

Diagnosing Physics-Informed Neural Networks for Motor Thermal Virtual Sensing: A Controlled Study on Single-Phase Induction Motors

Timothy A. Adeyi^{1,2}, Adrian D. Cheok¹, Steven Z. Zhou^{3,4}, Daniel O. Olasupo⁵,
Muhammed A. Makhdoom⁶, Samuel O. EnochOghene⁵

¹School of Automation, Nanjing University of Information Science and Technology, Nanjing, China

²Department of Mechanical Engineering, Lead City University, Ibadan, Nigeria

³NUS Suzhou Research Institute, Suzhou, China

⁴Department of Electrical Computer Engineering, National University of Singapore, Singapore

⁵Department of Electrical & Electronic Engineering, Lead City University, Ibadan, Nigeria

⁶School of Information & Communication Engineering, Nanjing University of Information Science and Technology, Nanjing, China

Email: adeyi.timothy@lcu.edu.ng, adrian@imagineeringinstitute.org, elezzy@nus.edu.sg, 202451520001@nuist.edu.cn, alimakhdoom968@gmail.com, enochoghene.samuel@lcu.edu.ng

How to cite this paper: Adeyi, T.A., Cheok, A.D., Zhou, S.Z., Olasupo, D.O., Makhdoom, M.A. and EnochOghene, S.O. (2026) Diagnosing Physics-Informed Neural Networks for Motor Thermal Virtual Sensing: A Controlled Study on Single-Phase Induction Motors. *Advances in Artificial Intelligence and Robotics Research*, 2, 127-146.

<https://doi.org/10.4236/airr.2026.23008>

Received: July 22, 2026

Accepted: August 22, 2026

Published: August 25, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Physics-informed neural networks (PINNs) have been proposed for motor thermal virtual sensing, yet the specific conditions under which embedded physics constraints outperform identical unconstrained networks remain unclear. This paper presents a controlled diagnostic study comparing a feedforward DNN and a PINN for winding temperature estimation in a 1 HP single-phase induction motor (SPIM). Unlike prior work, which focused on permanent magnet synchronous motors, this study addresses the distinct thermal modeling challenges of the SPIM's dual-winding structure by incorporating a first-order thermal ODE that accounts for main/auxiliary copper losses, bearing friction, and convective cooling. We systematically evaluate performance across five standard conditions: in-distribution operation, moderate load-profile shift, severe ambient temperature shift, cold-start transients, and sensor noise. Additionally, we conduct a self-supervised experiment and a systematic lambda sensitivity analysis across seven physics loss weights. To establish boundary conditions for physics prior effectiveness, we further evaluate both models under extreme distribution shifts (E1 - E3) and degraded data scenar-

ios (E4 - E7) encompassing sparse sampling, missing data, biased coverage, and noisy measurements. Results reveal a critical boundary condition: when temperature labels and temporal context are available, DNN and PINN achieve equivalent accuracy. However, when data are sparse, missing, biased, or noisy, the physics prior becomes valuable, reducing MAE by approximately 10% relative to DNN alone. The lambda analysis confirms that the DNN advantage is not an artifact of hyperparameter selection, while the self-supervised PINN fails ($\text{MAE} > 30^\circ\text{C}$), demonstrating that the simplified first-order prior is underdetermined without label supervision. These findings establish that physics-informed regularization using a simplified first-order prior is beneficial only under degraded data conditions, providing practitioners with clear decision rules for method selection in thermal virtual sensing.

Keywords

Diagnostic Study, Distribution Shift, Physics-Informed Neural Networks (PINNs), Single-Phase Induction Motor (SPIM), Thermal Virtual Sensing

1. Introduction

Electric motors are among the most widely deployed machines in modern engineering systems. They power industrial drives, household appliances, electric vehicles, pumps, fans, and compressors across every sector of the economy [1]-[3]. In all these applications, winding temperature is a critical operating variable. Elevated temperatures accelerate insulation degradation, reduce motor efficiency, and in severe cases, lead to premature failure [4]. Reliable temperature monitoring is therefore essential for condition monitoring, predictive maintenance, and protection system design.

Direct temperature measurement using embedded thermocouples or resistance temperature detectors is not always practical. In sealed motors, high-speed applications, and cost-sensitive consumer products, installing and maintaining physical temperature sensors introduces significant challenges in terms of cost, space, and long-term reliability [5]. This has motivated extensive research into virtual sensing, which involves estimating internal temperatures from electrical signals already available in the motor drive system, such as current, voltage, and speed.

Two broad classes of methods have been developed for motor thermal virtual sensing. Physics-based approaches, particularly lumped-parameter thermal networks (LPTNs), model the thermal behavior of the motor using electrical analogies of heat flow. These methods provide physically guaranteed estimates and generalize well across operating conditions, but they require detailed motor geometry, multi-node calibration, and access to thermal interface resistance data that is rarely available from manufacturers [6] [7]. Data-driven approaches, including support vector regression, convolutional neural networks, and long short-term memory networks, achieve high accuracy when sufficient labeled training data is

available but can produce thermodynamically inconsistent predictions when tested under conditions not represented in the training set [8].

Physics-informed neural networks (PINNs), introduced by Raissi *et al.* [9], offer a middle path between these two approaches. By embedding governing differential equations directly into the training loss function, PINNs impose physical consistency as a soft constraint during learning. Recent work has applied PINNs to motor thermal estimation with promising results. Wang *et al.* [10] combined a 10-node Motor-CAD-calibrated LPTN with neural network correction terms, achieving an MAE of 1.71°C on real PMSM data. Stensson [11] demonstrated PINN-based thermal modeling for electric motors, and Lim *et al.* [12] applied physics-informed learning to electric vehicle dynamics estimation.

Despite this progress, a fundamental question has not been answered: under what conditions does embedding a physics prior into the loss function actually improve performance compared to an identical unconstrained network? Existing studies propose PINN architectures and report performance on fixed test sets, but none isolates the marginal contribution of the physics constraint by holding everything else constant. If a physics prior only helps under conditions that are rarely encountered in practice, its added complexity may not be justified. If it provides robustness specifically where data-driven models fail, under distribution shift, sensor noise, or data scarcity, practitioners need to know this clearly.

A second gap exists in the motor type coverage of existing work. The vast majority of PINN-based thermal estimation studies focus on permanent magnet synchronous motors [10] [11] [13]. The single-phase induction motor, which is extensively used in household appliances, small pumps, and light industrial equipment worldwide, has received almost no attention in the neural network-based thermal estimation literature. The dual-winding structure of the SPIM, with separate main and auxiliary windings, each contributing to copper losses, presents a distinct thermal modeling challenge that the PMSM-oriented literature does not address.

This paper addresses both gaps through a controlled diagnostic study. A high-fidelity Simulink model of a 1 HP single-phase induction motor is used to generate five systematically varied datasets covering in-distribution operation, moderate load-profile shift, severe ambient temperature shift, cold-start transients, and noisy sensor measurements. An identical feedforward neural network architecture is trained with and without a first-order thermal ODE prior across all conditions. A self-supervised experiment, training without any temperature labels, is also conducted to establish the standalone capability of the physics prior.

The contributions of this paper are as follows:

- 1) The first controlled diagnostic comparison of DNN and PINN for single-phase induction motor thermal estimation uses identical architectures, varying only the loss function.
- 2) A dual-winding thermal ODE formulation for SPIM that explicitly accounts for main and auxiliary winding copper losses and bearing friction losses, following Calasan *et al.* [14].

3) A systematic lambda sensitivity analysis demonstrates that the DNN advantage over the PINN is not an artifact of hyperparameter selection but a consistent property across the full range of physics loss weights tested.

4) A self-supervised experiment to establish the minimum data requirements for physics-informed thermal estimation.

Data and Code Availability: To ensure reproducibility and facilitate further research, the Simulink model, Python training code, generated datasets, and all figures used in this study are publicly available at <https://github.com/Pstone123/SPIM-Thermal-PINN-v2>.

2. Related Work

2.1. Motor Thermal Estimation

The thermal monitoring of electric motors has been comprehensively reviewed by Wallscheid [5], who identified virtual sensing as a critical direction for condition monitoring in sealed and high-speed applications. Physics-based approaches have been the traditional standard. Wallscheid and Bocker [6] developed a globally identified four-node LPTN for PMSMs, demonstrating that accurate thermal modeling can be achieved with relatively few parameters when the network topology is well chosen. Kral *et al.* [7] proposed a practical two-node model for estimating permanent magnet and stator winding temperatures, which has since been widely adopted as a benchmark for lumped-parameter approaches. Chowdhury and Baski [15] demonstrated a simple lumped-parameter thermal model for totally enclosed fan-cooled machines, highlighting the trade-off between model simplicity and accuracy.

Data-driven approaches have advanced significantly over the past decade. The Paderborn PMSM benchmark introduced by Kirchgässner *et al.* [16] provided the research community with a publicly available real test bench dataset covering 139 hours of automotive PMSM operation, enabling reproducible comparison of data-driven methods. Jin *et al.* [8] proposed a model-based and data-driven integrated temperature estimation method that combines a simplified LPTN with a neural network correction layer, achieving strong performance on real motor data. These hybrid approaches represent an intermediate position between pure physics-based and pure data-driven methods.

2.2. Physics-Informed Neural Networks

The PINN framework, introduced by Raissi *et al.* [9], embeds governing differential equations into the loss function through an additional residual term computed via automatic differentiation. This approach has been applied across a wide range of engineering domains, from fluid dynamics to structural mechanics. In the motor thermal estimation domain, Wang *et al.* [10] demonstrated that combining a high-fidelity 10-node LPTN prior with neural network correction terms achieves an MAE of 1.71 °C on real PMSM data, substantially better than prior purely data-driven results on the same hardware. Stensson [11] showed that PINN-based ther-

mal modeling can be applied to electric motors using supervised training with temperature labels. Lim *et al.* [12] applied physics-informed learning to electric vehicle dynamics, demonstrating the potential of minimal-input PINN models for automotive applications.

Reinbold *et al.* [17] demonstrated that physically constrained learning can improve robustness on noisy, incomplete experimental data, which motivated the noise robustness experiment in the present study. Chen *et al.* [18] proposed gradient normalization as a technique for balancing competing loss terms in multi-task neural networks, which is directly relevant to the challenge of scaling the physics and data loss terms in PINN training, a challenge explicitly addressed in this paper through residual normalization.

2.3. Diagnostic Studies of PINN Performance

The broader PINN literature has begun to examine conditions under which physics-informed learning fails or underperforms. Krishnapriyan *et al.* [19] characterized failure modes in PINNs applied to stiff partial differential equations, identifying gradient pathologies and loss landscape properties that prevent convergence. These studies focus on forward and inverse PDE problems in computational physics rather than on engineering state estimation where the governing equation is a simplified approximation of the true system dynamics. No equivalent diagnostic study exists for motor thermal estimation. The present work fills this gap by providing the first systematic comparison of identical DNN and PINN architectures across controlled variations of data condition, distribution shift type, and physics loss weighting for motor thermal estimation.

3. Methodology

3.1. SPIM Simulation Model

A single-phase induction motor simulation model was developed in MATLAB/Simulink R2025b using the Simscape Electrical capacitor-start-run block. The motor was parameterized using manufacturer datasheet values for a 1 HP, 230 V, 50 Hz SPIM. The main winding stator resistance is $2.02\ \Omega$, the auxiliary winding stator resistance is $7.14\ \Omega$, and the rotor resistance referred to the stator is $4.12\ \Omega$. The rotor inertia is $0.0146\ \text{kg}\cdot\text{m}^2$. The motor was supplied through an H-bridge inverter configuration, and the mechanical load applied at the shaft port was controlled through a time-varying MATLAB function block.

The mechanical load torque was designed as a stepped time-varying profile to create rich, dynamic operating conditions suitable for neural network training. For the training dataset, the load torque alternated between 2.0 Nm and 3.5 Nm at 20-second intervals over a total simulation duration of 200 seconds. This represents a realistic scenario in which a motor experiences varying mechanical demand, such as a pump or fan operating under changing flow conditions. The load values were selected to maintain motor operation within the normal torque range, with rated torque estimated at approximately 4.75 Nm for a 1 HP, 50 Hz, four-

pole machine.

3.2. Thermal Subsystem

A first-order lumped thermal subsystem was coupled to the electrical model to compute winding temperature dynamics. The governing thermal equation is:

$$C_{th} \frac{dT}{dt} = I_{main}^2 R_{main}(T) + I_{aux}^2 R_{aux}(T) + B\omega - h(T - T_{amb}) \quad (1)$$

where T is the winding temperature in degrees Celsius, $C_{th} = 8000 \text{ J/}^\circ\text{C}$ is the thermal capacitance, $h = 28 \text{ W/}^\circ\text{C}$ is the convective heat transfer coefficient, T_{amb} is the ambient temperature, and ω is the rotor angular speed in rad/s. The first two terms represent copper losses in the main and auxiliary windings, respectively. The temperature dependence of copper resistance is modeled as:

$$R(T) = R_0 \left[1 + \alpha (T - T_{ref}) \right] \quad (2)$$

where R_0 is the reference resistance at T_{ref} , $\alpha = 0.00393/^\circ\text{C}$ is the temperature coefficient of copper, and $T_{ref} = 25^\circ\text{C}$ is the reference temperature. The third term, $B\omega$, represents friction and bearing losses following the viscous friction model proposed by Calasan *et al.* [14], where $B = 0.0005 \text{ Nm}\cdot\text{s/rad}$ is the bearing loss coefficient. The fourth term represents convective cooling to the ambient environment. Rotor copper losses were not included, as rotor current is not directly observable from stator measurements in this configuration. Core iron losses were excluded, as the iron loss coefficient is not available from the manufacturer datasheet. Both exclusions are consistent with virtual sensing approaches in the literature and are acknowledged as simplifying assumptions.

The model was solved using the variable-step ode45 solver with automatic step size control. Simulation outputs were sampled at 10 Hz (sample interval $\Delta t = 0.1 \text{ s}$), yielding 2001 samples per 200-second simulation. Current signals were processed through RMS blocks with a 0.1-second sliding window prior to logging, converting raw AC waveforms to effective heating values appropriate for thermal modeling.

3.3. Dataset Generation

Five datasets were generated by systematically varying the mechanical load profile and ambient temperature conditions. All datasets share the same input feature set: rotor speed (RPM), main winding RMS current (A), auxiliary winding RMS current (A), electromagnetic torque (Nm), elapsed time (s), and ambient temperature ($^\circ\text{C}$). The target variable is winding temperature ($^\circ\text{C}$). **Table 1** summarizes the dataset properties.

Dataset A served as the sole training source for all models. Datasets B1, B2, C, and D were used exclusively for evaluation; no model was retrained or fine-tuned on these sets. This design ensures that reported performance differences reflect genuine generalization capability rather than memorization.

For Dataset D, additive Gaussian noise with a standard deviation equal to 5%

of each signal's full-scale range was applied independently to speed, main current, and auxiliary current measurements after simulation. Temperature labels remained clean, reflecting the assumption that temperature sensors are more reliable than electrical measurement channels. The resulting temperature trajectories for these datasets are visualized in **Figure 1**.

Table 1. SPIM dataset summary.

Dataset	Purpose	Load Profile	Ambient (°C)	Duration (s)	Samples
A	Training	Stepped 2.0 - 3.5 Nm, 20 s intervals	25	200	2001
B1	Moderate shift	Different step patterns, varied magnitudes	25	200	2001
B2	Severe shift	Same as A, elevated ambient.	40	200	2001
C	Cold start	Same as A, early transient only.	25	100	1001
D	Sensor noise	Same as A, 5% Gaussian noise on inputs.	25	200	2001




Figure 1. Winding temperature trajectories for all four SPIM datasets.

To systematically investigate the boundary conditions of physics prior effectiveness, seven additional datasets (E1 - E7) were generated. Datasets E1 - E3 represent extreme distribution shifts: E1 (50°C ambient temperature), E2 (2× load scaling), and E3 (extended 500 s duration). Datasets E4 - E7 represent degraded data scenarios: E4 (sparse sampling at 0.5 s intervals), E5 (missing data with 20% of samples randomly removed), E6 (biased coverage with training restricted to high-torque operating regions), and E7 (noisy inputs with 10% Gaussian noise). These datasets are used exclusively for evaluation; no model is retrained on them.

3.4. Neural Network Architecture

Both the DNN and PINN share an identical feedforward architecture to ensure a fair, controlled comparison. The network takes five inputs: ω , I_{main} , I_{aux} , T_∞ , and t , which are rotor speed, main winding current, auxiliary winding current, torque, and time, respectively, with a single temperature output (T). Three hidden layers with 20, 20, and 10 neurons, respectively, use hyperbolic tangent (tanh) activation functions. The output layer is linear. This gives 761 trainable parameters, a deliberately compact architecture chosen to avoid overfitting on the relatively small training set of 2001 samples.

All inputs and the temperature output were normalized to $[0, 1]$ using the MinMaxScaler fitted exclusively on Dataset A. The same scaler was applied to all other datasets without refitting, preserving the distributional mismatch that defines the shift conditions.

3.5. Loss Functions

The DNN was trained using the mean squared error between predicted and true normalized temperature:

$$L_{DNN} = \frac{1}{N} \sum_{i=1}^N (\hat{T}_i - T_{true,i})^2 \quad (3)$$

The PINN was trained using a combined loss:

$$L_{PINN} = L_{data} + \lambda \cdot L_{physics} \quad (4)$$

where λ is the physics loss weight and $L_{physics}$ is the mean squared residual of the thermal ODE evaluated at each training point. The time derivative ($\frac{d\hat{T}}{dt}$) was computed via automatic differentiation with respect to the time input, following the standard PINN formulation of Raissi *et al.* [9]. The physics residual is:

$$L_{physics} = \frac{1}{N} \sum_{i=1}^N \left[\frac{C_{th} \frac{d\hat{T}}{dt} - (I_{main}^2 R_{main}(\hat{T}) + I_{aux}^2 R_{aux}(\hat{T}) + B\omega - h(\hat{T} - T_{amb}))}{1000} \right]^2 \quad (5)$$

The division by 1000^2 normalizes the residual, which operates in units of watts squared, to the same order of magnitude as the data loss, which operates in normalized temperature units. Without this scaling, the physics gradient overwhelms the data gradient at initialization, preventing convergence to a physically and empirically consistent solution.

The self-supervised PINN was trained with $\lambda \rightarrow \infty$ ($L_{data} = 0$), using only the physics residual as the training signal. This establishes the standalone estimation capability of the thermal prior without any temperature label supervision.

3.6. Training Protocol

All models were implemented in PyTorch and trained for 3000 epochs using the Adam optimizer with a fixed learning rate of 1×10^{-3} and full-batch gradient up-

dates. A fixed random seed (42) was used throughout for reproducibility. Training was performed on the CPU using Google Colab.

A systematic lambda sensitivity analysis was performed by training separate PINN instances with $\lambda \in \{1 \times 10^{-6}, 5 \times 10^{-6}, 1 \times 10^{-5}, 5 \times 10^{-5}, 1 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}\}$ and evaluating each on both Dataset A and Dataset B1. This analysis establishes whether the DNN superiority observed at any single λ value is an artifact of hyperparameter selection or a consistent property across the physics loss weighting range.

3.7. Evaluation Metrics

Model performance was quantified using three standard metrics for regression tasks in motor thermal estimation:

$$\text{Mean Absolute Error: } \text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{T}_i - T_{\text{true},i}| \quad (6)$$

$$\text{Root Mean Square Error: } \text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{T}_i - T_{\text{true},i})^2} \quad (7)$$

$$\text{Coefficient of Determination: } R^2 = 1 - \frac{\sum_{i=1}^N (T_{\text{true},i} - \hat{T}_i)^2}{\sum_{i=1}^N (T_{\text{true},i} - \bar{T})^2} \quad (8)$$

where $\bar{T}$ is the mean of the true temperature values over the evaluation set. MAE provides an interpretable error in degrees Celsius, while RMSE penalizes larger errors more heavily. R^2 indicates how well the predictions track the true temperature trajectory relative to the mean; a value below zero indicates performance worse than simply predicting the mean temperature.

4. Results

4.1. In-Distribution Performance (Dataset A)

Both models were trained exclusively on Dataset A and evaluated first on the same dataset to establish their baseline fitting capability. The DNN achieved an MAE of 0.066°C , RMSE of 0.077°C , and R^2 of 0.9999, indicating near-perfect reconstruction of the training temperature trajectory. The PINN achieved an MAE of 0.079°C , RMSE of 0.108°C , and R^2 of 0.9997. Both models fit the training data well, but the DNN fits more tightly, as expected. When a neural network is given sufficient clean data from the same distribution it was trained on, the physics constraint introduces a competing objective that prevents the model from fitting the data as closely as an unconstrained network. This small performance gap on in-distribution data is consistent with the general behavior of regularized versus unregularized models and does not indicate a problem with the PINN formulation.

4.2. Moderate Distribution Shift (Dataset B1)

Dataset B1 represents a change in load step timing and magnitude that the models were never exposed to during training. Under this condition, the DNN MAE in-

creased from 0.066°C to 0.434°C —a degradation of approximately 6.6 times relative to in-distribution performance. The R^2 remained at 0.9939, indicating that while the DNN’s absolute errors increased, it still tracked the general temperature trend reasonably well. The PINN produced an MAE of 0.471°C and R^2 of 0.9926 under the same condition, slightly worse than the DNN on both metrics.

A systematic sweep of the physics loss weight λ from 1×10^{-6} to 1×10^{-3} was conducted to determine whether any value of λ could produce a PINN advantage over the DNN. The results are presented in **Table 2**. Across all tested values, PINN performance degraded monotonically as λ increased on both Dataset A and Dataset B1. No value of λ produced a PINN that outperformed the DNN. At $\lambda = 1 \times 10^{-6}$, the physics constraint had negligible effect and the PINN matched the DNN almost exactly. At $\lambda = 1 \times 10^{-3}$, both in-distribution and shift performance degraded substantially. This confirms that the DNN superiority on Dataset B1 is not an artifact of suboptimal hyperparameter selection. This trend is visually summarized in **Figure 2**, which plots the MAE against the physics loss weight λ .

Table 2. Lambda sensitivity analysis—PINN performance vs. physics loss weight.

Lambda (λ)	A MAE ($^{\circ}\text{C}$)	A (R^2)	B1 MAE ($^{\circ}\text{C}$)	B1 (R^2)
DNN baseline	0.0660	0.9999	0.4338	0.9939
1×10^{-6}	0.0660	0.9999	0.4342	0.9939
5×10^{-6}	0.0660	0.9999	0.4356	0.9938
1×10^{-5}	0.0661	0.9999	0.4375	0.9938
5×10^{-5}	0.0694	0.9998	0.4524	0.9932
1×10^{-4}	0.0785	0.9997	0.4710	0.9926
5×10^{-4}	0.2230	0.9975	0.6102	0.9871
1×10^{-3}	0.3824	0.9925	0.7707	0.9792




Figure 2. PINN MAE as a function of physics loss weight λ on Dataset A and Dataset B1, with the DNN baseline shown as dashed lines.

4.3. Severe Distribution Shift (Dataset B2)

Dataset B2 introduced a more substantial distributional change by elevating the ambient temperature from 25°C to 40°C while keeping the load profile identical to Dataset A. Under this condition, the DNN MAE rose to 1.850°C and R^2 fell to 0.9270, a considerably larger degradation than observed under the moderate load shift of Dataset B1. This confirms that the ambient temperature shift created a more challenging test condition. The PINN produced an MAE of 1.871°C and R^2 of 0.9247. The gap between DNN and PINN remained small, with the DNN still marginally outperforming the PINN. Notably, R^2 values for both models dropped below 0.93, indicating meaningful degradation in prediction quality under this more severe shift. However, neither model showed a clear advantage, suggesting that under this type of thermal operating point shift, the simplified first-order prior does not provide sufficient constraint to counteract the degradation caused by the unseen ambient condition.

4.4. Cold Start Transient and Noisy Measurements

Dataset C captured the first 100 seconds of the same loading scenario as Dataset A, representing the early thermal transient from ambient temperature. The DNN achieved an MAE of 0.063°C and an R^2 of 0.9996, slightly better than its full-duration in-distribution performance. The PINN achieved an MAE of 0.079°C and an R^2 of 0.9991. The DNN's superior performance on Dataset C is consistent with the in-distribution finding; the early transient falls within the range of conditions seen during training, and the DNN interpolates accurately without requiring physical guidance. The physics prior does not improve performance when the test condition is well covered by the training data.

Dataset D introduced 5% Gaussian noise to the speed and current measurement channels, representing sensor uncertainty in a practical deployment scenario. The DNN MAE increased modestly from 0.066°C to 0.128°C, while R^2 remained at 0.9994. The PINN produced an MAE of 0.131°C and R^2 of 0.9994. Both models showed similar robustness to the noise level tested, with neither gaining a measurable advantage. The physics prior did not function as a noise filter at this noise level. This may be because the 5% noise, while detectable, did not substantially alter the underlying input-output mapping that the DNN had learned from clean data.

4.5. Hard Distribution Shift Evaluation

To test the limits of both models under extreme conditions not represented in the training data, three additional datasets (E1-E3) were evaluated. Dataset E1 elevated the ambient temperature to 50°C, a 25°C increase from the training ambient of 25°C. Dataset E2 applied a 2× load scaling, subjecting the motor to severe mechanical overload. Dataset E3 extended the simulation duration to 500 seconds, capturing cumulative thermal buildup beyond the 200-second training window.

The results are presented in **Table 3**. Under E1 (50°C ambient), the DNN

achieved an MAE of 13.63°C while the PINN achieved 13.77°C. Under E2 (2× load), both models produced MAEs of 13.63°C and 13.77°C, respectively. Under E3 (500 s extended run), the DNN achieved an MAE of 27.85°C and the PINN achieved 28.79°C.

Table 3. Hard distribution shift results.

Dataset	Condition	DNN MAE (°C)	PINN MAE (°C)
E1	50°C ambient	13.63	13.77
E2	2× load scaling	13.63	13.77
E3	500s extended run	27.85	28.79

Both models degraded substantially under these extreme shifts, with MAE increasing significantly relative to in-distribution performance. Critically, the physics prior provided no protective advantage; the DNN and PINN performed equivalently within the margin of measurement uncertainty.

4.6. Sampling Robustness and Data Degradation

To address the question of whether physics-informed regularization provides value under practical data limitations, four degraded-data scenarios (E4 - E7) were evaluated. Dataset E4 employed sparse sampling at 0.5 s intervals (2 Hz) instead of the original 0.1 s (10 Hz). Dataset E5 simulated missing data by randomly removing 20% of training samples. Dataset E6 introduced biased coverage by training exclusively on high-torque operating regions. Dataset E7 added 10% Gaussian noise to input measurements.

The results are presented in **Table 4**. The PINN outperformed the DNN across all four degraded-data conditions. Under E4 (sparse sampling), the PINN achieved an MAE of 15.44°C compared to the DNN MAE of 17.21°C, a 10.3% improvement. Under E5 (missing data), the PINN MAE was 15.43°C versus the DNN 17.22°C (10.4% improvement). Under E6 (biased coverage), the PINN MAE was 11.87°C versus the DNN 13.33°C (11.0% improvement). Under E7 (noisy inputs), the PINN MAE was 15.45°C versus the DNN 17.22°C (10.3% improvement).

Table 4. Sampling robustness results.

Condition	DNN MAE (°C)	PINN MAE (°C)	Ratio (PINN/DNN)
E4: Sparse sampling (0.5 s)	17.21	15.44	0.897
E5: Missing data (20%)	17.22	15.43	0.896
E6: Biased coverage	13.33	11.87	0.891
E7: Noisy inputs (10%)	17.22	15.45	0.897

This finding establishes a critical boundary condition: the physics prior provides measurable benefit (approximately 10% MAE reduction) precisely when

data quality degrades. Under conditions of complete, clean, and temporally dense data, the unconstrained DNN is sufficient. However, when data are sparse, missing, biased, or noisy, the physics-informed regularizer compensates for information loss.

4.7. Self-Supervised PINN

A self-supervised PINN was trained using only the physics residual loss, with no temperature labels provided during training. After 3000 epochs, the physics loss converged to near zero, indicating that the network found a solution that approximately satisfies the thermal ODE as shown in the training loss curves in **Figure 3**. However, evaluation against true temperature labels revealed MAE values of approximately 30°C across all datasets, with R^2 values below -18 . A negative R^2 of this magnitude indicates that the model performs far worse than simply predicting the mean temperature; the physically consistent solution found by the optimizer does not correspond to the actual thermal trajectory of the motor.

This result confirms that the first-order thermal ODE, while thermodynamically valid, is underdetermined as a standalone estimator. The ODE admits multiple solutions that satisfy the residual criterion; the network converges to one of these solutions, but without temperature labels to anchor it, there is no guarantee that this solution matches physical reality. Self-supervised thermal estimation using a simplified first-order prior is therefore not viable with the current model structure.





Figure 3. Training loss convergence curves for DNN, PINN, and self-supervised PINN over 3,000 epochs.

4.8. Summary of SPIM Results

Table 5 presents a complete summary of all SPIM results on datasets A-D and the self-supervised experiment. **Figure 4** provides a comparative bar chart of the MAE across these conditions.

Table 5. SPIM Results—DNN vs. PINN across all conditions.

Dataset	Condition	DNN (MAE)	DNN (R ²)	PINN (MAE)	PINN (R ²)	Winner
A	Training	0.0660	0.9999	0.0785	0.9997	DNN
B1	Moderate load shift	0.4338	0.9939	0.4710	0.9926	DNN
B2	Severe ambient shift	1.8500	0.9270	1.8705	0.9247	DNN
C	Cold start	0.0629	0.9996	0.0786	0.9991	DNN
D	Sensor noise	0.1276	0.9994	0.1313	0.9994	DNN
Self-supervised	No labels	—	—	~30	< -18	Fails

**Figure 4.** MAE comparison between DNN and PINN across all evaluation datasets ($\lambda = 1 \times 10^{-4}$).

The evaluation on datasets E1 - E7, summarized in **Table 4** and **Table 5**, further clarifies the boundary conditions of physics prior effectiveness. Under extreme distribution shifts (E1 - E3), both models degraded substantially and equally, with MAE increasing to 13.6°C - 28.8°C and the physics prior providing no protective advantage. However, under degraded data scenarios (E4 - E7), the PINN consistently outperformed the DNN, achieving approximately 10% lower MAE. This establishes a critical boundary condition: the physics prior becomes valuable when data quality degrades, while offering no benefit under complete data conditions or extreme distribution shifts.

5. Discussion

5.1. Boundary Conditions for Physics-Prior Effectiveness

The results, summarized in **Tables 3-5**, reveal a nuanced picture of physics-informed neural network effectiveness. The findings can be categorized into two

distinct regimes:

Regime 1—Complete Data with Temporal Context: When temperature labels are abundant, sampling is dense, data coverage is unbiased, and time is available as an input feature, DNN and PINN achieve equivalent accuracy (datasets A–D). The physics prior provides no measurable advantage because the data alone contain sufficient information to learn the thermal mapping.

Regime 2—Degraded Data: When data quality degrades (sparse sampling, missing observations, biased operating coverage, or noisy measurements—datasets E4–E7), the physics prior provides meaningful regularization, reducing MAE by approximately 10% relative to the DNN alone. This finding demonstrates that physics-informed learning is valuable not as a universal improvement but as a targeted tool for data-limited scenarios.

Under extreme distribution shifts (E1–E3), both models degraded substantially and equally. The physics prior offered no protection against extrapolation failure, consistent with the explanation that the first-order ODE, while encoding thermodynamic relationships, does not capture all relevant physical mechanisms (e.g., core iron losses, rotor copper losses, multi-dimensional heat paths).

5.2. Why the Physics Prior Did Not Help

There are three reasons why the simplified thermal ODE failed to provide a measurable benefit under any tested conditions.

First, the distribution shifts tested here, while genuine, did not degrade the DNN's interpolation capability to the point of failure. On Dataset B1, the DNN achieved an R^2 of 0.9939 even under the unseen load profile. On Dataset B2, R^2 fell to 0.9270, a meaningful degradation, but still far from the regime where a data-driven model breaks down entirely. The physics prior is most useful when a model is forced to extrapolate far outside its training distribution. In this study, both shift conditions were moderate enough that the DNN retained reasonable predictive capability without physical guidance.

Second, the simplified first-order ODE is an incomplete representation of the motor's actual thermal dynamics. The thermal prior used here accounts for copper losses in both windings, convective cooling, and bearing friction losses following Calasan *et al.* [14]. However, it excludes core iron losses, rotor copper losses, thermal interface resistances, and multi-dimensional heat flow paths. When the physics prior does not capture the dominant thermal mechanisms accurately, it does not guide the model toward better predictions; it guides it toward a thermodynamically approximate but empirically inaccurate solution. This is precisely what the lambda sensitivity analysis revealed: as the physics weight increased, both in-distribution and shift performance degraded monotonically. There was no operating point where physics helped.

Third, the self-supervised experiment confirmed that the ODE prior is under-determined. A network trained purely on the physics residual converged to a near-zero residual loss but produced temperature predictions with an MAE of approx-

imately 30°C. This demonstrates that multiple solutions satisfy the simplified ODE; the physics constraint alone cannot uniquely identify the motor's thermal trajectory. Temperature labels are necessary to anchor the solution, and when they are present in sufficient quantity, the data-driven model uses them effectively without needing physical guidance.

5.3. Comparison with Prior Work

The finding here contrasts with results reported by Wang *et al.* [10], who demonstrated a physics-informed approach achieving an MAE of 1.71°C on real PMSM data. However, there is an important difference in approach. Wang *et al.* used a 10-node Motor-CAD-calibrated lumped-parameter thermal network, a high-fidelity multi-node prior that captures inter-component heat transfer, not a single first-order ODE. A more faithful physics prior naturally provides stronger guidance and is more likely to yield measurable benefits. The present study tests the other end of this spectrum, a minimal first-order prior, and the results suggest that the benefit of physics-informed regularization is strongly dependent on the fidelity of the physics model used.

Stensson [11] and Lim *et al.* [12] similarly demonstrate PINN benefits under specific conditions, but neither study isolates the marginal contribution of the physics prior by comparing against an identical unconstrained baseline under the same conditions. The controlled experimental design in the present work, with identical architecture and data and only the loss function differing, allows a cleaner attribution of performance differences specifically to the physics constraint.

5.4. The Single-Phase Induction Motor—An Underexplored Domain

Prior PINN work on motor thermal estimation has focused almost exclusively on PMSMs [10] [11] [16]. The single-phase induction motor, despite its widespread use in household appliances, small pumps, and light industrial equipment, has received almost no attention in the neural network-based thermal estimation literature. The dual-winding structure of the SPIM, with separate main and auxiliary windings, each contributing to copper losses, requires a more nuanced thermal prior than the single-winding models commonly used for PMSMs. This study provides the first controlled diagnostic comparison of DNN and PINN for SPIM thermal estimation and establishes the thermal ODE formulation for this motor type with dual-winding copper loss terms.

5.5. Practical Implications

The results of this study convey a clear practical message for engineers deploying motor thermal estimators.

When training data are available and cover the intended operating range, a well-trained feedforward DNN is the preferred choice. It is simpler to implement, easier to train, requires no motor-specific thermal parameters, and consistently

achieves lower prediction error than a PINN using a simplified physics prior.

The physics prior adds value only when the simplified ODE is a faithful approximation of the dominant thermal dynamics and when the distribution shift is severe enough to substantially degrade data-driven interpolation. Neither condition was met in this study. For applications where these conditions are expected, such as motors operating under highly variable or unpredictable duty cycles with limited training data, a higher-fidelity prior, such as a multi-node LPTN calibrated to the specific motor, would be more appropriate than the first-order ODE used here.

Self-supervised thermal estimation, training without temperature labels using only the physics loss, is not viable with a simplified first-order prior. The under-determined nature of the ODE means that the network converges to physically plausible but empirically incorrect solutions. Applications that require label-free deployment will need either higher-fidelity physics models or alternative approaches, such as transfer learning or domain adaptation.

5.6. Limitations

Several limitations of this study should be noted.

First, the SPIM datasets were generated from a Simulink simulation model parameterized from manufacturer datasheet values. While the model produces physically plausible thermal trajectories, confirmed by steady-state temperature, speed, and torque values within expected ranges for a 1 HP motor, it has not been validated against direct thermocouple measurements from a physical motor. The absolute MAE values reported here reflect simulation-to-simulation performance and should not be compared directly to results from studies using real hardware measurements. Validation on real motor hardware remains an important direction for future work.

Second, the distribution shifts tested, while covering a range of operating conditions, were moderate in severity. The most severe shift tested, an ambient temperature change from 25°C to 40°C, degraded DNN R^2 to 0.927, which still represents reasonable predictive capability. Under more extreme shifts, such as a change in motor drive topology, significantly different ambient environments, or operation in completely different load regimes, the physics prior might provide a measurable advantage. This represents a direction for future investigation.

Third, the network architecture was fixed throughout the study (761 parameters). Architecture sensitivity was not explored; deeper or wider networks may alter the observed boundary conditions. Future work should examine whether network capacity affects the relative performance of DNN and PINN.

Fourth, the lambda sensitivity analysis fixed all other hyperparameters (optimizer, learning rate, loss scaling). Alternative optimizer settings and loss-balancing strategies such as GradNorm [18] merit future investigation.

Fifth, the study focused on a single motor type (SPIM). Generalization of these findings to three-phase induction motors, brushless DC motors, or switched reluctance motors requires additional studies.

6. Conclusions

This controlled diagnostic study demonstrates that the effectiveness of a first-order thermal physics prior for motor thermal virtual sensing is conditional, not universal. When temperature labels are abundant and temporal context is available, physics-informed and data-driven networks achieve equivalent accuracy. When data quality degrades, the physics prior provides meaningful regularization.

These findings establish a practical decision boundary for engineers deploying motor thermal estimators. When training data adequately covers the intended operating range, a well-trained DNN is the preferred choice. When data quality is compromised by sparse sampling, missing observations, biased coverage, or measurement noise, a physics-informed approach offers modest but meaningful improvements.

Future work should validate these findings on real motor hardware with direct thermocouple measurements, explore architecture sensitivity across network sizes and depths, and investigate whether higher-fidelity multi-node thermal priors provide greater benefits under extreme shifts.

Author Contributions

Conceptualization, T.A.A. and A.D.C.; methodology, T.A.A. and S.Z.Z.; software, T.A.A. and D.O.O.; validation, T.A.A., M.A.M., and E.S.O. formal analysis, T.A.A. and S.Z.Z.; investigation, D.O.O. and E.S.O.; resources, A.D.C. and S.Z.Z.; data curation, T.A.A. and D.O.O.; writing—original draft preparation, T.A.A.; writing—review and editing, A.D.C., S.Z.Z., and M.A.M.; visualization, T.A.A. and E.S.O.; supervision, A.D.C. and S.Z.Z.; project administration, A.D.C. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- [1] Mostajeran, E., Amiri, N., Ebrahimi, S. and Jatskevich, J. (2023) Electrical Machines in Electromagnetic Transient Simulations: Focusing on Efficient and Accurate Models. *IEEE Electrification Magazine*, **11**, 38-53.
<https://doi.org/10.1109/mele.2023.3320508>
- [2] Jing, H., Xiao, D., Wang, X., Chen, Z., Fang, G. and Guo, X. (2022) Temperature Estimation of Permanent Magnet Synchronous Motors Using Support Vector Regression. *2022 25th International Conference on Electrical Machines and Systems (ICEMS)*, Chiang Mai, 29 November-2 December 2022, 1-6.
<https://doi.org/10.1109/icems56177.2022.9983067>
- [3] Boldea, I. (2017) Electric Generators and Motors: An Overview. *CES Transactions on Electrical Machines and Systems*, **1**, 3-14.
<https://doi.org/10.23919/tems.2017.7911104>
- [4] Kim, D.W., Kang, D.H., Kim, C.H., Kim, J.S., Kim, Y.J. and Jung, S.Y. (2020) Operation Characteristic of IPMSM Considering PM Saturation Temperature. *IEEE Trans-*

- actions on Applied Superconductivity*, **30**, 1-4.
<https://doi.org/10.1109/tasc.2020.2989799>
- [5] Wallscheid, O. (2021) Thermal Monitoring of Electric Motors: State-of-the-Art Review and Future Challenges. *IEEE Open Journal of Industry Applications*, **2**, 204-223.
<https://doi.org/10.1109/ojia.2021.3091870>
 - [6] Wallscheid, O. and Bocker, J. (2016) Global Identification of a Low-Order Lumped-Parameter Thermal Network for Permanent Magnet Synchronous Motors. *IEEE Transactions on Energy Conversion*, **31**, 354-365.
<https://doi.org/10.1109/tec.2015.2473673>
 - [7] Kral, C., Haumer, A. and Lee, S.B. (2014) A Practical Thermal Model for the Estimation of Permanent Magnet and Stator Winding Temperatures. *IEEE Transactions on Power Electronics*, **29**, 455-464. <https://doi.org/10.1109/tpe.2013.2253128>
 - [8] Jin, L., Mao, Y., Wang, X., Lu, L. and Wang, Z. (2024) A Model-Based and Data-Driven Integrated Temperature Estimation Method for PMSM. *IEEE Transactions on Power Electronics*, **39**, 8553-8561. <https://doi.org/10.1109/tpe.2024.3382300>
 - [9] Raissi, M., Perdikaris, P. and Karniadakis, G.E. (2019) Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations. *Journal of Computational Physics*, **378**, 686-707. <https://doi.org/10.1016/j.jcp.2018.10.045>
 - [10] Wang, P., Wang, X. and Wang, Y. (2023) Physics-Informed Machine Learning Based Permanent Magnet Synchronous Motor Temperature Estimation. 2023 9th International Conference on Mechanical and Electronics Engineering (ICMEE), Xi'an, 17-19 November 2023, 582-587. <https://doi.org/10.1109/icmee59781.2023.10525605>
 - [11] Stensson, J. and Svantesson, K. (2023) Physics Informed Neural Network for Thermal Modeling of an Electric Motor. Master's Thesis, Chalmers University of Technology.
 - [12] Lim, H., Lee, J.W., Boyack, J. and Choi, J.B. (2024) EV-PINN: A Physics-Informed Neural Network for Predicting Electric Vehicle Dynamics. arXiv: 2411.14691.
 - [13] Raissi, M., Perdikaris, P. and Karniadakis, G.E. (2017) Physics Informed Deep Learning (Part I): Data-Driven Solutions of Nonlinear Partial Differential Equations. arXiv: 1711.10561.
 - [14] Čalasan, M., Alqarni, M., Rosić, M., Koljčević, N., Alamri, B. and Abdel Aleem, S.H.E. (2021) A Novel Exact Analytical Solution Based on Kloss Equation towards Accurate Speed-Time Characteristics Modeling of Induction Machines during No-Load Direct Startups. *Applied Sciences*, **11**, Article 5102. <https://doi.org/10.3390/app11115102>
 - [15] Chowdhury, S.K. and Baski, P.K. (2010) A Simple Lumped Parameter Thermal Model for Electrical Machine of TEFC Design. 2010 Joint International Conference on Power Electronics, Drives and Energy Systems & 2010 Power India, New Delhi, 20-23 December 2010, 1-7. <https://doi.org/10.1109/pedes.2010.5712385>
 - [16] Kirchgassner, W., Wallscheid, O. and Bocker, J. (2021) Data-Driven Permanent Magnet Temperature Estimation in Synchronous Motors with Supervised Machine Learning: A Benchmark. *IEEE Transactions on Energy Conversion*, **36**, 2059-2067.
<https://doi.org/10.1109/tec.2021.3052546>
 - [17] Reinbold, P.A.K., Kageorge, L.M., Schatz, M.F. and Grigoriev, R.O. (2021) Robust Learning from Noisy, Incomplete, High-Dimensional Experimental Data via Physically Constrained Symbolic Regression. *Nature Communications*, **12**, Article No. 3219.
<https://doi.org/10.1038/s41467-021-23479-0>
 - [18] Chen, Z., Badrinarayanan, V., Lee, C.Y. and Rabinovich, A. (2017) GradNorm: Gradient Normalization for Adaptive Loss Balancing in Deep Multitask Networks. arXiv:

1711.02257.

- [19] Krishnapriyan, A., Gholami, A., Zhe, S., Kirby, R. and Mahoney, M. (2021) Characterizing Possible Failure Modes in Physics-Informed Neural Networks. *Advances in Neural Information Processing Systems*, **34**, 26548-26560.