

Physics-Informed Recurrent Architecture with Embedded Thermodynamic Dynamics for Robust Sequence Modeling

Zafer Yiğit^{1 2}, Håkan Forsberg¹, and Masoud Daneshtalab^{1 3 *}

1- DeepHERO Lab - Mälardalen University - Sweden

2- Volvo Construction Equipment - Sweden

3- Tallinn University of Technology - Estonia

Abstract. Physics-informed machine learning has shown strong potential in improving generalisation under limited or noisy data, but most existing approaches treat physical priors only as soft regularisation terms on the loss. This work introduces a physics-structured recurrent architecture where thermodynamic differential equations are embedded directly into LSTM state updates. Adaptive physical parameters are learned through auxiliary multilayer perceptrons, forming a differentiable hybrid dynamical system that fuses physics priors with sequence learning. Experiments on industrial datasets show improved robustness under unseen fault conditions, outperforming conventional LSTMs and PINN-style models. The framework offers a scalable and generalizable approach to physics-aware recurrent modeling.

1 Introduction

Recent advances in physics-informed neural networks (PINNs) have shown that embedding physical priors into learning models can substantially improve sample efficiency, robustness, and interpretability [1, 2]. However, most existing implementations impose physical knowledge only as auxiliary loss regularizers, leaving the internal neural state evolution purely data driven. In contrast, this work proposes a physics-structured recurrent architecture in which thermodynamic differential equations are embedded directly into the recurrent state transition, effectively redefining the long short-term memory (LSTM) as a hybrid physics-driven dynamical system rather than a black box temporal approximator.

Conventional PINNs have demonstrated strong performance for static or steady state systems [3, 4]. However, they have difficulties maintaining the consistency for complex and time dependent dynamics such as thermal systems. Standard LSTM architecture provides strong capabilities for capturing long term dependencies but they are weak in predicting data noise and unseen scenarios since they do not have a built-in physical structure [5]. Some studies have introduced physics-informed loss functions or multi-fidelity schemes to enhance robustness to address these limitations [1, 2], while others have begun bridging PINNs and recurrent models [6, 7].

*This work is partially funded by the Swedish Knowledge Foundation within the RELIANT Research School. This study was co-funded by the European Union and Estonian Research Council via project TEM-TA138, the Swedish Research Council (VR) via project GreenDL.

The framework is domain-agnostic and applicable to any lumped thermal or flow system with known energy balance structure. In this study, it is evaluated on representative engine sensor datasets under normal and disturbed conditions, focusing on intake manifold boost temperature dynamics. The results show improved robustness and accuracy over conventional LSTM approach, confirming the benefit of embedding physics at the architectural level.

The main contributions of this paper are:

- A PI-LSTM architecture with physical dynamics in the recurrent transition
- Adaptive physical parameters through auxiliary multilayer perceptrons
- Improved robustness under fault and distribution shift scenarios

2 Methodology

The proposed architecture combines a standard LSTM structure with thermodynamic energy balance equations. This hybrid architecture preserves the ability of LSTM to capture temporal dependencies while enforcing physical consistency through embedded equations.

2.1 Baseline LSTM

A pure data driven LSTM network was trained to estimate the temperature signals in engine airpath systems. The architecture consists of two LSTM layers with 64 hidden units with a dropout rate of 0.2 and a fully connected output layer to estimate the final temperature values. The training objective minimizes a mean squared error between predicted and measured signals. This type of structure is commonly used in virtual sensor applications [5].

2.2 Physics Informed LSTM (PI-LSTM) Architecture

The proposed PI-LSTM architecture combines the baseline model with thermodynamic energy balance equations that govern the boost temperature dynamics. The model introduces fully connected multilayer perceptrons that estimate key thermal parameters used in physical equations. Temperature dynamics are represented using a lumped parameter dynamic thermal model, see Equation (1), where \dot{m}_a and \dot{m}_e are the mass flow values of the air and exhaust gas recirculation (EGR), T_a and T_{egr} are their transient temperatures, T_c represents the temperature of the wall (coolant), C_p is the specific heat of constant pressure, UA is the heat transfer coefficient and C_{eff} is the effective thermal capacity constant [8].

$$C_{eff} \frac{dT}{dt} = \dot{m}_a c_p (T_a - T) + \dot{m}_e c_p (T_{egr} - T) + UA(T_c - T) \quad (1)$$

The effective thermal response time constant $\tau(t)$ and the heat weight of the mixing gas $\varepsilon(t)$ can be expressed as in Equation (3), with $\dot{m}(t)$ as calculated in

Equation (2).

$$\dot{m}(t) = \dot{m}_a(t) + \dot{m}_e(t) \quad (2)$$

$$\tau(t) = \frac{C_{\text{eff}}}{\dot{m}(t) c_p + UA}, \quad \varepsilon(t) = \frac{\dot{m}(t) c_p}{\dot{m}(t) c_p + UA} \in (0, 1) \quad (3)$$

Equation (5) presents the equivalent equilibrium temperature T_{eq} is defined as a weighted average of the incoming boost temperature and the EGR gas mixture temperature T_{mix} shown in Equation (4) and the temperature of the wall T_c .

$$T_{\text{mix}}(t) = \frac{\dot{m}_a(t) T_a(t) + \dot{m}_e(t) T_{\text{egr}}(t)}{\dot{m}(t)} \quad (4)$$

$$T_{\text{eq}}(t) = \varepsilon(t) T_{\text{mix}}(t) + [1 - \varepsilon(t)] T_c(t) \quad (5)$$

The dynamics is discretised using the residual term $q(t)$, which represents unmodeled effects such as sensor biases; see Equation (6).

$$T_{t+1} = T_t + \Delta t \left[\frac{T_{\text{eq}}(t) - T_t}{\tau(t)} + q(t) \right] \quad (6)$$

The main contribution is the hybrid recurrent architecture that embeds physical equations directly into the state update.

2.3 Adaptive Parameter Learning

A single fully connected multilayer perceptron (MLP) estimates the target signal from the input data, and the result is used to compute the data loss during training. In addition, MLPs are introduced to estimate adaptive physical parameters selected according to the thermal characteristics of the target sensor. Each parameter estimated by MLP is used in the physical equations during training. Figure 1 shows one MLP used to estimate the target temperature sensor for data loss and two MLPs are employed to estimate UA and $q(t)$ for boost temperature sensor, capturing key thermal dynamics. This configuration allows the PI-LSTM to adapt its physical component to the specific dynamics of each sensor type.

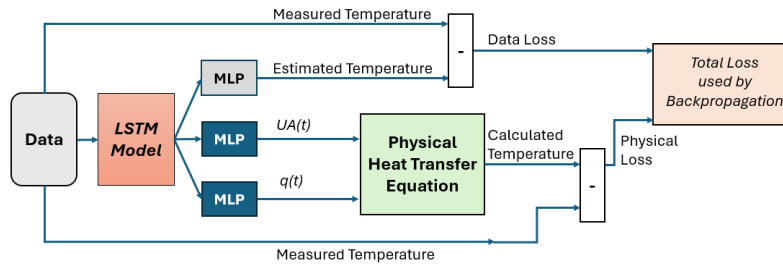


Fig. 1: PI-LSTM architecture for engine boost temperature sensor

3 Training and Experimental Setup

The baseline LSTM was trained using a standard data loss based on mean squared error (MSE), [9]. The training objective combines data loss with a physics-informed loss derived from the energy balance equations for the PI-LSTM, see Equation (7).

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}} + \lambda \mathcal{L}_{\text{physics}} \quad (7)$$

where λ balances data fidelity and physical consistency. Both terms are computed as MSE terms between predicted and reference trajectories. Both models were trained using the same dataset, optimizer, and hyperparameters to ensure a fair comparison. Training was performed in PyTorch on a workstation with an Intel(R) Core(TM) i7-12850HX CPU, 64 GB RAM, and an NVIDIA RTX A3000 GPU (12 GB). The baseline LSTM contained 53825 trainable parameters and required 1929.3 s for training. The PI-LSTM contained 64393 parameters and trained in 25604.8 s, reflecting a significant increase in complexity and runtime due to the embedded physical components.

4 Results and Comparison

The evaluation was performed on four independent datasets collected from diesel engine test rigs and real machines, covering both nominal and faulty operating conditions (boost leakage and EGR clogging). All test data were sampled at 100 ms and preprocessed identically to ensure compatibility.

4.1 Evaluation Metrics

Model performance was evaluated using standard regression metrics such as mean absolute error (MAE), root mean square error (RMSE), coefficient of determination (R^2), and the normalized RMSE (NRMSE) based on the interquartile range of the measured target values [9, 10]. This normalization makes the metric less sensitive to the absolute temperature scale and more representative of relative model accuracy.

4.2 Performance Comparison

Table 1 and Table 2 show that the PI-LSTM consistently achieves lower MAE and RMSE and higher R^2 across all test datasets compared to the baseline LSTM, demonstrating improved robustness under unseen fault conditions.

A visual comparison between the baseline LSTM and the proposed PI-LSTM model is shown in Figure 2. The baseline LSTM can follow general trends but fails to follow fast thermal transients under fault conditions. In contrast, the PI-LSTM model yields a closer match to the true temperature profile, confirming the advantage of embedding physical constraints into the network.

Test Dataset	Samples	Inf. Time (s)	MAE	RMSE	NRMSE	R2
Test 1 / Normal Condition	25040	2.411	0.6577	1.1786	0.2853	0.6978
Test 2 / Normal Condition	4838	0.467	0.8896	1.0130	0.2452	0.7957
Test 3 / Air-Leakage Failure	25180	2.256	1.2880	1.5748	0.3811	-0.0157
Test 4 / Clogged EGR Failure	4943	0.148	1.8203	2.4693	0.5976	0.2695

Table 1: Baseline LSTM performance for boost temperature estimation.

Test Dataset	Samples	Inf. Time (s)	MAE	RMSE	NRMSE	R2
Test 1 / Normal Condition	25040	6.606	0.3344	0.6069	0.1496	0.9199
Test 2 / Normal Condition	4838	1.231	0.5093	0.6237	0.1510	0.9226
Test 3 / Air-Leakage Failure	25180	7.178	0.9284	1.1628	0.2814	0.4462
Test 4 / Clogged EGR Failure	4943	1.184	0.8724	1.2988	0.3143	0.7979

Table 2: PI-LSTM performance for boost temperature estimation.

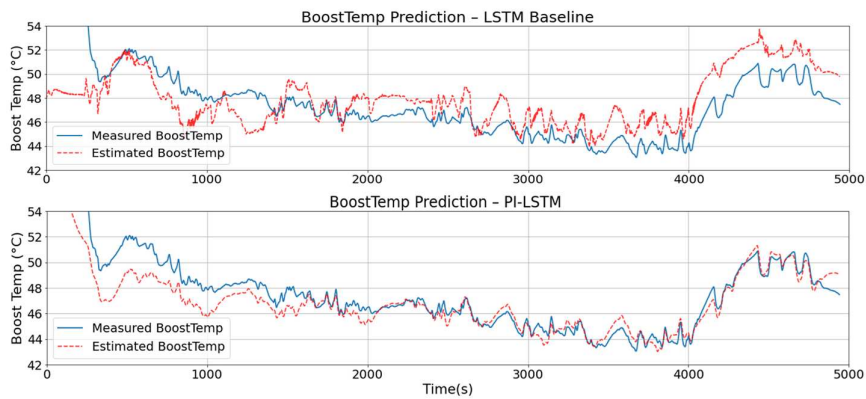


Fig. 2: Clogged EGR Condition - Boost Temperature Estimation of LSTM Models

5 Conclusion

This paper presented a physics-informed LSTM architecture that embeds thermodynamic equations directly into the recurrent state update. The approach improves prediction accuracy and robustness by combining data-driven learning with physical structure. Compared to a standard LSTM, the proposed PI-LSTM achieved lower prediction errors and higher R^2 scores across all test scenarios, including unseen fault conditions, with only a modest computational overhead. These results clearly demonstrate the benefit of embedding physical knowledge at the architectural level rather than treating it purely as a loss regularization. Future work will focus on integrating the PI-LSTM into the virtual engine rig shown in Figure 3 to build more robust engine model. The rig combines engine control software (vEMS), a GT-Suite based engine model, and an aftertreatment model consisting of a diesel oxidation catalyst (DOC), a diesel particulate filter (DPF), and a selective catalytic reduction (SCR) system. The framework will

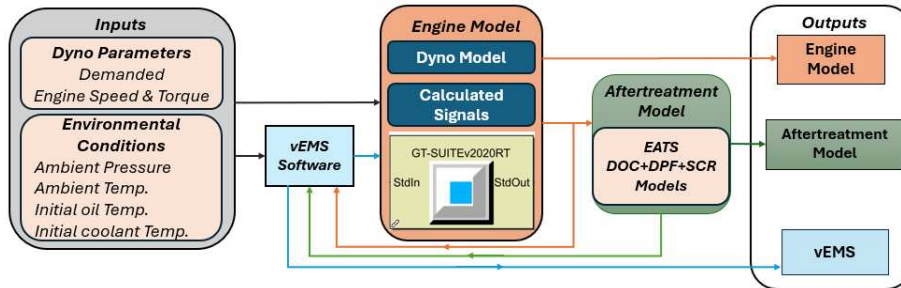


Fig. 3: Schematic of the virtual engine test rig

be expanded to include additional thermal and flow subsystems and enable more flexible physics data fusion strategies. Thus, the proposed framework is intended to be used as a general methodology for other applications using possibly other differential equations, but has yet to be confirmed. The virtual test rig will also be used to generate synthetic data for calibration, testing, fault analysis, machine learning, and digital twin applications.

References

- [1] Y. Wu, B. Sicard, S. A. Gadsden, Physics-informed machine learning: a comprehensive review on applications in anomaly detection and condition monitoring, *Expert Systems With Applications*, 255 (2024) 124678, Elsevier, 2024.
- [2] Y. Xu, P. Bansal, P. Wang, Y. Li, Physics-informed machine learning for system reliability analysis and design with partially observed information, *Reliability Engineering and System Safety*, 254 (2025) 110598, Elsevier, 2025.
- [3] S. Cuomo, V. S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, F. Piccialli, Scientific machine learning through physics-informed neural networks: where we are and what is next, *Journal of Scientific Computing*, Springer, 2022.
- [4] W. He, J. Li, X. Kong, and L. Deng, Multi-level physics informed deep learning for solving partial differential equations in computational structural mechanics, *Communications Engineering*,(2024) 3:151, 2024.
- [5] R. G. Herbon, G.G. Mateos, J. R. R. Ossorio, M. A. Prada, A. Moran, S. Alonso, J. J. Fuertes, M. Dominguez, Assessment and deployment of a LSTM-based virtual sensor in an industrial process control loop, *Neural Computing and Applications*, Springer, 2024.
- [6] X. Tang, Q. Hong, Y. Liu, B. Wang, Z. Cui, W. Shao, Physics-Informed LSTM based dynamic model of heat exchanger and application on the thermal management system considering time delay, *International Journal of Heat and Mass Transfer*, Elsevier, 2025.
- [7] K. Nath, X. Meng, D. J. Smith, G. E. Karniadakis, Physics-informed neural networks for predicting gas flow dynamics and unknown parameters in diesel engines, *Scientific reports*, Nature Portfolio, 2023.
- [8] L. Guzzella and C. H. Onder, editor. *Introduction to modeling and control of internal combustion engine systems*, Springer-Verlag Berlin Heidelberg, 2010.
- [9] G. James, D. Witten, T. Hastie, R. Tibshirani, editor. *An introduction to statistical learning with applications in python*, Springer, 2023.
- [10] D. C. Montgomery, G. C. Rungr, editor. *Applied statistics and probability for engineering*, Wiley, 7th Edition, 2018.