with hyperparameter optimization conducted via `Weights & Biases (WandB)`'s `sweep` functionality. A Bayesian optimization strategy was adopted, minimizing the root mean squared error (RMSE) on the test set as the objective metric. The search space for both NN and PINN models was centered around the optimal settings suggested in (Bourchas and Papalambrou 2025), exploring key parameters such as `learning rates` $1 \times 10^{-2}, 1 \times 10^{-3}, 3 \times 10^{-4}, 1 \times 10^{-4}$, `layer widths` of $64, 128, 256$ neurons, and `network depths` of $4, 6, 8$ layers. Each model was trained using the Adam optimizer with the selected learning rate for 1000 epochs, without learning rate scheduling or early stopping. For the PINN model specifically, the total loss was defined as in Eq. (6), with λ fixed to 100 as discussed in Section 2.3.

The final optimized hyperparameter values for all models (XGBoost, NN, and PINN), obtained through these comprehensive optimization procedures, are summarized in Tables 2, 3, 4.

3.3 Hardware Configuration

All computational experiments were conducted on a machine with Windows 10 Pro as the operating system, an AMD Ryzen 5500 CPU, an NVIDIA RTX 3060 GPU with 12 GB of VRAM, and 32 GB of RAM.

Table 3: Optimized hyperparameters for Neural Network (NN) models (Baseline vs. Hybrid).

Hyperparameter	Baseline	Hybrid
Learning rate	1×10^{-4}	3×10^{-4}
Hidden layers	6	4
Neurons per layer	256	64

Table 4: Optimized hyperparameters for Physics-Informed Neural Network (PINN) models (Baseline vs. Hybrid).

Hyperparameter	Baseline	Hybrid
Learning rate	1×10^{-4}	3×10^{-4}
Hidden layers	8	6
Neurons per layer	256	256

4 Results and Model Evaluation

This section evaluates the baseline and hybrid variants of the XGBoost, NN, and PINN models using both standard error metrics and a detailed analysis of their extrapolation behavior. Although the quantitative differences between baseline and hybrid models are modest, the qualitative analysis reveals clear advantages for the hybrid formulations in terms of physical consistency and robustness under unseen operating conditions.

4.1 Quantitative Analysis

Predictive accuracy was assessed using two standard metrics: Mean Absolute Error (MAE) and Root Mean Square Error (RMSE). After hyperparameter optimization, each model was evaluated on the training, validation, and test sets, with results summarized in Table 5.

Overall, the XGBoost baseline achieved slightly lower training errors than its hybrid counterpart, which is consistent with its strong capacity to fit the available data. Similar behavior was observed for the NN and PINN architectures, where baseline models occasionally achieved marginally lower MAE and RMSE values. However, across all architectures, these quantitative differences remained small and well within acceptable limits for operational maritime applications.

It is important to note that such global error metrics do not capture model behavior outside densely populated regions of the training dataset. As shown in the next subsection, baseline models display clear degradation in extrapolation scenarios, whereas the hybrid variants consistently deliver more stable and physically coherent predictions.

4.2 Qualitative Extrapolation Analysis

To assess robustness under unseen operating conditions, each model family was evaluated across speeds from 8–17 kn at ballast draft and 5-kn wind, with wind directions of $0°$, $90°$, and $180°$. Figures 3, 4, and 5 present representative results for XGBoost, NN, and PINN, respectively. Blue points denote the closest training samples, while yellow triangles represent nearby test samples, highlighting regions where the models must extrapolate.

Table 5: Performance comparison across all architectures.

Metric	XGBoost		Neural Network (NN)		Physics-Informed NN (PINN)	
	Baseline	**Hybrid**	**Baseline**	**Hybrid**	**Baseline**	**Hybrid**
Train MAE [kW]	114.2	143.3	181.01	235.30	147.03	214.73
Train RMSE [kW]	180.5	199.9	229.67	292.12	205.48	248.30
Validation MAE [kW]	114.2	143.3	156.29	214.08	135.52	167.17
Validation RMSE [kW]	180.5	199.9	212.94	272.38	204.21	219.51
Test MAE [kW]	122.2	148.8	162.66	219.32	144.30	171.19
Test RMSE [kW]	195.0	208.2	225.10	284.33	211.89	229.45

XGBoost (Figure 3). The baseline XGBoost model frequently exhibits non-monotonic or flattened power–speed behavior, especially at higher speeds where training data are not available. Such distortions violate the physically expected propeller-law scaling and lead to predictions that are not operationally credible. The hybrid XGBoost model corrects this issue by anchoring the prediction to the calm-water baseline, which preserves a smooth and monotonic increase in power even where training data are sparse.

Neural Network (Figure 4). The baseline NN displays similar extrapolation deficiencies, including overestimation at high speeds and noticeable sensitivity to wind direction in poorly sampled regions. The hybrid NN significantly reduces these artifacts. Because it learns only residual deviations around a physically consistent baseline, the resulting power–speed curves are smoother and exhibit improved generalization across the full speed range.

Physics-Informed Neural Network (Figure 5). The hybrid PINN demonstrates the strongest extrapolation performance across all architectures. While the baseline PINN benefits from its physics-informed loss, it can still deviate in sparsely represented regions. The hybrid formulation enforces consistency with both the sea-trial power law and the PINN derivative constraints, producing physically coherent predictions across all wind directions and throughout the 8–17 kn range.

Summary. Across all model families, the hybrid formulations consistently improve extrapolation stability while maintaining accuracy in dense data regions. These gains stem directly from embedding domain knowledge: by relieving the models from rediscovering the dominant power–speed law, the hybrid structure yields more reliable predictions in unseen or sparsely populated operating regimes. Among all architectures, the hybrid PINN demonstrates the best balance between numerical accuracy and physical consistency, making it particularly attractive for decision-support applications where extrapolation reliability is critical.

Taken together, these findings indicate that the main benefit of the hybrid formulations is not a reduction in error metrics, but a regularization of model behavior in sparsely sampled and extrapolative regimes. In the most demanding operating conditions examined, such as high speeds com-

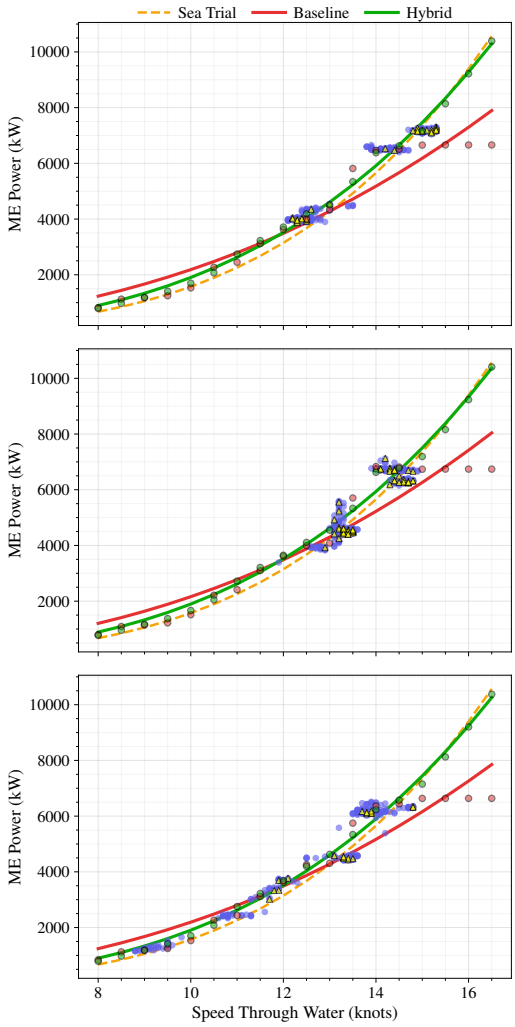

Figure 3: Extrapolation behavior of the XGBoost baseline and hybrid models at ballast draft and 5-kn wind. Rows correspond to wind directions of 0°, 90°, and 180°. Red and green curves show the baseline and hybrid predictions, respectively. Blue points denote the nearest training data, while yellow triangles indicate nearest test samples.

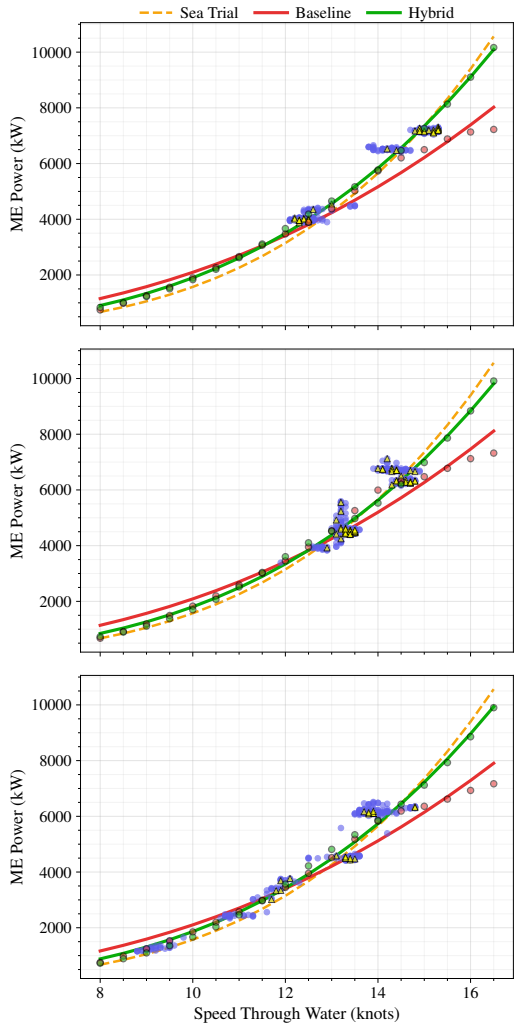

Figure 4: Extrapolation behavior of the NN baseline and hybrid models at ballast draft and 5-kn wind. Rows correspond to wind directions of $0°$, $90°$, and $180°$. Red and green curves show the baseline and hybrid predictions, respectively. Blue points denote the nearest training data, while yellow triangles indicate nearest test samples.

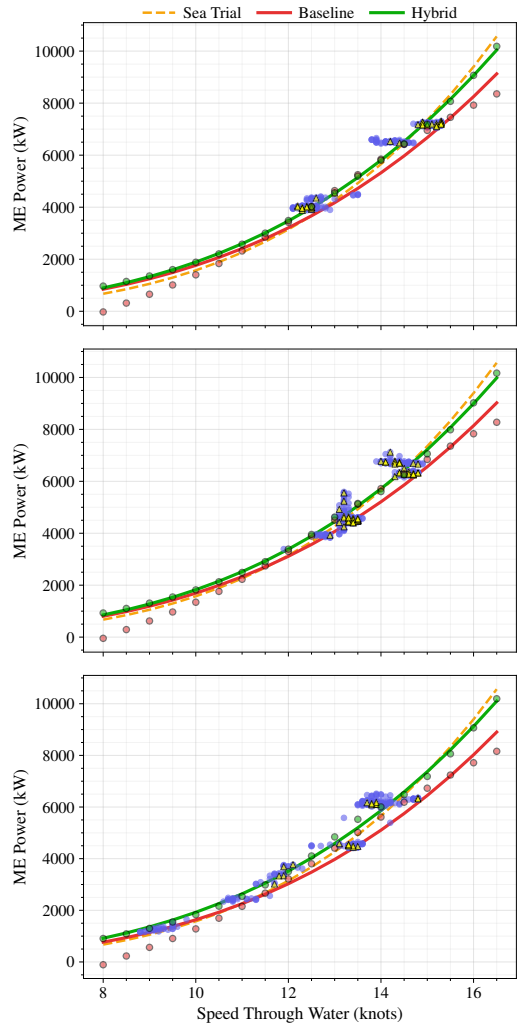

Figure 5: Extrapolation behavior of the PINN baseline and hybrid models at ballast draft and 5-kn wind. Rows correspond to wind directions of $0°$, $90°$, and $180°$. Red and green curves show the baseline and hybrid predictions, respectively. Blue points denote the nearest training data, while yellow triangles indicate nearest test samples.

bined with adverse winds, the hybrid models remain close to the physically expected power–speed trend, whereas the baseline models may flatten or deviate from the calm-water envelope. This distinction is particularly important for applications like weather routing and speed optimization, where decisions are often made precisely in regions with limited historical data coverage.

5 Conclusion

Across all model families, the baseline and hybrid variants showed small quantitative differences, well under 1% of the vessel's maximum power, indicating comparable accuracy in standard error metrics. However, these metrics do not capture a critical distinction: extrapolation performance. Hybrid models consistently produced smoother and more physically coherent power–speed curves outside the training envelope, while baseline models, especially XGBoost, often exhibited non-physical behavior at higher speeds or under extreme environmental conditions. Among all architectures, the hybrid PINN delivered the best overall performance, combining competitive accuracy with the strongest physical consistency in unseen regimes. These results demonstrate that incorporating physical structure through hybridization significantly improves reliability under extrapolation. The proposed framework is well suited for real-world maritime decision-support applications, such as weather routing and trim optimization, where accurate power prediction is required even when no historical data are available.

A Sea Trial Parameters

Sea trials are a set of standardized tests used to determine a vessel's performance characteristics under controlled, calm weather conditions. During these trials, a ship's speed (V) and main engine power (P) are systematically measured at specific draft conditions, typically in ballast (T_b) and design or laden (T_l) states.

The relationship between power and speed is commonly approximated using a power law:

$$P = cV^n$$

where c and n are coefficients specific to the vessel. To determine these coefficients, the power law is linearized by taking the natural logarithm of both sides:

$$\ln(P) = \ln(c) + n\ln(V)$$

This equation represents a linear relationship between $\ln(P)$ and $\ln(V)$. By performing a linear regression on the log-transformed sea trial data, the coefficients $\ln(c)$ and n can be determined, providing a robust power-speed curve for the vessel in calm water.

References

Agand, P.; Kennedy, A.; Harris, T.; Bae, C.; Chen, M.; and Park, E. J. 2023. Fuel consumption prediction for a passenger ferry using machine learning and in-service data: A comparative study. *Ocean Engineering*, 284: 115271.

Bassam, A. M.; Phillips, A. B.; Turnock, S. R.; and Wilson, P. A. 2023. Artificial neural network based prediction of ship speed under operating conditions for operational optimization. *Ocean Engineering*, 278: 114613.

Bourchas, O.; and Papalambrou, G. 2025. Physics informed neural networks for vessel main engine power prediction. *Ocean Engineering*, 332: 121344. In press.

Cai, Z.; Li, L.; Yu, L.; Li, C.; and Sun, M. 2024. Diversity, quality, and quantity of real ship data on the black-box and gray-box prediction models of ship fuel consumption. *Ocean Engineering*, 291: 116434.

Carlton, J. 2018. *Marine Propellers and Propulsion*. Butterworth-Heinemann.

Chen, T.; and Guestrin, C. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.

Chen, Y.; Sun, B.; Xie, X.; Li, X.; Li, Y.; and Zhao, Y. 2024. Short-term forecasting for ship fuel consumption based on deep learning. *Ocean Engineering*, 301: 117398.

Eshkofti, K.; and Barreau, M. 2025. Vanishing Stacked-Residual PINN for State Reconstruction of Hyperbolic Systems. *IEEE Control Systems Letters*, 9: 1417–1422. Early Access.

Fan, A.; Wang, Y.; Yang, L.; Tu, X.; Yang, J.; and Shu, Y. 2024. Comprehensive evaluation of machine learning models for predicting ship energy consumption based on onboard sensor data. *Ocean & Coastal Management*, 248: 106946.

Gkerekos, C.; Lazakis, I.; and Theotokatos, G. 2019. Machine learning models for predicting ship main engine Fuel Oil Consumption: A comparative study. *Ocean Engineering*, 188: 106282.

Guo, B.; Gupta, P.; Steen, S.; and Tvete, H. A. 2023. Evaluating vessel technical performance index using physics-based and data-driven approach. *Ocean Engineering*, 286: 115402.

Hornik, K.; Stinchcombe, M.; and White, H. 1989. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5): 359–366.

Howard, A. A.; Murphy, S. H.; Ahmed, S. E.; and Stinis, P. 2023. Stacked networks improve physics-informed training: applications to neural networks and deep operator networks. arXiv:2311.06483.

IMO. 2020. Fourth IMO GHG Study 2020. *International Maritime Organization, London*.

Kiouvrekis, Y.; Gkirtzou, K.; Zikas, S.; Kalatzis, D.; Panagiotakopoulos, T.; Lajic, Z.; Papathanasiou, D.; and Filippopoulos, I. 2025. An Explainable Machine Learning Approach for IoT-Supported Shaft Power Estimation and Performance Analysis for Marine Vessels. *Future Internet*, 17(6): 264. In press.

La Ferlita, A.; Qi, Y.; Di Nardo, E.; El Moctar, O.; Schellin, T. E.; and Ciaramella, A. 2024. A framework of a data-driven model for ship performance. *Ocean Engineering*, 309: 118486. In press.

Lang, X.; Wu, D.; and Mao, W. 2024. Physics-informed machine learning models for ship speed prediction. *Expert Systems with Applications*, 238: 121877. Published online 2023.

Mahmoodi, K.; Böling, J.; and Vettor, R. 2025. Evaluating the influence of marine weather parameters uncertainties on the ship fuel consumption with Monte Carlo analysis. *Ocean Engineering*, 341: 122531. In press.

Nguyen, S.; Fu, X.; Ogawa, D.; and Zheng, Q. 2023. An application-oriented testing regime and multi-ship predictive modeling for vessel fuel consumption prediction. *Transportation Research Part E*, 177: 103261.

Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2017a. Physics Informed Deep Learning (Part I): Data-driven Solutions of Nonlinear Partial Differential Equations. *arXiv preprint arXiv:1711.10561*.

Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2017b. Physics Informed Deep Learning (Part II): Data-driven Discovery of Nonlinear Partial Differential Equations. *arXiv preprint arXiv:1711.10566*.

Wang, K.; Liu, X.; Guo, X.; Wang, J.; Wang, Z.; and Huang, L. 2024. A novel high-precision and self-adaptive prediction method for ship energy consumption based on the multi-model fusion approach. *Energy*, 310: 133265.

Zhang, M.; Tsoulakos, N.; Kujala, P.; and 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.