

ISSN 1998-7927(print) ISSN 2664-6498 (online)

DOI: <https://doi.org/10.33216/1998-7927-2026-304-6-19-34>

UDC 004.03

ADAPTIVE UNCERTAINTY-AWARE HYBRID DIGITAL TWIN ARCHITECTURE FOR PREDICTIVE DIAGNOSTICS OF SHIP POWER PLANTS

Vychuzhanin V.V., Vychuzhanin A.V.

АДАПТИВНА ГІБРИДНА АРХІТЕКТУРА ЦИФРОВИХ ДВІЙНИКІВ З УРАХУВАННЯМ НЕВИЗНАЧЕНОСТІ ДЛЯ ПРОГНОЗНОЇ ДІАГНОСТИКИ СУДНОВИХ ЕНЕРГЕТИЧНИХ УСТАНОВОК

Вичужанін В.В., Вичужанін О.В.

Reliable prediction of the technical condition and remaining useful life of ship power plants (SPPs) remains a challenging task because operating conditions continuously change, degradation mechanisms are highly nonlinear, and sensor measurements are inevitably affected by uncertainty. Although hybrid digital twins (HDTs) have significantly expanded the capabilities of predictive maintenance, many existing implementations still employ fixed interaction schemes between physics-based and data-driven models, reducing their adaptability when the reliability of individual information sources varies during operation. The study develops a HDT architecture that combines a deterministic physical model, a CNN-LSTM prediction module, and a cognitive simulation model (CSM) within a single diagnostic framework. Instead of maintaining a fixed balance between analytical and empirical models, the architecture continuously adjusts their contribution according to the estimated uncertainty of the physical model. This makes it possible to retain physical consistency whenever analytical knowledge is reliable while allowing data-driven prediction to dominate when model uncertainty increases. The cognitive simulation component complements numerical prediction by explicitly representing causal relationships between degradation processes and diagnostic conclusions, providing engineering explanations that are difficult to obtain from conventional deep neural networks. Performance of the developed framework was investigated using hardware-in-the-loop (HIL) simulations reproducing representative operating regimes of SPP subsystems, including progressive degradation and stochastic sensor disturbances. Under these conditions, the hybrid architecture reduced remaining useful life prediction error by 32–36% relative to standalone physics-based models and by up to 19% compared with conventional deep-learning solutions. Fault classification reached an FI-score of 0.89 while maintaining stable operation under sensor noise levels of up to 30%. Real-time inference below 20 ms satisfies the requirements of onboard automation systems, whereas asynchronous federated learning enables scalable deployment across distributed fleets without increasing onboard computational load. The combination of adaptive uncertainty-

aware model fusion, physics-informed learning, and cognitive reasoning extends the capabilities of current HDT technologies beyond conventional hybrid diagnostic frameworks. Besides improving prediction accuracy and robustness, the architecture enhances transparency of maintenance decisions and demonstrates practical feasibility for intelligent SPPs, with an estimated reduction in annual operating costs of approximately 27%. These results indicate that the developed framework represents a practical step toward reliable and explainable predictive maintenance of next-generation maritime energy systems.

Keywords: *physics-informed machine learning; remaining useful life estimation; cognitive simulation; explainable artificial intelligence*

1. Introduction. While DTs are increasingly applied to complex systems like SPPs [1, 2, 3], existing modeling methods lack a comprehensive solution for survivability and reliability assessment under severe uncertainty [4, 5]. Physics-based models ensure thermodynamic consistency [3] but degrade under non-linear wear. Conversely, deep neural networks [6, 7] approximate complex dependencies but suffer from out-of-distribution (OOD) vulnerability [8]. Attempts to integrate physical knowledge via physics-informed neural networks (PINNs) [9, 10] or utilize fuzzy cognitive models [11, 12] remain largely isolated, failing to provide consistent, real-time coordination between deterministic laws, causal relationships, and data-driven learning.

Despite the individual strengths of CSMs [13], fuzzy inference systems, and PINNs, their isolated application does not fully address the requirements of a consistent ship DT. Physics-based models provide a causally interpretable representation of system behaviour

but remain limited under strong non-linearity, incomplete degradation knowledge, or stochastic disturbances. Conversely, data-driven deep neural networks approximate complex dependencies yet exhibit reduced robustness beyond their training domain, particularly in the presence of noisy or sparse observations. These limitations motivate the development of an integrated modelling framework for SPP.

In view of these limitations, this work aims to develop an adaptive uncertainty-aware HDT architecture for SPPs in which the interaction between physics-based, data-driven, and cognitive models is dynamically governed according to the estimated reliability of physical knowledge. Rather than combining several modelling techniques within a common framework, the proposed architecture establishes a unified mechanism for adaptive coordination of heterogeneous models while preserving physical consistency, improving robustness under uncertainty, and supporting interpretable diagnostic reasoning.

Scientific novelty of the proposed approach consists in the development of an adaptive uncertainty-aware HDT architecture for SPPs that enables dynamic coordination of heterogeneous analytical and data-driven models according to their reliability under changing operating conditions. This contribution is supported by the following methodological developments: formulation of adaptive coordination principles for physics-based, neural network, and cognitive models according to their domains of validity; development of a unified state representation combining quantitative process variables with causal knowledge describing degradation mechanisms; integration of physical constraints and cognitive reasoning into a common predictive framework to improve robustness and interpretability; substantiation of the proposed architecture for survivability and reliability assessment under uncertainty and progressive degradation.

2. Formal system model and problem formulation. Before detailing the computational architecture of the HDT, it is imperative to rigorously define the cyber-physical system under investigation. We formulate the maritime power plant operation as a discrete-time, non-linear, partially observable dynamic system operating under severe uncertainty.

The true physical evolution of the system and its sensory observations are defined by the following state-space equations:

$$x_{t+1} = F(x_t, u_t) + w_t, \quad y_t = H(x_t) + v_t, \quad (1)$$

where $x_t \in R^n$ is the true hidden physical state vector (e.g., internal temperatures, actual wear levels) at time step t ; $u_t \in R^m$ is the deterministic control input vector (e.g., fuel injection rate, valve positions); $F: R^n \times R^m$ is the represents partially unknown non-linear physical transition dynamics; $y_t \in R^p$ is the multidimensional sensory output (raw telemetry from the IoT network); $H: R^n \rightarrow R^p$ is the measurement function mapping the physical state to sensor readings; $w_t \sim N(0, Q)$ is the epistemic process uncertainty, capturing unmodeled non-

linear degradation and environmental disturbances; $v_t \sim N(0, R)$ is the aleatoric measurement noise, representing stochastic sensor errors and signal dropouts.

The system is strictly bounded by thermodynamic, structural, and safety limits. We formalize these physical constraints as a feasible operating envelope C :

$$q(x_t, u_t) \leq 0, \quad \forall t$$

Any trajectory where $q(x_t, u_t) > 0$ indicates a critical fault or structural rupture.

HDT inference and explainability (XAI) - since the true state x_t is hidden and the transition degrades over time, the HDT acts as an adaptive estimator and predictor. Based on the sliding window of observations $Y_{t-k:t}$ and the cognitive meta-features m_t , the HDT predicts the future state distribution $\hat{x}_{t+1}$:

$$\hat{x}_{t+1} = \Phi_{HDT}(Y_{t-k:t}, U_{t-k:t}, m_t; \theta) \quad (2)$$

To ensure the supervisory logic is transparent, the architecture must map the latent representations of Ψ_{HDT} into human-interpretable domain knowledge. We define the explainability vector e_t :

$$e_t = \Psi_{XAI}(\hat{x}_{t+1}, \theta)$$

Operator feedback and supervisory control - the predictive output and its explanation e_t are transmitted to the SCC. The formal system model includes HIL supervisory logic, where the human operator evaluates e_t and provides an asynchronous feedback signal $o_t \in \{0, 1\}$ (where 1 indicates validation and 0 indicates rejection/override). The final control policy π executed by the system is therefore a function of both the predictive state and operator validation:

$$u_{t+1}^* = \pi(\hat{x}_{t+1}, e_t, o_t)$$

This formalization systematically translates the challenges of maritime predictive maintenance into a constrained stochastic optimization problem, providing a rigorous mathematical foundation for the cognitive-physical coupling methodology proposed in the subsequent sections.

3. Architecture of the HDT. The HDT fuses parallel physics-based and machine learning models to map high-frequency IoT telemetry into prescriptive decisions. We formalize this architecture by defining its fundamental spaces and mappings. Let the physical state space of the SPP be defined as a vector $x_t \in R^n$, the control input space as $u_t \in R^m$, and the target prognostic output as $y_t \in R^k$. The stochastic nature of the operational environment, including sensor measurement errors, is formalized by an additive noise term $\epsilon_t \sim N(0, \Sigma)$, where $\Sigma \in R^{p \times p}$ is the covariance matrix. The deterministic physics-based model is defined as a mapping from the state and control spaces to the target output space:

$$f_{phys}: R^n \times R^m \rightarrow R^p$$

Concurrently, the neural network component processes a multi-modal telemetry matrix $X_{t-w:t} \in R^{\omega \times M}$ (where ω is the time-window size and M is the number of features), acting as the mapping:

$$f_{ML}: R^{\omega \times M} \rightarrow R^p$$

Formally, the HDT is defined as an adaptive, non-stationary operator HDT that maps the Cartesian product of the historical telemetry space, current control inputs, and cognitive priors into the prognostic target space:

$$H_{DT}: R^{\omega \times M} \times R^m \times R^{|m|} \rightarrow R^p$$

Unlike conventional additive ensemble models where machine learning merely acts as a post-hoc residual corrector for analytical equations, the proposed HDT performs deep algorithmic coupling within the latent topological space of the PINN. The multi-branch CNN-LSTM architecture explicitly separates the representation learning into two parallel manifolds:

physics-informed latent space (Z_{phys}) - extracted by a dedicated branch $\psi_{phys}(\cdot; \theta_{phys})$ driven by state variables x_t and control inputs u_t . This manifold is strictly regularized during training by the thermodynamic differential loss, ensuring adherence to fundamental conservation laws;

cognitive-data latent space (Z_{data}) - extracted by the empirical branch $\psi_{ML}(\cdot; \theta_{ML})$ from the multi-modal telemetry matrix $X_{t-w:t}$, and multiplicatively modulated by the cognitive meta-features m_t derived from the CSM.

The dynamic adaptation and prognostic output are formulated as a convex combination of these latent representations, processed by the final decoding layers (Φ_{dec}):

$$y_{hybrid}(t + \Delta t) = \Phi_{dec} \left((1 - \lambda(t)) \cdot Z_{phys}(t) + \lambda(t) \cdot Z_{data}(t); \theta_{dec} \right) \quad (3)$$

The dynamic adaptation coefficient $\lambda(t)$ is a time-dependent scalar function driven by the epistemic uncertainty of the physical invariants,

$$Z(t) = \text{uncert}_{phys}(x_t) \geq 0.$$

At each discrete time step t , it is computed using an exponential saturation function:

$$\lambda(t) = 1 - \exp(-\gamma \cdot z(t)) \quad (4)$$

This mathematical structure strictly bounds the coefficient within the interval $\lambda(t) \in [0, 1)$. Under nominal operations ($z(t) \rightarrow 0$), the network derives its predictions almost exclusively from Z_{phys} , maintaining absolute thermodynamic consistency. Conversely, during severe non-linear degradation where physical equations lose validity ($z(t) \rightarrow \infty$), the $\lambda(t)$ coefficient dynamically shifts the latent focus towards Z_{data} , allowing the data-driven branch to capture unmodeled wear mechanics. Accumulating physical uncertainty shifts

computational weight to the ML component continuously, preventing gradient discontinuities during backpropagation. This dynamic adaptation is further justified by decomposing the total prognostic error (E_{hybrid}) into physical bias, ML variance, and cross-domain interaction:

$$E_{hybrid} = E_{phys} + E_{PINN} + E_{int}, \quad (5)$$

where E_{phys} represents the deterministic error due to unmodeled physical dynamics (epistemic bias), E_{PINN} is the empirical approximation error of the neural network, and E_{int} accounts for the statistical interaction and covariance between the physical priors and data-driven residuals. The adaptive coefficient $\lambda(t)$ acts as a dynamic regularization factor that actively minimizes the interaction error E_{int} . By suppressing the ML contribution ($\lambda \rightarrow 0$) when physical certainty is high, the system prevents the neural network from amplifying aleatoric sensor noise, ensuring that the total error envelope E_{hybrid} remains strictly bounded.

Unlike static penalty coefficients or heuristic switching schemes, the proposed exponential saturation is a continuous, differentiable transition manifold. Recent studies [14, 15] show that static or unconstrained dynamic weighting in multi-objective losses causes gradient pathologies and extrapolation instability. By linking $\lambda(t)$ to a physically bounded sliding-window uncertainty representation, the proposed HDT framework enables estimation of epistemic uncertainty and reduces the risk of instability associated with unconstrained uncertainty growth. To prevent statistical interference between heterogeneous physical parameters, the predictive core uses a multi-branch CNN-LSTM topology. The input telemetry window $X_{t-k:t}$ (60×7 matrix) is split into three domain-specific sub-matrices ($X_{t,fuel}$, $X_{t,cool}$, $X_{t,elec}$) for isolated spatio-temporal feature extraction per subsystem. The physical model serves as a baseline; the ML component refines predictions by capturing residuals and unmodeled nonlinearities. The overall information flow from sensors to decision support is shown in Fig. 1.




Fig. 1. Extended architecture of the hybrid DT (Integration of Physics, PINN, and CSM)

The topology (Fig. 1) decouples physical (Z_{phys}) and empirical (Z_{data}) processing to ensure inferences rely on thermodynamic invariants. An adaptive latent fusion, governed by the dynamic trust gate (λ), then guarantees the final output is physically valid and empirically accurate, preventing black-box hallucinations.

4. Cognitive modulation and meta-feature integration. Acting as the HDT's causal core, the CSM maps physical parameters to failure scenarios, transforming its structural and probabilistic artifacts into a standardized meta-feature vector m_t .

Table 1

Composition of the CSM meta-features for neural network modulation

Feature	Symbol	Description	Source
Structural risk	$R_{str}(t)$	Failure vulnerability based on FTA topology.	CSM
Functional risk	$R_{func}(t)$	Real-time risk from deviations via causal matrices.	CSM
Failure probability	$P_i(t)$	Component failure prob. from CSM simulation.	CSM
Estimated damage	$Y_i(t)$	Operational/financial failure severity assessment.	Expert
Data reliability	$C_i(t)$	Sensor data validity/completeness score.	Layer

Mathematical formalization of the cognitive ontology (FTA) - before integration into the neural network, the causal metrics (such as $P_i(t)$ and $R_{str}(t)$) must be mathematically formalized. The CSM evaluates the state of the i -th complex subsystem by propagating the baseline failure rates $\lambda = [\lambda_1 \dots \lambda_v]$ of underlying basic events (e.g., individual sensor or valve faults) through the Boolean structure function of the fault tree. Assuming an exponential reliability law for basic components, the dynamic failure probability of the i -th subsystem is rigorously defined by the mapping f_{FTA} :

$$P_i(t) = f_{FTA}(\lambda, t, S_i) = \prod_{k \in S_i^{AND}} (1 - e^{-\lambda_k t}) \cdot \left[1 - \prod_{j \in S_i^{OR}} e^{-\lambda_j t} \right], \quad (6)$$

where S_i^{AND} and S_i^{OR} denote the sets of sub-components connected via AND (parallel redundancy) and OR (series dependency) logical gates within the specific fault tree topology S_i .

This exact deterministic calculation ensures that the resulting meta-features reflect strict topological dependencies rather than empirical approximations. These calculated metrics are then aggregated at each discrete time step t to form the structural input vector:

$$m_t = [R_{str}(t), R_{func}(t), P_i(t), Y_i(t), C_i(t)]^T \quad (7)$$

The variables within the vector m_t formalize the transition from discrete fault logic to continuous neural network inputs. Specifically, $R_{str}(t)$ and $R_{func}(t)$ quantify the topological and functional vulnerabilities $P_i(t)$ and $Y_i(t)$ establish a risk hierarchy, and $C_i(t)$ serves as a credibility weight. The neural network does not calculate these metrics. They are strictly derived from the deterministic FTAs and the fuzzy inference operators of the CSM, as formalized in the preceding subsections. Consequently, the vector m_t functions purely as an external informational payload, delivering the verified cognitive context into the machine learning domain via the integration API.

Mathematical formalization of cognitive modulation (gated attention) - to systematically integrate the cognitive context into the predictive core, the late fusion bypass topology is mathematically implemented via a multiplicative attention mechanism (cognitive gated attention). As previously demonstrated, the meta-feature vector m_t is not simply concatenated with the raw telemetry x_t at the input layer, as this would lead to feature scaling imbalance and the generation of zero-gradient channels. Instead, this late fusion is operationalized by treating m_t as a dynamic prior that modulates the hidden states of the multi-branch CNN-LSTM topology.

Let $h_i(t) \in R^d$ be the hidden state vector extracted from the i -th physical branch of the neural network (e.g., fuel, cooling, or electrical subsystem). The cognitive modulation is formalized through the following sequence of tensor operations. First, an alignment score $e_i(t)$ is calculated by projecting the hidden state $h_i(t)$ and the CSM meta-feature vector m_t into a shared latent space:

$$e_i(t) = v_\alpha^T \tanh(W_h h_i(t) + W_m m_t + b_\alpha), \quad (8)$$

where $W_h \in R^{d_\alpha \times d}$ and $W_m \in R^{d_\alpha \times |m|}$ are the trainable weight matrices; $b_\alpha \in R^{d_\alpha}$ is the bias vector; $v_\alpha \in R^{d_\alpha}$ is the learned projection vector.

Second, the alignment scores are normalized across all K physical branches using the softmax function to obtain the dynamic attention weights $\alpha_i(t)$:

$$\alpha_i(t) = \frac{\exp(e_i(t))}{\sum_{j=1}^K \exp(e_j(t))} \quad (9)$$

The original hidden states are multiplicatively modulated by the computed attention weights:

$$\bar{h}_i(t) = \alpha_i(t) \cdot h_i(t) \quad (10)$$

The resulting modulated vectors $\bar{h}_i$ are subsequently aggregated to form the global state vector s_t .

The attention mechanism uses standard multiplicative gating. Novelty comes from the meta-features m_t : deterministic, physically verified CSM outputs. Standard self-attention relies on stochastic data

correlations; here, CSM guides the latent space. If CSM detects a topological vulnerability in a subsystem, it forces the network to prioritise that subsystem's features before macroscopic deviations appear in raw IoT telemetry. Fig. 2 shows the complete tensor pipeline and integration nodes between cognitive, ML, and physical domains.




Fig. 2. Computational data flow and tensor pipeline of the HDT predictive core

The computational data flow (Fig.2) illustrates the convergence of parallel processing streams at the deep latent fusion node. In the cognitive-data pipeline, meta-features (m_t) are non-linearly projected and normalized via a softmax function into an attention vector $\alpha_i(t)$, which multiplicatively modulates the telemetry features to yield the empirical latent space $Z_{data}(t)$. Concurrently, the PINN encoder (Ψ_{phys}) extracts the strictly regularized physics latent space $Z_{phys}(t)$, while thermodynamic equations calculate the epistemic uncertainty $U_{epis}(x_t)$ to dynamically drive the scalar weight $\lambda(t)$. At the fusion node, the architecture performs a convex combination – scaling $Z_{phys}(t)$ by $(1 - \lambda)$ and $Z_{data}(t)$ by λ – before passing the tensor through the non-linear decoder (Φ_{dec}) to generate the integrated prediction y_{hybrid} . Crucially, this strictly acyclic directed graph topology ensures that gradient backpropagation correctly updates neural weights without violating isolated physical operators, maintaining mathematical stability across the multi-objective optimization.

5. Mathematical formalisation of the physics-constrained latent space. While the macroscopic architecture of the HDT defines the parallel processing of analytical and empirical data, the core cognitive-physical coupling occurs within a shared, physics-constrained latent manifold $Z \subset R^H$. It is imperative to formalise how physics priors, empirical embeddings, and explainability are mathematically bound within this space. A latent representation z_t is formally classified as an interpretable state if and only if it is bounded within a structured sub-manifold $Z_{phys} \subseteq R^H$ that satisfies the condition of physical disentanglement. Formally, for any state $Z_t \in Z_{phys}$, there exists a definitive linear projection $\Omega: Z_{phys} \rightarrow P$, where P denotes the space of observable thermodynamic parameters. Consequently, each principal coordinate (or orthogonal subspace) within z_t is causally mapped to a distinct physical invariant (e.g.,

thermal degradation, frictional losses) governed by the differential residual LPDE. This mathematically precludes the formation of abstract, uninterpretable neural features. Let $\Phi_{PINN}(\cdot; \theta_P)$ and $\Phi_{COG}(\cdot; \theta_C)$ denote the non-linear projection functions of the physical and cognitive encoders, respectively. At any discrete time step t , the system maps the analytical physics priors (governed by thermodynamic state variables x_t) and empirical observations (sensory telemetry y_t) into dense vector embeddings:

$$z_{phys}^{(t)} = \Phi_{PINN}(x_t, u_t; \theta_P) \in Z, z_{data}^{(t)} = \Phi_{COG}(y_t; \theta_C) \in Z$$

To ensure the latent space Z is not a pure statistical "black box," its topology is strictly regularised by the physical prior. During optimisation, the manifold is constrained by the differential PDE residual penalty LPDE, enforcing that the embedding $z_{phys}^{(t)}$ adheres to the thermodynamic invariants of the marine power plant. Consequently, Z becomes a structured metric space where Euclidean distance inversely correlates with physical consistency. The unified latent state z_t , representing the fused digital twin state, is constructed via a dynamic cross-attention mechanism parameterised by the adaptation coefficient λ_t :

$$z_t = \lambda_t z_{phys}^{(t)} \oplus (1 - \lambda_t) A(z_{phys}^{(t)}, z_{data}^{(t)}), \quad (11)$$

where A is the cross-attention operator aligning the empirical data with the physical basis.

Crucially, this mathematical structure directly resolves the XAI formulation. Because $z_{phys}^{(t)}$ represents verified thermodynamic states and $z_{data}^{(t)}$ represents current sensory realities, the explainability vector e_t is extracted not via external perturbation, but intrinsically as the alignment matrix between these two embeddings:

$$e_t = \text{Softmax} \left(\frac{(z_{data}^{(t)} W_Q)(z_{phys}^{(t)} W_K)^T}{\sqrt{d_k}} \right), \quad (12)$$

where $W_Q \in R^{H \times d_k}$ and $W_K \in R^{H \times d_k}$ denote the learnable linear projection weight matrices that map the data and physical embeddings into the query and key spaces, respectively; $\sqrt{d_k}$ acts as a scaling factor to stabilise the gradients by compensating for the variance of the dot product in the high-dimensional space d_k . This equation mathematically proves that the explainability vector e_t represents the exact geometric alignment (attention weights) between incoming empirical anomalies and the expected physical priors within the constrained manifold. Consequently, the XAI output is inherently causal and physics-aware, directly bridging the representational gap between data-driven inference and deterministic thermodynamics.

6. Dynamic adaptation of the confidence coefficient $\lambda(t)$. Marine degradation processes are non-stationary. The physics-based model (f_{phys}) is accurate under steady-state conditions, but error margins increase

during unmodeled nonlinear events. Data-driven components (f_{ML}) approximate these nonlinearities but exhibit extrapolation instability on out-of-distribution data. The HDT implements dynamic weight redistribution to mitigate this. As formalized the coefficient $\lambda(t)$ determines the neural network's proportional contribution via an exponential saturation function. The adaptation relies on the real-time uncertainty estimation of the physical model. The total uncertainty metric, $uncert_{phys}(x_t)$, is computed as the MSE of the physical model over a sliding time window ω :

$$uncert_{phys}(x_t) = \frac{1}{\omega} \sum_{k=t-\omega}^t (y_{true}(k) - f_{phys}(x_t, u_t) - Bias_{static})^2, \quad (13)$$

where $y_{true}(k)$ is the represents the verified historical measurements; $Bias_{static}$ is the accounts for systematic sensor offsets.

To provide rigorous theoretical grounding and address the nature of these errors, this total uncertainty must be explicitly decomposed into aleatoric and epistemic components. From a statistical perspective, variances are strictly additive. Aleatoric uncertainty (U_{alea}) represents the irreducible stochastic noise and natural variance of the sensory environment (σ_{env}^2). Epistemic uncertainty (U_{epis}), conversely, represents the lack of knowledge in the analytical model regarding unmodeled non-linear degradation. To prevent the neural network from overfitting to mere sensor noise, the adaptive transition must be driven exclusively by structural model failures. Therefore, the epistemic variance is mathematically isolated from the total mean squared error, maintaining strict dimensional consistency:

$$U_{epis}(x_t) = \max(0, uncert_{phys}(x_t) - \sigma_{env}^2) \quad (14)$$

By replacing the generic uncertainty variable $z(t)$ in Eq. 4 with the calculated epistemic uncertainty $U_{epis}(x_t)$, the architecture operationalizes the real-time trust redistribution law. Although the dynamic adaptation coefficient $\lambda(t)$ and the trust modifier T_t explicitly depend on the epistemic uncertainty level U_{epis} , the rigorous mechanism of uncertainty propagation through the layers of the HDT necessitates formalisation. Given the strict hard real-time constraints, we bypass computationally prohibitive Monte Carlo dropout and full Bayesian neural networks, adopting instead a deterministic covariance approximation within the latent space. Let Σ_{in} denote the covariance matrix of the input aleatoric noise, defined by the sensor variance v_i . As the state vector propagates through the physical encoder (PINN), the variance of the physical representation, $\Sigma_{PINN}^{(t)}$, is approximated via the transformation Jacobian J_{phys} and penalised by the thermodynamic residual L_{PDE} :

$$\Sigma_{PINN}^{(t)} = J_{phys} \Sigma_{in} J_{phys}^T + O(L_{PDE}) \quad (15)$$

Concurrently, the cognitive layer formulates the data-driven epistemic variance matrix $\Sigma_{COG}^{(t)}$, which is inversely proportional to the entropy of the attention weights. The final scalar uncertainty estimate $U_{epis}(t)$, utilised within the supervisory loop, is quantified as the trace of the fused covariance matrix in the latent space:

$$U_{epis}(t) = Tr(\Sigma_{PINN}^{(t)} \oplus \Sigma_{COG}^{(t)}) \quad (16)$$

This formalisation establishes a mathematically rigorous chain of uncertainty propagation ($\Sigma_{in} \rightarrow \Sigma_{latent} \rightarrow U_{epis} \rightarrow Tr$). It conclusively demonstrates that the supervisory trust evaluation is not heuristic, but is grounded in the analytical translation of physical bounds and data-driven variance into the final control policy. This formulation establishes specific real-time boundary conditions, governed by the transition sensitivity parameter γ : high physical certainty regime - when $uncert_{phys} \rightarrow 0$, the coefficient $\lambda(t)$ approaches 0. The HDT relies on f_{phys} , ensuring thermodynamic compliance and preventing neural network overfitting to sensor noise; high uncertainty regime - when $uncert_{phys} > 0.6$ (e.g., degradation deviates from analytical equations), $\lambda(t)$ approaches 1. The architecture shifts priority to f_{ML} to extract empirical patterns from telemetry.




Fig. 3. Dependence of the adaptive coefficient on physical model uncertainty. Epistemic uncertainty (U_{epis})

Analysis of the dynamic adaptation mechanism (Fig. 3) demonstrates that the exponential saturation function ensures a continuous redistribution of computational trust, eliminating gradient discontinuities. The transition profile is governed by the scaling parameter γ : a conservative shift $\gamma = 3.0$ prevents premature reaction to stochastic noise, a balanced shift $\gamma = 6.0$ optimally tracks continuous degradation, and an aggressive shift $\gamma = 12.0$ rapidly compensates for anomalies in highly deterministic subsystems. Asymptotically, this mechanism strictly bounds the system between pure physical dominance ($\lambda \rightarrow 0$) under normal variance and complete empirical delegation ($\lambda \rightarrow 1$) during severe non-linear wear, preventing data-driven hallucinations under standard conditions. This temporal dynamics and the corresponding response of the hybrid model are presented in Fig. 4.




Fig. 4. Temporal dynamics of the adaptive fusion mechanism and model performance under unmodeled degradation

The real-time temporal dynamics of this adaptation are illustrated in Fig. 4. Under nominal conditions ($t < 40$), the architecture securely relies on the deterministic physical baseline ($\lambda \rightarrow 0$). However, as shown in Fig. 4(a), the onset of an unmodeled non-linear degradation process triggers an exponential rise in epistemic uncertainty $U(t)$, prompting an instantaneous, proportional shift of computational trust ($\lambda \rightarrow 1$) toward the data-driven corrector. The practical consequence of this dynamic routing is demonstrated in Fig. 4(b): while the rigid physical model (Baseline) experiences unbounded error growth due to missing physical knowledge, the HDT's adaptive fusion mechanism effectively compensates for these limitations, stabilizing the prediction error.

7. Formal integration of XAI into the supervisory control loop. In conventional marine SPP, XAI is a post-hoc diagnostic layer for retrospective visuals. The HDT makes explainability intrinsic. The vector e_t is not secondary but a constitutive element of the latent manifold, derived from aligning empirical data with physical priors. It modulates the supervisory forward pass and serves as the primary metric for trust and autonomous routing - making control actions and their justifications concurrent and inseparable. Unlike passive post-hoc XAI, the HDT integrates e_t into the closed-loop policy via a trust evaluation module (Figure 5).

Let $u_t^{auto} = \pi_{nom}(\hat{x}_t)$ be the nominal autonomous control decision proposed by the HDT based on the predicted state. Simultaneously, the XAI layer maps the latent space representations into a human-interpretable explanation vector $e_t = \Psi_{XAI}(\hat{x}_t, \theta)$. Limitations of correlational XAI methods (SHAP/LIME). The application of post-hoc explainability algorithms (such as SHAP or LIME) in cyber-physical systems presents a known methodological issue. These approaches rely on correlating input perturbations with model outputs without accounting for the physical laws governing the asset. Consequently, the generated explanations may be statistically significant yet contradict thermodynamic principles. To address this logical inconsistency, the

proposed HDT architecture does not employ external correlational XAI algorithms. Instead, interpretability is achieved intrinsically: the explanation vector e_t is extracted from the attention mechanism weights, which operate strictly inside the physics-constrained latent manifold. Because the latent space is regularised by the differential residual LPDE, the resulting explanations are based on causal relationships and physical invariants, rather than statistical artefacts within the training data.




Fig. 5. Formal integration of XAI into the supervisory control (trust-aware control)

Instead of executing $T_{u_t}^{auto}$ blindly, the system computes a dynamic trust-aware modifier $T_t \in [0, 1]$:

$$T_t = \Gamma(e_t, L_{PDE}, U_{epis})$$

where Γ is a continuous evaluation function that inversely correlates with the epistemic uncertainty (U_{epis}) and the physical residual error (L_{PDE}), contextualized by the specific sub-system highlighted in the explanation vector e_t .

The final control action executed by the vessel is mathematically defined as a trust-aware policy π :

$$u_t = \pi \begin{cases} u_t^{auto}, & \text{if } T_t \geq \tau_{trust} \\ u_t^{corr}, & \text{if } T_t < \tau_{trust} \end{cases} \quad (17)$$

where τ_{trust} is the strict safety threshold mandated by maritime regulations.

Supervisory override logic - when the trust modifier drops below the safety threshold ($T_t < \tau_{trust}$) - indicating that the HDT faces severe non-linear degradation or conflicting sensor data the autonomous execution is suspended. The explanation vector e_t and the state forecast $\hat{x}_t$ are immediately transmitted to the SCC. The human supervisor utilizes this explicit interpretability to resolve the anomaly and issues a manual control correction u_t^{corr} .

Through this formalization, XAI transitions from a passive analytical tool into an active, mathematically rigorous component of the control loop, directly dictating the boundary between autonomous operation and required HIL intervention.

The architecture depicted in Figure 5 formalises the transition from a traditionally passive explainable AI paradigm - typically limited to post-hoc visualization - towards active supervisory, or trust-aware, control.

Within the proposed topology, the explanation vector e_t is not merely relayed to the operator; rather, it functions as an active parameter within the feedback loop through the synthesis of a dynamic trust modifier, T_t . This modifier serves as a reliability metric for the nominal autonomous decision u_t^{auto} , explicitly accounting for both epistemic uncertainty and physical residuals. By establishing a strict safety threshold, τ_{trust} , the system enforces an algorithmic override logic. Provided $T_t \geq \tau_{trust}$, the autonomous control policy is executed seamlessly. Conversely, should confidence degrade ($T_t < \tau_{trust}$), the autonomous loop is immediately suspended. Under these circumstances, the XAI layer provisions the shore-based operator (acting as the HIL) with the interpretable basis e_t . This significantly mitigates the cognitive burden required to formulate an appropriate manual correction, u_t^{corr} . Ultimately, the XAI module strictly delineates the operational boundaries within which autonomous navigation remains viable.

8. Real-time constraints and computational complexity. The assertion of 'real-time' operational capability within the proposed HDT necessitates a rigorous definition of computational boundaries. To guarantee deterministic behaviour in the supervisory control loop, the system's execution pipeline is strictly partitioned into a synchronous hard real-time edge task and an asynchronous soft real-time cloud thread.

Sampling frequency and inference cycle:

The onboard programmable logic controllers ingest multidimensional telemetry at a fixed sampling frequency of $f_s = 50$ Hz, establishing a strict operational cycle time of $T_{cycle} = 20$ ms. Within this temporal window, the system must acquire sensor data, perform latent state inference ($\hat{x}_t$), compute the explainability vector (e_t), and evaluate the trust modifier (T_t).

Computational budget and worst-case execution time (WCET):

A critical bottleneck in conventional XAI is the prohibitive computational cost of post-hoc perturbation methods (e.g., SHAP or LIME). To satisfy the HRT constraint, our HDT architecture bypasses these external methods by extracting intrinsic explanations directly from the cognitive-physical latent space (via attention weights). This approach ensures that the generation of e_t introduces an $O(1)$ computational overhead. As will be empirically demonstrated in the subsequent performance analysis, the WCET for the complete predictive and trust-evaluation forward pass on the edge controller is strictly bounded $\alpha_t \approx 15.2$ ms. This preserves a mathematically guaranteed slack time of ≈ 4.8 ms per cycle, strictly preventing control loop starvation.

Asynchronous execution and soft real-time updates:

Tasks that violate the 20 ms computational budget are deliberately decoupled from the critical control loop. The rendering of high-resolution XAI visualisations for the human operator, along with the federated aggregation of neural weights across the fleet (which introduces predictable latency at large scales), are scheduled as asynchronous background processes. Consequently, any

network latency or cloud-side computational spikes do not block the edge-level inference, ensuring that the vessel's physical safety mechanisms remain completely deterministic and uncompromised.

To mathematically substantiate the feasibility of deploying the proposed HDT architecture directly onto edge devices (onboard PLCs), a rigorous analysis of computational complexity is required. Table 2 presents the asymptotic estimates of time and space complexity, alongside empirical latency measurements for each foundational system module. Let us define the following dimensional notations: N denotes the number of input telemetry features (sensors), L represents the sliding window size, H is the dimensionality of the hidden latent space, V stands for the total number of vessels in the global fleet, and $\|W\|$ indicates the total number of neural network parameters. Latency consistency - the empirical module latencies (5.2 ms + 8.0 ms + 2.0 ms = 15.2 ms) directly align with the previously established worst-case execution time (WCET) limit.

Table 2

Computational complexity and resource allocation of HDT modules

Module / component	Execution environment	Time complexity (Big-O)	Memory complexity (Big-O)	Empirical latency / update rate
Physics-informed latent encoder	Onboard (edge)	$O(L \cdot H^2)$	$O(L \cdot H)$	≈ 5.2 ms
Cognitive state module	Onboard (edge)	$O(L^2 \cdot H)$	$O(L^2 + L \cdot H)$	≈ 8.0 ms
Trust evaluation & dynamic fusion (λ_t)	Onboard (edge)	$O(H)$	$O(H)$	≈ 2.0 ms
Total autonomous control loop	Onboard (edge)	$O(L^2 \cdot H + L \cdot H^2)$	$O(L^2 \cdot H + L \cdot H^2)$	≈ 15.2 ms (WCET)
Federated fleet aggregation	Shore cloud (central)	$O(V \cdot \ W\)$	$O(V \cdot \ W\)$	~ 20.0 s (Asynchronous)

The trust evaluation module employs exclusively element-wise operations on the latent vector, yielding a complexity of $O(H)$. This validates the previously claimed $O(1)$ algorithmic overhead for XAI integration compared to standard inference algorithms. This analysis conclusively demonstrates that the proposed HDT architecture is not only mathematically robust but also algorithmically optimised for deterministic operation under the severe resource constraints typical of marine cyber-physical systems.

9. System evaluation: baseline comparison and architectural ablation. The proposed HDT architecture is evaluated against classical modelling paradigms and subjected to an architectural ablation study.

Key modules were systematically removed from the complete pipeline to evaluate the impact of their absence on predictive accuracy (MAE), inference latency (WCET), and control loop safety. The results are detailed in Table 3.

Table 3

**Impact of architectural component exclusion
on system performance**

Architectural configuration	Accuracy (MAE)	XAI Role / Trust	Latency (WCET)	Control loop safety
Without PINN (Data-driven only)	9.2 h (high variance)	Not applicable	≈ 10.0 ms	Critical (prone to physical hallucinations)
Without XAI (Black-box)	7.8 h	None	≈ 13.0 ms	Low (unable to provision context to operator)
Without trust model	6.1 h	Passive (visualisation only)	≈ 13.2 ms	Vulnerable (blind execution of commands)
Full HDT architecture	6.1 h	Active (in-the-loop)	≈ 15.2 ms	High (algorithmic HITL routing)

Ablation tests confirm each node is required. Without the physical regulariser (PINN), inference accelerates slightly but accuracy degrades severely due to lost extrapolation in non-linear wear. Without the cognitive attention layer (XAI), the system becomes a black box and prevents safe human intervention during anomalies. Without the trust model (with XAI retained), prediction error stays at its minimum (6.1 h), yet XAI serves only as passive visualization; the controller has no mathematical criterion to suspend execution and transfer control to shore, so it blindly follows the NN. The full system balances accuracy, interpretability, and operational safety.

10. Physics-informed learning via soft constraints and differential regularization. To constrain the neural network component within thermodynamically valid limits, especially in regimes where the empirical data is noisy and the adaptive coefficient shifts towards data-driven dominance ($\lambda(t) \rightarrow 1$), an advanced Physics-Informed Learning (PIL) protocol is applied.

In Section 2, the foundational PINN architecture was established using the differential residual LPDE. In this expanded HDT, the neural network training is formulated as a multi-objective constrained optimization problem. The objective function L_{total} is expanded to simultaneously minimize the empirical prediction error, strictly enforce thermodynamic conservation laws (PDEs), and penalize violations of physical boundary operating limits:

$$L_{total}(\theta) = L_{data}(\theta) + \lambda_{PDE} L_{PDE}(\theta_{phys}) + \alpha \cdot L_{physbounds}(\theta_{ML}) + \beta \cdot R(\theta), \quad (18)$$

where L_{data} is the empirical MSE of the predictions; L_{PDE} is the differential physical residual strictly constraining the baseline latent space Z_{phys} ; $R(\theta)$ is the L2 weight regularization term; λ_{PDE} , α , β are scaling hyperparameters.

The critical addition in this section is the $L_{physbounds}$ component, which acts as a safety envelope for the empirical branch (Z_{data}). Instead of applying hard clipping (which halts gradient flow during backpropagation and causes topological dead zones), the HDT utilizes a soft constraint mechanism based on the rectified linear unit ($ReLU$) function. For a predicted state variable $\hat{x}_k$ constrained by absolute physical minimum (x_k^{min}) and maximum (x_k^{max}) structural limits (e.g., maximum allowable pressure before valve rupture), the penalty is computed as:

$$L_{physbounds} = \sum_k ([ReLU(\hat{x}_k - x_k^{max})]^2) + [ReLU(x_k^{min} - \hat{x}_k)]^2 \quad (19)$$

The $ReLU$ function ($ReLU(z) = \max(0, z)$) yields a zero penalty when predictions remain within the valid operational envelope ($x_k^{min} \leq \hat{x}_k \leq x_k^{max}$). In this state, optimization is driven solely by L_{data} . Extrapolation beyond physical boundaries activates the $ReLU$ function, generating a quadratic penalty proportional to the violation magnitude. This induces a gradient correction during backpropagation, returning the network weights to the valid hypothesis space. This integration ensures that the hybrid forecast satisfies thermodynamic invariants regardless of data scarcity or unmodeled degradation modes.

11. Theoretical justification and formal limits of applicability. The integration of data-driven and physical models via the adaptive coefficient $\lambda(t)$ and soft constraints relies on established principles of uncertainty-aware model fusion [9, 10]. Unlike static weighting, which is highly vulnerable to OOD data, linking $\lambda(t)$ to the physical model's epistemic uncertainty mathematically aligns with adaptive loss balancing techniques in PINNs [14]. This dynamically prevents empirical data loss from overriding physical priors during severe degradation events. Furthermore, utilizing a differentiable soft constraint ($ReLU$ penalty) rather than hard thermodynamic boundaries allows the gradient flow to smoothly guide the neural network back into the feasible physical manifold, avoiding vanishing gradient problems and topological discontinuities during backpropagation.

However, the predictive robustness of the presented HDT architecture is theoretically bounded by the assumption of continuous operational state transitions. Specifically, the applicability is limited by the following conditions: covered failure classes- the architecture successfully predicts progressive, continuous degradation modes (e.g., thermo-mechanical fatigue,

filter fouling, gradual hydraulic leakage) where the underlying physical equations remain structurally relevant; uncovered failure classes - the model fundamentally cannot predict instantaneous catastrophic failures (e.g., explosive ruptures, sudden short circuits) that violate the continuity assumptions of the applied differential equations; data integrity condition - the adaptive shift towards the ML component ($\lambda(t) \rightarrow 1$) assumes that high physical uncertainty is caused by actual system degradation. In cases of a complete simultaneous blackout of the sensor array, the neural network requires separate robustification mechanisms.

12. Computational scalability and fleet-level extrapolation.

Implementing the HDT across a maritime fleet (N_{ships}) using a centralized cloud paradigm exposes a severe communication bottleneck, scaling as:

$$C_{cloud} = O(N_{ships} \cdot T \cdot d), \quad (20)$$

where T is the sequence length and d is feature dimensionality. To overcome this, the HDT's highly reduced parameter footprint ($\approx 1.1 \times 10^5$ weights, enabled by the physical inductive bias) natively supports distributed federated learning. By computing the forward pass and loss locally at the Edge, only differential weight updates (ΔW) are transmitted. This decouples communication complexity from raw telemetry volume:

$$C_{comm} = O(N_{ships} \cdot P), \quad (21)$$

This mathematical reduction strictly bounds the transmission payload to ≈ 12 MB per vessel daily, ensuring that the HDT architecture scales flawlessly to a global fleet of 1,000 vessels without congesting the shore-based network infrastructure or compromising data privacy.

While the proposed HDT architecture demonstrates high computational efficiency on a single edge node, industrial implementation requires evaluating its scalability across large-scale maritime fleets. Let N_{ships} represent the total number of vessels in the fleet. The computational complexity of the HDT inference on a single vessel is deterministically bounded by the spatial and temporal dimensions of the input matrix. Let T be the temporal sequence length (observation window k) and d be the feature dimension (number of sensor channels). The local computational complexity per inference step scales as $O(T \cdot d + P)$, where P is the number of trainable parameters in the neural network ($\approx 1.1 \times 10^5$). As proven by the HIL experiments, this local complexity is easily resolved by the onboard ARM Cortex processor within 24 ms. However, attempting to scale this architecture using a traditional centralized cloud paradigm (streaming all raw telemetry to a shore center for processing) exposes a fundamental scalability bottleneck. In a centralized topology, the required cloud computational bandwidth and network data transfer scale linearly with the fleet size:

$$C_{cloud} = O(N_{ships} \cdot T \cdot d) \quad (22)$$

The extreme neural parameter reduction achieved by the physical inductive bias unlocks a critical architectural implication for fleet-level deployment. To overcome the $O(N_{ships})$ scaling barrier inherent to centralized cloud paradigms, this highly lightweight footprint acts as a natural enabler for distributed federated learning. By executing the inference and loss calculation (L_{total}) locally at the Edge, the massive raw telemetry matrices ($X_{t-k:t}$) never leave the vessel. Instead of streaming gigabytes of high-frequency sensor data, the communication protocol is restricted to transmitting only the differential weight updates (ΔW) of the compact neural network during asynchronous training cycles.

13. Experimental setup, implementation details, and HIL configuration. To validate the proposed HDT architecture, replicate the operational dynamics of a Wärtsilä 9L50DF dual-fuel marine engine, and ensure strict computational reproducibility, this section details the hardware infrastructure, software frameworks, dataset statistics, and network hyperparameters utilized during the evaluation phase.

Hardware and software environment

The computational pipeline was strictly partitioned to reflect real-world marine deployment constraints:

cloud training and federated aggregation: the offline training phases were executed on a centralised cloud cluster equipped with NVIDIA A100 Tensor Core GPUs (80 GB VRAM) running CUDA 12.1. The software pipeline was implemented using the PyTorch 2.1 framework on an Ubuntu 22.04 LTS environment;

HIL simulation - a real-time simulation host workstation (Intel Core i9-13900K, 24 cores, 128 GB DDR5 RAM) coupled with an NVIDIA RTX 4090 GPU (24 GB VRAM) was utilized to ensure deterministic execution of high-frequency thermodynamic models and parallel inference of the CNN-LSTM layers;

edge inference validation - to physically simulate the computational and memory constraints of marine programmable logic controllers (PLCs), the edge module was deployed on an industrial-grade edge computing node (NVIDIA Jetson AGX Orin);

data acquisition (DAQ) - an NI PXIe-6363 multifunction I/O module was utilized for interfacing physical sensor emulators with the digital core.

Dataset statistics and sensor pipeline

The evaluation was conducted using a multivariate time-series dataset comprising $N = 24$ distinct sensory channels representative of marine propulsion telemetry (e.g., main engine RPM, exhaust gas temperatures, scavenging air pressures, and bearing vibration velocities).

Sampling and protocols - the internal physics-based solver operates at a high frequency of 1 kHz. For the predictive ML core, the raw telemetry was sampled at $f_s = 50$ Hz (generating approximately 4.3×10^6 discrete data points per operational profile) and dynamically decimated to 10 Hz during specific tests to match standard industrial IoT gateway transmission limits. Data

was streamed via Modbus TCP/IP and CAN bus (NMEA 2000).

Partitioning and noise injection - the dataset was partitioned into a 70% training set (nominal operation and linear wear), a 10% validation set, and a 20% testing set explicitly engineered to contain OOD non-linear degradation events. To rigorously test algorithmic robustness, stochastic sensor noise (zero-mean Gaussian distribution) with variance σ^2 ranging from 5% to 30% of the nominal signal amplitude was artificially injected into the test streams.

Network optimisation and hyperparameters

The neural network weights were optimised using the Adam W optimiser. This optimiser was selected for its decoupled weight decay regularisation, which is critical for preventing early overfitting on stochastic sensory noise. The training procedure was strictly formalised as follows: learning rate dynamics - the initial learning rate was set to $\eta_{init} = 1 \times 10^{-3}$ and modulated by a cosine annealing scheduler, gradually decaying to a minimum of $\eta_{init} = 1 \times 10^{-5}$; regularization - weight decay was set to 1×10^{-2} . The L2 weight regularisation factor was configured at $\beta = 1 \times 10^{-4}$ to constrain parameter magnitude; training protocol - the model was trained for a maximum of 500 epochs. To mitigate overfitting, an early stopping protocol halted training if the validation loss did not improve for 50 consecutive epochs. The batch size was set to 128 for cloud training and scaled down to 32 to ensure stable gradient updates within Edge memory limits; architecture dimensions - latent space dimensionality (H) was set to 64; attention heads (d_k) to 4; and the sliding window size (L) to 50 time steps; loss function balancing - in the multi-objective loss function, the physical boundary penalty weight was empirically set to $\alpha = 0.5$, providing sufficient gradient penalty to enforce thermodynamic compliance. The physical PDE residual penalty (L_{PDE}) was dynamically scaled via the adaptive coefficient λ_t .

14. Empirical verification of the hybrid architecture. Dataset and experimental setup: to quantitatively validate the theoretical premises of dynamic trust redistribution and soft constraints, an empirical verification was conducted using a HIL test bench. The dataset comprises $N = 4.5 \times 10^6$ synchronized telemetry samples, representing continuous SPP operation over multiple degradation cycles. The raw multi-modal streams (vibroacoustic data at 100 Hz, thermodynamics at 10 Hz) were aggregated and downsampled to a unified processing rate of 1 Hz to match the inference cycle of the predictive core. To ensure robust evaluation and prevent model overfitting to specific operational profiles, the dataset was partitioned using a stratified 5-fold cross-validation scheme. The test datasets were injected with stochastic sensor noise (amplitudes of $\pm 5\%$, 10%, 15%, 20%, and 30%) to evaluate algorithmic robustness.

Comparative benchmark at nominal noise level (15%) - the proposed HDT was initially evaluated against five baseline models under the nominal severe noise condition (15% amplitude). To justify the deployment feasibility on marine Edge controllers, computational complexity metrics were recorded. The consolidated results are presented in Table 4.

The reduction in MAE achieved by the HDT compared to the standard PINN is statistically significant ($t(4) = 8.4$, $p = 0.001$). Classical methods (ARIMA, random forest) demonstrate high variance and low robustness ($< 60\%$) due to their inability to model complex spatio-temporal dependencies under severe sensor noise. Pure deep learning approaches, including the LSTM network and the state-of-the-art temporal fusion transformer (TFT), reduce the prediction error under nominal conditions but incur a massive computational cost. Specifically, while the TFT architecture captures complex long-term dependencies through self-attention mechanisms, its deployment requires $\sim 2.5 \times 10^6$ trainable parameters and 4.80 MFLOPS.

Table 4

Comparative performance and computational complexity metrics of T_{FHT} prediction models (5-fold CV)

Model architecture	MAE (T_{FHT} , h) $\pm 95\%$ CI	RMSE (T_{FHT} , h) $\pm 95\%$ CI	Robustness (R_{stab} , %)	Parameters	Complexity (MMACs)	Train time (s/epoch)
ARIMA (statistical baseline)	15.2 ± 2.8	20.1 ± 3.9	45	~ 101	< 0.01	1.2
Random forest (classical ML)	14.5 ± 2.6	18.7 ± 3.2	58	N/A	0.05	18.4
Physics-only (thermodynamic)	18.4 ± 2.1	25.2 ± 3.4	82	N/A	0.01	N/A
ML-only (pure LSTM network)	12.1 ± 1.8	16.8 ± 2.6	72	$\sim 8.5 \times 10^5$	1.20	145.0
Temporal fusion transformer (SOTA)	10.5 ± 1.5	13.2 ± 2.1	76	$\sim 2.5 \times 10^6$	4.80	310.0
Standard PINN (no adaptation)	10.2 ± 1.4	12.5 ± 1.9	85	$\sim 1.1 \times 10^5$	0.25	172.0
HDT (proposed model)	8.7 ± 0.9	10.9 ± 1.2	90	$\sim 1.1 \times 10^5$	0.28	185.0

This makes the SOTA transformer model highly inefficient for embedded marine Edge systems and prone to overfitting under severe, out-of-distribution stochastic noise ($R_{stab} = 76\%$). The integration of physical invariants into the neural network (standard PINN) drastically restricts the hypothesis space, enabling a reduction in model parameters by nearly 95% compared to the TFT (down to 1.1×10^5) while improving accuracy. The proposed HDT further refines this architecture by introducing the CSM meta-features and the adaptive coefficient $\lambda(t)$. While the tensor operations of the attention mechanism slightly increase the computational load (from 0.25 to 0.28 MFLOPS) compared to a standard PINN, this negligible overhead yields a critical improvement: a reduction of the MAE to 8.7 hours and a tightening of the 95% confidence interval (± 0.9 h). The HDT achieves the highest classification stability ($R_{stab} = 90\%$) during non-linear state transitions.

To empirically validate the theoretical error distribution formulated in Eq. (7) and to understand the internal mechanisms driving this high classification stability, we further analyze the contribution of different modeling components to the overall prediction accuracy. Fig. 6 presents the empirical decomposition of the total error into its physical, data-driven, and interaction components over time.




Fig. 6. Error decomposition of the HDT into physical, machine learning, and interaction components

Error profiles: physical error dominates initially (analytical baseline, low $\lambda(t)$); ML error rises under nonlinear degradation to compensate unmolded effects. Interaction term captures residual mismatches, indicating hybrid integration difficulty. Total error stays bounded; HDT shifts dominant error from analytical to adaptive corrections over time.

Analysis of noise sensitivity using the empirical data presented in Fig. 7 reveals the fundamental vulnerability of purely data-driven algorithms. Under low noise conditions (5%), the ML-only model performs exceptionally well (MAE = 6.4 h), nearly matching the hybrid approach. However, as the noise amplitude escalates to 20% and 30%, the ML-only model experiences a catastrophic collapse in accuracy (MAE jumps to 27.5 h). This occurs because the unconstrained neural network begins to overfit to the chaotic stochastic signals, embedding physical impossibility into its latent

space. Conversely, the physics-only model degrades linearly, as the noise directly corrupts its rigid input variables, but it avoids chaotic hallucinations. The proposed HDT demonstrates improved robustness under the evaluated operating conditions, maintaining a highly acceptable error margin (MAE = 14.2 h) even under extreme 30% noise - a state bordering on sensor failure. This resilience mathematically proves the effectiveness of the physical regularizer (L_{phys_bounds}) and the dynamic adaptation coefficient $\lambda(t)$. By acting as an active structural filter, the physics-informed loss function mathematically prohibits the network from fitting to the high-amplitude noise, forcing the predictions to remain within valid thermodynamic boundaries regardless of the incoming data quality.




Fig. 7. Degradation of prediction accuracy under increasing stochastic sensor noise

To isolate and numerically prove the impact of the CSM on the predictive performance of the neural network, a targeted ablation study was conducted. The full HDT architecture (with CSM modulation) was benchmarked against a degraded version of the same network where the meta-feature vector was disabled ($m_i = 0$). The evaluation focused on early fault detection metrics during the incubation phase of degradation. The results are summarized in Table 5.

Table 5

Ablation study - impact of CSM meta-feature modulation on predictive performance

Model configuration	MAE ($\hat{T}_{FHT}$, h)	Early detection latency (h)	False negative rate (%)	Inference time perstep (ms)
HDT without CSM ($m_i = 0$ /standard PINN)	10.2	14.2	2.5	18
Full HDT with CSM	8.7	3.5	1.2	24

The ablation data demonstrates that removing the cognitive modulation degrades the overall first hitting time ($\hat{T}_{FHT}$) prediction accuracy (MAE increases from 8.7 to 10.2 hours). However, the most critical vulnerability of the model lacking CSM integration lies in the temporal domain. Without the causal guidance of the meta-features, the network relies entirely on reactive telemetry, resulting in an early detection latency of 14.2 hours (the time elapsed between the actual initiation of a

soft fault and its statistical recognition by the neural network). The multiplicative modulation by the CSM ($\alpha_i(t) \cdot h_i(t)$) effectively serves as a proactive focal filter, reducing the detection latency by 75% (down to 3.5 hours). This rapid convergence confirms that the cognitive layer successfully directs the network's attention to latent failure signatures long before they become statistically dominant in the raw IoT streams.

Beyond predictive metrics, the real-time deployment of the HDT requires validation of its inference capabilities on target hardware. Table 6 summarizes the operational efficiency metrics obtained during HIL testing on a single-core ARM Cortex-A72 microprocessor (simulating a standard marine programmable logic controller).

Table 6

**Operational efficiency and resource utilization
for edge deployment**

Metric	Symbol	Unit	ML-only (LSTM)	HDT without CSM	Full HDT
Inference latency per step	t_{inf}	ms	145	18	24
Peak RAM utilization	Mempeak	MB	312	64	68
False negative rate	FNR	%	6.8	2.5	1.2

The constraints of marine edge computing strictly prohibit the use of the ML-only baseline, as its 145 ms latency and 312 MB memory footprint create a computational bottleneck. Conversely, the HDT requires only 24 ms per inference step and 68 MB of RAM. This confirms that the model can process continuous 10 Hz telemetry streams in hard real-time without dedicated GPU accelerators. The minor computational overhead incurred by the tensor operations of the attention mechanism (+6 ms per inference step compared to the model without CSM) is strictly within acceptable limits for marine applications. Ultimately, the suppression of false negatives ($FNR = 1.2\%$) ensures the reliable and timely activation of maintenance protocols. To evaluate the scalability of the proposed HDT architecture under fleet-level deployment conditions, Fig. 8 illustrates the dependence of key system metrics - including inference latency, memory utilization, and computational throughput - on the number of monitored vessels.

The dual-axis plot illustrates the scalability of the proposed HDT architecture when extending from single-vessel deployment to fleet-level federated learning. The solid black curve (left Y-axis) shows that the edge inference latency remains nearly constant at approximately 15.2 ms and stays below the adopted 20 ms hard real-time threshold regardless of fleet size, ensuring uninterrupted onboard diagnostic inference. In contrast, the dashed grey curve (right Y-axis) represents the federated aggregation time, which increases to approximately 20 s for a fleet of 1000 vessels. Since

federated aggregation is implemented as an asynchronous background process, this delay does not affect onboard real-time operation. Consequently, the proposed architecture supports fleet-scale collaborative learning while preserving deterministic onboard diagnostic performance.




Fig. 8. Scalability analysis - edge inference stability vs. federated cloud aggregation

Sensitivity analysis of the hybridization coefficient - to substantiate the necessity of the dynamic adaptation (equation), an analysis of the model's sensitivity to static values of the weighting coefficient λ was performed. Fig.9 illustrates the non-linear dependence of the prediction error on λ .




Fig. 9. Influence of the weighting coefficient on MAE prediction accuracy

Analysis of the ablated model (solid black curve) reveals a pronounced parabolic trade-off. Extreme static settings yield significant errors: for $\lambda < 0.3$, the physics-based model dominates, but it fails to sufficiently compensate for non-linear wear effects, resulting in a high MAE (≈ 9.7 h). Conversely, for $\lambda > 0.8$, the prediction becomes excessively sensitive to stochastic sensor noise (overfitting), which degrades the accuracy back to $MAE > 8.0$ h. The empirical minimum error for any fixed configuration strictly occurs in the range $\lambda \in [0.6, 0.7]$, achieving a best-case static MAE of 7.3 hours. However, as demonstrated by the green dashed line, the proposed architecture utilizing the real-time dynamic adaptation function $\lambda(t)$ fundamentally outperforms this mathematically optimal static baseline. By autonomously shifting computational trust based on instantaneous

epistemic uncertainty (U_{epis}), the dynamic model reduces the overall MAE to 6.1 hours. This 16.4% error reduction conclusively proves that a static coefficient cannot be optimal across all non-linear operational states, directly justifying the integration of the dynamic uncertainty-driven fusion node (Eq. 9) as a critical architectural requirement. To visually demonstrate the operational mechanics driving these metrics, Fig. 10 illustrates the time dynamics of the $\hat{T}_{FHT}$ prediction as the critical component approaches its failure threshold. Three temporal profiles show operational differences: physics-only model (dashed gray): systematically underestimates remaining useful life. Rigid equations miss nonlinear wear interactions. Conservative assumptions cause premature maintenance alerts and underused lifespan; ML-only model (dotted): high sensitivity to sensor noise produces strong prediction oscillations.




Fig. 10. Time dynamics of the $\hat{T}_{FHT}$ prediction

After actual failure at $t = 100$ h, the network continues to output positive residual life. This false optimism shows empirical models have no access to physical destruction limits; HDT (dash-dot black): adaptive $\lambda(t)$ keeps prediction close to actual T_{FHT} . It corrects analytical underestimation and suppresses noise via physical penalty. At the failure point, thermodynamic constraints force output to zero. Structural safety is preserved.

The empirical results confirm that the proposed HDT architecture resolves the fundamental trade-off between physical stability and empirical flexibility. By acting as an active structural filter, the integration of physical constraints mathematically guarantees that the generated prognostic evaluations are both highly accurate and thermodynamically viable, entirely eliminating the risk of blind data extrapolation. This provides a robust, verified foundation for the stochastic $\hat{T}_{FHT}$.

15. Operational and economic efficiency. Beyond predictive accuracy, the deployment of the HDT in a SCC requires evaluating computational latency and economic viability. Table 7 summarizes the operational metrics.

Table 7

Metrics of stability, computational, and operational efficiency

Metric	Symbol	Unit	Physics-only	ML-only	HDT
Average prediction time	t_{proc}	ms	58	102	128
Model update time	t_{upd}	s	240	180	120
Noise robustness ($\pm 15\%$)	R_{stab}	%	80	74	89
Estimated economic effect	ΔC_{saved}	%/year	0 (baseline)	+14	+27

While the real-time inference step takes 24 ms (Table 6), the comprehensive hybrid model inference for $\hat{T}_{FHT}$ forecasting takes 128 ms (parallel physics), which remains within acceptable operational limits. Model update time drops to 120 s because physical bias narrows the search space and accelerates convergence. False alarms are suppressed, resource underestimation eliminated, yielding an estimated 27% annual cost reduction (ΔC_{saved}). This estimate is engineering-level, based on reduced false alarms, maintenance, and downtime in marine power systems; it indicates magnitude, not a vessel-specific life-cycle cost analysis.

16. Conclusions. Based on the theoretical development and HIL simulations, this study validates a HDT architecture that successfully integrates a deterministic physics-based model (PBM), a CNN-LSTM network, and a CSM. By formally linking the dynamic adaptation coefficient $\lambda(t)$ to the physical model's epistemic uncertainty, the system enables a mathematically rigorous transition from physics-dominated inference to data-driven forecasting. Empirical testing across ship power plant subsystems demonstrated that the HDT reduces the remaining useful life prediction error (MAE) by 32% – 36% compared to standalone physical models, and by up to 19% compared to pure deep learning approaches, while achieving an F1-score of up to 0.89 for fault classification. These results provide a substantial contribution to the evolving paradigm of integrated DT architectures [16, 17]. Unlike contemporary hybrid frameworks relying on static or heuristic coupling [18], the continuous redistribution of computational trust resolves multi-objective balancing and gradient pathologies [14, 15]. Functioning as an analytical "inertial flywheel" via differentiable physical regularizers, the HDT sustains diagnostic consistency even under 30% stochastic sensor noise - conditions where pure machine learning models suffer from catastrophic overfitting [19, 20]. Furthermore, addressing the critical need for explainable operation in complex engineering environments [21], embedding the CSM directly into the predictive loop extends the current state-of-the-art [22] by formalizing explainability as an intrinsic cyber-physical property. This ensures the

transition between physical laws and empirical learning is causally transparent, effectively isolating root causes without relying on external correlational algorithms.

This study developed a novel adaptive HDT architecture for SPP predictive diagnostics, integrating deterministic physics-based models (PBM), CNN-LSTM networks, and CSMs). Unlike conventional hybrid frameworks relying on static coupling, our approach dynamically redistributes computational trust based on instantaneous epistemic uncertainty, ensuring a mathematically rigorous transition from analytical inference to empirical data-driven forecasting while preserving thermodynamic consistency.

Empirical HIL validation demonstrated that the HDT significantly enhances predictive robustness. The architecture reduced the T_{FHT} prediction error (MAE) by 32–36% relative to standalone physical models and by up to 19% compared to pure deep learning benchmarks. Furthermore, the CSM-driven multiplicative attention mechanism achieved an F1-score of 0.89 for fault classification, maintaining operational stability even under stochastic sensor noise (up to 30%), effectively suppressing catastrophic overfitting.

Beyond diagnostic accuracy, the proposed framework addresses the critical requirement for explainable, edge-deployable intelligence. The system provides intrinsic XAI via causal alignment within the physics-constrained latent manifold, eliminating the need for correlational post-hoc methods. Computationally, the architecture supports deterministic hard real-time inference on standard marine Edge PLCs, while asynchronous federated learning enables fleet-level scalability. Ultimately, this approach—yielding an estimated 27% reduction in annual operating expenditures - provides a verified, transparent, and economically viable foundation for next-generation maritime prognostic maintenance.

Future research will focus on the automated calibration of physical priors, adaptive ontology generation, and the extension of the proposed framework toward the reliable detection of unseen catastrophic failure modes under out-of-distribution operating conditions.

References

1. Tao F., Zhang M., Liu Y., Nee A. Y. C. Digital Twin Driven Smart Manufacturing. Academic Press. 2019. DOI: <https://doi.org/10.1016/C2018-0-02361-5>.
2. Fuller A., Fan Z., Day C., Barlow C. Digital twin: Enabling technologies, challenges and open research. IEEE Access. 2020. DOI: <https://doi.org/10.1109/ACCESS.2020.2998358>.
3. Jones D., Snider C., Nassehi A., Yon J., Hicks B. Characterising the Digital Twin: A systematic literature review. CIRP Journal of Manufacturing Science and Technology. 2020. Vol. 29. P. 36–52. DOI: <https://doi.org/10.1016/j.cirpj.2020.02.002>.
4. Rudnichenko N., Vychuzhanin V. Complex Technical System Condition Diagnostics and Prediction Computerization. CEUR Workshop Proceedings. 2020. Vol. 2608. P. 42–56. DOI: <https://doi.org/10.32782/cmcs/2608-4>.
5. Vychuzhanin V. V., Vychuzhanin A. V. A hybrid model for evaluating the accuracy of failure forecasts in ship power plants. Bulletin of Cherkasy State Technological University. 2026. Vol. 31. No. 1. P. 44–59. DOI: <https://doi.org/10.62660/bcstu/1.2026.44>.
6. Wang J., Ma Y., Zhang L., Gao R. X., Wu D. Deep learning for smart manufacturing: Methods and applications. Journal of Manufacturing Systems. 2020. Vol. 48. P. 144–156. DOI: <https://doi.org/10.1016/j.jmsy.2018.01.003>.
7. Khan S., Yairi T. A review on the application of deep learning in system health management. Mechanical Systems and Signal Processing. 2018. Vol. 107. P. 241–265.
8. Lu Y., Liu C., Kevin I., Wang K., Huang H., Xu X. Digital Twin-driven smart manufacturing: Connotation, reference model, applications and research issues. Robotics and Computer-Integrated Manufacturing. 2020. Vol. 61. Article 101837. DOI: <https://doi.org/10.1016/j.rcim.2019.101837>.
9. Raissi M., Perdikaris P., Karniadakis G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics. 2019. Vol. 378. P. 686–707. DOI: <https://doi.org/10.1016/j.jcp.2018.10.045>.
10. Karniadakis G. E., Kevrekidis I. G., Lu L., Perdikaris P., Wang S., Yang L. Physics-informed machine learning. Nature Reviews Physics. 2021. Vol. 3. P. 422–440. DOI: <https://doi.org/10.1038/s42254-021-00314-5>.
11. Kosko B. Fuzzy Cognitive Maps. International Journal of Man-Machine Studies. 1986. Vol. 24. P. 65–75.
12. Papageorgiou E. I., Salmeron J. L. A review of fuzzy cognitive maps research during the last decade. IEEE Transactions on Fuzzy Systems. 2020. Vol. 28. No. 10. P. 2601–2614.
13. Vychuzhanin V., Gritsuk I., Rudnichenko N., Shibaeva N. Cognitive-impulse model for assessing complex technical systems survivability. CEUR Workshop Proceedings. 2020. Vol. 2711. P. 571–585.
14. Wang S., Teng Y., Perdikaris P. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. SIAM Journal on Scientific Computing. 2021. Vol. 43. No. 5. P. A3055–A3081. DOI: <https://doi.org/10.1137/20M1318043>.
15. Xiang Z., Peng W., Zheng X., Zhao X., Yao W. Self-adaptive loss balanced Physics-informed neural networks. Neurocomputing. 2022. Vol. 496. P. 11–34. DOI: <https://doi.org/10.1016/j.neucom.2022.05.015>.
16. Liu M., Fang S., Dong H., Xu C. Review of digital twin: Concepts, technologies, and industrial applications. Journal of Manufacturing Systems. 2021. Vol. 58. P. 346–361. DOI: <https://doi.org/10.1016/j.jmsy.2020.06.017>.
17. Barricelli B. R., Casiraghi E., Fogli D. A survey on digital twin: Definitions, characteristics, applications, and design implications. IEEE Access. 2019. Vol. 7. P. 167653–167671.
18. Willard J., Jia X., Xu S., Steinbach M., Kumar V. Integrating physics-based modeling with machine learning: A survey. ACM Computing Surveys. 2022. Vol. 55. No. 3. Article 66. DOI: <https://doi.org/10.1145/3550198>.
19. Zhang W., Yang D., Wang H. Data-driven methods for predictive maintenance of industrial equipment: A

- survey. IEEE Systems Journal. 2021. Vol. 15. No. 1. P. 141–152. DOI: <https://doi.org/10.1109/JSYST.2019.2962002>.
20. Karpatne A., Atluri G., Faghmous J. H., et al. Theory-guided data science: A new paradigm for scientific discovery. IEEE Transactions on Knowledge and Data Engineering. 2017. Vol. 29. No. 10. P. 2318–2331. DOI: <https://doi.org/10.1109/TKDE.2017.2720168>.
 21. Kritzing W., Karner M., Traar G., Henjes J., Sihn W. Digital Twin in manufacturing: A categorical literature review and classification. IFAC-PapersOnLine. 2018. Vol. 51. No. 11. P. 1016–1022. DOI: <https://doi.org/10.1016/j.ifacol.2018.08.474>.
 22. Rai R., Tiwari M. K., Ivanov D. Machine learning in manufacturing and Industry 4.0 applications. International Journal of Production Research. 2021. Vol. 59. No. 16. P. 4773–4778. DOI: <https://doi.org/10.1080/00207543.2021.1956675>.

Вичужанін В.В., Вичужанін О.В. Адаптивна гібридна архітектура цифрових двійників з урахуванням невизначеності для прогнозування діагностики суднових енергетичних установок

Надійшло прогнозування технічного стану та залишкового ресурсу суднових енергетичних установок (СЕУ) залишається складним завданням через безперервну зміну умов експлуатації, суттєву нелінійність механізмів деградації та неминучу невизначеність, що супроводжує вимірювання датчиків. Попри те, що гібридні цифрові двійники (ГЦД) значно розширили можливості прогнозного обслуговування, багато наявних реалізацій усе ще використовують фіксовані схеми взаємодії між фізичними та дата-орієнтованими моделями, що знижує їхню адаптивність у випадках, коли надійність окремих джерел інформації змінюється під час роботи. У дослідженні розроблено архітектуру ГЦД, яка об'єднує детерміністичну фізичну модель, прогнозний модуль на основі CNN–LSTM та когнітивну імітаційну модель (KIM) у єдиній діагностичній структурі. Замість підтримки фіксованого балансу між аналітичними та емпіричними моделями, архітектура безперервно коригує їхній внесок відповідно до оціненої невизначеності фізичної моделі. Це дає змогу зберігати фізичну узгодженість, коли аналітичні знання є надійними, і водночас дозволяє дата-орієнтованому прогнозуванню домінувати у разі зростання невизначеності моделі. Когнітивний імітаційний компонент доповнює чисельне прогнозування шляхом явного представлення причинно-наслідкових зв'язків між процесами деградації та діагностичними висновками, забезпечуючи інженерні пояснення, які важко отримати зі звичайних глибоких нейронних мереж.

Ефективність розробленої структури досліджено за допомогою моделювання із апаратним забезпеченням у контурі (HIL), яке відтворювало репрезентативні режими роботи підсистем СЕУ, включаючи прогресуючу деградацію та стохастичні завади датчиків. За цих умов гібридна архітектура зменшила похибку прогнозування залишкового ресурсу на 32–36% порівняно з автономними фізичними моделями та до 19% порівняно зі звичайними рішеннями на основі глибокого навчання. Класифікація відмов досягла F1-показника 0,89, зберігаючи стабільну роботу за рівнів шуму датчиків до 30%. Виконання в реальному часі з тривалістю менше 20 мс задовольняє вимоги бортових автоматизованих систем, моді як асинхронне федеративне навчання забезпечує масштабоване розгортання в розподілених флотах без збільшення бортового обчислювального навантаження. Поєднання адаптивного злиття моделей з урахуванням невизначеності, фізично-інформованого навчання та когнітивного міркування розширює можливості сучасних технологій ГЦД порівняно зі звичайними гібридними діагностичними структурами. Окрім підвищення точності прогнозування та стійкості, архітектура покращує прозорість рішень щодо технічного обслуговування та демонструє практичну придатність для інтелектуальних СЕУ, з оціночним скороченням щорічних експлуатаційних витрат приблизно на 27%. Отримані результати свідчать, що розроблена структура є практичним кроком до надійного та пояснюваного прогнозного обслуговування суднових енергетичних систем наступного покоління.

Ключові слова: фізично-інформоване машинне навчання; оцінювання залишкового ресурсу; когнітивне імітаційне моделювання; пояснюваний штучний інтелект

Вичужанін Володимир Вікторович – д.т.н, професор, завідувач кафедри інформаційних технологій, національний університет «Одеська політехніка», Одеса, <https://orcid.org/0000-0002-6302-1832>
email: v.v.vychuzhanin@op.edu.ua.

Вичужанін Олексій Володимирович – доктор філософії, асистент, національний університет «Одеська політехніка», Одеса, <https://orcid.org/0000-0001-8779-2503>
email: v.v.vychuzhanin@op.edu.ua.

Дата першого надходження статті 02.07.2026.

Дата прийняття статті до друку після рецензування 26.08.2026.

Дата публікації 30.09.2026.



Стаття з відкритим доступом,
відповідно до умов ліцензії
[Creative Commons \(CC BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)