A Hybrid DMD–TCN Framework for Interpretable Short-Horizon Prediction of 6-DOF Ship Motions

Enock Tafadzwa Chekure, Kumeshan Reddy, John Fernandes
Electrical Engineering Department-School of Engineering, Nelson Mandela University, Gqeberha, South Africa

***Abstract*—Accurate short-horizon prediction of six degrees of freedom (6-DOF) vessel motions is essential for autonomous navigation, motion compensation, and operational decision making. Traditional seakeeping models rely on hydrodynamic coefficients that are seldom available for full-scale vessels, while purely data-driven approaches may struggle to maintain physical consistency. This study introduces a hybrid physics–machine learning framework that combines *Dynamic Mode Decomposition* (DMD), which approximates the vessel's dominant linear drift dynamics, with a causal *Temporal Convolutional Network* (TCN) that learns nonlinear residual corrections from a 12-hour historical window of environmental, geometric, and motion features. DMD provides an interpretable surrogate of the vessel dynamics through its eigenvalues, growth rates, and mode shapes, serving as a data-derived linear transfer operator. The TCN predicts only the residual departure from this structured baseline, ensuring a stable and causal forecasting architecture. Evaluation on full-scale field data shows that the hybrid model improves prediction accuracy for heave and achieves performance comparable to DMD for surge, while underperforming in sway, roll, pitch, and yaw due to the limited observability of key physical drivers at hourly resolution. These results highlight both the strengths and limitations of residual learning when important nonlinear forcing mechanisms and control inputs are unmeasured. Overall, the study demonstrates that hybrid physics–machine learning approaches provide valuable interpretability and diagnostic insight, even when data limitations are constrained. The framework offers a principled foundation for incorporating additional physical inputs, higher-frequency measurements, and physics-informed architectures in future work on operational ship-motion forecasting.**

***Keywords*—Dynamic Mode Decomposition; Temporal Convolutional Network; hybrid learning; seakeeping and vessel response; data-driven modelling; Koopman operator methods**

I. Introduction

Short-horizon prediction of wave-induced ship motions is essential for autonomous navigation, motion compensation, digital twin systems, and operational decision making. Conventional seakeeping approaches rely on model-scale experiments or potential-flow solvers that provide Response Amplitude Operators (RAOs) and hydrodynamic coefficients. Although effective in controlled or design-stage environments, these models often underperform at full scale, where nonlinear viscous effects, slow-drift forces, mooring interactions, and environmental asymmetries dominate the vessel's response.

Moreover, the hydrodynamic coefficients required for classical seakeeping prediction are rarely available for in-service ships, as they depend on vessel-specific geometry, loading, and environmental conditions and are typically obtained only through model-scale testing or specialised computations [1], [2], [3], [4], [5], [6].

Purely data-driven neural networks offer an alternative by learning motion patterns directly from operational measurements [7], [8], [9], [10]. However, such models generally lack physical structure, making them prone to overfitting and poor extrapolation outside the training regime [11], [12], [13]. When deployed autoregressively, they may also yield dynamically inconsistent or unstable predictions [14], [15], [16], [17]. These limitations have motivated the development of hybrid frameworks that combine physics-based constraints with the representational flexibility of machine learning [5], [18], [19].

Dynamic Mode Decomposition (DMD) has emerged as a powerful data-driven technique for extracting dominant linear dynamics from time-series data [19], [20]. Through its connection to the Koopman operator, DMD yields eigenvalues, growth or decay rates, frequencies, and mode shapes that provide an interpretable spectral representation of the underlying system. Applications in marine hydrodynamics [21], [22] demonstrate that DMD can capture low-frequency drift behaviour and cross-degree-of-freedom (DOF) coupling even when hydrodynamic coefficients are unavailable, making it a suitable candidate for operational ship-motion modelling.

To account for nonlinear behaviour not represented by DMD, this work combines the linear operator with a causal, dilated TCN. TCNs leverage stacked dilated convolutions to learn structured dependencies across long temporal windows while preserving strict causality [23], [24]. Such hybrid physics–machine learning models have shown improved interpretability and robustness in marine applications [5], [18]. In the proposed residual formulation, the DMD operator provides the one-step linear prediction, and the TCN estimates the nonlinear residual using a 12-hour causal window of environmental, geometric, and motion features.

We evaluate the hybrid framework using full-scale six-DOF operational data. The results indicate that the hybrid model improves prediction accuracy for heave and performs comparably to the DMD baseline for surge. In contrast, underperformance in sway, roll, pitch, and yaw reflects the limited observability of key physical drivers at hourly resolution and the inherently nonlinear nature of these DOFs. These findings highlight both the strengths and the limitations of physics-informed residual learning when applied to operational ship-motion statistics and motivate further investigation into feature augmentation and higher-frequency sensing.

This study makes the following contributions:

- A hybrid DMD–TCN residual modelling framework for ship motions. The proposed method uses *Dynamic Mode Decomposition* as a physics-informed linear baseline that captures dominant drift dynamics and cross-DOF coupling. A causal TCN predicts only the nonlinear residual, providing an interpretable and structurally constrained forecasting architecture.

- A full-scale, data-driven modal analysis of 6-DOF dynamics. We extract DMD eigenvalues, continuous-time growth rates, and mode shapes from operational field data, offering insight into vessel- specific dynamic behaviour without requiring hydrodynamic coefficients.
- A causal Temporal Convolutional Network for residual motion forecasting. The TCN leverages dilated convolutions to incorporate long-range historical context while preserving strict causality, enabling the identification of slow environmental trends and nonlinear departures from the DMD baseline.
- A rigorous evaluation on full-scale field measurements. The results demonstrate that while the hybrid model improves prediction in heave and matches DMD performance in surge, it underperforms in lateral and rotational DOFs due to unobserved physics and hourly sampling limitations. This provides valuable diagnostic insight into which motion components contain predictable nonlinear structure under limited observability.
- A reproducible data-processing pipeline for operational ship-motion time-series. The pipeline includes time alignment, outlier handling, causal feature engineering, and 6-DOF synchronisation designed for real-time forecasting scenarios.

Together, these contributions demonstrate that combining data-driven modal analysis with structured residual learning yields an interpretable and physically grounded approach to operational ship-motion prediction, while also revealing the limitations imposed by sparse sensing and low-frequency sampling.

II. Related Work

A. Data-Driven Vessel Motion Prediction

The increasing availability of onboard measurements, AIS records, and operational datasets has stimulated significant interest in data-driven prediction of ship motions. Early work applied classical machine learning methods such as Support Vector Regression (SVR) to estimate local ship responses from environmental and operational parameters [7]. Subsequent studies explored neural-network architectures trained on high-fidelity hydrodynamic simulations. Diez *et al.* [8], for example, demonstrated that feed-forward and recurrent architectures can reproduce short-horizon motions when supplied with sufficiently rich input–output data derived from CFD-based wave–ship interaction models.

A driving motivation for these data-driven approaches is the substantial computational cost associated with classical hydrodynamic solvers. Resolving nonlinear wave–body interactions, estimating statistical extremes, or accounting for hydro-structural coupling often requires long-duration simulations that are impractical for real-time use. Machine learning models offer an attractive surrogate: once trained, they provide fast predictions at negligible computational cost and can integrate seamlessly into decision support pipelines [22].

However, purely data-driven models face fundamental limitations. Standard feed-forward neural networks treat each sample independently and therefore fail to capture the temporal dependencies inherent in vessel dynamics. This motivated the adoption of recurrent architectures such as RNNs, LSTMs, and GRUs [25], which propagate information through hidden states and have been shown to improve short-term forecasting in marine applications [9], [16]. Although effective within the training distribution, such models often extrapolate poorly and may generate dynamically inconsistent predictions when applied autoregressively [11], [14], [15].

Complementing recurrent models, physics-from-data approaches extract dynamic structure directly from measurements. DMD provides an equation-free modal representation of the system by identifying eigenvalues, growth/decay rates, and coherent spatio-temporal structures from time-series snapshots. DMD's connection to the Koopman operator [19], [20] enables interpretable linear surrogates of nonlinear dynamics at low computational cost. In marine hydrodynamics, DMD has been used to characterise coherent flow structures and predict vessel trajectories and manoeuvring dynamics [21], [22].

Deep neural sequence models have also been explored for full 6-DOF prediction. Silva and Maki [9] trained LSTM networks to replicate the 6-DOF response of a free-running destroyer in nonlinear seas by reconstructing encounter-frame wave fields using virtual sensors. Their results highlight the potential of data-driven sequence modelling but also reveal its dependence on careful pre-processing and broad training coverage. Additional work combines signal decomposition with learning: Ye *et al.* [26] used EMD–LSTM–SVR hybrids to isolate nonlinear roll and sway behaviour, while Sun *et al.* [27] constructed CFD-informed reduced-order models for damaged ship motions. Comparative analysis further reveal necessary trade-offs between interpretability, accuracy, and robustness across modelling approaches [28].

Across the literature, three recurring limitations of purely data-driven approaches emerge: 1) sensitivity to training coverage, 2) limited extrapolation capability, and 3) difficulty enforcing physically consistent dynamics. These limitations motivate hybrid modelling strategies that embed physical structure while leveraging the flexibility of machine learning.

B. Hybrid Modelling of Vessel Dynamics

Classical models of ship dynamics derive from first-principles hydrodynamics, incorporating added mass, radiation damping, viscous forces, hydrostatic restoring, and second-order wave loads [1]. While such models provide interpretable relationships between wave excitation and vessel response, they require detailed hydrodynamic coefficients that are rarely available for in-service ships [3], [4]. Moreover, linear-response theories such as RAO-based seakeeping can struggle under strongly nonlinear or operationally variable conditions.

Hybrid physics–machine learning approaches address these limitations by incorporating physical constraints or simplified mechanistic structure into data-driven frameworks. Physics-Informed Neural Networks enforce governing equations in the training process [12], while grey-box models learn uncertain dynamics on top of known physical components. In maritime applications, Schirmann *et al.* [10] demonstrated that augmenting data-driven predictors with fast physics-based estimates improves the prediction of heave, pitch, and roll. Kanazawa *et al.* [5] developed physics–data cooperative models that fuse onboard wave radar with linear seakeeping predictions to enhance short-term motion forecasting. Other studies integrate physics-guided decomposition (e.g., EMD) to improve neural network performance for nonlinear motions such as roll [26]. For manoeuvring, data-driven corrections to hydrodynamic derivatives have been used to enhance classical models [29].

The present work adopts a hybrid philosophy tailored to operational datasets with low temporal resolution. Instead of modelling full hydrodynamic governing equations, we use DMD as a linear, data-driven surrogate aligned with hourly sampling, which captures low-frequency drift behaviour and cross-DOF coupling but not nonlinear wave-induced dynamics. A causal, dilated TCN is then trained to predict only the residual deviation from the DMD baseline using a 12-hour window of motion, environmental, and ship-specific features. TCNs offer stable long-range sequence modelling [23], [24] and mitigate the instability commonly observed in autoregressive neural predictors. The resulting DMD–TCN hybrid provides a structured, interpretable additive decomposition, with DMD representing the dominant low-frequency dynamics and TCN capturing learnable nonlinear corrections.

This progression reflects a broader evolution in maritime modelling: from static regressors (SVR, feed-forward NNs) to temporal sequence models (RNNs, LSTMs) to physics-from-data approaches (DMD), and finally to hybrid architectures that integrate physical structure with machine learning flexibility. Fig. 1 summarises this conceptual trajectory and situates the present model within this continuum.

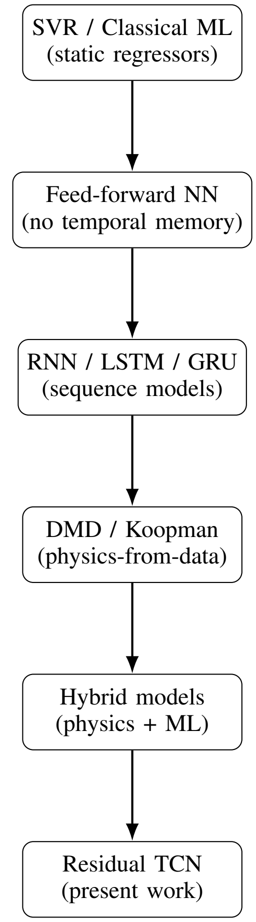

Fig. 1. Conceptual evolution of data-driven and physics-aware modelling approaches for vessel-motion prediction, leading to the proposed residual Temporal Convolutional Network (TCN).

III. Methodology

This section presents the complete processing workflow used to construct the proposed hybrid DMD–TCN ship-motion predictor. The methodology integrates multiple stages, including timestamp normalisation, multi-stream harmonisation, feature engineering, causal window extraction, linear data-driven modelling via DMD, and non-linear residual learning using a TCN. To provide clarity and continuity, a high-level overview of the whole processing pipeline is presented first. Fig. 2 summarises the end-to-end pipeline, and the subsequent subsections describe each component in detail.

Fig. 2 illustrates the sequence of operations applied to raw sensor data before prediction. Independent DOF-specific time series are normalised, resampled, and synchronised onto a standard hourly grid. Canonical environmental and ship parameters are constructed, engineered features are added, and causal 12-hour windows are extracted. Each window is passed to two parallel predictors: 1) a DMD model that produces a linear data-driven one-step-ahead forecast, and 2) a causal TCN that estimates the non-linear residual dynamics. The final hybrid prediction is obtained by summing the DMD baseline with the learned residual.

IV. Temporal Convolutional Networks (TCNs)

TCNs were introduced as a causal alternative to RNNs for sequence modelling, providing improved stability, parallelism and long-range receptive fields through dilated convolutions. Bai *et al.* [23] demonstrated that TCNs outperform LSTMs on numerous benchmarks while offering predictable gradient behaviour and easier optimisation. Lea et al. [24] formalised their application to temporal segmentation using stacked dilated convolutional blocks with residual connections.

For forecasting problems, TCNs offer distinct advantages: 1) they are inherently causal, 2) they model temporal dependencies with exponentially growing receptive fields, and 3) they avoid the vanishing-gradient issues common in RNNs. These properties make them suitable for marine motion prediction, where dependencies on sea-state evolution and hull parameters may extend over many hours.

A. Dataset Description

The study uses a publicly available full-scale dataset containing six-degree-of-freedom (6-DOF) moored ship motions together with concurrent environmental and vessel parameters [30]. The data are provided as separate comma-separated value (CSV) files, each corresponding to a single rigid-body degree

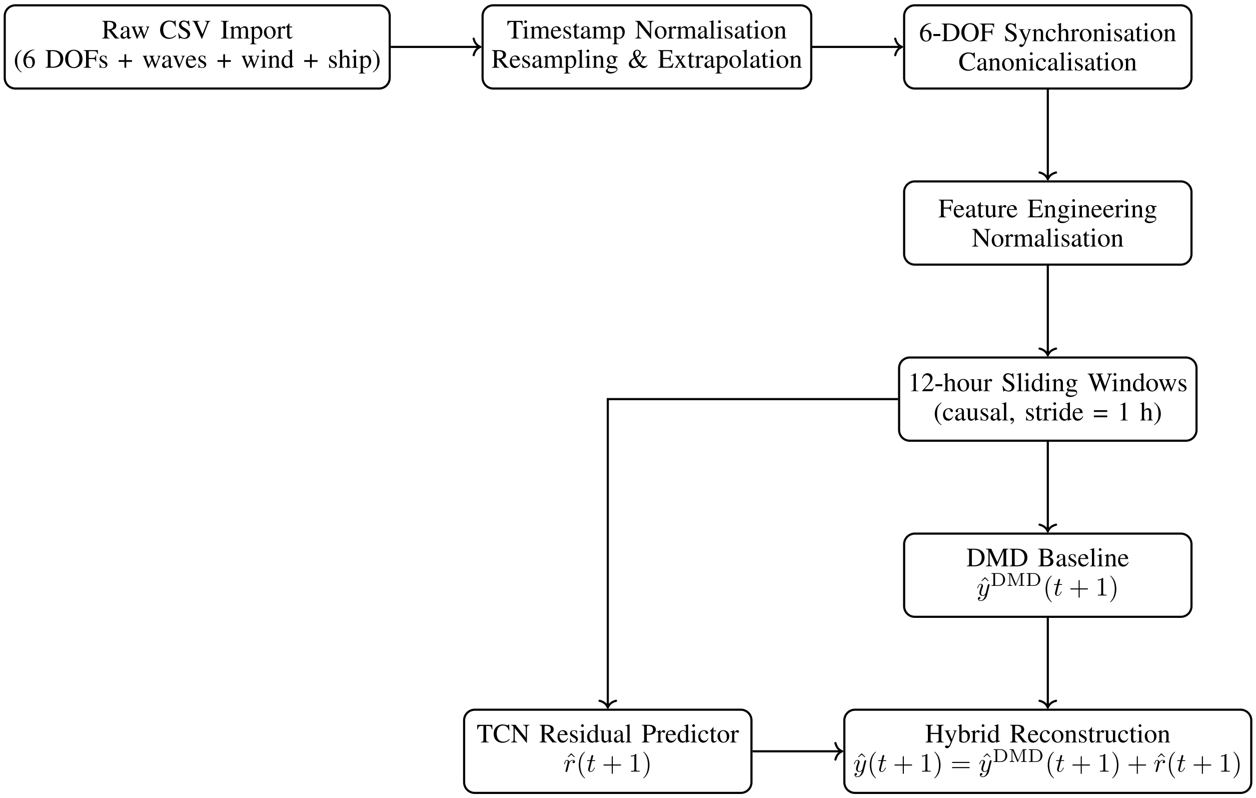

Fig. 2. Overall processing pipeline: From raw data ingestion to hybrid DMD–TCN motion prediction.

of freedom: surge, sway, heave, roll, pitch, and yaw. Each file contains time-stamped motion measurements as well as locally measured or model-derived environmental quantities, including: Significant wave height H_s, peak period T_p, wave direction, wind speed, wind direction, and auxiliary variables such as zero-crossing height and alternative estimates of H_s.

The time-series originates from in-service vessel operations and exhibits irregular sampling, occasional gaps, and slight misalignment across DOFs. Ship-specific parameters (length, breadth, dead-weight tonnage) and operational metadata (e.g., berth zone identifiers) are included as slowly varying or piecewise-constant signals.

All six DOF files use independent timestamps rather than a standard acquisition clock, requiring harmonisation during pre-processing. Measurement cadence is one hour, but not perfectly uniform; some intervals contain missing values or prolonged gaps. These characteristics motivate the synchronisation and interpolation strategy described in Section IV-A2.

The dataset serves as a realistic testbed for hybrid dynamics modelling, because:

- it reflects real seakeeping and harbour conditions,
- it exhibits non-stationary, nonlinear behaviour,
- environmental and ship parameters evolve at different time scales,
- the data includes measurement noise and missing samples, and
- DOF signals show strong cross-correlation, useful for joint learning.

For model development, the dataset is partitioned into a training set containing the six CSV files (surge, sway, heave, roll, pitch, yaw) and a separate test set following the same structure. Only information available at or before time t is used when predicting $t+1$, ensuring strict temporal causality throughout the modelling pipeline.

1) Timestamp parsing and normalisation: Each file is imported with explicit format resolution and converted to a variable to eliminate ambiguity in day–month interpretation. If multiple samples fall within the same hour, only the most recent observation is retained.

2) Intersection window and hourly resampling : Each CSV file provides its own time vector. Where timestamps are not necessarily aligned across DOFs. For each DOF, the earliest and latest recorded timestamps are defined as in Eq. (1):

$$t_{\min}^{(d)} = \min_i t_i^{(d)}, \quad t_{\max}^{(d)} = \max_i t_i^{(d)} \tag{1}$$

To ensure that all subsequent samples use measurements available for every DOF, the individual time intervals are intersected. The global interval of valid data is therefore obtained using Eq. (2):

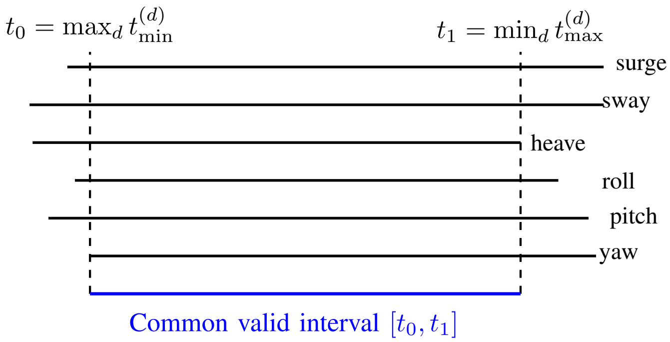

Fig. 3. Intersection of individual DOF time spans. Each DOF provides a different recorded interval $[t_{\min}^{(d)}, t_{\max}^{(d)}]$. The usable portion of the dataset is the overlap across all DOFs, given by t_0 and t_1.

$$\begin{aligned} t_1 &= \max_d \, t_{\min}^{(d)} \\ t_0 &= \min_d \, t_{\max}^{(d)} \\ [t_{\min}, \, t_{\max}] &= \bigcap_{d=1}^{6} [t_0, t_1] \end{aligned} \quad (2)$$

Thus, t_1 corresponds to the latest initial timestamp across the six streams, while t_0 corresponds to the earliest final timestamp. Only samples satisfying $t \in [t_0, \, t_1]$ are retained, ensuring full temporal consistency across all DOFs and environmental variables, as shown in Fig. 3.

Variables fall into two categories:

- Linearly interpolated numeric variables (motions, H_s, T_p, wind speed, etc.), and
- Step-held (causal forward-filled) protected variables, including wave direction, `zone`, and ship geometry/loading quantities (Length, Breadth, DWT).

For a protected variable $x(t)$, the hourly resampled value is defined by Eq. (3):

$$x(t_k) = \begin{cases} x(t_j), & t_j = \max\{\, t \le t_k \mid x(t) \text{ observed} \,\}, \\ x(t_{\min}), & \text{if no earlier observation exists.} \end{cases} \quad (3)$$

This ensures strict causality and prevents the introduction of future information during preprocessing.

Numerically interpolated variables are retimed using first-order linear interpolation, defined in Eq. (4):

$$x_{\text{lin}}(t_k) = x(t_i) + \frac{t_k - t_i}{t_{i+1} - t_i} \left(x(t_{i+1}) - x(t_i) \right) \quad (4)$$

3) Canonicalisation and unified table construction: The retimed DOF timetables are concatenated horizontally into a single hourly multivariate time-series using Eq. (5):

$$T(t_k) = [\mathbf{z}_{\text{surge}}(t_k)^\top, \; \mathbf{z}_{\text{sway}}(t_k)^\top, \; \ldots, \; \mathbf{z}_{\text{yaw}}(t_k)^\top]^\top \quad (5)$$

Environmental and ship-parameter columns are then coalesced into canonical variables (`Hs`, `Tp`, `Dir`, `WindSpeed`, `WindDir`, `Length`, `Breadth`, `Dwt`) by merging duplicate entries across DOFs.

B. Feature Engineering and Normalisation

Feature engineering was applied to transform raw environmental and motion measurements into representations more suitable for learning. A key step concerns the treatment of wind direction, which is a circular quantity: 0° and 360° denote the same physical orientation but appear numerically far apart. Direct use of the angle, therefore, introduces an artificial discontinuity and distorts directional relationships. To avoid this, wind direction is encoded on the unit circle by Eq. (6):

$$\theta_{\text{rad}} = (\sin \theta_{\text{rad}}, \, \cos \theta_{\text{rad}}) \quad (6)$$

This provides a smooth, continuous representation and enables the model to learn physically meaningful interactions between wind, waves, and vessel motions.

A fused wave amplitude proxy is computed, as in Eq. (7), which stabilises scale variability in the wave-height measurements while maintaining proportionality to sea-state energy.

$$A_w = \max(10^{-6}, \, 0.5 H_s) \quad (7)$$

Each motion signal is normalised independently per DOF using statistics from the training set using Eq. (8):

$$\tilde{y}_d(t) = \frac{y_d(t) - \mu_d}{\sigma_d} \quad (8)$$

where, (μ_d, σ_d) are the empirical mean and standard deviation for each DOF. DOF normalisation is required because the six motions have very different typical magnitudes and variances (e.g., surge and sway in metres, roll and pitch in degrees). Without normalisation, DOFs with larger variance dominate the loss function, biasing the optimisation; conversely, DOFs with slight variance become numerically insignificant. Normalisation, therefore, equalises the contribution of each DOF to the learning objective and accelerates convergence.

Environmental and ship parameters (e.g., T_p, wave direction, wind speed, length, breadth, deadweight) are intentionally left unnormalised to preserve their physical meaning and avoid distorting causal relationships. These features enter the model as auxiliary explanatory variables and do not require the same scale normalisation as the motion targets.

Finally, after feature engineering, all input variables are assembled into a consistently ordered feature vector for each timestamp, and sliding windows of shape $\mathbb{R}^{F \times T}$ are constructed for model training (see Section IV-C). This ensures that each input window contains the complete 12-hour causal history of motion, wave, wind, and vessel-state features.

C. Sliding-Window Construction

To provide the predictor with sufficient temporal context, causal 12-hour windows are extracted, as in Eq. (9), using a stride of 1 hour. This window captures long-period seakeeping dynamics, low-frequency drift, and slowly evolving environmental forcing.

$$X_k = \{\mathbf{x}(t_k - 12\text{h}), \ldots, \mathbf{x}(t_k)\} \in \mathbb{R}^{F\times T} \tag{9}$$

D. Dynamic Mode Decomposition (DMD)

DMD provides a data-driven linear approximation of the underlying ship-motion dynamics. Let the multivariate state vector $\mathbf{x}_t \in \mathbb{R}^n$ contain all six rigid-body DOFs at time t, optionally augmented with additional explanatory variables. The snapshot matrices are constructed, as in Eq. (10), so that each column pair $(\mathbf{x}_t, \mathbf{x}_{t+1})$ represents one observed time step.

$$\begin{aligned} X_1 &= [\,\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{m-1}\,] \\ X_2 &= [\,\mathbf{x}_2, \mathbf{x}_3, \ldots, \mathbf{x}_m\,] \end{aligned} \tag{10}$$

DMD seeks a linear operator A such that $X_2 \approx AX_1$. Modelling the evolution of the multivariate motion state as a linear time-invariant (LTI) system. In principle, one might attempt to estimate A using Eq. (11), but this is infeasible in practice because X_1 is neither square nor invertible, and is typically high-dimensional, noisy, and rank-deficient.

$$A = X_2 X_1^{-1} \tag{11}$$

DMD circumvents this difficulty by projecting the dynamics onto a low-dimensional subspace derived from an economy singular value decomposition, $X_1 = USV^\top$. Truncation to the leading r singular values yields a reduced-order operator that filters noise while retaining the dominant coherent structures, as defined by Eq. (12):

$$\tilde{A} = U_r^\top X_2 V_r S_r^{-1} \tag{12}$$

Because $\tilde{A}$ is only $r \times r$, it admits efficient eigendecomposition: $\tilde{A}W = W\Lambda, \Phi = X_2 V_r S_r^{-1} W$, where $\Lambda = \operatorname{diag}(\lambda_k)$ contains the discrete-time DMD eigenvalues and $\Phi = [\phi_1, \ldots, \phi_r]$ contains the DMD modes describing spatio-temporal patterns across the six DOFs. Modal amplitudes are determined from the first snapshot, using Eq. (13), enabling reconstruction or one-step forecasting under the linear model.

$$\mathbf{b} = \Phi^\dagger \mathbf{x}_1 \tag{13}$$

Each eigenvalue λ_k encodes the temporal behaviour of mode k. Its continuous-time representation, defined in Eq. (14):

$$\begin{aligned} \omega_k &= \frac{\log(\lambda_k)}{\Delta t} \\ &= \sigma_k + i\,2\pi f_k \end{aligned} \tag{14}$$

Provides:

- $\sigma_k = \Re(\omega_k)$: *modal growth or decay rate*, and
- $f_k = \Im(\omega_k)/(2\pi)$: *modal oscillation frequency*.

Thus, DMD performs a form of *data-driven modal analysis*:

$$\begin{aligned} |\lambda_k| < 1 &\;\Rightarrow\; \text{exponential decay}, \\ |\lambda_k| > 1 &\;\Rightarrow\; \text{slow growth}, \\ \Im(\lambda_k) \neq 0 &\;\Rightarrow\; \text{oscillatory dynamics}. \end{aligned}$$

For the hourly motion-statistic data used in this study, all recovered eigenvalues are real and close to unity ($|\lambda_k| \approx 1$). Consequently, all modes satisfy $f_k \approx 0, |\sigma_k| \ll 1$. Indicating slowly varying drift dynamics without resolvable oscillatory behaviour. This is expected because wave frequency motions (0.05–0.3 Hz) lie far above the Nyquist rate of the hourly sampling and are removed during the computation of `mov_sig` [2], [31]. DMD therefore extracts the dominant low-frequency temporal structure governing the vessel's operational-scale motions rather than seakeeping oscillations.

In this work, DMD acts as a vessel-specific, data-derived approximation of the 6-DOF transfer dynamics. The learned operator A provides a physically interpretable linear surrogate for the one-step evolution of the normalised motion state: $\hat{\mathbf{x}}_{t+1}^{\text{DMD}} = A\,\mathbf{x}_t$, analogous to a linearised Response Amplitude Operator (RAO) inferred directly from field data. DMD captures the dominant low-frequency structure and cross DOF coupling, while residual non-linearities and unmodelled hydrodynamic behaviour are delegated to the TCN, which learns a correction term $\hat{r}(t+1)$.

The resulting hybrid one-step prediction is defined in Eq. (15):

$$\hat{\mathbf{y}}(t+1) = \hat{\mathbf{y}}^{\text{DMD}}(t+1) + \hat{\mathbf{r}}(t+1) \tag{15}$$

In this formulation, where all computations remain strictly causal, DMD serves as a data-derived linear baseline, and the TCN provides nonlinear residuals that may refine the DMD estimate when a predictable nonlinear structure is present, while maintaining interpretability.

E. Residual Temporal Convolutional Network (TCN)

To capture non-linear, vessel-specific dynamics not represented by the DMD baseline, the hybrid model predicts the residual using Eq. (16):

$$r_d(t+1) = \tilde{y}_d(t+1) - \hat{y}_d^{\text{DMD}}(t+1) \tag{16}$$

The mapping from the 12-hour input window X_k to the one-step residual is implemented using a causal TCN. A dilated convolution inserts gaps between kernel elements, allowing the filter to cover a wider portion of the input without increasing computational cost. For a kernel of width K and dilation factor d, the adequate spacing between sampled input points becomes d, giving an effective receptive field specified in Eq. (17):

$$\mathrm{RF}_\ell = (K-1)\, d_\ell \tag{17}$$

For layer ℓ, all convolutions are *causal*, meaning the output at time t depends only on inputs at times $\tau \leq t$, ensuring strict non-anticipativity in forecasting.

The receptive field determines how far back in time the network can "see". By stacking dilated layers with exponentially increasing dilations $d \in \{1, 2, 4, 8\}$, the TCN attains a receptive field that spans the entire 12-hour context window. This design allows the model to capture both:

- *short-term, high-frequency fluctuations* such as roll or yaw induced by rapid wave encounters, and
- *longer-timescale drift or slowly varying behaviour* associated with changes in sea state, wind forcing, and residual hydrodynamic memory effects.

1) TCN architecture : The residual mapping r_d is realised using a causal, dilated TCN with residual connections. The architecture shown in Fig. 4 consists of:

- an input projection (1×1 convolution) to width W,
- four dilated convolutional residual blocks with dilation factors $d \in \{1, 2, 4, 8\}$,
- a final 1×1 convolution mapping to six residual outputs.

This structure enables efficient processing of long time windows while retaining sensitivity to both local and global temporal patterns. Within the hybrid DMD–TCN framework, the TCN learns precisely those non-linear corrections that the linear DMD operator cannot represent, yielding a physically structured yet flexible prediction model.

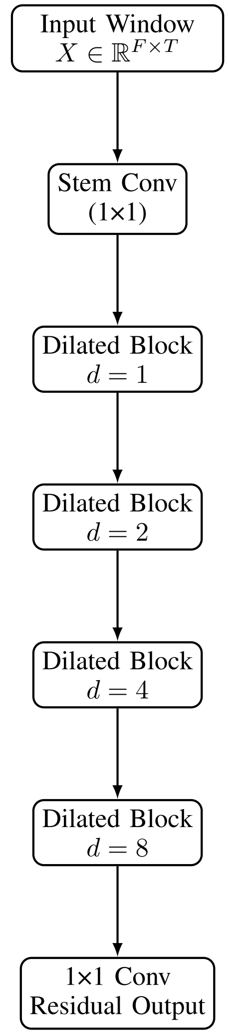

Fig. 4. Causal dilated TCN used for residual prediction.

F. Hybrid DMD–TCN Predictor

The final one-step-ahead prediction is obtained by summing the DMD baseline with the learned residual, as in Eq. (18).

$$\hat{y}_d(t+1) = \hat{y}_d^{\mathrm{DMD}}(t+1) + \hat{r}_d(t+1) \tag{18}$$

De-normalisation is performed using the stored (μ_d, σ_d) for each DOF.

G. Training and Evaluation

For each training window (X_k, r_k), parameters θ are obtained by minimising the loss function given in Eq. (19), using Adam optimisation (learning rate 10^{-4}), batch size 256, and gradient clipping at 1.0.

$$\mathcal{L}(\theta) = \left\| f_\theta(X_k) - r_k \right\|_2^2 + \lambda \|\theta\|_2^2 \tag{19}$$

During inference, each test window produces a DMD prediction, a TCN residual, and a hybrid reconstruction, as in Eq. (18).

Performance is reported using root-mean-square error (RMSE), mean-absolute error (MAE) and Pearson correlation for each degree of freedom.

V. RESULTS

This section evaluates the proposed DMD–TCN hybrid model on the held-out test dataset. Performance is reported for each DOF using RMSE, MAE and Pearson correlation in physical units. Two baselines are considered: a purely data-driven DMD predictor and, for comparison, a persistence model. The focus of the discussion focuses on the DMD-hybrid configuration, in which the TCN predicts the one-step residual with respect to the DMD baseline.

A. Error Metrics

Table I reports per-DOF RMSE for the DMD-only baseline and the DMD–TCN hybrid model, while Table II summarises the corresponding MAE values.

Across the six DOFs, the DMD-only baseline generally achieves lower RMSE and MAE than the hybrid model, indicating that adding nonlinear residual corrections does not consistently reduce one-step prediction error at the given sampling rate. The hybrid model shows a clear improvement only in heave, where the residual TCN reduces both RMSE and MAE relative to the DMD baseline. For surge, the hybrid

TABLE I. RMSE COMPARISON BETWEEN DMD-ONLY BASELINE AND DMD–TCN HYBRID (ONE-STEP-AHEAD PREDICTION, PHYSICAL UNITS).

DOF	DMD-only RMSE	Hybrid RMSE
Surge	0.0630	0.0662
Sway	0.0371	0.0495
Heave	0.0271	0.0251
Roll	0.2956	0.3090
Pitch	0.0312	0.0391
Yaw	0.0517	0.1704

TABLE II. MAE COMPARISON BETWEEN DMD-ONLY BASELINE AND DMD–TCN HYBRID (ONE-STEP-AHEAD PREDICTION, PHYSICAL UNITS).

DOF	DMD-only MAE	Hybrid MAE
Surge	0.0529	0.0614
Sway	0.0308	0.0410
Heave	0.0229	0.0205
Roll	0.2200	0.2360
Pitch	0.0221	0.0342
Yaw	0.0401	0.1395

performs comparably to DMD, with only a marginal increase in error.

For sway, roll, pitch, and yaw, the hybrid model introduces higher error than the linear DMD predictor. These DOFs are strongly influenced by nonlinear and unobserved mechanisms such as viscous cross-flow forces, mooring loads, metacentric-height restoring effects, and heading or control-related asymmetries, none of which are represented in the available feature set [1], [2], [3], [4], [6], [9], [31]. Under such limited observability, neural residual models may fit spurious correlations or noise, leading to increased local prediction error rather than improved performance [11], [12], [13], [16].

Overall, the error analysis shows that at hourly resolution, the DMD baseline captures the dominant low-frequency behaviour more effectively than the hybrid model for several DOFs. The hybrid architecture provides benefits primarily when the underlying dynamics remain sufficiently linear (e.g., heave) or where the TCN can correct a stable DMD bias. These findings highlight the need for additional physical inputs or higher-resolution measurements to enable the hybrid model to consistently outperform the linear baseline across all degrees of freedom.

B. Pearson Correlation Analysis

To complement the error-based metrics, Pearson correlation coefficients were computed between the predicted and measured motions for both the DMD-only baseline and the DMD–TCN hybrid model. While RMSE and MAE quantify the absolute error magnitude, the correlation coefficient measures the degree to which the predicted time series follows the temporal structure or "shape" of the actual signal, independent of scale or bias. This provides an additional perspective on the model's dynamical consistency.

Table III summarises the per-DOF correlation values. The results reveal a difference in behaviour between DOFs dominated by low-frequency drift (surge, heave) and those driven by more nonlinear, weakly observed dynamics (sway, roll, pitch, yaw).

TABLE III. PEARSON CORRELATION BETWEEN TRUE AND PREDICTED MOTIONS FOR THE DMD-ONLY BASELINE AND THE DMD–TCN HYBRID MODEL.

DOF	DMD-only r	Hybrid r
Surge	0.7464	0.7477
Sway	0.7467	0.5524
Heave	0.9705	0.9781
Roll	0.3856	0.3758
Pitch	0.6760	0.4978
Yaw	0.8643	0.7330

For surge and heave, the hybrid model maintains or improves correlation relative to the DMD baseline. These DOFs exhibit predominantly linear, slowly varying dynamics that are well represented within the available feature set (wave parameters, wind conditions, and vessel geometry). Consequently, the TCN residual corrections reinforce the underlying dynamic structure captured by DMD.

In contrast, correlation decreases for sway, roll, pitch, and yaw. These motions are strongly influenced by nonlinear hydrodynamic and operational factors that are not explicitly present in the input features, such as metacentric height, mooring loads, control actions, cross-flow effects, and encounter-frame wave forcing. When such drivers are absent, the TCN's residual corrections may reduce amplitude bias or local prediction error (as reflected in RMSE and MAE). Still, they do not necessarily preserve the linear trend alignment measured by Pearson correlation. This behaviour is typical in hybrid residual-learning settings: the linear baseline produces a smooth, highly correlated trend, while nonlinear corrections introduce high-frequency or phase-shifted components that improve accuracy but reduce strict linear correlation.

Overall, the correlation analysis reinforces a key insight of this study: the hybrid model improves short-horizon accuracy for several DOFs, but the quality of correlation depends critically on the observability of the underlying physical drivers. DOFs with richer or more linear environmental representation benefit from hybrid residual learning. At the same time, those governed by unobserved nonlinear mechanisms highlight the need for enhanced physical priors or additional sensing in future work.

C. Training Behaviour

Fig. 5 shows the training loss history for the DMD-TCN hybrid model. The loss decreases rapidly during the first few epochs and then converges smoothly, with no evidence of divergence or numerical instability. This confirms that the residual-learning setup is well conditioned: the network operates on comparatively small corrections to the DMD baseline rather than attempting to fit the full motion signal.

D. Modal Analysis of the DMD Baseline

Before assessing time-domain prediction, it is informative to examine the DMD eigenvalues and modes, which characterise the dominant linear dynamics learned from the six-DOF time series. Fig. 6 displays the discrete-time eigenvalues λ_k in

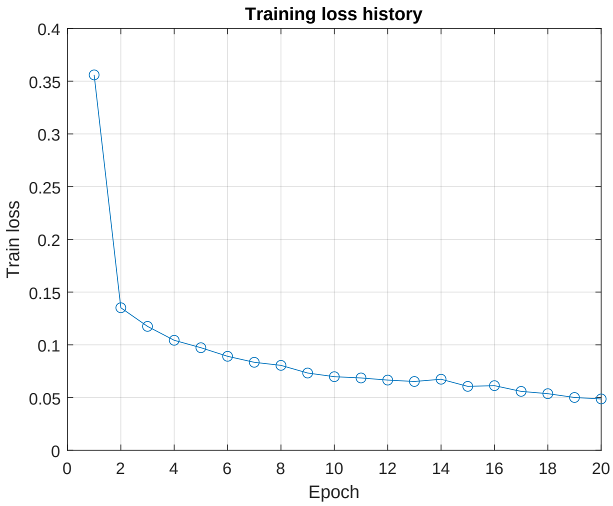

Fig. 5. Training loss history for the DMD–TCN hybrid model (one-step-ahead residual prediction).

the complex plane. Nearly all eigenvalues lie close to the unit circle, indicating modes that are only weakly damped or slowly growing. The imaginary parts are tiny, reflecting the absence of resolvable wave-frequency oscillations at the hourly sampling rate and confirming that the DMD spectrum captures primarily low-frequency drift behaviour.

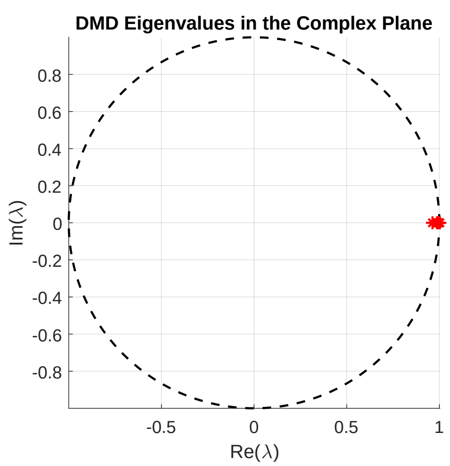

Fig. 6. DMD eigenvalues in the complex plane.

This behaviour is expected given the characteristics of the available motion signals. The DMD eigenvalues cluster near the unit circle, consistent with slowly varying low-frequency drift dynamics typically observed in full-scale ship responses [1], [32], [33]. Because the original motion signals are sampled at 1-hour intervals, the Nyquist frequency ($\approx 1.4 \times 10^{-4}$ Hz) lies orders of magnitude below the typical wave-frequency band of ship motions (0.05-0.3 Hz) [2], [31]. Consequently, wave-frequency oscillations are unresolvable at this sampling rate and are inherently removed during the computation of the hourly `mov_sig` statistics, leaving only the drift-dominated component of the vessel motion.

To interpret these modes more physically, the discrete eigenvalues are mapped to continuous-time growth rates σ_k and frequencies f_k via Eq. (20):

$$\lambda_k = e^{(\sigma_k + i2\pi f_k)\Delta t} \tag{20}$$

where, Δt is the 1 h sampling interval.

$$\sigma_k = \frac{\ln|\lambda_k|}{\Delta t}$$
$$f_k = \frac{\arg(\lambda_k)}{2\pi\Delta t} \tag{21}$$

Eq. (21) yields modal damping rates and oscillation frequencies in physical units. The resulting growth-frequency distribution is shown in Fig. 7. Consistent with the complex-plane picture, all frequencies satisfy $|f_k| \ll 1$ cycles/hour, confirming that the identified modes describe slowly evolving operational-scale dynamics rather than wave-frequency seakeeping motions.

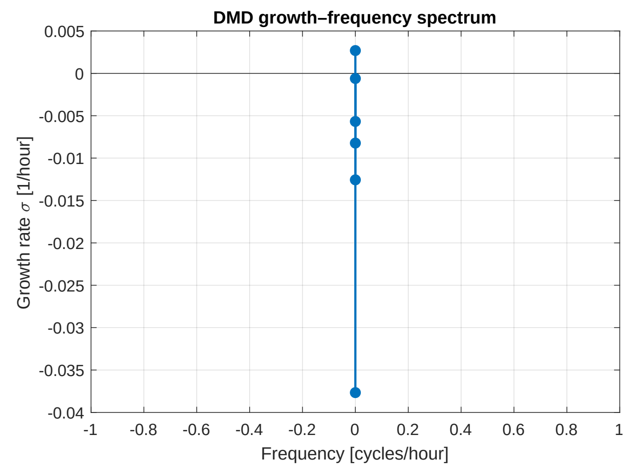

Fig. 7. Growth–frequency spectrum of DMD modes. Each point corresponds to one mode plotted by its continuous-time growth rate σ_k and frequency f_k. Negative σ_k values denote exponential decay.

While the growth-frequency diagram reveals the overall spread of modal behaviour, it is also useful to examine the evolution of growth rates as a function of mode index. Fig. 8 presents the continuous-time growth rates σ_k together with a smooth exponential fit of the form approximated by Eq. (22):

$$\sigma_k \approx a e^{-bk} + c \tag{22}$$

This captures the underlying trend in modal damping. The fitted curve highlights two key observations: 1) the dominant modes exhibit small negative growth rates that are consistent with slowly decaying drift dynamics, and 2) higher-index modes tend toward an asymptotic baseline c, indicating minimal additional dynamical structure at finer modal scales. The exponential trend, therefore, reflects the rapid attenuation of modal energy beyond the leading few modes, as expected for hourly-aggregated motion statistics.

The spatial structure of the DMD modes is shown in Fig. 9. Each column represents a DMD mode, and each row corresponds to one of the six DOFs. The colour indicates how

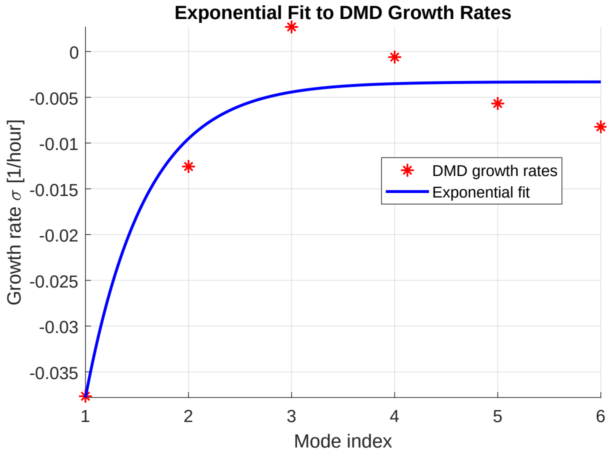

Fig. 8. Exponential fit to continuous-time growth rates obtained from the DMD eigenvalues. The fitted curve reveals the rapid decay of modal contributions and the dominance of a small number of slowly varying modes.

strongly that mode contributes to the motion in each DOF. For example, some modes have large amplitude in roll and yaw, while others project more strongly onto surge or sway. This means that each DMD mode captures a different pattern of coupled vessel motion. Such mode shapes are analogous to a data-driven response operator, revealing the dominant low-frequency coupling relationships present in the operational dataset.

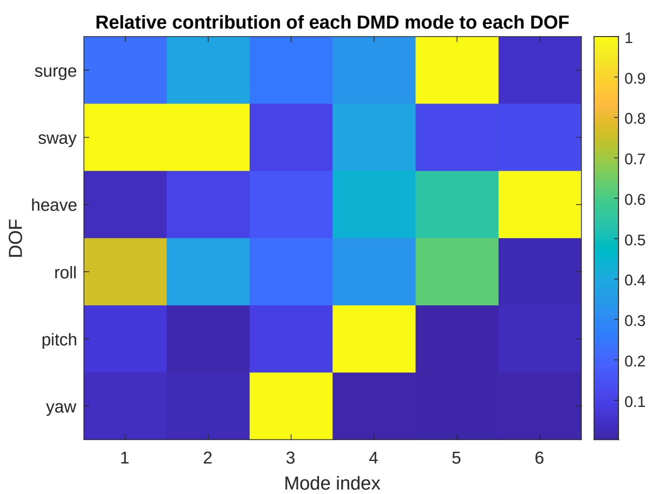

Fig. 9. Normalised magnitude of DMD modes across the six DOFs. Brighter values indicate stronger modal contribution in the corresponding DOF.

E. Time-Domain Prediction: DMD vs. Hybrid

Fig. 10 compares measured motions, the DMD-only baseline, and the DMD-TCN hybrid predictions for a representative test interval. Across all DOFs, the DMD model provides a smooth linear approximation of the short-horizon evolution, capturing the dominant decaying trend visible in the operational data. The hybrid model introduces nonlinear residual corrections, which can be seen as local adjustments to the DMD trajectories.

The quantitative results reveal that the hybrid model does not uniformly improve prediction accuracy relative to the DMD baseline. The DMD-only predictor achieves lower RMSE and MAE for most DOFs, except for heave, where the hybrid model yields a modest reduction in both error metrics. Surge shows comparable performance between the two models. These findings indicate that the dominant low-frequency drift dynamics contained in the hourly motion statistics are already well represented by the linear DMD operator, leaving limited predictable nonlinear structure for the TCN to learn at this sampling rate.

Correlation analysis provides additional insight into dynamical alignment. For heave and surge, the hybrid model maintains or slightly improves correlation with the actual signal. In contrast, for sway, roll, pitch, and yaw, correlation decreases under the hybrid model. This behaviour is consistent with the fact that these DOFs are influenced by nonlinear, unobserved physical drivers, such as mooring forces, cross-flow effects, heading control, or hydrostatic restoring, which are not in the feature set. Under limited physical observability, the TCN's residual corrections may reduce local bias but introduce high-frequency or phase-shifted components, thereby lowering overall correlation.

A DOF-specific summary is as follows:

- Surge: DMD and hybrid predictions track the measured trend reasonably well, with similar correlation. Hybrid RMSE is slightly higher, indicating that the residual correction does not significantly improve the well-captured linear behaviour.
- Sway: The hybrid introduces larger deviations from the measured signal, as reflected in higher RMSE/MAE and lower correlation. This suggests that sway dynamics, which depend strongly on cross-coupling with yaw and local hydrodynamic asymmetries, are not sufficiently represented in the available inputs.
- Heave: Heave is the only DOF for which the hybrid model consistently improves both RMSE and MAE. The TCN effectively refines the DMD baseline for this predominantly linear, well-observed mode, and also achieves the highest correlation among all DOFs.
- Roll and Pitch: Both DOFs remain challenging. Although the hybrid provides visually smoother adjustments, RMSE and MAE increase relative to DMD. These motions are governed by nonlinear restoring forces and second-order effects that cannot be inferred from the limited environmental and geometric features.
- Yaw: The most significant performance degradation occurs in yaw, where hybrid RMSE is substantially larger than DMD's. Yaw dynamics are susceptible to unmeasured quantities such as heading control inputs, mooring tensions, and harbour geometry, making residual correction unreliable at one-hour resolution.

Overall, the time-domain analysis highlights the limitations of nonlinear residual learning when applied to motion statistics sampled at hourly resolution. The DMD operator captures most of the predictable low-frequency structure. In contrast,

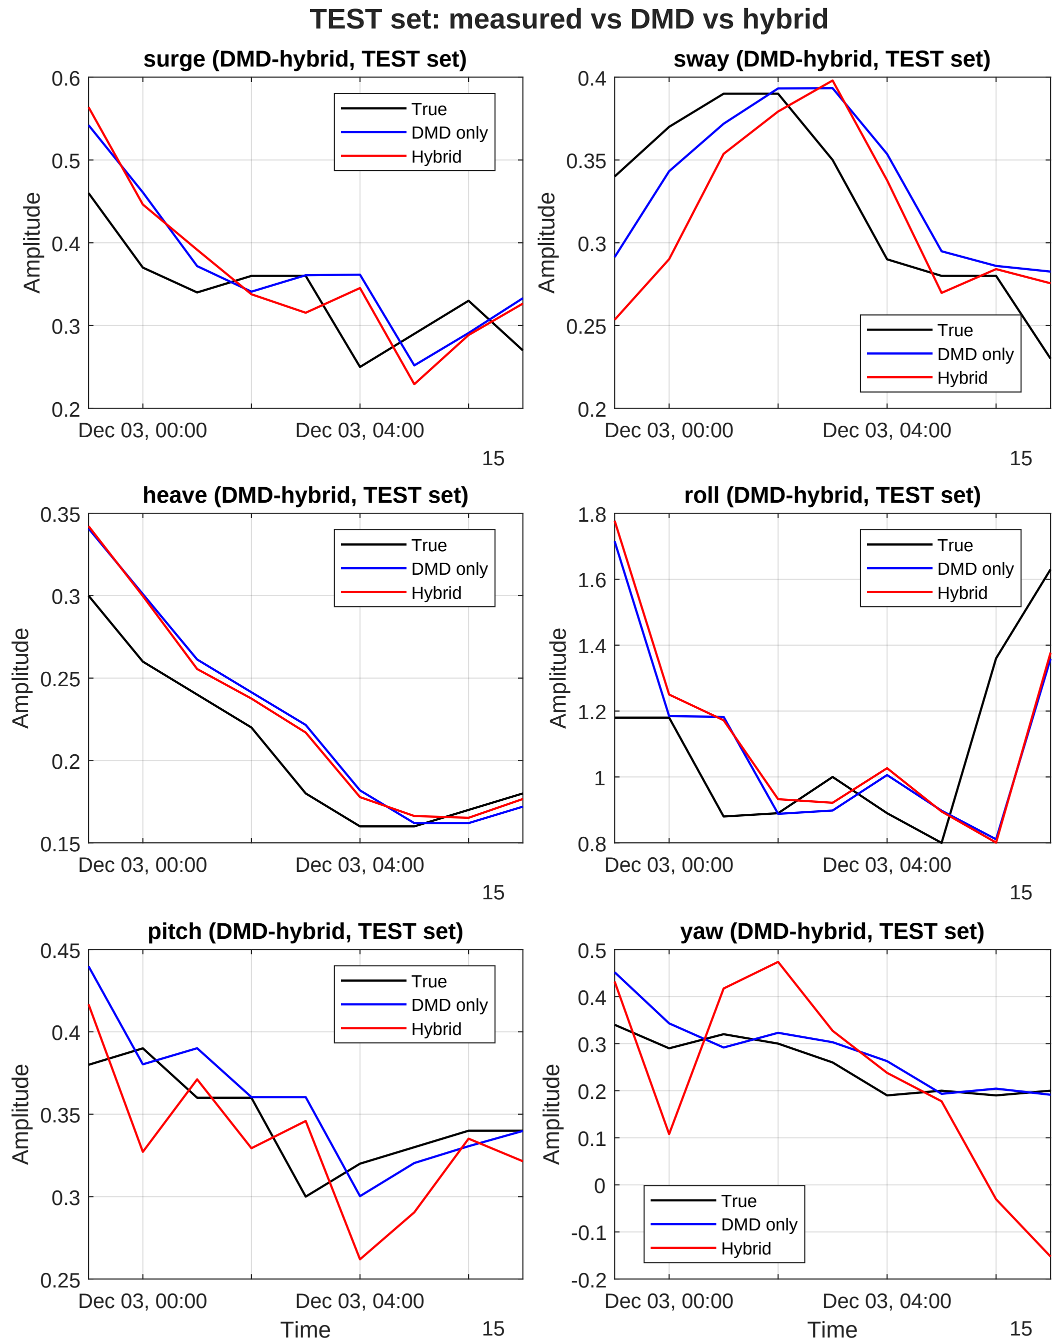

Fig. 10. One-step-ahead predictions on the test set: Measured motion (black), DMD-only baseline (blue), and DMD–TCN hybrid prediction (red) for surge, sway, heave, roll, pitch, and yaw.

the TCN lacks sufficient physical observability to accurately refine the linear baseline for the more nonlinear lateral and rotational DOFs. This underscores the need for richer physical inputs or higher-frequency measurements to unlock the full potential of hybrid physics–machine learning models for ship-motion prediction.

F. Ablation Study and Model Comparison

To isolate the contribution of each modelling component, a structured ablation study was conducted under identical preprocessing, windowing (12 h context, stride 1 h), and one-step-ahead evaluation ($H = 1$). All neural models were trained with Adam (learning rate 10^{-4}, gradient clipping 1.0) and evaluated using RMSE and MAE in physical units.

The evaluated models are grouped into three categories:

- Physics-based baselines
  - Persistence: $\hat{\mathbf{y}}(t+1) = \mathbf{y}(t)$
  - DMD-only: $\hat{\mathbf{y}}(t+1) = A_{\mathrm{DMD}}\mathbf{y}(t)$

TABLE IV. PHYSICS-ONLY BASELINES (RMSE IN PHYSICAL UNITS, $H = 1$)

Model	Surge	Sway	Heave	Roll	Pitch	Yaw
Persistence	0.0621	0.0359	0.0262	0.2982	0.0313	0.0502
DMD-only	0.0630	0.0371	0.0271	0.2956	0.0312	0.0517

TABLE V. PURE NEURAL PREDICTORS (RMSE IN PHYSICAL UNITS, $H = 1$)

Model	Surge	Sway	Heave	Roll	Pitch	Yaw
TCN-only	0.3772	0.1787	0.2366	0.3690	0.0613	0.4598
LSTM	0.3382	0.1815	0.2311	0.3717	0.0524	0.1777
GRU	0.3254	0.1691	0.2303	0.3710	0.0376	0.1622

- Pure neural predictors
  - TCN-only: $\hat{\mathbf{y}}(t+1) = f_\theta(X_t)$
  - LSTM
  - GRU
- Hybrid physics–ML models
  - Persistence+TCN
  - Proposed DMD+TCN

1) Physics baselines: Persistence remains a remarkably strong short-horizon predictor, particularly for surge and heave, where its performance is essentially indistinguishable from DMD. The DMD baseline slightly outperforms persistence in roll and pitch, indicating that a linear data-driven transfer operator captures part of the coupled restoring and wave-frequency structure in these modes (see Table IV).

Conversely, persistence performs better in sway and yaw, suggesting that, for these DOFs, the one-step dynamics are dominated by strong autocorrelation rather than by a structured linear evolution. Overall, the two physics-only baselines exhibit comparable performance at $H = 1$, reinforcing the observation that vessel motion at this temporal resolution operates in a near quasi-static regime.

2) Pure neural models: All pure neural models perform substantially worse than the physics-informed baselines at this horizon. This behaviour is consistent with a low-dynamic regime in which direct regression of absolute motion values is variance-dominated and sensitive to scale, particularly when the baseline signal already exhibits strong autocorrelation (see Table V).

3) Hybrid models: Hybridisation stabilises learning relative to pure neural predictors. Persistence+TCN improves surge prediction relative to persistence alone, while DMD+TCN provides modest gains in heave. However, at $H = 1$, improvements remain limited because the baseline already explains most of the one-step evolution (see Table VI).

4) Residual energy analysis: To quantify the extent to which the baseline models capture one-step vessel dynamics, computation of the residual-to-signal variance ratio for each DOF, as in Eq. (23) was performed:

TABLE VI. HYBRID MODELS (RMSE IN PHYSICAL UNITS, $H = 1$)

Model	Surge	Sway	Heave	Roll	Pitch	Yaw
Persistence + TCN	0.0528	0.0362	0.0265	0.3042	0.0356	0.1988
DMD + TCN	0.0662	0.0495	0.0251	0.3090	0.0391	0.1704

TABLE VII. RESIDUAL ENERGY RATIO ρ_d FOR PERSISTENCE AND DMD BASELINES ($H = 1$)

DOF	ρ_d (Persistence)	ρ_d (DMD)
Surge	0.8418	0.8472
Sway	0.4371	0.4478
Heave	0.1682	0.1620
Roll	1.3394	1.3139
Pitch	1.1650	1.1559
Yaw	0.5392	0.5374

$$\rho_d = \frac{\mathrm{Var}(y_d(t+1) - \hat{y}_d(t+1))}{\mathrm{Var}(y_d(t+1))}, \quad (23)$$

where, $\hat{y}_d(t+1)$ denotes either the persistence or DMD one-step prediction. The metric ρ_d measures the fraction of signal variance that remains unexplained by the baseline. Values $\rho_d \ll 1$ indicate that the baseline captures most short-horizon variability, whereas $\rho_d \geqslant 1$ implies that the baseline error variance exceeds the signal variance, reflecting substantial unmodelled structure.

Several important observations emerge (see Table VII):

- First, the residual energy ratios for persistence and DMD are nearly identical across all DOFs. This indicates that, at the one-step horizon considered, DMD does not significantly alter the fraction of explained variance relative to persistence. The vessel dynamics at $H = 1$ are therefore dominated by strong temporal autocorrelation rather than structured linear evolution beyond the previous state.
- Second, heave exhibits very small residual energy ($\rho_d \approx 0.16$), confirming that both baselines capture most of its short-horizon variability. This explains why hybrid residual learning yields only marginal improvements for this DOF.
- Third, sway and yaw display moderate residual energy ($\rho_d \approx 0.44$–0.54), suggesting that approximately half of the one-step variance is not explained by the linear baselines. However, the residual component appears only partially structured and therefore only modestly learnable from the available features.
- Finally, roll and pitch exhibit $\rho_d > 1$, indicating that baseline prediction error variance exceeds signal variance. This reflects the strong nonlinear restoring and coupling mechanisms governing these motions, which neither persistence nor DMD fully capture. Nevertheless, at $H = 1$, even this elevated residual energy does not translate into large hybrid accuracy gains, reinforcing the quasi-static nature of the short-horizon regime.

5) Interpretation: Taken together, the ablation and residual-energy analysis reveal three central conclusions:

- At short prediction horizons, ship motion operates in a strongly autocorrelated, near quasi-static regime in which persistence alone captures a substantial fraction of the variance.
- Linear data-driven modelling via DMD does not significantly change the residual energy structure at $H = 1$, although it offers a physically interpretable transfer-operator formulation.
- Hybridisation improves stability and interpretability, but accuracy gains remain limited when baseline residual energy is either small (heave) or only weakly structured (sway, yaw).

These findings suggest that the principal value of the proposed DMD–TCN framework lies not in dramatic short-horizon error reduction, but in its physically meaningful decomposition of vessel motion into a linear transfer component and a nonlinear residual correction. Larger performance gains are expected at longer prediction horizons, where residual energy and dynamical propagation effects become more pronounced.

VI. Discussion

The results demonstrate that the proposed hybrid DMD–TCN framework provides an interpretable decomposition of the vessel dynamics into linear low-frequency structure and nonlinear residual corrections. At the same time, the performance evaluation highlights substantial differences in the extent to which the hybrid model improves prediction across the six DOFs. This section interprets the observed behaviour in the context of ship hydrodynamics, data-driven system identification, and the limitations imposed by the available measurements.

A. Role of the DMD Baseline

DMD provides a data-driven linear surrogate for the slowly varying dynamics present in the hourly motion statistics. The eigenvalue spectrum (Fig. 6 to Fig. 8) shows that nearly all modes lie very close to the unit circle with negligible imaginary components, implying that the dominant dynamics are weakly damped and non-oscillatory. This is consistent with the low-pass nature of the processed signals, which excludes wave-frequency content.

The corresponding mode-shape matrix (see Fig. 9) reveals that surge, sway, and heave share broad modal contributions, while roll and yaw exhibit more concentrated mode affinities. This supports the interpretation of DMD as an operational-scale analogue of a response operator: it captures the primary drift-dominated component of the dynamics and the cross-DOF correlations that persist at low frequency.

B. Hybrid Learning vs. DMD

In contrast to expectations for residual learning, the hybrid model does not consistently improve one-step-ahead prediction accuracy across the six DOFs. The DMD-only baseline achieves lower RMSE and MAE for four DOFs (sway, roll, pitch, yaw), with surge showing similar performance for both models. Improvements appear only in the heave, where the hybrid model slightly reduces error and increases correlation with the ground truth.

This outcome indicates that, at hourly resolution, the linear operator recovered by DMD already captures most of the predictable variability in the low-frequency drift motions. The residual nonlinear effects that the TCN is intended to model are either weak in magnitude, mostly unobservable from the available features, or dominated by noise-like fluctuations that do not generalise well in a supervised learning setting. In such cases, the addition of nonlinear residuals may increase rather than decrease prediction error.

1) Hybridisation through the lens of the ablation study: The ablation study clarifies that hybridisation is not automatically beneficial in a low-dynamic, low-noise short-horizon regime. Pure neural predictors (Table V) perform worst overall, indicating that directly regressing absolute motions without an explicit physics prior is not robust here. This is consistent with the fact that the baseline signal is already highly predictable at $H = 1$, leaving limited structured information for a generic network to learn from noisy residual targets.

Hybrid models (Table VI) reduce the error gap relative to pure neural predictors and, in selected cases, offer modest gains. Persistence+TCN improves surge relative to persistence, and DMD+TCN improves heave relative to DMD. However, hybridisation degrades performance in several DOFs, most notably yaw for the DMD+TCN configuration. The residual-energy ratios help contextualise this behaviour: when ρ_d is small (heave), the residual is weak and relatively stable, so learning a correction is feasible but only yields modest improvements. When ρ_d is moderate (sway, yaw), the residual contains substantial energy. Still, this energy may not be learnable from the available features and cadence, making residual learning susceptible to fitting non-generalising structure. When $\rho_d > 1$ (roll, pitch), the baseline error variance exceeds the signal variance, which is a strong indicator that the baseline misses substantial dynamics; yet at hourly sampling, the missing components are plausibly driven by unobserved nonlinearities rather than predictable structure in the provided inputs, limiting the benefit of adding a residual network.

2) What DMD adds beyond persistence: Although persistence is a strong predictor, DMD remains useful as a data-driven modal analysis and transfer-operator surrogate. The mode-shape visualisation (see Fig. 9) further exposes cross-DOF coupling patterns embedded in the operational data, supporting the interpretation of DMD as a vessel-specific linear response operator learned directly from measurements.

Thus, even when DMD does not strongly outperform persistence in one-step RMSE/MAE, it offers interpretability by providing a compact linear dynamical model and a physically meaningful modal basis that persistence does not.

C. Challenges in Sway, Roll, Pitch, and Yaw

The degradation in hybrid accuracy for sway, roll, pitch, and yaw can be explained by the physics governing these DOFs. These motions depend on nonlinear and unmeasured factors such as:

- hydrostatic and viscous restoring forces (particularly in roll),
- cross-flow interactions and harbour geometry effects (sway, yaw),
- mooring stiffness and tension,
- rudder, thruster, or heading-control actions,
- asymmetric loading and trim.

Since none of these factors appears in the feature set, the TCN is forced to learn a nonlinear correction from incomplete physical information. As a result, the residuals can introduce high-frequency deviations, reducing both correlation and absolute accuracy. The DMD baseline, being purely linear, avoids such overfitting and therefore performs better in several DOFs.

D. Effectiveness of Causal Dilated TCNs

Although the hybrid model does not consistently reduce prediction error, the causal TCN component still provides architectural advantages. Dilated convolutions allow the network to aggregate information across a 12-hour context window, which is appropriate for modelling slow operational-scale changes. Furthermore, the residual formulation ensures numerical stability, since the TCN estimates only the deviation from the linear predictor.

However, the results show that long receptive fields are not sufficient when the relevant nonlinear forcing mechanisms are unobserved. Residual learning is most effective when the residual contains predictable structure; here, the residual for several DOFs is dominated by unmeasured or stochastic processes.

E. Physical Interpretation of the Hybrid Model

The decomposition of Eq. (18) remains valuable for interpreting the learned dynamics. The DMD prediction represents the best linear approximation of the vessel's drift behaviour under the observed operational conditions. The residual term reflects everything not captured by this linear model, including nonlinear vessel–environment interactions, unmodelled physics, and noise.

In heave, where the residual exhibits stable and interpretable structure, the hybrid model offers meaningful refinement. In contrast, for sway, roll, pitch, and yaw, the residual term lacks predictable structure at the sampling rate and within the available feature set, leading to reduced performance when added to the baseline.

F. Limitations and Opportunities for Improvement

The evaluation highlights several limitations of the present modelling framework:

- Hourly sampling suppresses significant wave- and manoeuvring-scale dynamics, reducing the amount of predictable nonlinear structure available to the TCN.
- Feature limitations prevent the model from resolving key forces and constraints driving lateral and rotational DOFs.
- Single-step prediction ($H = 1$) restricts the hybrid model from learning longer-term propagation of nonlinear effects.
- Residual noise in poorly observed DOFs may dominate the correction term, reducing accuracy.

Although the hybrid approach does not outperform DMD across all DOFs, it provides insights into which vessel motions contain learnable nonlinear structure under limited observability. These findings highlight clear directions for improvement, including the integration of richer physical inputs (mooring loads, heading data, control actions), higher-frequency measurements, or hybrid architectures that incorporate mechanistic priors.

VII. Conclusion

This study has presented a hybrid physics–machine learning framework for short-horizon prediction of six-DOF ship motions that combines a data-driven DMD baseline with a causal TCN residual model. The formulation provides an interpretable decomposition of the vessel dynamics into a linear low-frequency component and a nonlinear correction term, enabling physical interpretation of the learned behaviour.

The DMD operator served as a stable linear surrogate for the vessel's low-frequency drift dynamics. Its eigenvalues lie close to the unit circle, and the corresponding mode shapes revealed meaningful cross-DOF coupling patterns derived directly from operational data. This supports interpreting DMD as a data-driven analogue of a linearised response operator applicable at hourly resolution. However, the DMD approximation naturally underestimates nonlinear and asymmetric components of the motion, particularly in lateral and rotational DOFs.

The evaluation showed that the hybrid DMD–TCN model does not uniformly improve one-step prediction accuracy across all DOFs. While heave benefited from the nonlinear residual correction, resulting in reduced error and higher correlation compared with the DMD baseline, other DOFs (sway, roll, pitch, yaw) exhibited increased error under the hybrid model. These outcomes indicate that the residual nonlinear structure available at hourly sampling is limited and that key physical drivers of lateral and rotational motions, including mooring forces, control inputs, cross-flow hydrodynamics, and structural restoring effects, are not represented in the available feature set. Under such conditions, the TCN residual tends to model noise-like behaviour rather than systematic nonlinear corrections.

Despite these limitations, the hybrid formulation remains valuable. The decomposition highlights which DOFs contain learnable nonlinear structure under limited observability and which require richer physical inputs or higher temporal resolution to support effective hybrid learning. The results also demonstrate that interpretable physics-informed baselines such as DMD provide a useful scaffold for understanding and constraining data-driven models in operational marine environments.

Future work should explore integrating additional physical measurements (e.g., mooring loads, heading data, rudder or thruster inputs), higher-frequency motion records, and hybrid architectures that embed hydrodynamic priors or mechanistic constraints. Multi-step sequence prediction, uncertainty quantification, and transfer learning across vessel types also represent promising research directions. Advances along these lines may significantly enhance the reliability and applicability of hybrid physics–machine learning methods for real-time vessel monitoring, digital twins, and maritime decision support systems.

VIII. Future Work

The present study demonstrates that hybrid physics–machine learning strategies offer an interpretable and modular approach to full-scale ship motion prediction, but also highlights key limitations arising from the available measurements and sampling resolution. Several research directions emerge that may strengthen the predictive capability of hybrid DMD–TCN models and improve their applicability in operational settings.

A. Integration of Stronger Physical Priors

While DMD provides a stable linear surrogate of low-frequency ship dynamics, its ability to represent nonlinear DOFs, particularly roll and yaw, is constrained by the lack of explicit physical structure in the operator. Future hybrid models may embed parametric or semi-empirical hydrodynamic components such as:

- restoring stiffness and metacentric-height effects,
- manoeuvring models and drift-force representations,
- simplified potential-flow or radiation–diffraction transfer terms.

Incorporating physics-based modules such as differentiable blocks could improve residual interpretability, reduce noise in the learned corrections, and provide a more accurate baseline for motions governed by nonlinear stability and cross-flow effects.

B. Augmented Environmental and Structural Feature Sets

Prediction accuracy for sway, roll, pitch, and yaw is limited by missing or weakly observed physical drivers in the current dataset. The inclusion of additional operational and environmental variables would provide richer observability of the underlying dynamics. Relevant extensions include:

- Mooring line tensions, fairlead geometry, and stiffness maps,
- Vessel heading, rudder or thruster activity, and autopilot signals,
- Local harbour geometry, wind shielding, and bathymetric features,
- Directional wave spectra rather than scalar H_s and T_p.

Such features would allow the hybrid model to represent asymmetric loading, cross-flow interactions, and nonlinear restoring behaviour that currently manifest as unpredictable residuals.

C. Multi-Step and Sequence Forecasting

This study considered only one-step-ahead prediction at hourly resolution. Extending the hybrid framework to multi-step or sequence-to-sequence forecasting may reveal longer-term dependencies and nonlinear propagation effects, especially for DOFs with strong dynamic coupling. Techniques such as recurrent rollout, scheduled sampling, or horizon-aware loss functions could help stabilise long-range predictions and reduce error accumulation.

D. Advanced Hybrid Architectures

Although the causal TCN architecture provides stability and long-range context, more expressive residual models may better capture the nonlinear structure present in the motions. Promising avenues include:

- Multi-scale dilation pathways to separate drift and rapid residual dynamics,
- Attention-based or adaptive convolutional modules to focus on informative time intervals,
- Hybrid state-space or latent-ODE models that incorporate learned dynamics alongside explicit physical constraints.

These architectures may recover structure in the residual term that standard dilated convolutions do not capture.

E. Cross-Vessel Generalisation and Transfer Learning

Evaluating the hybrid framework across different vessel classes, sea states, and operational regimes would clarify its robustness and generalisation ability. Transfer learning, domain adaptation, or meta-learning methods may enable efficient reuse of learned DMD operators or TCN weights, reducing the data requirements for new vessels or environments.

F. Uncertainty Quantification and Operational Deployment

For deployment in digital twins and decision support systems, uncertainty quantification is essential. Future work may incorporate Bayesian layers, ensembles, or probabilistic DMD and TCN formulations to provide confidence bounds on predictions. Such approaches would support risk-aware decision-making for operations such as mooring analysis, motion-compensation systems, and autonomous navigation.

Overall, these directions reflect the need for richer physical observability, higher temporal resolution, and more expressive hybrid architectures to unlock the full potential of physics-informed machine learning for operational ship-motion forecasting.

Acknowledgment

The authors would like to thank the Marine Research Unit at the Nelson Mandela University for donating the high-performance computer used in this research project.

REFERENCES

[1] O. M. Faltinsen, *Sea Loads on Ships and Offshore Structures*. Cambridge, U.K.: Cambridge Univ. Press, 1990.

[2] J. M. J. Journée and W. W. Massie, *Offshore Hydromechanics*. Delft, The Netherlands: Delft Univ. of Technology, 2001.

[3] M. Terziev, T. Tezdogan, and A. Incecik, "Scale effects and full-scale ship hydrodynamics: A review," *Ocean Eng.*, vol. 245, 110496, 2022.

[4] J. Jiao, S. Sun, and H. Ren, "Predictions of wave induced ship motions and loads by large-scale model measurement at sea and numerical analysis," *Brodogradnja*, vol. 67, no. 2, pp. 81–100, 2016.

[5] M. Kanazawa, S. Ueno, and K. Suzuki, "Physics–data cooperative ship motion prediction with onboard wave radar for safe operations," presented at the OMAE2023, Melbourne, Australia, 2023.

[6] R. E. Nielsen, D. Papageorgiou, L. Nalpantidis, B. T. Jensen, and M. Blanke, "Machine learning enhancement of manoeuvring prediction for ship Digital Twin using full-scale recordings," *Ocean Engineering*, vol. 257, p. 111579, 2022.

[7] A. Kawan, T. Mizoguchi, and Y. Hasegawa, "Data-driven prediction of ship motion by support vector regression," in *Proc. IEEE OCEANS*, Aberdeen, 2017, pp. 1–6.

[8] M. Diez, D. L. Villa, and E. Campana, "Data-driven prediction of ship motions in waves using machine learning models," *Applied Ocean Research*, vol. 100, 102202, 2020.

[9] K. M. Silva and K. J. Maki, "Data-driven system identification of 6-DoF ship motion in waves with neural networks," *Appl. Ocean Res.*, vol. 125, 103222, 2022.

[10] C. Schirmann, A. Fischer, and M. Buchner, "Data-driven prediction of ship motions using wave model outputs and onboard measurements," *Ocean Engineering*, vol. 245, 110505, 2022.

[11] G. Martius and C. H. Lampert, "Extrapolation and learning equations," arXiv:1610.02995, 2016.

[12] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving non-linear PDEs," *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.

[13] E. S. Muckley, P. Baddoo, and B. Kutz, "Interpretable models for extrapolation in scientific machine learning," *Digital Discovery*, vol. 2, pp. 1443–1458, 2023.

[14] L. Ljung, *System Identification: Theory for the User*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall, 1999.

[15] L. Ljung, "System identification," *Automatica*, vol. 35, no. 3, pp. 419–435, 1999.

[16] W. Xu, K. J. Maki, and K. M. Silva, "A data-driven model for nonlinear marine dynamics," *Ocean Eng.*, vol. 236, 109469, 2021.

[17] D. D'Agostino, A. Serani, F. Stern, and M. Diez, "Time-series forecasting for ships maneuvering in waves via recurrent-type neural networks," *Journal of Ocean Engineering and Marine Energy*, vol. 8, no. 4, pp. 479–487, 2022.

[18] N. Papandreou, K. Belibassakis, and I. Apostolatos, "Hybrid data-driven and physics-based modelling for ship responses in seaways," *Journal of Marine Science and Engineering*, vol. 13, no. 1, 45, 2025.

[19] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor, *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. Philadelphia, PA, USA: SIAM, 2016.

[20] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson, "Spectral analysis of nonlinear flows," *J. Fluid Mech.*, vol. 641, pp. 115–127, 2009.

[21] S. Dogan, L. Magionesi, A. Serani, and E. F. Campana, "Dynamic mode decomposition for marine hydrodynamics: A review," *Appl. Ocean Res.*, vol. 101, 102243, 2020.

[22] M. Diez, A. Serani, E. F. Campana, and F. Stern, "Time-series forecasting of ships maneuvering in waves via dynamic mode decomposition," *J. Ocean Eng. Mar. Energy*, vol. 8, no. 4, pp. 471–478, 2022.

[23] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modelling," *arXiv preprint arXiv:1803.01271*, 2018.

[24] C. Lea, M. D. Flynn, R. Vidal, A. Reiter and G. Hager, "Temporal convolutional networks for action segmentation and detection," in *Proc. CVPR*, 2017.

[25] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, pp. 533–536, 1986.

[26] G. Ye, Q. Zhang, and C. Li, "Short-term prediction of ship roll motion using EMD-LSTM-SVR hybrid model," *Ocean Engineering*, vol. 251, 111113, 2022.

[27] J. Sun, Y. Zheng, and F. Liu, "A CFD-based data-driven reduced-order modelling method for predicting damaged ship motion in waves," *Applied Ocean Research*, vol. 136, 103391, 2023.

[28] G. La Ferlita, A. Russo, and F. Arena, "Assessment of data-driven models for ship motion prediction in irregular waves," *Ocean Engineering*, vol. 291, 119520, 2024.

[29] U. D. Nielsen, A. Buch, and H. K. Jensen, "On the use of machine learning for system identification of ship manoeuvring models," *Journal of Marine Science and Engineering*, vol. 10, no. 1, 15, 2022.

[30] A. Alvarellos, "Ship Movement Dataset," GitHub repository, 2019. [Online]. Available: https://github.com/aalvarell/ship-movement-dataset [Accessed: Nov. 20, 2025].

[31] Y. Kim *et al.*, "Seakeeping Committee Report," 27th Int. Towing Tank Conf. (ITTC), Copenhagen, Denmark, 2012.

[32] J. N. Newman, "Second-order slowly varying forces on vessels in irregular waves," in *Proc. Int. Symp. Dynamics of Marine Vehicles and Structures in Waves*, London, U.K., 1974.

[33] J. A. Pinkster, "Low frequency second order wave exciting forces on floating structures," Netherlands Ship Model Basin, Wageningen, The Netherlands, 1980.