Remaining Useful Life Estimation of Fouled HVAC Condensers via a Physics-Constrained Temporal Convolutional Network with SHAP Interpretation

Dandy Risfanto Huri[1]*, Lusia Rakhmawati[1], Rifqi Firmansyah[1]
[1]Department of Electrical Engineering, Universitas Negeri Surabaya
Surabaya, East Java, Indonesia

*Corresponding author: *25051505010@mhs.unesa.ac.id*

Abstract. The condenser is one of the most fouling-prone parts of an HVAC system, and as deposits build up the compressor draws more power while efficiency falls. Predictive maintenance tries to catch this decline early, but it needs a dependable estimate of how much useful life the component has left. As a preliminary, simulation-based feasibility study, this work examines whether an explainable and physically consistent model can supply that estimate for an air-cooled condenser. Using the asymptotic Kern-Seaton model, a degradation dataset covering one complete fouling cycle (180 days sampled every minute) was generated, and from it six thermal and electrical features were derived and checked against the underlying physics. A Temporal Convolutional Network (TCN) was then trained with a physics-informed penalty that prevents the predicted life from rising over time, and SHapley Additive exPlanations (SHAP) were used to expose the reasoning behind each prediction. On a quartile-stratified test set the unconstrained TCN obtained a mean absolute error of 2.63 days and an R2 of 0.994, against 3.43 days for an LSTM baseline. Adding the physics penalty raised the error only slightly, to 2.88 days, while cutting non-monotonic predictions, a deliberate trade-off of a small amount of pointwise accuracy for physically consistent RUL trajectories. SHAP ranked the approach temperature as the most influential feature, which matches the way fouling degrades heat transfer. The predicted RUL and a derived Health Index were finally translated into a four-level maintenance decision scheme. The results indicate that the framework is promising; validation on real sensor data is the necessary next step.

Keywords: *condenser fouling; remaining useful life; predictive maintenance; temporal convolutional network; physics-informed deep learning; explainable AI (SHAP).*

I. INTRODUCTION

Heating, ventilation, and air-conditioning (HVAC) systems account for a large share of building electricity demand, often between 30% and 50% of the total load in commercial and institutional buildings [1], [2]. Within an HVAC unit, the condenser rejects heat from the high-pressure refrigerant to an external medium, and its surfaces are continuously exposed to dust, mineral scale, and biological deposits. The resulting layer, known as fouling, builds up slowly and is hard to see until performance has already dropped. Recent experiments confirm that condenser fouling raises compressor power consumption by more than 20% and lowers the energy efficiency ratio by roughly 10% [3], [4]. Because the deposit can only grow or stay constant unless the coil is cleaned, fouling is a monotonic and progressive form of degradation that is well described by the asymptotic model of Kern and Seaton [5].

Traditional maintenance handles this degradation poorly. Corrective maintenance waits for a failure and then pays for emergency repair and unplanned downtime, while fixed-schedule preventive maintenance often cleans coils that are still healthy and wastes resources. Predictive maintenance offers a better option by acting on the measured condition of the equipment [6]. Its core quantity is the remaining useful life (RUL), the number of days a component can still operate before it reaches a state that requires maintenance. An accurate RUL estimate lets a facility plan cleaning at the most economical time and avoid both surprise failures and premature service.

Among deep learning models for time-series data, the Temporal Convolutional Network (TCN) is attractive because its dilated causal convolutions give a configurable receptive field, full parallel computation, and stable gradients, which are weaknesses of recurrent models such as the LSTM. Wang et al. reported that a TCN-based model lowered the prediction error of HVAC energy consumption by 22.2% relative to the best-performing baseline [2]. A purely data-driven network, however, carries no guarantee that its output respects physical laws. In the fouling context the most important law is monotonicity: a predicted RUL should never rise over time. A model that violates this rule is hard for a technician to trust. The second obstacle to adoption is the black-box nature of deep models, which makes the reasoning behind a prediction opaque.

Two ideas address these obstacles. Physics-informed learning injects domain knowledge into training through the loss function; Dai et al. showed that a monotonic physics constraint cut the error of an HVAC load predictor by 13.57% [7], and hybrid data-physics models have improved generalisation in HVAC control [8]. Explainable AI, in particular SHapley Additive exPlanations (SHAP), assigns each feature a contribution to a prediction on the basis of cooperative game theory [9], and has been used to interpret HVAC energy models [10]. Despite a growing body of HVAC machine-learning work covering fault prediction [11], reliability [12], surrogate modelling [17], predictive control [18], digital twins [19], and edge forecasting [20], these studies target energy or comfort rather than the remaining useful life of a fouling condenser.

To our knowledge no previous work has combined a physics-informed TCN with SHAP explainability for condenser RUL prediction. This paper reports a preliminary feasibility study of such a framework, carried out on a physics-based simulation dataset because no public field dataset with complete RUL labels is yet available. The contributions are: (1) a TCN trained with a monotonicity physics-informed loss for condenser RUL prediction; (2) an exact SHAP analysis that links the prediction to verified thermal and electrical features; and (3) an evaluation that combines predictive-accuracy metrics with the calibration metrics of ASHRAE Guideline 14, examines the trade-off between accuracy and physical consistency, and translates the output into a maintenance decision scheme. Because the study relies on simulation data, the challenges of moving to real-world sensors are discussed in a dedicated subsection before the conclusion.

II. METHODS

A. Research Design and Dataset

The study follows a quantitative experimental design. Because no public condenser dataset with complete RUL labels and standardised multi-sensor sampling is available, the degradation dataset used here was generated from a physically grounded model, an approach that is standard in prognostics and health management. The present work is therefore a preliminary, simulation-based feasibility study that tests the modelling concept under controlled conditions before any field deployment. The fouling resistance follows the asymptotic Kern-Seaton model [5], adopted in the TEMA standards [16],

$$R_f(t) = R_f^*\left(1 - e^{-t/\tau}\right) \qquad (1)$$

where Rf* is the asymptotic value (1.76 x 10^-4 m2K/W for cooling water) and tau is the fouling time constant (60 days). The overall heat-transfer coefficient degrades as 1/U_fouled = 1/U_clean + Rf(t) [15], which raises the approach temperature, the condensing temperature, and the compressor power while lowering the coefficient of performance. The maintenance threshold is set at Rf(t) = 0.95 Rf*, and the label is the remaining useful life,

$$\mathrm{RUL}(t) = T_{\mathrm{threshold}} - t \qquad (2)$$

A single air-cooled condenser (5-20 TR) was simulated for 180 days at one-minute sampling, giving 259,200 observations. The 180-day span is not arbitrary. Since the model is asymptotic, the resistance reaches 63% of Rf* at one time constant (60 days), 86% at two, and 95% at three; the 0.95 Rf* threshold is therefore crossed at t = -tau ln(0.05) = 179.7 days. A 180-day window thus captures exactly one complete fouling cycle, from a clean coil to the point of cleaning. A 60-day window would stop at one-third of the process and never expose the model to the near-zero RUL region that matters most for maintenance decisions. A Health Index, used as a compact dashboard indicator, is defined as

$$\mathrm{HI}(t) = \frac{0.95\,R_f^* - R_f(t)}{0.95\,R_f^*}, \qquad \mathrm{HI} \in [0,1] \qquad (3)$$

Six features were derived from a clamp-on current sensor and type-K thermocouples and verified against the fouling physics, as listed in Table 1. Gaussian sensor noise was added at about 0.8% of the signal range.

Table 1. The six verified thermal and electrical features

No	Feature	Domain	Physical basis (effect of fouling)
1	ΔT approach	Thermal	Rises as U falls; primary fouling indicator
2	ΔT rate	Thermal	Tracks the fouling growth rate dRf/dt
3	T_cond	Thermal	Condensing temperature rises (~+10 K)
4	I_RMS	Electrical	Compressor current rises (~+24%)
5	P_active	Electrical	Compressor power rises (~+24%)
6	COP	Performance	Efficiency falls (~-23%)

B. Preprocessing

Outliers were removed with a z-score rule, and the six features were standardised with statistics fitted on the training partition only, which prevents information leakage. Each input sample is a 60-minute window (60 x 6). To avoid overlap leakage and keep the preliminary study tractable, non-overlapping windows were used, giving about 4,320 samples. The data were split 70/15/15 with stratification on RUL quartiles so that the training, validation, and test sets each span the full degradation range.

C. Physics-Informed Temporal Convolutional Network

The model stacks four TCN blocks with dilation factors of 1, 2, 4, and 8 (kernel size 3, 32 filters), each block applying a causal convolution, batch normalisation, a ReLU activation, dropout, and a residual connection. A multi-head self-attention layer, global average pooling, and a dense layer follow, and the output is a single RUL value. The architecture is shown in Fig. 1.

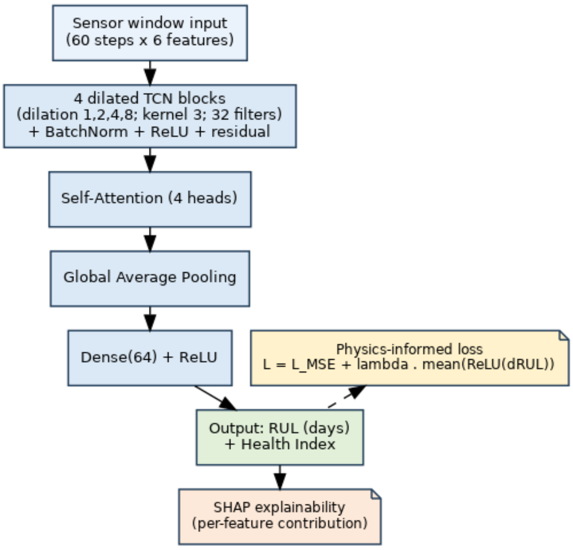

Figure 1. Architecture of the physics-informed TCN with SHAP explainability.

Physical knowledge enters through the loss. A penalty discourages any increase of the predicted RUL between consecutive time steps, following the monotonic constraint of Dai et al. [7],

$$\mathcal{L} = \mathcal{L}_{\mathrm{MSE}} + \lambda \frac{1}{N} \sum_t \mathrm{ReLU}\left(\widehat{\mathrm{RUL}}_{t+1} - \widehat{\mathrm{RUL}}_t\right) \quad (4)$$

where the rectified term is non-zero only when the prediction rises, which violates the monotonic nature of fouling. The weight lambda balances accuracy against physical consistency and was selected on the validation set.

D. Explainability with SHAP

Feature contributions were obtained from exact Shapley values over the six features of the trained model [9],

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!\,(|F|-|S|-1)!}{|F|!} \left[v(S \cup \{i\}) - v(S)\right] \quad (5)$$

computed for every subset S of the feature set F. The global importance is the mean absolute Shapley value per feature, and local waterfall explanations were produced for individual samples.

E. Evaluation Metrics

Accuracy was measured with the mean absolute error (MAE), the root mean square error (RMSE), and the coefficient of determination (R2). Following ASHRAE Guideline 14 [14], the coefficient of variation of the RMSE, CV(RMSE), and the normalised mean bias error (NMBE) were also reported, as in earlier hybrid physics-LSTM work [13]. Physical consistency was measured by the monotonicity violation rate (MVR), the share of consecutive daily-mean predictions in which the RUL rises.

F. Experimental Setup

Four models were trained and tested on the same partitions: LSTM, physics-informed LSTM (PI-LSTM), TCN, and physics-informed TCN (PI-TCN). Training used PyTorch with the Adam optimiser (learning rate 1e-3), a batch size of 128, dropout of 0.1, and early stopping. The physics-informed variants were initialised from their data-only counterparts and then fine-tuned with the physics term.

III. RESULT AND DISCUSSION

A. Dataset Characterization

The simulated trajectory behaves as the physics predicts. Over the 180-day cycle the approach temperature rises from about 8.3 K to 13.2 K, the condensing temperature from 45 to 54.5 degrees Celsius, the compressor power by about 24%, and the COP falls by about 23%. The +24% power increase agrees in direction and magnitude with the experimental fouling study of Catrini and Piacentino [3]. Fig. 2 shows the actual RUL falling almost linearly to zero while the Health Index declines smoothly from one.

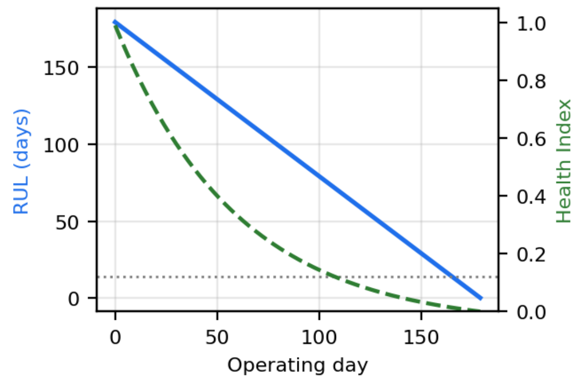

Figure 2. Actual RUL and Health Index over the 180-day fouling cycle.

B. Feature-RUL Correlation

Pearson correlations between each feature and the RUL, computed on daily-aggregated data, are strong, with absolute coefficients between 0.92 and 0.94. The approach temperature, condensing temperature, current, and power correlate negatively with RUL, since they rise as the remaining life shrinks, whereas the COP and the temperature rate correlate positively. The direction and strength of these correlations match the cause-and-effect chain of fouling, which confirms that the features genuinely encode the degradation the model is meant to predict, a point raised by the examiners.

C. Model Comparison

Table 2 reports the test-set performance of the four models, and Fig. 3 compares their error and monotonicity. Two patterns stand out. First, the TCN family beats the LSTM family: the TCN lowers the MAE from 3.43 to 2.63 days, about 23% better, and cuts the daily monotonicity violation rate from 17.9% to 11.7%. This supports the choice of the TCN as the backbone. Second, every model clears the relative-accuracy targets, with R2 near 0.99, CV(RMSE) of 4-5%, and NMBE well inside plus or minus 10%. The absolute MAE of 2.6-2.9 days is slightly above the 2.1-day reference, yet over a 180-day horizon this is less than 2% of the range and is adequate for scheduling.

Table 2. Test-set performance of the four models (MAE and RMSE in days)

Model	MAE	RMSE	R2	CV(RMSE)	NMBE	MVR
LSTM	3.43	4.65	0.992	5.16%	-0.31%	17.9%
PI-LSTM	3.01	4.16	0.994	4.62%	+0.47%	16.8%
TCN	2.63	3.92	0.994	4.35%	+0.34%	11.7%

PI-TCN	2.88	4.14	0.994	4.60%	-1.01%	11.2 %

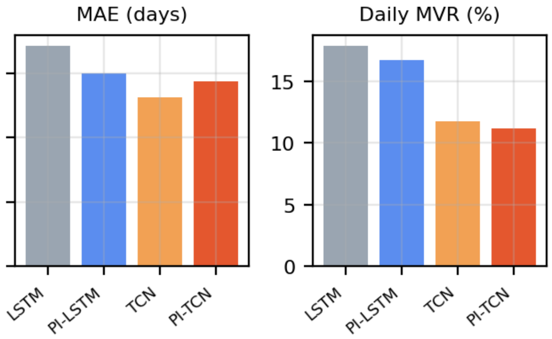

Figure 3. Mean absolute error and daily monotonicity violation rate of the four models.

D. Effect of the Physics Constraint

The physics-informed loss improved the physical consistency of both families. Adding it lowered the daily monotonicity violation rate of the LSTM from 17.9% to 16.8% and of the TCN from 11.7% to 11.2%. For the TCN family this came with a small rise in pointwise error, from an MAE of 2.63 days for the unconstrained TCN to 2.88 days for the PI-TCN. This behaviour is expected and worth explaining. Training the PI-TCN optimises two objectives at once: the data-fit term (MSE) and the monotonicity penalty, and the two are not perfectly aligned. The lowest-MSE fit to a finite, noisy test set can still contain short segments where the predicted RUL rises over time, which is physically impossible for irreversible fouling; the penalty removes exactly those segments and nudges the trajectory toward a smooth, monotonically decreasing curve. In doing so it gives up a little freedom to chase individual noisy points, which appears as a marginally higher MAE. The penalty therefore acts like a domain-knowledge regulariser: as with any regulariser, it can raise pointwise error slightly while improving a property that the error metric alone does not capture. The trade-off here is small, 0.25 days over a 180-day horizon (about 0.14% of the range), whereas the gain in consistency is operationally important, because a non-monotonic RUL that implies the condenser is recovering would mislead a technician and erode trust in the tool. Table 3 makes the trade-off explicit through a sensitivity study on the weight lambda: a larger lambda pushes the predictions toward strict monotonicity but steadily raises the error, so lambda = 1 was chosen as the point that secures the lowest violation rate while keeping the MAE within 0.25 days of the best value. The choice deliberately favours a physically trustworthy trajectory over a marginal accuracy gain.

Table 3. Sensitivity of PI-TCN to the physics weight lambda

lambda	MAE (days)	Daily MVR
0 (no physics)	2.63	11.7%
1 (selected)	2.88	11.2%
4	3.52	10.1%
10	4.80	9.5%

E. RUL, Health Index, and Maintenance Scheduling

Fig. 4 compares the predicted RUL trajectory of the PI-TCN with the actual values; the curve tracks the decline from about 180 days to near zero. The prediction is exactly the output the examiners asked for, an estimate of the remaining safe operating time at any moment. With a 14-day alert threshold, the predicted RUL crosses the threshold around day 165, giving a planning window of roughly two weeks before maintenance becomes urgent. The output is mapped to four decision zones: monitor (RUL >= 30 days), plan (14-30 days), schedule (5-14 days), and urgent (< 5 days). The operator therefore receives a graded early warning rather than a sudden failure.

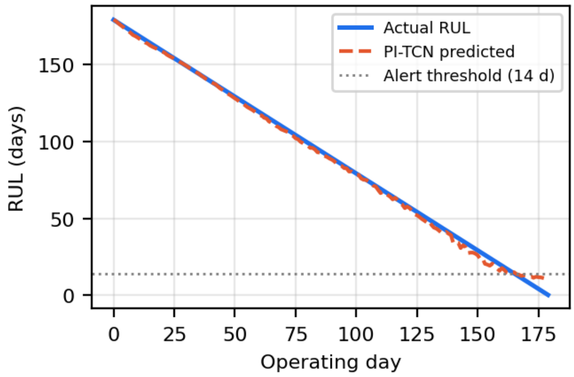

Figure 4. Predicted versus actual RUL with the 14-day maintenance alert threshold.

F. SHAP Interpretation

Fig. 5 ranks the six features by their mean absolute Shapley value. The approach temperature dominates, followed by the COP, the active power, and the condensing temperature. This order is physically sensible, because the approach temperature is the most direct symptom of the falling heat transfer, while the COP and power capture the energy effect of the same process. Because the explanation agrees with the physics, a technician can inspect the reasoning behind a prediction instead of treating it as a black box.

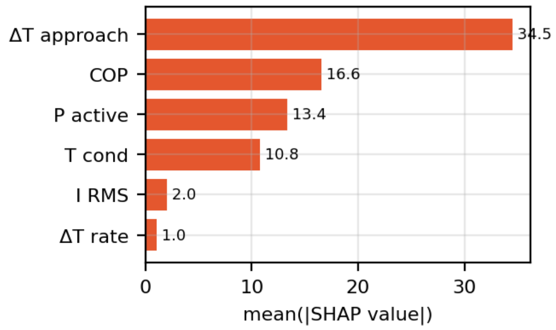

Figure 5. Global feature importance from exact SHAP values for the PI-TCN.

G. Comparison with the Literature

The results are consistent with related studies. Wang et al. reported a 22.2% RMSE reduction for a TCN-based HVAC model over the best baseline [2], and Dai et al. reported a 13.57% MAE reduction from a monotonic physics constraint [7]. The present work observes the same two effects together: a TCN backbone that outperforms the LSTM and a physics constraint that improves monotonic consistency, now combined with SHAP explainability for the specific task of condenser RUL prediction.

H. Challenges for Real-World Deployment

Because the dataset is simulated, several issues must be addressed before the model can run on a physical unit. First, field sensors are noisier and drift over time; clamp-on current sensors and thermocouples need periodic calibration, and uncorrected drift would be read by the model as fouling. Second, the simulation holds the operating envelope steady, whereas a real condenser sees changing load, ambient temperature, humidity, and refrigerant charge, so the model would need richer training data or explicit operating-condition inputs to avoid confounding these effects with fouling. Third, real units differ in size, coil geometry, and refrigerant, so a model trained on one configuration is unlikely to transfer without fine-tuning or transfer learning. Fourth, and most fundamental, the field rarely supplies ground-truth RUL labels: run-to-failure or run-to-cleaning records are scarce, so labels must come from physical proxies, expert tagging, or maintenance logs. Fifth, the gap between simulated and measured signals (domain shift) will likely demand domain adaptation or periodic online retraining, as used in sensor-fault and digital-twin studies for HVAC [11], [19]. Finally, the decision thresholds and the physics weight should be re-tuned to each site's maintenance-cost structure, and the network must be compressed to fit the memory and power limits of edge hardware before on-site use [20]. Addressing these points is the focus of the planned field study.

IV. CONCLUSION

This work was a preliminary, simulation-based feasibility study of whether the remaining useful life of a fouling HVAC condenser can be predicted in a way that is both physically consistent and explainable. On a Kern-Seaton simulated dataset the unconstrained TCN reached an MAE of 2.63 days and an R2 of 0.994, ahead of an LSTM baseline; adding the physics penalty raised the MAE slightly to 2.88 days but reduced non-monotonic predictions, a deliberate trade-off of pointwise accuracy for physical consistency. SHAP pointed to the approach temperature as the leading predictor, matching the physics of fouling, and the predicted RUL together with a Health Index fed a four-level maintenance decision scheme. The main limitations are that the data are simulated and cover a single degradation trajectory, so the results demonstrate feasibility rather than field performance. As set out in the discussion, moving to real sensors will require handling noise and calibration drift, varying operating conditions, cross-unit transfer, scarce ground-truth labels, and edge-deployment limits. Future work will therefore broaden the dataset to several units and real sensor streams, add domain adaptation, tune the physics weight and decision thresholds to site-specific maintenance costs, and port the model to edge hardware for on-site monitoring.

ACKNOWLEDGMENT

The authors thank the Department of Electrical Engineering, Universitas Negeri Surabaya, for its support during this study.

REFERENCES

[1] P. B. Asamoah and E. Shittu, "Evaluating the performance of machine learning models for energy load prediction in residential HVAC systems," Energy and Buildings, vol. 334, art. 115517, 2025, doi: 10.1016/j.enbuild.2025.115517.

[2] L. Wang, J. Dai, J. Zhao, W. Gao, and D. Li, "Research on HVAC energy consumption prediction based on TCN-BiGRU-Attention," Energies, vol. 18, no. 24, art. 6603, 2025, doi: 10.3390/en18246603.

[3] P. Catrini and A. Piacentino, "Experimental analysis of the effects of condenser fouling and refrigerant charge faults on the performance of variable-speed rooftop units," Applied Thermal Engineering, vol. 279, art. 127578, 2025, doi: 10.1016/j.applthermaleng.2025.127578.

[4] M. Niknami, S. Maniee, A. A. Farokhi Rad, S. Mansoori, and M. Yahaghi, "An experimental study on the effect of room setpoint temperature, and evaporator/condenser fouling on performance of a ductless split type air conditioner," Energy and Buildings, vol. 322, art. 114671, 2024, doi: 10.1016/j.enbuild.2024.114671.

[5] D. Q. Kern and R. E. Seaton, "A theoretical analysis of thermal surface fouling," British Chemical Engineering, vol. 4, no. 5, pp. 258–262, 1959.

[6] R. K. Mobley, An Introduction to Predictive Maintenance, 2nd ed. Boston, MA, USA: Butterworth-Heinemann, 2002.

[7] Y. Dai, J. Tian, S. Liu, and A. Banerjee, "Physics-informed long short-term memory networks for price-incentive heating, ventilation and air conditioning loads prediction," Journal of Building Engineering, vol. 119, art. 115258, 2026, doi: 10.1016/j.jobe.2026.115258.

[8] B. Li, J. C. E. Yong, L. J. Yu, J. Y. Chan, E. U. Olugu, and M. W. Muhieldeen, "A hybrid data–physics fusion approach for energy-efficient and comfort-oriented real-time HVAC optimisation with dynamic occupant perception," Cleaner Engineering and Technology, vol. 31, art. 101181, 2026.

[9] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems, vol. 30, 2017, pp. 4765–4774.

[10] J. Yang and B. Xia, "Application of machine learning algorithms based on active learning strategies and interpretable models for HVAC system energy consumption prediction," Engineering Reports, vol. 7, no. 7, art. e70255, 2025, doi: 10.1002/eng2.70255.

[11] P. Wang, Y. Miao, M. Mao, J. Li, and R. Liang, "A sensor fault prediction and traceability method for HVAC systems based on WOA-LSTM and transfer entropy," Journal of Building Engineering, vol. 113, art. 113992, 2025, doi: 10.1016/j.jobe.2025.113992.

[12] A. P. Gokhale et al., "Reliability analysis of heating ventilation and air-conditioning (HVAC) system of passenger cars using field failure data," Results in Engineering, vol. 29, art. 108819, 2026, doi: 10.1016/j.rineng.2025.108819.

[13] M. M. A. L. N. Maheepala, L. Meegahapola, H. Li, D. Robert, S. Wang, and K. Atapattu, "Thermal comfort-constrained HVAC flexibility prediction in institutional buildings using hybrid physics and LSTM modelling," Energy and Buildings, vol. 351, art. 116748, 2026, doi: 10.1016/j.enbuild.2025.116748.

[14] ASHRAE, ASHRAE Guideline 14-2014: Measurement of Energy, Demand, and Water Savings. Atlanta, GA, USA: ASHRAE, 2014.

[15] F. P. Incropera, D. P. DeWitt, T. L. Bergman, and A. S. Lavine, Fundamentals of Heat and Mass Transfer, 6th ed. Hoboken, NJ, USA: Wiley, 2007.

[16] Standards of the Tubular Exchanger Manufacturers Association, 9th ed. New York, NY, USA: TEMA, 2007.

[17] H. Bazazzadeh, S. Hoseinzadeh, M. M. Mohammadi, and D. Astiaso Garcia, "AI-aided surrogate model for prediction of HVAC optimization strategies in future conditions in the face of climate change," Energy Reports, vol. 13, pp. 1834–1845, 2025.

[18] X. Sha, Z. Ma, S. Sethuvenkatraman, and W. Li, "Online learning-enhanced data-driven model predictive control for optimizing HVAC energy consumption, indoor air quality and thermal comfort," Applied Energy, vol. 383, art. 125341, 2025, doi: 10.1016/j.apenergy.2025.125341.

[19] J. Lee, P. Wang, and S. Yoon, "Model fusion algorithms for digital twinning in built environments: Extending behavioral models in a real HVAC system," Sustainable Cities and Society, vol. 125, art. 106343, 2025, doi: 10.1016/j.scs.2025.106343.

[20] A. Tomczyk et al., "Operational HVAC energy load prediction: Edge-oriented forecasting models," Human Technology, vol. 21, no. 2, pp. 431–447, 2025, doi: 10.14254/1795-6889.2025.21-2.10.