

TCN-FECAM-iTransformer: A Hybrid Framework With Frequency-Enhanced Channel Attention And Inverted Transformer For Ship Speed Prediction

Jie Zhang¹, Yanghui Tan^{1*}, Jundong Zhang², and Enzhi Zhu¹

¹ Maritime College, Tianjin University of Technology, Tianjin 300384, China

² Marine Engineering College, Dalian Maritime University, Dalian 116026, China

* Corresponding author. E-mail: tyanghui@email.tjut.edu.cn (T.Y.)

Received: May 12, 2026; Accepted: Jun. 17, 2026

Accurate ship speed prediction in real-world marine operations remains challenging due to noise interference, strong multivariate coupling, non-stationarity, and data distribution shifts. This study proposes TCN-FECAM-iTransformer, a novel hybrid deep learning framework that synergistically integrates Temporal Convolutional Network (TCN), Frequency-Enhanced Channel Attention Mechanism (FECAM), and iTransformer for high-precision one-step-ahead ship speed prediction. To address the critical challenges of strong marine environmental noise, non-stationarity, and complex multivariate coupling in real-world ship operations, the framework incorporates feature selection and carefully designed component collaboration. Specifically, Pearson correlation and mutual information (MI)-based feature selection are first applied to improve input quality by capturing linear and nonlinear dependencies, respectively.

Subsequently, the TCN module extracts multi-scale local temporal features, the FECAM module suppresses noise and enhances channel-wise representations through frequency-domain analysis, and the iTransformer effectively models long-range dependencies across multiple variables. On real voyage data collected from an LPG carrier, the proposed framework achieves the best overall prediction performance, attaining an RMSE of **0.1952 kn** and an R^2 of **0.9707** under the Pearson Top-12 feature selection strategy. Compared with the strongest baseline model (iTransformer), the proposed approach reduces the RMSE by approximately **13.5%**. Ablation studies further verify the effectiveness and complementary roles of the TCN and FECAM modules. Overall, this study provides an efficient and robust solution for ship speed prediction and demonstrates that the combination of appropriate feature selection and hybrid deep learning architectures can significantly improve the modeling of complex multivariate maritime time series.

Keywords: Ship Speed Prediction; Frequency-Enhanced Channel Attention; Inverted Transformer; Mutual Information

© The Author(s). This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY 4.0\)](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are cited.

http://dx.doi.org/10.6180/jase.202610_33.034

1. Introduction

With the rapid development of the global shipping industry and the increasingly stringent requirements of the International Maritime Organization (IMO) on ship energy efficiency and carbon emissions, accurate ship speed prediction has become an important enabling technology for intelligent, energy-efficient, and low-carbon maritime trans-

portation [1]. According to the 2023 IMO Strategy on Reduction of Greenhouse Gas (GHG) Emissions from Ships, the carbon intensity of international shipping must be reduced by at least 40% by 2030 compared to 2008 levels, while pursuing net-zero GHG emissions by or around 2050 [2]. These ambitious targets impose higher demands on ship energy consumption optimization. Accurate ship

speed prediction serves as core technical support for precise estimated time of arrival calculation, fuel consumption minimization, dynamic route optimization, and overall carbon emission control [3].

Ship speed is not only directly linked to navigation safety, port scheduling efficiency, and emission performance, but is also governed by multiple highly coupled environmental and operational factors under complex sea conditions, including wind speed and direction, wave height and period, sea current, cargo load, draft/trim, and main engine operating parameters. These factors exhibit strong nonlinearity, non-stationarity, and multivariate coupling characteristics, along with significant marine environmental noise. Traditional physics-based models often rely on empirical assumptions and handcrafted parameters, which may limit their adaptability and prediction accuracy under highly dynamic marine environments [4].

The accumulation of vast navigation data from the Automatic Identification System (AIS), and high-frequency onboard sensors has enabled data-driven approaches to become a research hotspot [5]. Machine learning and deep learning methods can directly mine hidden patterns from historical data without complex physical modeling, offering greater flexibility and generalization capability [6]. As shown in the overall technical framework (Fig. 1 Overall framework of the proposed method.), a typical data-driven ship speed prediction pipeline generally involves raw ship operation data collection, sensor monitoring, data preprocessing, feature selection, advanced modeling, and final performance evaluation.

Existing studies on ship speed prediction mainly fall into traditional machine learning methods, deep learning time-series models, and advanced Transformer-based architectures [7]. Traditional machine learning approaches, such as Random Forest(RF) [8], eXtreme Gradient Boosting(XGBoost) [9], Categorical Boosting (CatBoost) [10], and Light Gradient-Boosting Machine (LightGBM) [11], have demonstrated strong nonlinear fitting ability on ship observation data and are widely used for energy-efficient navigation and fuel consumption prediction [12]. However, these models show limitations in capturing long-range temporal dependencies and exhibit reduced robustness under drastic sea condition changes. Moreover, feature selection strategies play a pivotal role in these methods: Pearson correlation excels at identifying linear relationships [13], while mutual information (MI) effectively captures nonlinear dependencies. The differential suitability of these strategies across model types has emerged as an important research focus [14].

In recent years, Transformer variants—including In-




Fig. 1. Overall framework of the proposed method.

former [15], Autoformer [16], and iTransformer [17]—have achieved remarkable progress in long-sequence time-series forecasting by improving computational efficiency and multivariate modeling capability. Nevertheless, in the specific scenario of ship speed prediction characterized by strong noise, non-stationarity, and complex multivariate coupling, existing models still face key challenges, including insufficient extraction of multi-scale temporal features, inadequate suppression of marine environmental noise, and limited modeling of long-range inter-variate dependencies.

To address these limitations, this paper proposed a novel hybrid deep learning framework, TCN-FECAM-iTransformer, for one-step-ahead ship speed prediction, which tightly integrates adaptive feature selection with multi-scale temporal modeling, frequency-enhanced channel attention, and inverted multivariate attention. As illustrated in the technical framework (Fig. 1), the method first applies Pearson correlation and mutual information-based feature selection to optimize input quality, followed by a TCN [18] encoder for local temporal feature extraction, a FECAM [19] that leverages 1D Discrete Cosine Trans-

form (DCT) to suppress noise and strengthen channel-wise dependencies in the frequency domain, and an iTransformer to efficiently capture long-range interactions along the variate dimension. This design achieves effective local-to-global feature fusion for marine operational data characterized by high noise and strong coupling. Furthermore, extensive comparative experiments were conducted to investigate the impact of different feature selection strategies on the predictive model, providing additional insights into the role of feature engineering in ship speed prediction.

2. Materials and methods

2.1. Data

The dataset used in this study is derived from real voyage data of a liquefied petroleum gas (LPG) carrier, originating from onboard automation system records, noon reports, and associated meteorological and oceanographic data [20]. A sampling interval of 5 minutes is adopted for the monitoring data. Compared with higher-frequency sampling, this relatively low sampling frequency effectively filters short-term random noise while preserving key dynamic variations in ship operations, thereby capturing long-term trends and stable characteristics of the marine environment and vessel state. This provides more stable and reliable input for time-series models. Meanwhile, the larger sampling interval imposes greater challenges on capturing long-range dependencies and generalizing predictions. This setting enables a more realistic evaluation of model robustness under real operational conditions. The main technical parameters of the ship are presented in Table 1.

Table 1. Specifications of the ship

Parameter	Value	Units
Length	225	m
Breadth	37	m
Capacity	54,340	DWT
Maximum output power	12,400	kw
Service speed	16.8	kn
Main engines	MAN B&W 6G60ME-C9.2	—

The dataset contains 25 input feature variables, mainly covering main engine operating parameters, ship attitude, wind-wave-current environmental factors, and meteorological conditions, as detailed in Table 2. The target variable is ship speed over ground, in knots.

Data preprocessing only includes imputation of a very small number of missing values (using linear interpolation) to maximize preservation of the original data's true dis-

tribution characteristics. After preprocessing, the dataset contains 21,053 valid samples. To ensure stable training and generalization, the dataset is split chronologically in an 8:1:1 ratio based on the temporal order of voyage records: the first 80% (16,842 samples) for training, the next 10% (2,105 samples) for validation, and the last 10% (2,106 samples) for testing. This division ensures consistent distribution across subsets, with the validation set used for hyperparameter tuning and early stopping, and the test set for final performance evaluation. The authenticity, multivariate nature, and reasonable sample size of this dataset provide a reliable foundation for verifying model performance in real complex marine environments [21]. The target variable is ship speed over ground (Speed_Ground). To avoid target leakage, the model adopts a single-step ahead sliding window prediction scheme. At each prediction step, the input consists of the past 24 time steps of all selected features (including historical Speed_Ground), while the output is the Speed_Ground value at the next time step.

Note: The model uses only historical values of Speed_Ground with other input features to predict the future value.

2.2. Feature Selection Strategies

To investigate the impact of different feature selection strategies on model performance, this paper designs three input schemes: no feature selection (Raw data, using all 25 features), Pearson correlation coefficient-based selection, and mutual information-based selection.

The Pearson correlation coefficient measures the linear correlation between features and the target variable (speed over ground, SOG), calculated as:

$$r = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^n (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^n (y_i - \bar{y})^2}} \quad (1)$$

Mutual information quantifies the amount of shared information between an input feature X and the target variable Y , and is defined as:

$$MI(X; Y) = \sum p(x, y) \log \left(\frac{p(x, y)}{p(x)p(y)} \right) \quad (2)$$

where $p(x, y)$ denotes the joint probability density function, and $p(x)$ and $p(y)$ are the corresponding marginal probability density functions.

In this study, MI scores are estimated using the k-nearest neighbor estimator implemented in the `mutual_info_regression` function of Scikit-learn. To prevent data leakage, feature selection is performed exclusively on the training subset obtained by chronological splitting, and the selected feature subset is consistently applied to the validation and test sets.

Table 2. Description of the dataset.

No	Feature	Description	Units
1	ME_RPM	Main Engine (ME) revolutions per minute	min^{-1}
2	SFOC	Specific Fuel Oil Consumption	g/kWh
3	ME_Fuel_Total	Main Engine total Fuel Oil consumption	mT
4	Speed_Ground	Ship's speed over ground	knots
5	Heading	Ship's heading	$^{\circ}$
6	Wind_Speed	Wind speed from anemometer	knots
7	Fuel_Efficiency	Fuel efficiency index	—
8	Busbar_Load	Total load on busbars	kW
9	Air_Temp	Ambient air temperature	$^{\circ}\text{C}$
10	Sea_Temp	Sea water temperature	$^{\circ}\text{C}$
11	Trim	Trim in meters (+aft, -fore)	m
12	List	List in degrees	$^{\circ}$
13	COG	ECDIS Course Over Ground	deg
14	Wind_U	U-component of wind	m/s
15	Wind_V	V-component of wind	m/s
16	Wind_Gust	Wind gust speed	m/s
17	Wave_Height	Significant wave height	m
18	Wave_Direction	Wave direction	$^{\circ}$
19	Wave_Period	Wave period	s
20	Current_Direction	Sea current direction (Struja smjer)	deg
21	Current_Speed	Sea current speed (Struja brzina)	m/s
22	Douglas_Sea	Douglas sea scale (DSS, 0-4)	—
23	Douglas_Swell	Douglas swell scale (DSS, 0-4)	—
24	Beaufort_Scale	Beaufort wind scale (Bft, 0-12)	—
25	Wind_Direction	ECDIS wind direction	deg

Table 3. Input feature configurations

Strategy	Input Features (Historical Sequence: $\mathbf{X}_{t-23:t}$)
Raw Data	Historical Speed_Ground and all 24 operational/environmental variables (25 variables in total).
Pearson	Historical Speed_Ground, ME_RPM, Fuel_Efficiency, ME_Fuel_Total, Douglas_Sea, Beaufort_Scale, Douglas_Swell, Wind_Speed, Heading, Wave_Height, Wind_Gust, Wave_Period, and COG.
MI	Historical Speed_Ground, Wave_Period, Wind_U, Wave_Direction, Wave_Height, Wind_V, Wind_Gust, Current_Speed, Current_Direction, ME_RPM, Heading, COG, and ME_Fuel_Total.
Output	Speed_Ground_{t+1}

All features are ranked according to their Pearson correlation coefficients (absolute values) or MI scores, and the Top- K most informative features are selected. A sensitivity analysis (See Section 3.2) with $K = \{6, 12, 18\}$ is further conducted to determine an appropriate feature subset size. The Pearson correlation coefficients and MI scores of all candidate features are illustrated in Figs. 2 and 3. Both metrics are calculated exclusively on the training subset obtained by chronological splitting.

For all feature selection, the forecasting task is formulated as a one-step-ahead prediction problem, where the historical observations within the previous 24 time steps are used to predict the ship speed at the next time step ($t + 1$). The corresponding input configurations are sum-

marized in Table 3.

2.3. TCN-FECAM-iTransformer Model

Building upon the feature selection stage described in the previous section, the subsequent modeling stage aims to learn informative temporal patterns from the selected variables and to capture the underlying relationships among them for accurate ship speed prediction. To achieve this, this study proposes the TCN-FECAM-iTransformer model, a hybrid architecture that integrates temporal convolution, frequency-enhanced channel attention, and inverted Transformer-based global dependency modeling. The detailed structure of the proposed model is presented in the following.




Fig. 2. Pearson correlation




Fig. 3. Mutual information values between features and speed (five-round average)

As shown in Fig. 4, the model consists of four main components that operate in a sequential and synergistic manner:

(1) TCN Encoder: The raw multivariate input [B, C, T] is first processed by the TCN Encoder. This module comprises multiple stacked residual blocks, each containing grouped dilated causal convolutions with exponentially increasing dilation rates, followed by weight normalization, ReLU activation, dropout, and a 1×1 convolution for residual connection. By stacking residual blocks with different dilation factors, the encoder achieves a large receptive field while maintaining causality, effectively extracting multi-scale local temporal dependencies from the ship operation time series. Independent convolution is performed along the variable dimension to avoid early inter-variate mixing. The output of the TCN Encoder retains the original shape [B, C, T] and is subsequently passed to the FECAM module.

(2) FECAM Module: To suppress marine environmental noise and enhance inter-variable dependencies, the TCN-encoded features [B, C, T] is fed into the FECAM. As depicted in the architecture, the module first applies a 1D DCT to transform the time-domain signals into the frequency domain. Frequency-domain attention is then modeled through a fully connected neural network. The enhanced frequency representations are fused with the original features via a scaling operation and residual addition. This process jointly strengthens channel-wise dependencies and important semantic patterns while effectively attenuating noise. After FECAM processing, the refined features with dimension [B, C, d_model] are delivered to the iTransformer module.

(3) iTransformer Encoder: The FECAM-refined features are embedded into tokens and then processed by the iTransformer. In this stage, each variate is treated as a token, and an inverted multi-head self-attention mechanism is applied along the variate dimension. Specifically, Query, Key, and Value projections are computed from the embedded tokens to construct a multivariate correlations map through matrix multiplication (MatMul). This inverted attention efficiently captures long-range and complex interactions among different variables (such as engine parameters, wind, waves, and trim). Each iTransformer layer further includes a position-wise feed-forward network, layer normalization, and residual connections. Multiple such layers are stacked to progressively model global dependencies. Finally, a linear projector maps the output of the last layer to the predicted ship speed value.

By organically combining the TCN Encoder's multi-




Fig. 4. TCN-FECAM-iTransformer Model Architecture

scale local temporal extraction, the FECAM module's frequency-domain noise suppression and channel refinement, and the iTransformer's powerful global multivariate modeling capability, the proposed framework achieves effective local-to-global feature fusion, making it particularly suitable for noisy, non-stationary, and strongly coupled ship speed prediction tasks. Model parameters and training hyperparameters are shown in Table 4.

2.4. Baseline Models

To comprehensively evaluate the proposed TCN-FECAM-iTransformer framework, nine mainstream baseline models are selected, covering both traditional machine learning and deep learning time-series approaches. All baselines are implemented using identical data preprocessing, feature selection strategies, and training-validation-test splits to ensure a fair and reproducible comparison.

Traditional machine learning models include Random RF, XGBoost, CatBoost, and LightGBM. These tree-based ensemble methods are known for their strong nonlinear fitting capability, feature importance interpretability, train-

ing efficiency, and robustness to noise. Previous studies have demonstrated their effectiveness in ship speed prediction. For instance, Bassam et al. systematically compared various regression algorithms on real high-frequency ship operational data, highlighting the efficiency and robustness of tree ensemble models. Gan et al. evaluated multiple models on shipboard observation data from a similar LPG carrier and reported that LightGBM achieved the highest accuracy. CatBoost is additionally included as a competitive baseline due to its superior handling of mixed-type features and built-in mechanisms to prevent overfitting.

Despite these advantages, traditional machine learning models have inherent limitations in capturing long-range temporal dependencies and modeling complex dynamic changes under varying sea conditions.

Deep learning time-series models include Long Short-Term Memory (LSTM) [22], TCN, Informer, Autoformer, and iTransformer. LSTM, a classic recurrent architecture, captures short-term dependencies through gating mechanisms and has been widely applied in early ship trajectory and speed prediction tasks. TCN employs causal dilated convolutions to achieve large receptive fields with parallel computation and stable gradients. Informer introduces ProbSparse self-attention to improve computational efficiency for long sequences [23]. Autoformer enhances trend and seasonality modeling via decomposition blocks and auto-correlation mechanisms. iTransformer innovatively inverts the variate dimension as the attention sequence, enabling efficient modeling of long-range multivariate interactions.

3. Results and discussion

3.1. Comparative Experiment

To comprehensively evaluate model prediction performance, this paper selected four commonly used regression evaluation metrics: Root Mean Square Error (RMSE), Coefficient of Determination (R^2), Mean Square Error (MSE), and Mean Absolute Error (MAE). These metrics quantify model accuracy from aspects of error magnitude, fitting goodness, and robustness widely used in ship speed prediction research. The calculation formulas for each metric are as follows:

Root Mean Square Error (RMSE):

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2} \quad (3)$$

Coefficient of Determination (R^2):

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (4)$$

Table 4. Parameters and hyperparameters

Parameter	Value	Description
Input Sequence Length	24	Historical input time steps
Prediction Length	1	Future steps to predict
Embedding Dimension	512	Dimension of token embeddings
Number of Attention Heads	8	Multi-head attention heads in iTransformer
TCN Kernel Sizes	2, 3, 5, 3	Kernel sizes in MultiScaleTCN branches
TCN Dilations	1, 2, 4, 8	Dilation rates for multi-scale receptive field
iTransformer Layers	3	Number of encoder layers
Feed-Forward Dimension	1024	Dimension of FFN in transformer blocks
Batch Size	128	Training batch size
Learning Rate	1.00E – 04	Initial learning rate
Optimizer	AdamW	With weight decay = 1e – 5
Loss Function	MAE (L1 Loss)	Used for both training and validation
Training Epochs	200	Fixed 200 epochs
Random Seed	42	For reproducibility

Mean Square Error (MSE):

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2 \quad (5)$$

Mean Absolute Error (MAE):


$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (6)$$

where n is the number of samples, y_i is the true speed of the i -th sample, $\hat{y}_i$ is the corresponding predicted value, and $\bar{y}_i$ is the mean of the true values. This paper uses RMSE as the primary metric, as it directly reflects the magnitude of prediction error (in knots) and is sensitive to large errors; R^2 , MAE, and MSE serve as auxiliary metrics to assess overall fitting goodness, robustness to outliers, and squared error distribution, respectively.

To evaluate the effectiveness of the proposed TCN-FECAM-iTransformer framework, comprehensive comparative experiments were conducted under three feature selection strategies: Raw data, Pearson correlation, and Mutual Information. All models, including nine base-lines, were implemented with identical data preprocessing, training-validation-test splits, and hyperparameter tuning procedures. The quantitative results are summarized in Table 5.

Red circles represent Raw data, blue squares represent Pearson correlation selection, and orange triangles represent Mutual Information (MI) selection. The proposed TCN-FECAMiTransformer (Ours) achieves the best performance across all metrics and strategies.

As shown in Table 5 and Fig. 5, feature selection strategies substantially influence prediction performance, although their effects vary across different model architectures. Overall, the Pearson correlation-based strategy consistently improves the predictive accuracy of most models.

**Fig. 5.** Model performance comparison under different feature selection strategies

Compared with the raw-data setting, the RMSE of Random Forest decreases from 0.4139 kn to 0.3365 kn, XGBoost improves from 0.3168 kn to 0.2520 kn, and CatBoost decreases from 0.5460 kn to 0.3076 kn. Similarly, LightGBM achieves its best performance under Pearson selection with an RMSE of 0.2288 kn and an R^2 of 0.9598. These results suggest that linear correlation filtering effectively removes redundant variables and enhances the robustness of conventional machine learning models.

For deep learning models, the influence of feature selection is more architecture-dependent. Pearson selection noticeably improves the performance of LSTM, TCN, Informer, and Autoformer, while the iTransformer exhibits only marginal differences among the three input configu-


Table 5. Model performance under different feature selection strategies

Strategy	Model	RMSE	R ²	MAE	MSE
Raw data	RF	0.4139	0.8684	0.2086	0.1713
	XGBoost	0.3168	0.9229	0.1403	0.1003
	Catboost	0.5460	0.7710	0.2451	0.2981
	LightGBM	0.2798	0.9398	0.1594	0.0783
	LSTM	0.6817	0.6430	0.4056	0.4647
	TCN	0.4974	0.8099	0.2622	0.2474
	Informer	0.3705	0.8945	0.2559	0.1373
	Autoformer	0.5161	0.7953	0.3743	0.2664
	iTransformer	0.2352	0.9575	0.1047	0.0553
	Ours	0.2177	0.9636	0.0766	0.0474
Pearson	RF	0.3365	0.9130	0.1595	0.1132
	XGBoost	0.2520	0.9512	0.0949	0.0635
	Catboost	0.3076	0.9273	0.1364	0.0946
	LightGBM	0.2288	0.9598	0.1052	0.0523
	LSTM	0.5604	0.7587	0.3634	0.3141
	TCN	0.3301	0.9163	0.1559	0.1090
	Informer	0.2877	0.9364	0.1455	0.0828
	Autoformer	0.4386	0.8522	0.2804	0.1923
	iTransformer	0.2258	0.9608	0.0869	0.0510
	Ours	0.1952	0.9707	0.0749	0.0381
MI	RF	0.3801	0.8890	0.1811	0.1445
	XGBoost	0.2765	0.9413	0.1205	0.0765
	Catboost	0.3680	0.8959	0.1611	0.1354
	LightGBM	0.2722	0.9431	0.1121	0.0741
	LSTM	0.4988	0.8088	0.2785	0.2488
	TCN	0.3563	0.9025	0.1786	0.1270
	Informer	0.3129	0.9248	0.1322	0.0979
	Autoformer	0.6308	0.6943	0.4544	0.3979
	iTransformer	0.2329	0.9583	0.1020	0.0542
	Ours	0.2064	0.9673	0.0770	0.0426

rations. Specifically, the RMSE of iTransformer varies only from 0.2352 kn (Raw) to 0.2258 kn (Pearson) and 0.2329 kn (MI), indicating that its inverted attention mechanism possesses strong capability for handling redundant multi-variate inputs. In contrast, the performance degradation of several models under the MI strategy implies that preserving nonlinear dependencies alone may not always provide the optimal feature subset for ship speed prediction.

The proposed TCN-FECAM-iTransformer framework consistently achieves the best performance under all three feature selection strategies. Under the Pearson setting, the proposed model attains the lowest RMSE of 0.1952 kn, the lowest MAE of 0.0749 kn, and the highest R² of 0.9707. Compared with the strongest baseline, iTransformer under the Pearson strategy, the proposed method reduces the RMSE by approximately 13.5%. It also outperforms the best traditional machine learning model, LightGBM, by about 14.7% in terms of RMSE. Even without feature selection, the proposed framework maintains superior performance, achieving an RMSE of 0.2177 kn, which further demonstrates its robustness against redundant input information.

The superior performance of the proposed model can be attributed to the complementary integration of its three

**Fig. 6.** Prediction results of each model under its best-performing feature selection strategy

major components. The TCN encoder effectively captures local multiscale temporal patterns, the FECAM module suppresses marine environmental noise while enhancing informative channel representations in the frequency domain, and the iTransformer efficiently models long-range inter-variable dependencies through the inverted attention mechanism. Their combination enables effective local-to-global feature extraction for highly coupled and noisy ship operational data.

Overall, the comparative experiments demonstrate that

appropriate feature selection can further improve prediction accuracy, with the Pearson strategy providing the most balanced and effective feature subset for the investigated dataset. Meanwhile, the proposed TCN-FECAM-iTransformer framework consistently delivers the best overall performance, providing a robust and accurate solution for practical ship speed prediction and intelligent energy-efficiency management.

3.2. Sensitivity analysis experiments

Table 6. Sensitivity analysis experiment results

Model	RMSE	R2	MAE	MSE
Ours (Raw data)	0.2177	0.9636	0.0766	0.0474
Ours_Pearson (Top6)	0.2084	0.9666	0.0756	0.0434
Ours_Pearson (Top12)	0.1952	0.9707	0.0749	0.0381
Ours_Pearson (Top18)	0.2021	0.9686	0.0782	0.0409
Ours_MI(Top6)	0.2240	0.9615	0.0759	0.0502
Ours_MI(Top12)	0.2064	0.9673	0.0770	0.0426
Ours_MI(Top18)	0.2134	0.9650	0.0776	0.0456

The results indicate that Top-12 provides the best trade-off between information preservation and redundancy reduction. When only six features are retained, important environmental information may be discarded, resulting in degraded performance. Conversely, increasing the number of selected features to eighteen introduces redundant variables and slightly weakens the prediction accuracy. Therefore, Top-12 was selected as the default feature subset size for experiments, as it provides a balanced trade-off between information preservation and redundancy reduction for both Pearson correlation and mutual information based feature selection.

3.3. Ablation Experiments

To further verify the contribution and necessity of each component in the proposed TCN-FECAM-iTransformer framework, ablation studies were conducted under the Pearson feature selection strategy, which was identified as optimal in the comparative experiments.

Three ablated variants were constructed: (1) removing the FECAM module (w/o FECAM), (2) removing the TCN Encoder (w/o TCN), and (3) removing both TCN and FECAM modules simultaneously (w/o TCN&FECAM), which degenerates to the baseline iTransformer. The results are presented in Table 7.

As shown in Table 7, removing either the TCN encoder or the FECAM module leads to a noticeable decrease in

Table 7. Ablation experiment results

	RMSE	R2	MAE	MSE
Ours	0.1952	0.9707	0.0749	0.0381
w/o FECAM	0.2234	0.9617	0.0957	0.0499
w/o TCN	0.2154	0.9643	0.0870	0.0464
w/o TCN & FECAM	0.2258	0.9608	0.0869	0.0510

prediction accuracy, indicating that both components contribute positively to the overall framework. Specifically, removing FECAM increases the RMSE from 0.1952 to 0.2234, suggesting that frequency-domain channel enhancement effectively suppresses environmental noise and improves feature representation. Similarly, removing the TCN encoder degrades the RMSE to 0.2154, demonstrating the importance of multi-scale local temporal modeling. When both modules are removed, the architecture degenerates into the baseline iTransformer, whose RMSE further increases to 0.2258. These results confirm that the proposed TCN-FECAM-iTransformer benefits from the complementary effects of local temporal extraction, frequency-aware channel attention, and global multivariate dependency modeling.

4. Conclusions

This study proposes TCN-FECAM-iTransformer, a hybrid deep learning framework for accurately predicting ship speeds in real-world marine environments. The main contributions of this study can be summarized as follows:

- (1) A hybrid TCN-FECAM-iTransformer framework is proposed for ship speed prediction, effectively integrating local temporal modeling, frequency-domain channel attention, and long-range multivariate dependency learning.
- (2) A systematic comparison between Pearson correlation and mutual information based feature selection strategies is conducted. Experimental results indicate that feature selection has a significant impact on prediction accuracy; among the methods evaluated, the Top-K selection method based on the Pearson correlation coefficient demonstrated superior overall performance on the datasets studied.
- (3) Extensive experiments based on real voyage data demonstrate that the proposed framework achieves the best prediction accuracy among several representative machine learning and deep learning baselines.

In summary, this work provides an effective and high-precision data-driven solution for ship speed-over-ground prediction, offering practical value for ship energy efficiency, route optimization, and low-carbon maritime operations.

Data availability

Data will be made available on request.

Acknowledgements

The authors acknowledge funding support from High-technology Ship Research Program (CBG3N21-3-3).

References

- [1] I. C. on Clean Transportation. *New IMO study highlights sharp rise in short-lived climate pollution*. 2020. URL: <https://theicct.org/new-imo-study-highlights-sharp-rise-in-short-lived-climate-pollution/>.
- [2] I. M. Organization. *2023 IMO Strategy on Reduction of GHG Emissions from Ships*. 2023. URL: <https://www.imo.org/en/ourwork/environment/pages/2023-imo-strategy-on-reduction-of-ghg-emissions-from-ships.aspx>.
- [3] S. Nguyen, M. Gadel, K. Wang, J. Li, X. Zhang, S.-C. Kong, X. Fu, and Z. Qin, (2025) "Maritime decarbonization through machine learning: A critical systematic review of fuel and power prediction models" **Cleaner Logistics and Supply Chain** 14: 100210. DOI: 10.1016/j.clscn.2025.100210. URL: <https://doi.org/10.1016/j.clscn.2025.100210>.
- [4] Y. Kim, P. Gupta, and S. Steen, (2025) "A comprehensive review of data processing for ship performance analysis" **Applied Ocean Research** 162: 104737. DOI: 10.1016/j.apor.2025.104737. URL: <https://doi.org/10.1016/j.apor.2025.104737>.
- [5] S. Wang, Y. Li, and H. Xing, (2023) "A novel method for ship trajectory prediction in complex scenarios based on spatio-temporal features extraction of AIS data" **Ocean Engineering** 281: 114846. DOI: 10.1016/j.oceaneng.2023.114846. URL: <https://doi.org/10.1016/j.oceaneng.2023.114846>.
- [6] W. Gan, D. Ma, and Y. Duan, (2025) "Comparisons of Machine Learning Methods in Ship Speed Prediction Based on Shipboard Observation" **Journal of Marine Science and Engineering** 13(6): 1011. DOI: 10.3390/jmse13061011. URL: <https://doi.org/10.3390/jmse13061011>.
- [7] A. M. Bassam, A. B. Phillips, S. R. Turnock, and P. A. Wilson, (2023) "Artificial neural network based prediction of ship speed under operating conditions for operational optimization" **Ocean Engineering** 278: 114613. DOI: 10.1016/j.oceaneng.2023.114613. URL: <https://doi.org/10.1016/j.oceaneng.2023.114613>.
- [8] L. Breiman, (2001) "Random Forests" **Machine Learning** 45(1): 5–32. DOI: 10.1023/A:1010933404324. URL: <https://doi.org/10.1023/A:1010933404324>.
- [9] T. Chen and C. Guestrin. "XGBoost: A Scalable Tree Boosting System". In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2016, 785–794. DOI: 10.1145/2939672.2939785. URL: <https://doi.org/10.1145/2939672.2939785>.
- [10] L. Prokhorenkova et al. "CatBoost: unbiased boosting with categorical features". In: *Advances in Neural Information Processing Systems*. 2018. URL: https://papers.neurips.cc/paper_files/paper/2018/hash/14491b756b3a51daac41c24863285549-Abstract.html.
- [11] G. Ke et al. "LightGBM: A Highly Efficient Gradient Boosting Decision Tree". In: *Advances in Neural Information Processing Systems*. 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/hash/6449f44a102fde848669bdd9eb6b76fa-Abstract.html.
- [12] A. Fan, S. Sun, Z. Hu, N. Vladimir, and W. Mao, (2026) "Enhancing transfer learning strategies for ship fuel consumption prediction under data scarcity" **Ocean Engineering** 351: 124398. DOI: 10.1016/j.oceaneng.2026.124398. URL: <https://doi.org/10.1016/j.oceaneng.2026.124398>.
- [13] X. Jiang et al., (2025) "Research on ship speed prediction based on time series imaging and deep convolutional network fusion method" **Applied Ocean Research** 154: 104384. DOI: 10.1016/j.apor.2024.104384. URL: <https://doi.org/10.1016/j.apor.2024.104384>.
- [14] X. Zhang et al., (2025) "Ship manoeuvring model identification using mutual information clustering with stepwise regression structural optimisation under small-sample constraints" **Ocean Engineering** 341: 122594. DOI: 10.1016/j.oceaneng.2025.122594. URL: <https://doi.org/10.1016/j.oceaneng.2025.122594>.
- [15] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, (2021) "Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting" **Proceedings of AAAI** 35(12): 11106–11115. DOI: 10.1609/aaai.v35i12.17325. URL: <https://doi.org/10.1609/aaai.v35i12.17325>.
- [16] H. Wu et al. "Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting". In: *Advances in Neural Information Processing Systems*. 2021. URL: <https://proceedings.neurips.cc/paper/2021/hash/>

bcc0d400288793e8bdcd7c19a8ac0c2b - Abstract .
html.

- [17] Y. Liu, T. Hu, H. Zhang, H. Wu, S. Wang, L. Ma, and M. Long. "iTransformer: Inverted Transformers Are Effective for Time Series Forecasting". In: *International Conference on Learning Representations*. 2024. DOI: [10.48550/arXiv.2310.06625](https://doi.org/10.48550/arXiv.2310.06625). URL: <http://arxiv.org/abs/2310.06625>.
- [18] S. Bai, J. Z. Kolter, and V. Koltun. *An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling*. 2018. DOI: [10.48550/arXiv.1803.01271](https://doi.org/10.48550/arXiv.1803.01271). URL: <http://arxiv.org/abs/1803.01271>.
- [19] M. Jiang, P. Zeng, K. Wang, H. Liu, W. Chen, and H. Liu, (2023) "FECAM: Frequency enhanced channel attention mechanism for time series forecasting" **Advanced Engineering Informatics** 58: 102158. DOI: [10.1016/j.aei.2023.102158](https://doi.org/10.1016/j.aei.2023.102158). URL: <https://doi.org/10.1016/j.aei.2023.102158>.
- [20] A. Vorkapić, R. Radonja, and S. Martinčić-Ipšić, (2021) "Predicting Seagoing Ship Energy Efficiency from the Operational Data" **Sensors** 21(8): 2832. DOI: [10.3390/s21082832](https://doi.org/10.3390/s21082832). URL: <https://doi.org/10.3390/s21082832>.
- [21] P. Karagiannidis and N. Themelis, (2021) "Data-driven modelling of ship propulsion and the effect of data pre-processing on the prediction of ship fuel consumption and speed loss" **Ocean Engineering** 222: 108616. DOI: [10.1016/j.oceaneng.2021.108616](https://doi.org/10.1016/j.oceaneng.2021.108616). URL: <https://doi.org/10.1016/j.oceaneng.2021.108616>.
- [22] S. Hochreiter and J. Schmidhuber, (1997) "Long Short-Term Memory" **Neural Computation** 9(8): 1735–1780. DOI: [10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735). URL: <https://doi.org/10.1162/neco.1997.9.8.1735>.
- [23] M.-W. Li et al., (2025) "A self-optimised informer-based model for nonlinear ship roll motion forecasting" **Transportmetrica A: Transport Science**: DOI: [10.1080/23249935.2025.2606152](https://doi.org/10.1080/23249935.2025.2606152). URL: <https://doi.org/10.1080/23249935.2025.2606152>.