machines

MDPI

Article

Adaptive Digital Twin Framework for PMSM Thermal Safety Monitoring: Integrating Bayesian Self-Calibration with Hierarchical Physics-Aware Network

Jinqiu Gao [1], Junze Luo [2], Shicai Yin [3,*], Chao Gong [4], Saibo Wang [4] and Gerui Zhang [4]

[1] School of Water Resources and Hydro-Electric Engineering, Xi'an University of Technology, Xi'an 710048, China; gaojinqiu@xaut.edu.cn
[2] China Energy Engineering Group, Hunan Electric Power Design Institute Co., Ltd. (CEEC-HEPDI), Changsha 410007, China; carlfeynman@163.com
[3] School of Electrical Engineering and Automation, Nantong University, Nantong 226019, China
[4] School of Automation, Northwestern Polytechnical University, Xi'an 710072, China; chaogong@nwpu.edu.cn (C.G.); wangsaibo@mail.nwpu.edu.cn (S.W.); zhanggerui@mail.nwpu.edu.cn (G.Z.)
* Correspondence: yinsci@ntu.edu.cn

Abstract

To address the limitations of parameter drift in physical models and poor generalization in data-driven methods, this paper proposes a self-evolving digital twin framework for PMSM thermal safety. The framework integrates a dynamic-batch Bayesian calibration (DBBC) algorithm and a hierarchical physics-aware network (HPA-Net). First, the DBBC eliminates plant–model mismatch by robustly identifying stochastic parameters from operational data. Subsequently, the HPA-Net adopts a "physics-augmented" strategy, utilizing the calibrated physical model as a dynamic prior to directly infer high-fidelity temperature via a hierarchical training scheme. Furthermore, a real-time demagnetization safety margin (DSM) monitoring strategy is integrated to eliminate "false safe" zones. Experimental validation on a PMSM test bench confirms the superior performance of the proposed framework, which achieves a Root Mean Square Error (RMSE) of 0.919 °C for the stator winding and 1.603 °C for the permanent magnets. The proposed digital twin ensures robust thermal safety even under unseen operating conditions, transforming the monitoring system into a proactive safety guardian.

Keywords: digital twin; PMSM; thermal monitoring; Bayesian calibration; physics-aware machine learning

Check for updates

Academic Editors: Pengfei Liang, Shaogan Ye and Qun Chao

Received: 5 January 2026
Revised: 15 January 2026
Accepted: 21 January 2026
Published: 24 January 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

With the global consensus on carbon neutrality and sustainable development, the electrification of transportation and the expansion of renewable energy generation have become pivotal strategies for decarbonization [1–3]. In this context, permanent magnet synchronous machines (PMSMs) have emerged as the core energy conversion devices across various sectors. Due to their superior power density, high efficiency, and reliability, PMSMs are not only the dominant traction motors in electric vehicles (EVs) but also widely adopted as direct-drive generators in wind power turbines and key actuators in hydroelectric governor systems [4–6]. Whether driving a high-speed train or harvesting energy from turbulent winds, the operational reliability of PMSMs is critical for the safety and stability of the entire energy system [7–9].

However, these applications inevitably expose PMSMs to harsh and time-varying operating conditions [10,11]. For EVs, frequent acceleration and hill-climbing impose severe thermal overloads. For wind turbines, unpredictable wind speed fluctuations and grid faults can cause transient current surges. Similarly, hydro-generators often face long-term continuous operation under variable load demands. In all these scenarios, excessive heat accumulation is a common threat. Elevated temperatures accelerate the aging of stator winding insulation and, more critically, degrade the performance of neodymium iron boron (NdFeB) permanent magnets (PMs). As the remanence of magnets decreases linearly with temperature, overheating can trigger irreversible demagnetization (ID) faults, leading to efficiency drops, increased vibration, or even catastrophic system failures [12,13]. Therefore, accurate real-time monitoring of the motor's internal thermal state is a prerequisite for active safety protection, regardless of the application field.

However, the limitations of the sensing devices in existing systems still pose certain challenges to current monitoring technologies [14,15]. Direct temperature measurement of critical internal components (especially the rotating rotor) is technically challenging and cost-prohibitive due to the limitations of embedded sensors in high-voltage or enclosed environments [16,17]. Consequently, "soft-sensing" technology, which estimates internal temperatures based on accessible electrical signals, has become the standard solution. Existing approaches generally fall into model-based methods (e.g., simplified thermal networks, simplified TNs) and data-driven methods (e.g., deep neural networks) [18–22]. While simplified TNs offer physical interpretability, their accuracy relies heavily on precise parameter specification.

This reliance on parameters poses a significant challenge in practical engineering, particularly for long-lifespan assets like wind turbines or hydro-generators. During years of operation, thermal parameters (resistances and capacitances) drift significantly due to material aging, fouling of cooling channels, or varying ambient conditions [23–25]. Similarly, in mass-produced EVs, manufacturing tolerances lead to parameter dispersion. Such "parameter uncertainty" causes traditional TNs to deviate from reality over time. Although data-driven methods can capture complex non-linearities without physical parameters, they act as "black boxes" and often lack the robustness required for safety-critical energy infrastructure—a model trained on one wind site may fail when deployed to another [26,27].

To address the limitations of static physical models and pure data-driven algorithms, the Digital Twin (DT) offers a transformative solution. However, existing frameworks often struggle to simultaneously resolve the parameter mismatch caused by aging and the inherent limitations of simplified physical models [28,29]. Standard adaptive methods frequently attempt to force linear parameters to fit unmodeled non-linear dynamics (e.g., complex convection), resulting in physically biased estimates. Consequently, a critical gap remains for a systematic framework that can effectively fuse the interpretable physical baseline with the learning capacity of neural networks.

In response, this paper proposes a self-evolving digital twin framework for PMSM thermal safety monitoring, as illustrated in Figure 1, designed to be adaptable across different industrial applications. Specifically, to reconcile "model reliability" with "environmental adaptability," this study establishes a closed-loop evolution mechanism. It leverages physical laws to mitigate the uncertainty of pure data-driven methods while utilizing neural networks to capture complex non-linear dynamics. By synergizing DBBC for parameter self-identification and HPA-Net for direct high-fidelity inference, the framework evolves from a static model into a proactive dynamic guardian. The main contributions are summarized as follows:

- Adaptive physical twin construction: A dynamic-batch Bayesian calibration (DBBC) algorithm is proposed to enable the "self-calibration" of the physical twin (simplified

TN). By adaptively adjusting the sampling strategy, this algorithm robustly identifies uncertain thermal parameters from operational data, ensuring the model adapts to the specific conditions of the individual motor, whether it is in a vehicle or a wind turbine.

- Fidelity enhancement via HPA-Net: To bridge the "reality gap" caused by model simplifications, a hierarchical physics-aware network (HPA-Net) is developed. Through a hierarchical training strategy (collaborative physics embedding followed by reality-gap fine-tuning), the network internalizes the physical laws from the simulation domain and adapts to the reality gap in the experimental domain.
- Active safety margin monitoring: Moving beyond temperature estimation, a real-time demagnetization safety margin (DSM) monitoring strategy is introduced. By dynamically calculating the safety boundary of the magnets based on real-time temperature feedback, this method provides early warnings for potential thermal failures, enhancing the resilience of critical electric drive systems.

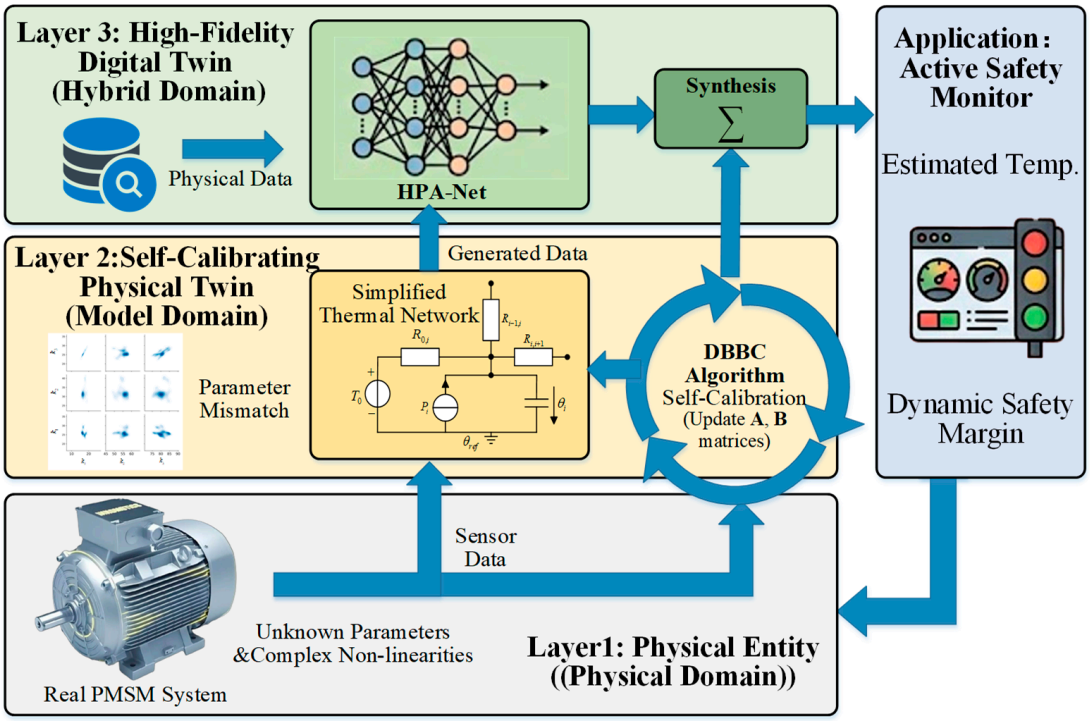

Figure 1. Schematic diagram of the proposed self-evolving digital twin framework.

The remainder of this paper is organized as follows: Section 2 establishes the physical twin model. Section 3 introduces the self-calibration via DBBC. Section 4 details the HPA-Net and its training strategy. Section 5 presents experimental validation and safety monitoring applications. Finally, Section 6 concludes the paper.

2. State-Space Modeling of the Reduced-Order Physical Twin

The construction of a reliable digital twin necessitates a robust physical baseline capable of capturing the dominant thermal dynamics of the PMSM. While finite element analysis (FEA) offers high spatial resolution, its computational prohibitiveness renders it unsuitable for real-time monitoring and digital twin applications. Consequently, this study adopts a TN approach. This section details the derivation of a reduced-order physical twin, focusing on the trade-off between model fidelity and computational efficiency.

2.1. Thermal Dynamics Formulation

The thermal behavior of a PMSM is a complex multi-physics phenomenon involving electromagnetic losses converted into heat, which is then dissipated through conduction, convection, and radiation. Based on the law of conservation of energy and Fourier's law of heat conduction, the thermal equilibrium for any given volumetric node i within the motor can be described by the following first-order differential equation [30]:

$$C_i \frac{d\vartheta_i(t)}{dt} = Q_{gen,i}(t) - Q_{out,i}(t) \tag{1}$$

where $\vartheta_i(t)$ represents the temperature rise in node i relative to the ambient environment, C_i denotes the thermal capacitance (heat capacity) of the node, $Q_{gen,i}(t)$ is the internal heat generation rate caused by electromagnetic losses, and $Q_{out,i}(t)$ represents the net heat transfer rate to adjacent nodes or the environment.

In Figure 2, substituting the conductive heat transfer equation, the heat exchange term can be expanded as a summation of fluxes driven by temperature gradients:

$$Q_{out,i}(t) = \sum_{j \in \Omega_i} \frac{\vartheta_i(t) - \vartheta_j(t)}{R_{ij}} \tag{2}$$

where Ω_i represents the set of neighboring nodes thermally coupled to node i, and R_{ij} denotes the equivalent thermal resistance between node i and node j. Combining Equations (1) and (2), the governing equation for the nodal temperature dynamics is derived as:

$$\frac{d\vartheta_i(t)}{dt} = \frac{1}{C_i} Q_{gen,i}(t) - \sum_{j \in \Omega_i} \frac{1}{R_{ij} C_i} (\vartheta_i(t) - \vartheta_j(t)) \tag{3}$$

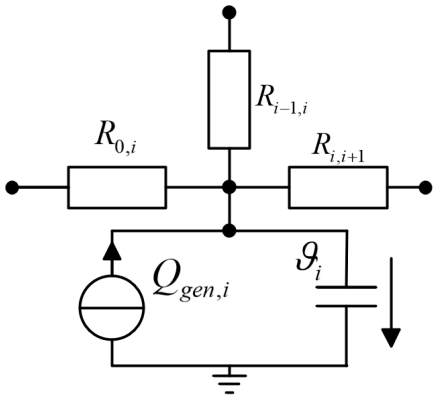

Figure 2. Equivalent thermal circuit diagram of a generic volumetric node i.

To construct a reduced-order model suitable for the digital twin, the motor geometry is discretized into a few critical lumped nodes: the stator tooth, the stator winding, the PM and etc. These nodes are selected because the winding insulation and PMs are the most thermally vulnerable components, and the stator tooth serves as a critical heat conduction bridge.

2.2. Continuous State-Space Representation

For the purpose of advanced control and observation, the coupled differential equations for the these nodes are reformulated into a standard continuous-time state-space model. Let the state vector $\mathbf{x}(t)$ be defined as the temperature rise vector:

$$\mathbf{x}(t) = [\vartheta_{st}(t), \vartheta_{wdg}(t), \vartheta_{pm}(t)]^T \tag{4}$$

The input vector $\mathbf{u}(t)$ incorporates the system's active heat sources and boundary conditions. In high-speed PMSMs, the heat generation is primarily distributed among the windings, the stator core, and the rotor magnets. To ensure the mathematical completeness of the state-space model, the input vector is defined to include the stator copper loss (Q_{cu}), the stator iron loss (Q_{iron}), the PM eddy current loss (Q_{pm}), and the ambient temperature (ϑ_{amb}):

$$\mathbf{u}(t) = [Q_{cu}(t), Q_{iron}(t), Q_{pm}(t), \vartheta_{amb}(t)]^T \tag{5}$$

Correspondingly, the input matrix $\mathbf{B}(\boldsymbol{\psi}) \in \mathbb{R}^{3\times 4}$ maps these power losses and boundary conditions to the temperature rate of change for each node. Specifically, Q_{cu} directly drives the winding temperature dynamics, Q_{iron} acts on the stator tooth node, while Q_{pm} serves as the direct heat source for the PM node.

Accordingly, the system dynamics can be expressed in the matrix form:

$$\dot{\mathbf{x}}(t) = \mathbf{A}(\boldsymbol{\psi})\mathbf{x}(t) + \mathbf{B}(\boldsymbol{\psi})\mathbf{u}(t) \tag{6}$$

where the system matrix $\mathbf{A}(\boldsymbol{\psi}) \in \mathbb{R}^{3\times 3}$ describes the internal thermal coupling. These matrices are parameterized by the thermal parameter set $\boldsymbol{\psi} = \left\{R_{st,wdg}, R_{st,pm}, R_{st,amb}, \ldots, C_{st}, C_{wdg}, C_{pm}\right\}$.

Explicitly, the diagonal elements of $\mathbf{A}(\boldsymbol{\psi})$ represent the self-dissipation rates, while the off-diagonal elements represent the inter-nodal coupling. For instance, the dynamics of the winding node are governed by:

$$\frac{d\vartheta_{wdg}}{dt} = \frac{Q_{cu}}{C_{wdg}} - \frac{\vartheta_{wdg} - \vartheta_{st}}{R_{wdg,st}C_{wdg}} - \frac{\vartheta_{wdg} - \vartheta_{amb}}{R_{wdg,amb}C_{wdg}} \tag{7}$$

This formulation highlights the physical interpretability of the model: every element in the state matrices corresponds to a specific physical parameter.

2.3. Discretization and Model Uncertainty

To implement the digital twin on a digital processor, the continuous model (6) is discretized using the Euler method with a sampling interval Δt. The discrete-time evolution equation is given by:

$$\mathbf{x}_{k+1} = (\mathbf{I} + \mathbf{A}(\boldsymbol{\psi})\Delta t)\mathbf{x}_k + (\mathbf{B}(\boldsymbol{\psi})\Delta t)\mathbf{u}_k \tag{8}$$

While this physical twin captures the fundamental heat transfer mechanisms, its fidelity relies heavily on the accuracy of the parameter set $\boldsymbol{\psi}$. In practice, determining these parameters is fraught with challenges:

- Geometric simplification: The lumping of complex 3D geometries into 1D nodes introduces inherent structural errors.
- Material uncertainty: Thermal properties (e.g., conductivity, specific heat) vary with manufacturing tolerances and material grades.
- Aging effects: Thermal interface materials degrade over time, altering contact resistances ($R_{contact}$).

Consequently, a nominal set of parameters $\boldsymbol{\psi}_{nom}$ derived from datasheets or rough geometric calculations is insufficient for high-fidelity monitoring. This necessitates a mechanism for the digital twin to "self-calibrate" against the actual motor, which leads to the proposed Bayesian calibration framework in the next section.

3. Digital Twin Self-Calibration via DBBC

To mitigate the parameter uncertainties identified in Section 2, this paper proposes a probabilistic calibration framework. Unlike deterministic optimization methods (e.g., genetic algorithms or particle swarm optimization) that provide point estimates without

confidence intervals, Bayesian inference treats the parameters as random variables, estimating their full posterior probability distributions. To address the computational inefficiency of traditional Bayesian methods on large-scale monitoring data, we introduce a novel DBBC algorithm.

3.1. Probabilistic Formulation of Parameter Inversion

The goal of self-calibration is to infer the unknown thermal parameter set $\boldsymbol{\psi}$ given a sequence of observed operational data $\mathcal{D} = \{(\mathbf{u}_k, \mathbf{y}_k)\}_{k=1}^{N}$, where y_k represents the measured temperatures from sensors (e.g., stator winding sensors). According to Bayes' theorem, the posterior distribution $P(\boldsymbol{\psi}|\mathcal{D})$ is proportional to the likelihood of the data given the parameters multiplied by the prior belief:

$$P(\boldsymbol{\psi}|\mathcal{D}) \propto \mathcal{L}(\mathcal{D}|\boldsymbol{\psi}) \cdot \pi(\boldsymbol{\psi}) \tag{9}$$

Here, $\pi(\boldsymbol{\psi})$ encodes prior knowledge (e.g., physical bounds of thermal resistance), and $\mathcal{L}(\mathcal{D}|\boldsymbol{\psi})$ is the likelihood function. Assuming the measurement noise and model error follow a Gaussian distribution with zero mean and variance σ^2, the log-likelihood function for the dataset can be expressed as:

$$\ln \mathcal{L}(\mathcal{D}|\boldsymbol{\psi}) = -\frac{N}{2}\ln(2\pi\sigma^2) - \frac{1}{2\sigma^2}\sum_{k=1}^{N} ||\mathbf{y}_k - \hat{\mathbf{y}}_k(\boldsymbol{\psi})||^2 \tag{10}$$

where $\hat{\mathbf{y}}_k(\boldsymbol{\psi})$ is the output predicted by the physical twin using parameters $\boldsymbol{\psi}$. σ represents the standard deviation of the measurement noise, which is determined a priori based on the precision of the temperature sensors.

3.2. The Dynamic-Batch Bayesian Calibration (DBBC) Algorithm

Standard Markov chain Monte Carlo (MCMC) methods, such as the metropolis-hastings (MH) algorithm, require calculating the likelihood over the entire dataset $\mathcal{D}$ (size N) at every iteration to accept or reject a candidate parameter. For digital twins accumulating massive amounts of operational data, this full-batch evaluation becomes computationally prohibitive [31,32]. Furthermore, standard MH often suffers from slow mixing, trapping in local modes of the posterior distribution.

The proposed DBBC algorithm overcomes these limitations by introducing a dynamic sub-sampling strategy. Inspired by the efficiency of stochastic gradient descent (SGD) in deep learning, DBBC utilizes a subset (batch) of data $\mathcal{B}_t \subset \mathcal{D}$ at iteration t to estimate the likelihood. The innovation lies in the adaptive scheduling of the batch size $|\mathcal{B}_t|$.

3.2.1. The Acceptance Ratio with Dynamic Batches

Let $\boldsymbol{\psi}_t$ be the current parameter state and $\boldsymbol{\psi}'$ be a candidate sampled from a proposal distribution $q(\boldsymbol{\psi}'|\boldsymbol{\psi}_t)$. The acceptance probability α in DBBC is modified to rely on the dynamic batch $\mathcal{B}_t$:

$$\alpha(\boldsymbol{\psi}_t, \boldsymbol{\psi}') = \min\left(1, \frac{\mathcal{L}(\mathcal{B}_t|\boldsymbol{\psi}')\pi(\boldsymbol{\psi}')q(\boldsymbol{\psi}_t|\boldsymbol{\psi}')}{\mathcal{L}(\mathcal{B}_t|\boldsymbol{\psi}_t)\pi(\boldsymbol{\psi}_t)q(\boldsymbol{\psi}'|\boldsymbol{\psi}_t)}\right) \tag{11}$$

The likelihood ratio is approximated using the batch data, which introduces stochastic noise into the acceptance decision. The DBBC algorithm strategically manipulates this noise through a two-phase process:

Phase 1: Stochastic Exploration (Warm-Up)

In the initial iterations, a small fixed batch size B_{small} is employed.

- Mechanism: Small batches introduce significant variance into the likelihood estimation. This "gradient noise" acts as a thermal agitation, allowing the Markov chain to traverse the high-dimensional parameter space rapidly.
- Benefit: This phase prevents the algorithm from prematurely converging to local optima, effectively exploring the global landscape of the parameter space. It mimics the "simulated annealing" process where high temperature facilitates exploration.

Phase 2: Precision Convergence (Refinement)

Once the chain stabilizes (detected via convergence diagnostics), the algorithm transitions to the convergence phase. The batch size is adaptively increased according to a schedule function $\Gamma(t)$, eventually reaching the full dataset size or a sufficiently large B_{large}.

- Mechanism: Increasing the batch size suppresses the estimation noise, sharpening the likelihood surface.
- Benefit: This ensures that the samples settle into the mode of the true posterior distribution, providing high-precision parameter estimates required for the physical twin.

3.2.2. Algorithmic Implementation

The execution of the DBBC algorithm involves the following steps:

1. Initialization: Initialize parameters $\boldsymbol{\psi}_0$ from prior distributions (e.g., rough geometric calculations).
2. Proposal Generation: At step t, generate a candidate $\boldsymbol{\psi}' = \boldsymbol{\psi}_t + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \Sigma_{prop})$
3. Batch Selection: Select a random mini-batch $\mathcal{B}_t$ of size $|\mathcal{B}_t|$. The size is determined by the adaptive schedule:

$$|\mathcal{B}_t| = \begin{cases} B_{start} & \text{if } t < T_{warmup} \\ \min(N, B_{start} + \gamma(t - T_{warmup})) & \text{if } t \geq T_{warmup} \end{cases} \tag{12}$$

   where γ is the growth rate.
4. Evaluation and Decision: Calculate the acceptance ratio α using Equation (11) and accept $\boldsymbol{\psi}'$ with probability α.
5. Constraint Handling: Physical constraints (e.g., $R > 0$, $C > 0$) are enforced via the prior $\pi(\boldsymbol{\psi})$, which returns zero probability for infeasible regions.

3.3. Convergence Analysis

The convergence of the DBBC algorithm is monitored using the Gelman-Rubin statistic ($\hat{R}$) and the Geweke diagnostic. The Gelman-Rubin test runs multiple parallel chains and compares the inter-chain variance to the intra-chain variance. Convergence is declared when $\hat{R} < 1.1$ for all parameters. Upon convergence, the empirical mean of the posterior samples is extracted as the optimal parameter set $\boldsymbol{\psi}^*$. These calibrated parameters are then injected back into the state-space model (Equation (6)), completing the "self-calibration" of the physical twin.

4. High-Fidelity Temperature Estimation via Hierarchical Physics-Aware Network

While the DBBC algorithm significantly improves model accuracy by calibrating the linear parameters, the physical twin derived in Section 2 remains a reduced-order approximation. It relies on the linear assumption, presupposing that thermal resistances

and capacitances are constant. However, in reality, thermal dynamics in PMSMs exhibit strong non-linearities:

- Convection Non-linearity: The convective heat transfer coefficient is highly dependent on rotor speed and air gap fluid dynamics.
- Loss Distribution: Eddy current losses are non-uniformly distributed and temperature-dependent.

These unmodeled dynamics create a reality gap—a persistent estimation error $\mathbf{e}(t)$ between the calibrated physical twin output and the true motor temperature. To bridge this gap without discarding the physical model, this section proposes a HPA-Net.

4.1. The Physics-Augmented Modeling Strategy

Instead of replacing the physical model with a black-box neural network, HPA-Net is designed as an "add-on" compensator. The digital twin's final output is synthesized as:

To achieve high-fidelity monitoring under complex dynamic conditions, this paper proposes an HPA-Net. The core philosophy of HPA-Net is to utilize the calibrated physical twin (from Section 3) as a "dynamic knowledge generator" that provides real-time physical guidance to the neural network. Mathematically, the Digital Twin's final output y_{DT} is formulated as a deep non-linear mapping $\mathcal{F}$ driven by physics-augmented features:

$$y_{DT} = \mathcal{F}(\mathbf{x}_{raw}, y_{TN}, \theta) \tag{13}$$

where $\mathbf{x}_{raw}$ represents the measurable operation variables (currents i_d, i_q, rotor speed ω, etc.). y_{TN} is the theoretical temperature estimate calculated by the Calibrated TN. Crucially, y_{TN} is embedded into the network's input layer as a Dynamic Physical Prior. It informs the data-driven model about the fundamental thermal trends derived from Fourier's law. θ denotes the learnable weights of the network.

By incorporating y_{TN} directly into the feature space, the HPA-Net transforms the problem from "learning physics ab initio" to "physics-guided refinement." The network leverages the physical model's output as a robust baseline and learns to map this coarse estimate—along with other operational variables—to the precise ground truth temperature. This physics-augmented architecture ensures that the estimation is strictly anchored by physical laws, thereby significantly improving convergence speed and generalization capability compared to pure black-box models.

4.2. Architecture of HPA-Net

To effectively leverage the physical prior while avoiding overfitting to noise, HPA-Net employs a specialized multi-branch autoencoder architecture [33,34]. Unlike standard feedforward networks, this topology forces the network to learn a robust, physically meaningful latent representation of the thermal state.

4.2.1. Shared Feature Encoder

The encoder acts as a feature extractor. It takes a multi-physics input vector containing:

- Operational Variables: *d*-*q* axis currents (i_d, i_q), rotor speed (ω), *d*-*q* axis voltages (u_d, u_q), etc.
- Physical Estimates: The temperature states predicted by the physical twin $\mathbf{x}_{phys}$. This physics-augmented input provides the network with a strong initial guess.

Before feeding into the network, all input variables are normalized to the range [0, 1] using Min-Max scaling to eliminate dimensional heterogeneity and accelerate convergence. The encoder maps these inputs into a low-dimensional latent space $\mathbf{h}$. This compres-

sion forces the network to filter out high-frequency measurement noise and capture the underlying thermal manifold of the system.

4.2.2. Physics Reconstruction Branch

The first decoder branch is the reconstruction branch. Its objective is to reconstruct the original physical inputs from the latent code **h**.

- Function: This acts as a physics-based regularization mechanism. By forcing the latent code to retain enough information to reproduce the physical states, it ensures that the learned features are physically meaningful and not just arbitrary numerical mappings, thereby preventing overfitting on small datasets.

4.2.3. Physics Consistency Branch

The second decoder branch (Yellow module, Bottom) is the Physics Consistency Branch.

- Function: This branch maps the latent code h to known physical correlations (e.g., theoretical power losses derived from currents).
- Objective: The associated loss $\mathcal{J}_{phy}$ enforces the network to respect fundamental energy conservation laws. This prevents the latent representation from learning unphysical patterns during the pre-training phase, ensuring the gray-box nature of the model.

4.2.4. Specialized Task Branch

The third decoder branch is the specialized task branch, which serves as the core estimation engine.

- Function: Instead of predicting an error term, this branch maps the latent code **h** directly to the final high-fidelity temperature vector y_{DT}.
- Mechanism: Leveraging the physical features extracted by the encoder, this branch functions as a refinement engine. It preserves the dominant physical laws while non-linearly fine-tuning the magnitude to match the ground truth, effectively compensating for the structural limitations of the linear physical model.

4.3. Physics-Constrained Hierarchical Training

Training neural networks on limited experimental data often leads to overfitting. To address this, HPA-Net employs a novel hierarchical training strategy that leverages the concept of transfer learning. The schematic of the proposed HPA-Net is illustrated in Figure 3.

Phase 1: Physics-Constrained Pre-training (Simulation Domain)

In this phase, the network is trained using a massive dataset generated by the calibrated physical twin (from Section 3) under a wide range of operating conditions (e.g., extreme speeds, overload currents). The goal is to embed the physical mechanisms into the neural network weights.

- Objective: The optimization objective is a compound loss function consisting of three components:

$$\mathcal{J}_{total} = \mathcal{J}_{rec} + \omega_{task} \cdot \mathcal{J}_{task} + \omega_{phy} \cdot \mathcal{J}_{phy} \tag{14}$$

where ω_{task}, ω_{phy} are weighting coefficients balancing the tasks.

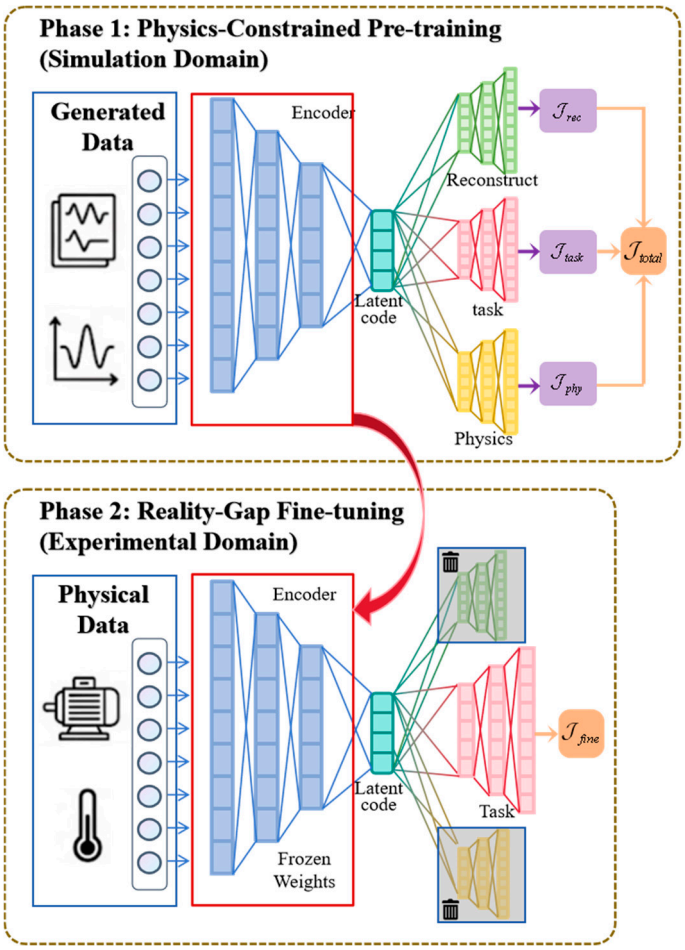

Figure 3. Schematic representation of the proposed HPA-Net.

1. Reconstruction loss ($\mathcal{J}_{rec}$): Ensures the encoder captures comprehensive features from the inputs $\mathbf{x}_{in}$.

$$\mathcal{J}_{rec} = \frac{1}{N}\sum ||\hat{\mathbf{x}}_{rec} - \mathbf{x}_{in}||^2 \tag{15}$$

2. Specialized task loss ($\mathcal{J}_{task}$): The core task branch is trained to mimic the output of the physical twin. We force the network to reproduce the physical model's prediction $\mathbf{y}_{TN}$ (e.g., temperature change rates). This creates a strong "physical prior" in the network.

$$\mathcal{J}_{task} = \frac{1}{N}\sum ||\hat{\mathbf{y}}_{pred} - \mathbf{y}_{TN}||^2 \tag{16}$$

3. Physics consistency loss ($\mathcal{J}_{phy}$): Corresponds to the "empirical feature branch". It enforces the network to learn known physical correlations (e.g., ensuring the latent features correlate with theoretical losses).

$$\mathcal{J}_{phy} = \frac{1}{N}\sum ||\hat{\mathbf{f}}_{phy} - \mathbf{f}_{theory}||^2 \tag{17}$$

where $\hat{\mathbf{f}}_{phy}$ represents the feature extraction expectation value computed from the input based on empirical knowledge, and $\mathbf{f}_{theory}$ denotes the output of the branch decoder.

- Outcome: The encoder learns to extract robust features related to power losses and thermal dynamics. The network effectively learns physics from the simulation data, establishing a stable initialization for the weights.

Phase 2: Reality-Gap Fine-tuning (Experimental Domain)

In this phase, the network is exposed to real experimental data collected from the test bench.

- Strategy: The weights of the shared encoder are frozen (or updated with a very low learning rate) to preserve the learned physical features. Only the weights of the specialized task branch are actively optimized.
- Objective: The loss function targets the prediction of the reality gap:

$$\mathcal{J}_{fine} = \left|\left|\mathbf{y}_{sensor} - \hat{\mathbf{y}}_{pred}\right|\right|^2 + \xi\left|\left|\Theta\right|\right|^2 \tag{18}$$

  where ξ is the regularization coefficient.
- Outcome: The network focuses solely on learning the specific non-linearities (the reality gap) that the TN missed. This drastically reduces the amount of experimental data required for convergence and ensures excellent generalization to unseen operating points.

Through this two-step process, HPA-Net evolves from a purely mathematical observer to a physics-aware compensator, completing the construction of a self-evolving, high-fidelity digital twin.

4.4. Real-Time Thermal Safety Assessment Strategy

Accurate temperature estimation is the prerequisite for safety monitoring. Based on the estimated PM temperature ($\hat{\vartheta}_{pm}$) from HPA-Net, the motor's Demagnetization Safety Margin (DSM) is calculated in real-time. The knee point of the demagnetization curve (H_{knee}) for NdFeB magnets degrades linearly with temperature:

$$B_{limit}(t) = B_{r0} \cdot (1 + \alpha_{\mathrm{Br}}(\hat{\vartheta}_{pm}(t) - 20)) \tag{19}$$

By monitoring the distance between the actual operating point and this temperature-dependent B_{limit}, the system triggers an early warning when the safety margin drops below 10%, effectively preventing irreversible demagnetization faults.

5. Verification Results

To comprehensively validate the fidelity and adaptability of the proposed self-evolving digital twin framework, experimental tests are conducted on a specialized test bench driven by a 24-slot/4-pole interior PMSM. Crucially, to assess the framework's "self-calibration" capability, the experiment is designed to verify if the proposed DBBC algorithm can robustly identify the exact thermal parameters of this specific physical motor from operational data. This process is essential to eliminate the inherent uncertainties caused by manufacturing tolerances and material variations, ensuring the physical twin perfectly matches the real plant before the neural network compensation is applied.

5.1. Experimental Setup

5.1.1. Hardware Configuration

The overall hardware architecture is visualized in Figure 4. Regarding the control and loading system, the digital twin algorithms run in parallel with the Field-Oriented Control (FOC) strategy on a floating-point DSP TMS320F28335 controller board (Texas Instruments, Dallas, TX, USA). Real-time data monitoring and interaction occur through a host console on the primary PC, with a high-bandwidth oscilloscope capturing transient signal waveforms. To provide dynamic load profiles for thermal testing, a programmable magnetic powder brake is mechanically coupled to the motor under test.

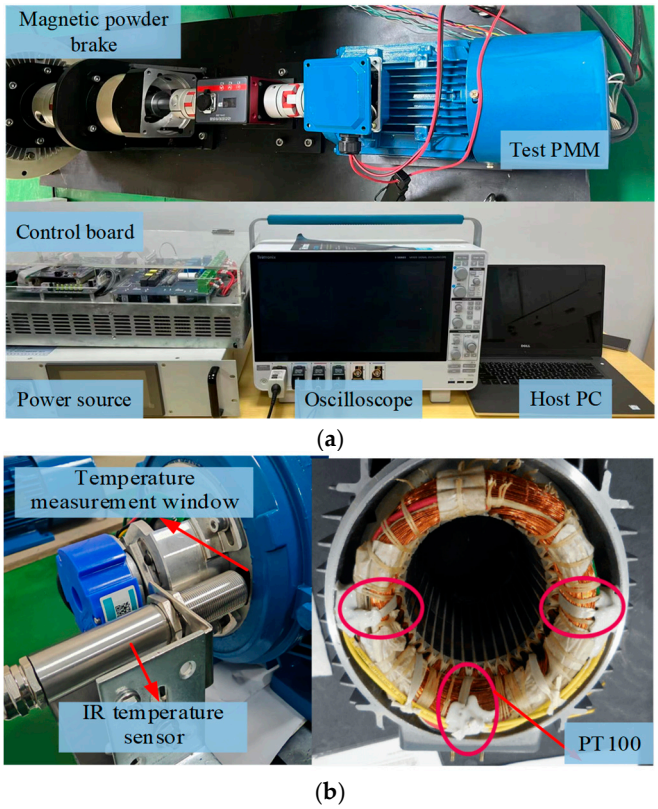

Figure 4. Illustration of the verification test bench. (**a**) Overall hardware architecture and control interface; (**b**) Detailed view of the motor temperature measurement.

In order to acquire high-precision thermal "ground truth" to evaluate the estimation accuracy of the digital twin, the platform employs a comprehensive independent temperature sensing system. For the stationary components, the platform uses high-precision PT100 resistance temperature detectors (Dwyer Instruments, Michigan City, IN, USA, IEC 60751 Class A accuracy) to obtain the actual temperature of the coil and stator tooth. A number of PT100 sensors are evenly distributed at critical hotspots within the stator slots and winding end-turns to fully capture the stationary thermal field distribution. It is worth mentioning that accurately validating the rotor magnet temperature is vital for the proposed safety margin monitoring strategy. Therefore, to measure the temperature of the rotating rotor without altering thermal boundary conditions, as shown in Figure 4b, a specialized observation window was machined on the motor's rear cover. Through this aperture, a non-contact infrared radiometer is positioned to measure the rotor's real-time surface temperature. This setup serves as the "Reference Reality," providing the benchmark data required to calculate the reality gap and train the HPA-Net. To coordinate the data acquisition of heterogeneous sensors, a multi-rate processing scheme is implemented. The electrical signals are sampled at a frequency of 10 kHz. To ensure stable torque regulation, the control algorithm is executed at a control frequency of 5 kHz. In contrast, considering the large thermal inertia of the motor components, the temperature data from the thermocouples are acquired at a lower sampling rate of 1 Hz.

5.1.2. Digital Twin Visualization Interface

To realize the comprehensive interaction between the physical asset and the virtual model, a dedicated human–machine interface (HMI) system was developed, as illustrated in Figure 5. This software platform serves as the visual interaction layer of the DT, communicating with the DSP controller via a high-speed serial communication interface.

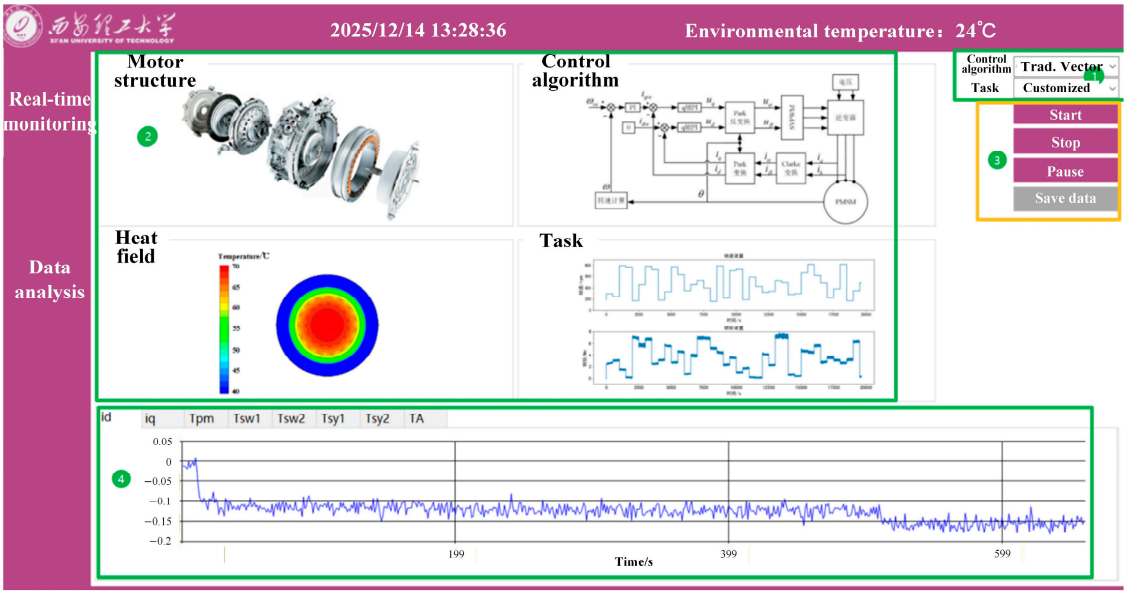

Figure 5. Interface of the developed DT monitoring software.

The visualization platform is designed with three core functional modules to ensure lifecycle manageability:

(1) Operational Control and Task Loading: Operators can define dynamic "Task Profiles" (e.g., customizable speed/load cycles) through the interface, which are transmitted to the controller to execute complex testing scenarios. Control strategies (e.g., vector control) can also be switched online.

(2) Multi-Physics Visualization: This module renders the "virtual entity" of the motor. Crucially, it maps the nodal temperatures estimated by the HPA-Net onto the motor geometry in real-time, generating a dynamic thermal heatmap. This provides operators with intuitive insight into the internal thermal distribution, revealing hotspots that physical sensors cannot cover.

(3) Transient Analysis: A virtual oscilloscope monitors critical state variables, including *d-q* axis currents and real-time safety margins, facilitating immediate fault diagnosis and performance evaluation.

5.2. Validation of Physical Twin Self-Calibration (DBBC Performance)

The first stage of validation aims to verify if the proposed DBBC algorithm can autonomously identify the accurate thermal parameters of the specific test motor from scratch, thereby eliminating the "Parameter Uncertainty" caused by manufacturing tolerances. Moreover, as an offline calibration tool, DBBC eliminates real-time updates given the quasi-static nature of thermal parameters. This ensures the online digital twin operates with fixed parameters, maintaining a minimal computational load ideal for embedded implementation.

5.2.1. Experimental Protocol

To rigorously evaluate the digital twin's performance under realistic, non-stationary conditions, a stochastic drive cycle was employed as the validation benchmark. Unlike standard steady-state tests, this cycle introduces rapid fluctuations in both speed and load, mimicking the harsh thermal transients encountered in actual industrial applications (e.g., EV acceleration or wind gusts). Figure 6 illustrates the distribution of the operating points covered during the experiment. As shown, the test envelope spans a wide range of speed (0–800 rpm) and torque (0–10 Nm) combinations, ensuring that both the DBBC

and the HPA-Net are validated against global dynamic behaviors rather than limited local operating points.

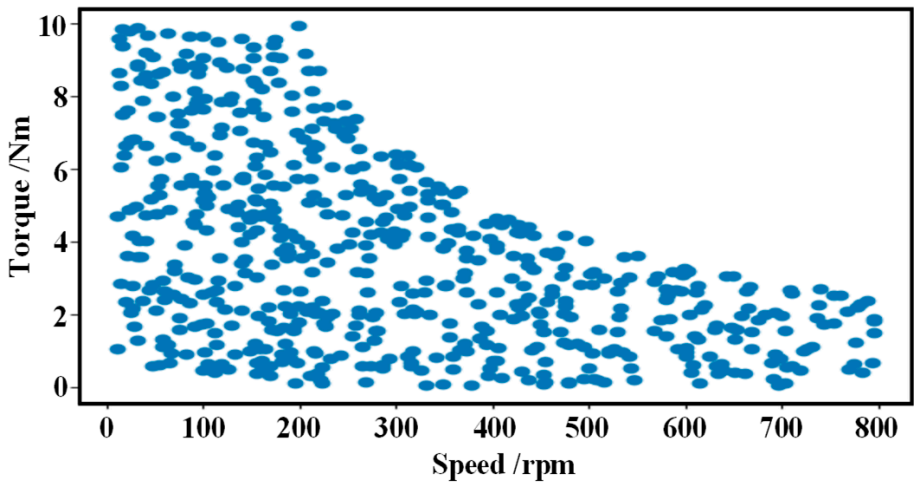

Figure 6. Distribution of operating points tested during the experimental validation.

To simulate a realistic scenario of "Parameter Uncertainty," the seven key thermal parameters (scaling factors k_1 to k_7, representing thermal resistances and capacitances relative to their nominal values) were initialized with random deviations of up to ±30%.

5.2.2. Convergence Analysis

Figure 7 illustrates the optimization trajectory of the mean squared error (MSE) for ten parallel Markov chains (MC_1–MC_{10}) during the calibration process.

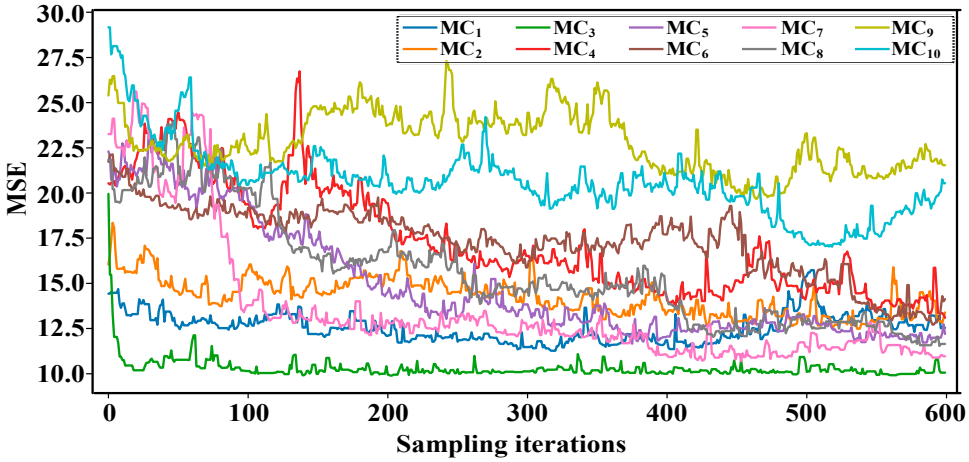

Figure 7. Convergence traces of the Mean Squared Error (MSE) during the DBBC pre-heating phase.

Stochastic Exploration: In the initial "pre-heating" iterations (0–100), the MSE curves exhibit sharp declines. The algorithm utilizes the gradient noise from small batches to drive the chains rapidly out of high-error regions, avoiding local minima.

Stability: After approx. 200 iterations, all chains converge to a stable low-error floor (MSE ≈ 10.0), indicating that the physical twin has successfully aligned its dynamics with the actual motor measurements.

5.2.3. Posterior Parameter Identification

Unlike deterministic methods that yield only point estimates, the DBBC framework quantifies the uncertainty of the identified parameters. Figure 8 presents the matrix of the posterior distributions for the seven calibrated parameters (k_1–k_7).

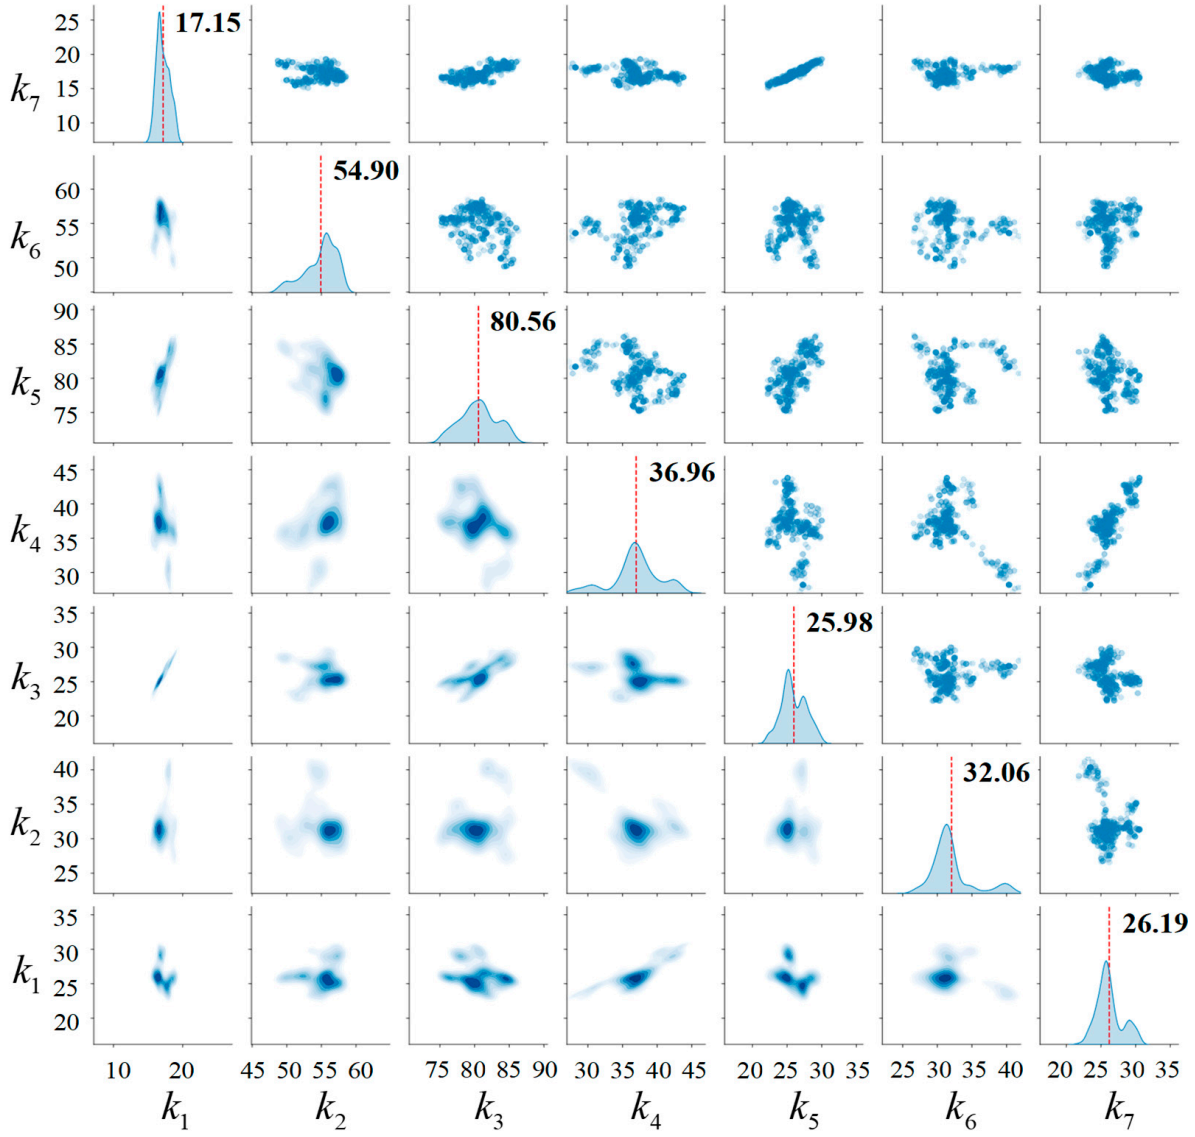

Figure 8. Posterior distribution matrix of the identified thermal parameters.

Marginal Distributions (Diagonal): As shown in the diagonal density plots, all parameters exhibit clear, unimodal peaks. For instance, parameter k_7 (associated with the critical winding heat capacity) shows a sharp convergence around its true physical value (approx. 17.15), confirming high identification precision.

Correlation Analysis (Off-diagonal): The scatter plots reveal the physical coupling between parameters. The compact clustering of samples in the high-probability regions demonstrates that the algorithm has effectively captured the "thermal manifold" of the motor, ensuring a robust physical baseline for the subsequent digital twin operations.

5.3. Fidelity Enhancement via HPA-Net

While the DBBC-calibrated physical twin (Section 5.1) captures the dominant linear thermal dynamics, it inevitably suffers from errors under highly dynamic operating

conditions due to model simplifications (e.g., constant thermal resistances ignoring speed-dependent convection). This section validates the capability of the proposed HPA-Net to bridge this "reality gap."

Before presenting the comparative results, the implementation details of the proposed network are clarified. The specific hyperparameters of the implemented HPA-Net are detailed in Table 1. Structurally, the central encoder employs an overcomplete architecture with a latent dimension of 32 to capture rich feature representations, and the ReLU activation function is utilized for all hidden layers to ensure efficient gradient propagation. In addition, these coefficients in (14) are determined via a grid search method to ensure that the gradients from the reconstruction branch, task branch, and physics branch share the same order of magnitude. Based on the experimental tuning, the optimal values are set as $\omega_{task} = 10$ and $\omega_{phy} = 0.1$.

Table 1. Detailed architecture parameters of HPA-Net.

Component	Input Dimension	Output Dimension	Hidden Layers	Neurons per Layer
Central Encoder	12	32	2	32
Reconstruction Branch	32	12	1	32
Specialized Task Branch	32	3	1	16
Physics Consistency Branch	32	2	1	16

To strictly evaluate the generalization performance, the validation was conducted on a dynamic "stop-and-go" drive cycle that was not included in the training dataset. The proposed framework was compared against two baseline methods: Calibrated TN: The physics-only model optimized in Section 3; GRU: A Gated Recurrent Unit network representing standard RNN variants.

A. Temperature Estimation Accuracy

Figure 9 presents the time-domain comparison of temperature estimation for both the Stator Winding and the PM.

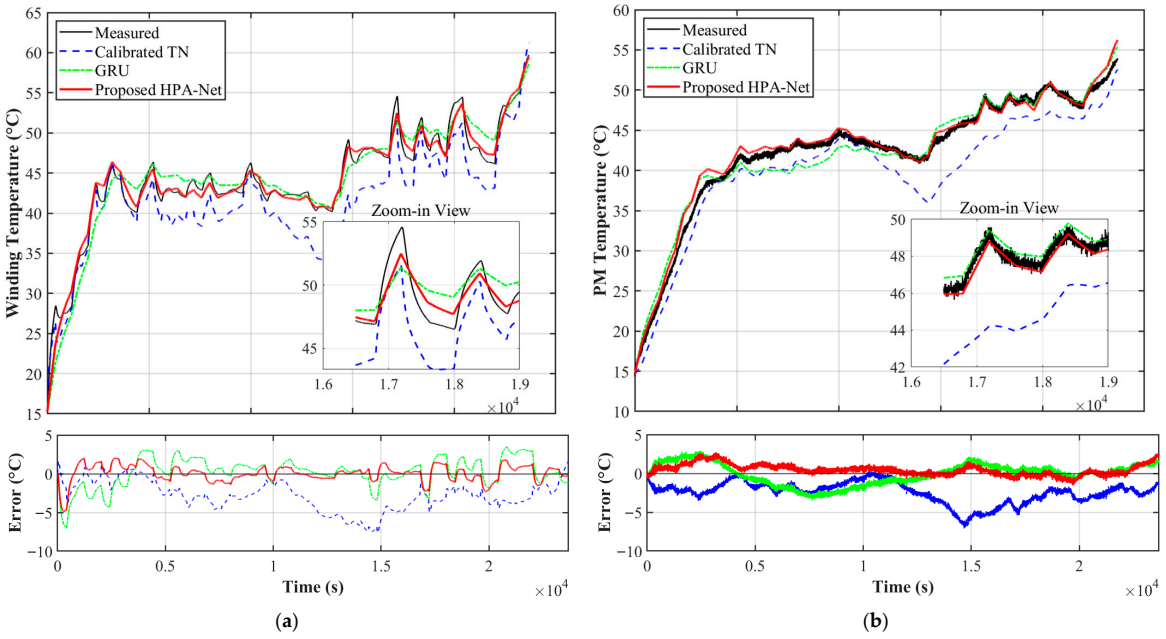

Figure 9. Comparative experimental results of temperature estimation under dynamic loading. (**a**) Estimation temperature for the stator winding; (**b**) Estimation temperature for the PM.

Stator Winding Analysis (Figure 9a): As observed in the "Zoom-in View" and the error plot, the Calibrated TN exhibits a persistent steady-state bias and fails to capture the peak temperatures during rapid loading, reaching a maximum deviation of over 5 °C. This is attributed to the linear model's inability to account for the non-linear variations in iron loss and convection heat transfer. The Standard GRU, while tracking the trend, shows noticeable fluctuations and instability. In contrast, the proposed HPA-Net tightly hugs the measured curve. By learning the non-linear term, it effectively eliminates the bias of the TN while avoiding the overfitting noise of the GRU.

PM Temperature Analysis (Figure 9b): Accurate estimation of the rotor PM temperature is critical for safety but challenging due to the lack of direct physical connection. The Calibrated TN shows a significant drift over time, underestimating the rotor temperature by nearly 4 °C at the later stages of the cycle. The HPA-Net demonstrates superior generalization capability, maintaining high accuracy even for this difficult-to-observe node. The zoom-in view and error curves confirm that the proposed method captures the subtle thermal transients of the rotor that the physical model misses.

Statistical Performance Summary: To further verify the generalization capability of the proposed framework, statistical evaluations were conducted across different dynamic tests. Table 2 summarizes the quantitative performance metrics (MAE, RMSE, MAPE, and R^2) for the Stator Winding (SW) and PM.

Table 2. Quantitative performance comparison across different tests.

Method	Comp.	MAE	RMSE	MAPE	R^2
HPA-Net	SW	0.787	0.919	1.792%	0.982
	PM	1.169	1.603	1.890%	0.9667
GRU	SW	2.056	2.731	4.525%	0.8552
	PM	2.971	3.943	5.416%	0.8383
Calibrated TN	SW	3.334	3.945	5.669%	0.755
	PM	4.169	4.621	7.698%	0.665

As indicated in Table 2, the proposed HPA-Net achieves the lowest average RMSE (0.919 °C for SW and 1.603 °C for PM) across all tests, significantly outperforming the GRU and calibrated TN baselines. This confirms that the superior performance observed in Figure 9 is consistent and robust against different operating conditions.

B. Dynamic Consistency and Physical Interpretability

A key advantage of the proposed "Physics-Aware" training strategy is that the network learns to respect physical laws rather than merely fitting data. To verify this, we analyzed the temperature rate of change ($\hat{\vartheta}$), which represents the thermal dynamics governing the system.

Rate-of-Change Analysis: Figure 10 compares the per-second temperature change rate predicted by the HPA-Net versus the calibrated TN. The TN output exhibits significant jagged noise and lag, which is typical for simplified differential equations under discrete sampling. Conversely, the HPA-Net output is smoother and more stable. Crucially, the HPA-Net does not generate non-physical spikes, proving that the physics consistency loss ($\mathcal{J}_{phy}$) successfully regularized the network to follow the underlying thermal inertia of the motor.

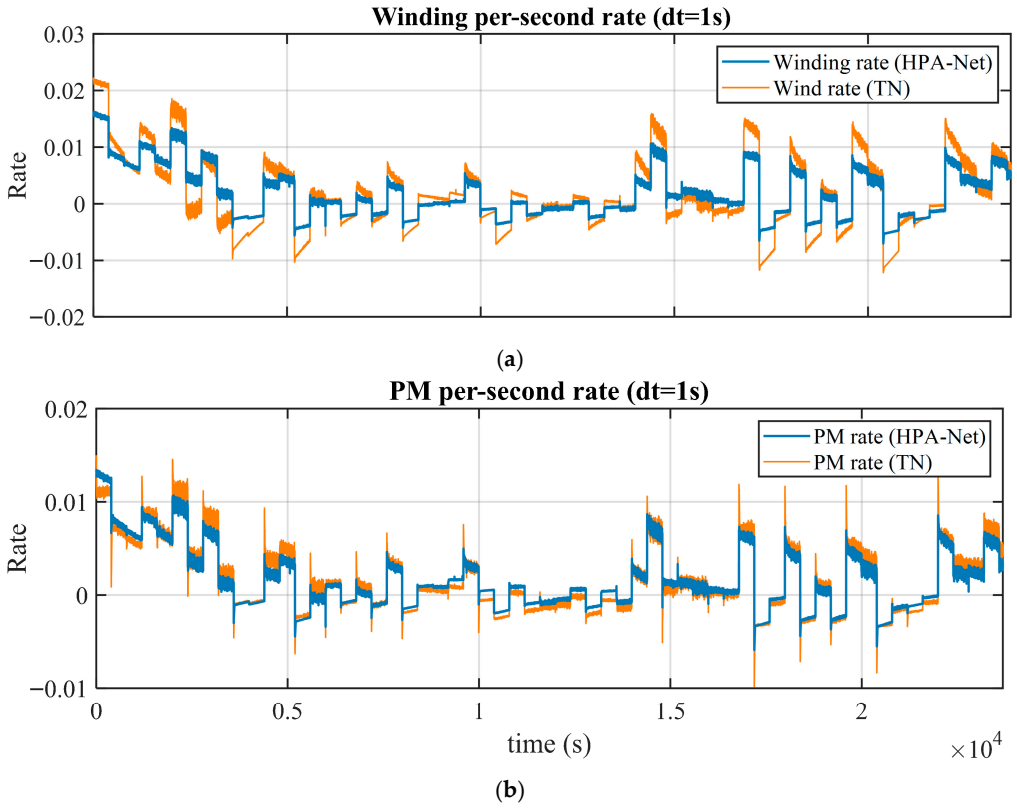

Figure 10. Comparison of the temperature rate of change between the HPA-Net and the TN. (**a**) Dynamics of the stator winding temperature rate; (**b**) Dynamics of the PM temperature rate.

Correlation Verification: Figure 11 further quantifies this physical consistency through scatter plots correlating the dynamic rates of the HPA-Net and the TN. For the winding rate, the correlation coefficient is $r = 0.953$. For the PM rate, the correlation coefficient reaches $r = 0.978$. The strong linearity in these plots confirms that the HPA-Net retains the physical directions derived from the thermal network while applying necessary magnitude corrections. This proves that the digital twin is not a Black Box but a physically interpretable Gray Box model.

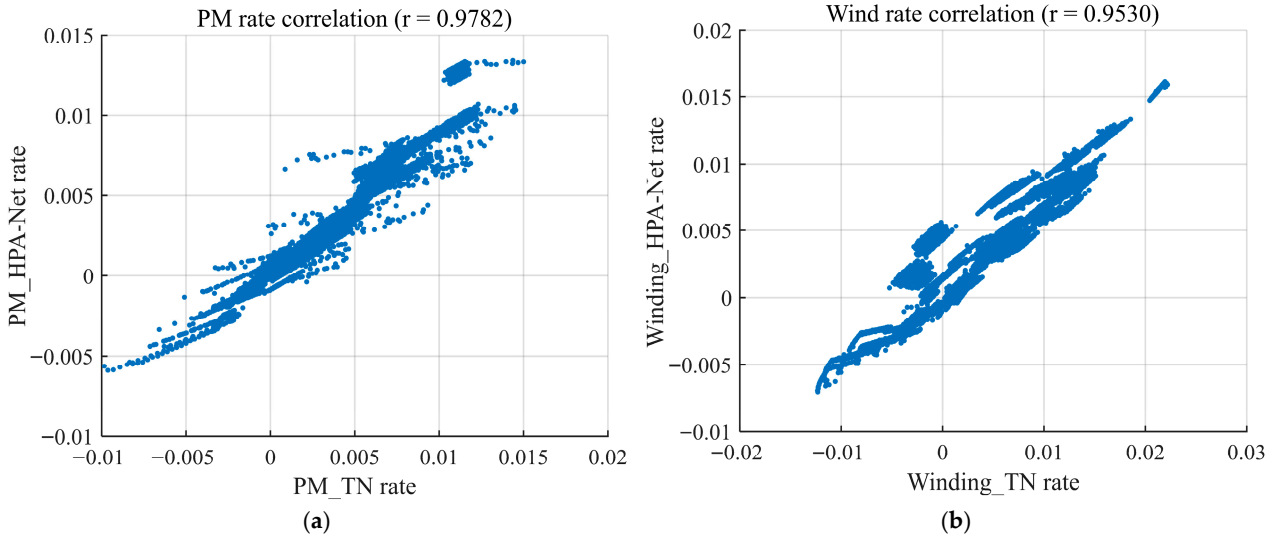

Figure 11. Correlation analysis of the thermal dynamics between the Physics Model and the Neural Network. (**a**) Linear correlation of stator winding temperature rates; (**b**) Linear correlation of PM temperature rates.

5.4. Real-Time Demagnetization Safety Assessment

The ultimate objective of the digital twin is to prevent irreversible demagnetization of the PM. While Section 5.3 demonstrated the temperature tracking accuracy, this section quantifies how that accuracy translates into reliable safety monitoring. Since the demagnetization knee point (H_{knee}) is a direct function of the PM temperature, any estimation error in the thermal model propagates directly into the safety margin calculation.

A. Safety Margin Formulation

The demagnetization safety margin (DSM) is defined as the distance between the current operating point and the temperature-dependent knee point limit. A negative error leads to an overly optimistic DSM, creating a "false safe" zone where the monitoring system fails to trigger an alarm despite the motor being at risk.

B. Quantitative Risk Analysis at Peak Load

To evaluate the safety reliability, we analyzed the critical high-temperature transient at $t \approx 1.72 \times 10^4$ s from Figure 9b. As shown in the zoom-in view of Figure 9b, the rapid loading caused a sharp temperature rise.

Failure of physical model: The Calibrated TN predicted a PM temperature of only 44.1 °C, whereas the actual temperature reached 49.5 °C. This underestimation of 5.4 °C is dangerous. For N35 grade NdFeB magnets, a 5 °C difference can degrade the intrinsic coercivity by approximately 30 kA/m. In a critical overload scenario (e.g., approaching 120 °C), such an error would cause the control system to overestimate the safety margin by roughly 12%, potentially allowing the motor to operate beyond its physical limit. Reliability of HPA-Net: In contrast, the HPA-Net estimated the peak at 48.9 °C, with a negligible error of 0.6 °C. This high-fidelity tracking ensures that the calculated DSM reflects the true thermal state of the rotor.

C. Implications for Active Protection

Furthermore, the dynamic analysis in Figure 10 confirms that the HPA-Net captures temperature changes with minimal delay. In practical engineering, this early detection capability is crucial. The traditional TN model typically exhibits a thermal delay due to the filtering effect of lumped capacitances. By compensating for the transient errors, the proposed DT can trigger thermal derating (current reduction) seconds earlier than traditional methods, providing a wider time window for the cooling system to react.

In summary, although the Calibrated TN provides a baseline, only the HPA-Net ensures the precision required for boundary-critical operations, effectively converting the gray box model into a robust safety guardian.

6. Conclusions

This paper proposes a self-evolving digital twin framework that integrates the DBBC and a HPA-Net. The DBBC eliminates the initial plant–model mismatch by robustly identifying stochastic parameters, while the HPA-Net adopts a physics-augmented strategy to directly infer high-fidelity temperature distributions. Additionally, the integrated DSM monitoring strategy effectively eliminates false safe zones, transforming the digital twin into a proactive safety guardian.

Experimental validation on a PMSM test bench confirms the superior performance of the proposed framework. Specifically, it achieves a RMSE of 0.919 °C for the stator winding and 1.603 °C for the critical permanent magnets, significantly outperforming calibrated TN and GRU baselines under stochastic dynamic loads.

Furthermore, the proposed methodology demonstrates strong generalizability. It is designed to be transferable: it adapts to homologous machines (e.g., brushless DC

electric motor) via topological updates, and extends to heterologous machines (e.g., switch reluctance machine) by reconstructing the physical baseline. In both cases, the core strategy of using HPA-Net to compensate for model limitations remains universally effective.

Despite these achievements, a limitation of the current framework lies in the hyperparameter optimization. The weighting coefficients in the loss function are currently determined via grid search, which may not be optimal for all operating phases. Future work will aim to eliminate this dependency on manual tuning. Inspired by multi-objective balancing, we plan to investigate adaptive weighting mechanisms (e.g., uncertainty-weighted loss) to dynamically adjust the focus between physical constraints and data-driven tasks.

Author Contributions: J.G.—Conceptualization, review and editing; J.L.—Original draft, Resources, Methodology; S.Y.—review and editing, Methodology; C.G.—Formal analysis, Investigation; S.W.—Data Curation, Software; G.Z.—Experimental Validation. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported in part by the National Natural Science Foundation of China under Grant 62503388, 62503250, the Fundamental Research Funds for the Central Universities under Grant G2024KY05104 and Basic Research Program of Jiangsu BK20250957.

Data Availability Statement: Data are contained within the article.

Conflicts of Interest: Author Junze Luo was employed by the company Hunan Electric Power Design Institute Co., Ltd. (CEEC-HEPDI). The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Ahmad, H.; Yaqub, M.; Lee, S.H. Global trends in carbon neutrality: A scientometric review on energy transition challenges, practices, policies, and opportunities. *Environ. Dev. Sustain.* **2025**, 1–26. [CrossRef]
2. Soni, R.; Dvivedi, A.; Kumar, P. Carbon neutrality in transportation: In the context of renewable sources. *Int. J. Sustain. Transp.* **2025**, *19*, 1–15. [CrossRef]
3. Lucchini, F.; Torchio, R.; Bianchi, N. A Survey on the Sustainability of Traditional and Emerging Materials for Next-Generation EV Motors. *Energies* **2024**, *17*, 5861. [CrossRef]
4. Kumari, R.; Prabhakaran, K.K.; Desingu, K.; Chelliah, T.R.; Sarma, S.V.A. Improved hydroturbine control and future prospects of variable speed hydropower plant. *IEEE Trans. Ind. Appl.* **2020**, *57*, 941–952. [CrossRef]
5. Nouraldin, N.A.; Chebaani, M.; Számel, L.; Abdelwahab, S.A.M.; Abdellatif, W.S. Experimental investigation of predictive control for PMSM-based wind turbine generation system. *Comput. Electr. Eng.* **2024**, *119*, 109554. [CrossRef]
6. Özçiflikçi, O.E.; Koç, M.; Bahçeci, S.; Selçuk, E. Overview of PMSM control strategies in electric vehicles: A review. *Int. J. Dyn. Control* **2024**, *12*, 2093–2107. [CrossRef]
7. Du, B.; Huang, W.; Cheng, Y.; Chen, J.; Tao, R.; Cui, S. Fault diagnosis and separation of PMSM rotor faults using search coil based on MVSA and random forests. *IEEE Trans. Ind. Electron.* **2024**, *71*, 15089–15099. [CrossRef]
8. De Donato, G.; Scelba, G.; Pulvirenti, M.; Scarcella, G.; Capponi, F.G. Low-cost, high-resolution, fault-robust position and speed estimation for PMSM drives operating in safety-critical systems. *IEEE Trans. Power Electron.* **2018**, *34*, 550–564. [CrossRef]
9. Liu, X.; Liu, S.; Fu, Z.; Ge, L. A Critical Review of Fault-Tolerant Control for Multiphase PMSM Drive Systems. *IEEE Trans. Transp. Electrif.* **2025**, *11*, 13684–13704. [CrossRef]
10. Zhang, Y.; Wang, Q.G.; Liu, J.; Yu, J. Command Filtered Fuzzy Adaptive Control for PMSMs with Time-Varying State Constraints and Stochastic Disturbances. *IEEE Trans. Transp. Electrif.* **2023**, *10*, 5614–5623. [CrossRef]
11. Wireko-Brobby, A.; Hu, Y.; Wang, G.; Gong, C.; Lang, W.; Zhang, Z. Analysis of the sources of error within PMSM-based electric powertrains—A review. *IEEE Trans. Transp. Electrif.* **2023**, *10*, 6370–6406. [CrossRef]
12. Li, X.; Wang, X.; Xing, Z. Analysis and research on local irreversible demagnetization of inverter-powered PM motor. *IEEE Trans. Ind. Electron.* **2023**, *71*, 5618–5629. [CrossRef]
13. Gao, J.; Gui, W.; Yang, C.; Peng, T.; Luo, J.; Yang, C. Demagnetization modeling and analysis for MTPA-driven permanent magnet motors with distributed lap windings. In *Proceedings of the IEEE Applied Power Electronics Conference and Exposition (APEC), Orlando, FL, USA, 19–23 March 2023*; IEEE: New York, NY, USA, 2023; pp. 1352–1357.

14. Melluso, F.; Spirto, M.; Nicolella, A.; Malfi, P.; Tordela, C.; Cosenza, C.; Savino, S.; Niola, V. Torque fault signal extraction in hybrid electric powertrains through a wavelet-supported processing of residuals. *Mech. Syst. Signal Process.* **2026**, *242*, 113652. [CrossRef]
15. Fan, Z.; Gao, R.X.; He, Q.; Huang, Y.; Jiang, T.; Peng, Z.; Thévenaz, L.; Xiong, Y.; Zhong, S. New sensing technologies for monitoring machinery, structures, and manufacturing processes. *J. Dyn. Monit. Diagn.* **2023**, *2*, 69–88.
16. Hu, J.; Lv, K.; Huang, W.; Lin, H.; Wang, D. Temperature measurement of PMSM rotor based on Wound-Rotor resolver reuse technique. *IEEE Trans. Transp. Electrif.* **2024**, *11*, 2415–2424. [CrossRef]
17. Gao, J.; Gui, W.; Xie, X.; Wan, R.; Yin, S.; Yang, C. Temperature-Insensitive Irreversible Demagnetization Fault Diagnosis Based on Dual-State Flux Estimation for PMSM. *IEEE Trans. Power Electron.* **2025**, *40*, 13324–13336. [CrossRef]
18. Abubakar, U.; Wang, X.; Shah, S.H.; Sheng, Y.; Maina, A.D. Coupled electromagnetic-LPTN model of high speed PMSM for mechanical vapor recompression applications. In *Proceedings of the International Conference on Electrical Machines and Systems (ICEMS), Chiang Mai, Thailand, 29 November–2 December 2022*; IEEE: New York, NY, USA, 2022; pp. 1–6.
19. Kirchgässner, W.; Wallscheid, O.; Böcker, J. Data-driven permanent magnet temperature estimation in synchronous motors with supervised machine learning: A benchmark. *IEEE Trans. Energy Convers.* **2021**, *36*, 2059–2067. [CrossRef]
20. Feng, G.; Lai, C.; Tjong, J.; Kar, N.C. Noninvasive Kalman filter based permanent magnet temperature estimation for permanent magnet synchronous machines. *IEEE Trans. Power Electron.* **2018**, *33*, 10673–10682. [CrossRef]
21. Kirchgässner, W.; Wallscheid, O.; Böcker, J. Thermal neural networks: Lumped-parameter thermal modeling with state-space machine learning. *Eng. Appl. Artif. Intell.* **2023**, *117*, 105537. [CrossRef]
22. Feng, J.; Liang, D.; Zhu, Z.Q.; Guo, S.; Li, Y.; Zhao, A. Improved low-order thermal model for critical temperature estimation of PMSM. *IEEE Trans. Energy Convers.* **2021**, *37*, 413–423. [CrossRef]
23. Zhang, S.; Zhou, Z.; Pu, Y.; Li, Y.; Xu, Y. Parameter Identification of Permanent Magnet Synchronous Motor Based on LSOSMO Algorithm. *Sensors* **2025**, *25*, 2648. [CrossRef]
24. Ahn, H.; Park, H.; Kim, C.; Lee, H. A review of state-of-the-art techniques for PMSM parameter identification. *J. Electr. Eng. Technol.* **2020**, *15*, 1177–1187. [CrossRef]
25. Wang, P.; Zhu, Z.Q.; Liang, D. Improved position-offset based online parameter estimation of PMSMs under constant and variable speed operations. *IEEE Trans. Energy Convers.* **2024**, *39*, 1325–1340. [CrossRef]
26. Huang, B.; Wang, J. Applications of physics-informed neural networks in power systems-a review. *IEEE Trans. Power Syst.* **2022**, *38*, 572–588. [CrossRef]
27. Wang, J.; Li, Y.; Gao, R.X.; Zhang, F. Hybrid physics-based and data-driven models for smart manufacturing: Modelling, simulation, and explainability. *J. Manuf. Syst.* **2022**, *63*, 381–391. [CrossRef]
28. Iliuţă, M.E.; Moisescu, M.A.; Pop, E.; Ionita, A.-D.; Caramihai, S.-I.; Mitulescu, T.-C. Digital twin—A review of the evolution from concept to technology and its analytical perspectives on applications in various fields. *Appl. Sci.* **2024**, *14*, 5454. [CrossRef]
29. Chen, S.; Bekar, E.T.; Bokrantz, J.; Skoogh, A. AI-enhanced digital twins in maintenance: Systematic review, industrial challenges, and bridging research–practice gaps. *J. Manuf. Syst.* **2025**, *82*, 678–699. [CrossRef]
30. Wallscheid, O.; Böcker, J. Global identification of a low-order lumped-parameter thermal network for permanent magnet synchronous motors. *IEEE Trans. Energy Convers.* **2016**, *31*, 354–365. [CrossRef]
31. Kruschke, J.K. Bayesian data analysis. *Wiley Interdiscip. Rev. Cogn. Sci.* **2010**, *1*, 658–676. [CrossRef]
32. Bi, S.; Beer, M.; Cogan, S.; Mottershead, J. Stochastic model updating with uncertainty quantification: An overview and tutorial. *Mech. Syst. Signal Process.* **2023**, *204*, 110784. [CrossRef]
33. Spigler, G. Denoising autoencoders for overgeneralization in neural networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **2019**, *42*, 998–1004. [CrossRef] [PubMed]
34. Li, W.; Huang, R.; Li, J.; Liao, Y.; Chen, Z.; He, G.; Yan, R.; Gryllias, K. A perspective survey on deep transfer learning for fault diagnosis in industrial scenarios: Theories, applications and challenges. *Mech. Syst. Signal Process.* **2022**, *167*, 108487. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.