

Development of an Intelligent Predictive Maintenance and Energy Optimisation Framework for HVAC Systems Using Machine Learning and Real-Time Equipment Performance Data

Tyler Leow Pei Liang¹, Nick Chong Woei Chyang², Anson Leow Pei Iu³, Andy Chan Hsi Ming⁴,
V.B. Kumaresan⁵

^{1&5}European International University, Paris, France

²University of Wilmington, New Castle, United Kingdom

^{3 &4}MAIWP International University, Kuala Lumpur, Malaysia

Abstract - This study develops an intelligent predictive maintenance and energy optimisation framework for heating, ventilation and air-conditioning (HVAC) systems by integrating machine learning, real-time equipment condition monitoring, failure-risk estimation and energy-aware control. The study adopts a design-science orientation supported by an integrative literature synthesis of ten studies published between 2023 and 2025. The evidence was coded according to application domain, data inputs, machine-learning function, maintenance outcome, energy outcome and implementation contribution. The synthesis identifies four interdependent capabilities: real-time sensing and data quality management; machine-learning-based anomaly detection, fault diagnosis and failure-risk prediction; maintenance decision support; and energy optimisation through intelligent or model-predictive control. The evidence indicates that predictive analytics can reduce unplanned downtime and maintenance burden, while energy-aware HVAC control can produce material energy savings. However, outcomes vary according to building type, baseline condition, control strategy and data quality. The study contributes a research-ready framework, testable propositions, performance indicators and an implementation pathway for future validation using operational building data.

Key Words: HVAC, predictive maintenance, machine learning, energy optimisation, real-time monitoring, fault detection and diagnosis.

1. INTRODUCTION

Heating, ventilation and air-conditioning (HVAC) systems are among the most energy-intensive technical systems in commercial and institutional buildings. The source manuscript notes that HVAC can account for more than 40% of building energy consumption and argues that maintenance condition and energy performance are closely connected. Conventional reactive maintenance addresses faults after performance has deteriorated or equipment has failed, while time-based preventive maintenance may replace or service components regardless of their actual condition. Both approaches can create avoidable cost, downtime and energy waste.

Machine learning (ML) and real-time sensing offer an alternative. Sensor streams such as temperature, pressure, flow, vibration, current, power, valve position, compressor state and indoor environmental variables can be transformed into indicators of equipment health. ML models can then detect anomalies, classify faults, forecast degradation or estimate failure risk. When those predictions are connected to an energy-management layer, the same information can also support operating-point optimisation, scheduling and model-predictive control.

The principal weakness of the original article is that it reads mainly as a narrative overview. It discusses predictive

maintenance and energy optimisation but does not specify a research problem, research questions, methodology, analytical procedure, validation logic or a coherent set of findings. This revision addresses those gaps while retaining the original topic and evidence base.

1.1 Research Problem

Existing HVAC studies frequently examine fault detection, predictive maintenance or energy optimisation as separate technical problems. In practice, however, maintenance condition affects energy performance and energy-control actions can affect equipment stress. A fragmented approach may therefore identify faults without translating them into optimal maintenance actions, or optimise energy without adequately accounting for equipment degradation. The research problem is the absence of an integrated, research-testable framework that links real-time equipment data, ML-based health prediction, maintenance decision support and energy optimisation in a single closed-loop architecture.

1.2 Research Objectives

RO1: To identify the principal data, machine-learning and decision-support components required for intelligent HVAC predictive maintenance.

RO2: To examine how predictive maintenance information can be integrated with energy optimisation strategies.

RO3: To synthesise the expected operational, maintenance and energy-performance outcomes reported in the source literature.

RO4: To develop an integrated framework and measurable performance indicators for future empirical validation using real-time HVAC data.

1.3 Research Questions

RQ1: What real-time equipment and operational data are most relevant to ML-enabled HVAC predictive maintenance?

RQ2: What ML functions are required to convert HVAC performance data into actionable maintenance intelligence?

RQ3: How can predictive maintenance outputs be linked to HVAC energy optimisation?

RQ4: What technical and operational indicators should be used to evaluate the proposed integrated framework?

1.4 Contribution of the Study

The study contributes an integrated architecture rather than treating predictive maintenance and energy optimisation independently. It also provides an explicit measurement model for later empirical work, including prediction metrics (precision, recall, F1-score, MAE/RMSE where applicable), maintenance metrics (downtime, mean time between failures, maintenance cost and false-alarm rate), and energy metrics (kWh, peak demand, coefficient of performance and energy-cost reduction).

2. LITERATURE REVIEW

2.1 HVAC Maintenance and Energy Efficiency

The literature cited in the source manuscript consistently positions maintenance quality as an important determinant of HVAC reliability and energy performance. Degraded heat-transfer surfaces, refrigerant problems, sensor drift, airflow restrictions and control faults can increase energy consumption before a complete failure becomes visible. This makes condition monitoring valuable not only for reliability but also for energy management.

2.2 Machine Learning for Predictive Maintenance

ML-based predictive maintenance converts historical and streaming equipment data into estimates of abnormality, fault class, degradation state or future failure. The source manuscript refers to support vector regression and broader AI/ML approaches. Across the cited studies, the common analytical logic is: acquire equipment data, clean and transform the signals, extract informative features, train or update a predictive model, evaluate model performance and translate the model output into a maintenance action.

2.3 Energy Optimisation and Intelligent Control

Energy optimisation extends the maintenance layer by using equipment-health information together with demand, weather, occupancy and comfort constraints. Intelligent control and model-predictive control can modify set points, schedules, staging and load distribution while respecting thermal comfort and equipment limitations. The source manuscript reports examples of energy-cost reductions above 20% and, in some cited intelligent-control contexts, energy-use reductions exceeding 40%. These values should be interpreted as study-specific reported outcomes rather than universal effects.

2.4 Research Gap and Conceptual Integration

The reviewed material supports a closed-loop view in which equipment condition and energy performance are jointly managed. The proposed framework therefore integrates four layers: sensing and data management; ML analytics; maintenance decision support; and energy optimisation/control. Feedback from maintenance actions and control responses is returned to the data layer so that model performance can be monitored and updated.

2.5 Proposed Research Framework

Table 1: Proposed intelligent HVAC framework architecture

Framework Layer	Main Inputs	Analytical Function	Primary Outputs
1. Real-time data layer	Temperature, pressure, flow, vibration, current, power, equipment state, weather/occupancy where available	Data validation, synchronisation, missing-value handling, feature engineering	Clean time-series dataset and health/energy features
2. ML condition-intelligence layer	Engineered features and historical fault/maintenance records	Anomaly detection, fault classification, regression/forecasting, failure-risk estimation	Anomaly score, fault class, predicted degradation/failure risk
3. Maintenance decision layer	Predicted condition, risk, criticality and maintenance constraints	Priority scoring and maintenance scheduling	Work-order priority, recommended inspection/service window
4. Energy optimisation layer	Equipment health, load, comfort limits, tariffs and operating constraints	Optimisation or model-predictive control	Energy-efficient set points, schedules and equipment staging

2.6 Testable Propositions for Future Empirical Validation

P1: Higher quality and higher-frequency equipment-condition data improve HVAC fault-detection and failure-prediction performance.

P2: ML-based condition intelligence reduces unplanned HVAC downtime relative to reactive maintenance.

P3: Integrating predicted equipment health into maintenance scheduling reduces maintenance cost and increases equipment availability.

P4: Integrating predictive-maintenance information into energy optimisation reduces HVAC energy consumption relative to conventional rule-based operation.

P5: The effect of the integrated framework on energy performance is mediated by improved equipment condition and more efficient operating decisions.

3. METHODOLOGY

3.1 Research Design

This revision uses a design-science research orientation supported by a structured integrative literature synthesis. Design science is appropriate because the principal research output is an artefact: an integrated predictive-maintenance and energy-optimisation framework. The evidence synthesis is used to identify required components, relationships, performance measures and implementation conditions.

3.2 Evidence Base and Unit of Analysis

The unit of analysis is the published study or technical paper. The core evidence base consists of the ten references already included in the source manuscript, spanning HVAC energy optimisation, intelligent control, predictive maintenance, AI-enabled maintenance and related industrial applications. The publication years range from 2023 to 2025. Studies outside HVAC were retained only when they contributed transferable predictive-maintenance methods or implementation logic.

3.3 Inclusion and Exclusion Logic

A source was included when it addressed at least one of the following: (i) machine learning or AI for predictive maintenance; (ii) HVAC energy efficiency or intelligent control; (iii) real-time or condition-monitoring data; or (iv) integration of predictive information with operational decision-making. Purely descriptive material without relevance to the framework architecture was not used as evidence for a specific causal claim. Because the original manuscript did not document a database search protocol, this revision is an integrative synthesis rather than a PRISMA-compliant systematic review.

3.4 Coding and Analytical Procedure

Each source was coded across six dimensions: application domain; data/inputs; ML or AI function; maintenance contribution; energy contribution; and framework implication. The analysis proceeded in four stages. First, recurring technical functions were extracted. Second, similar functions were grouped into higher-order themes. Third, relationships among the themes were mapped into a closed-loop architecture. Fourth, measurable indicators were assigned to each framework layer to support future validation.

3.5 Future Empirical Validation Protocol

For a field validation, real-time data should be collected from one or more HVAC assets such as chillers, air-handling units, pumps, cooling towers or packaged units. A time-stamped dataset should combine sensor readings with alarm logs, work orders, maintenance events and energy-meter data. The observation window should be long enough to include normal operation, seasonal/load variation and a sufficient number of fault or degradation events.

Data preparation should include timestamp alignment, sensor-range checks, missing-value treatment, outlier review, resampling, lagged features and rolling statistics. To avoid temporal leakage, model development should use chronological training, validation and test partitions rather than random splitting where future observations can enter the training set. Where fault classes are imbalanced, class-weighting or carefully validated resampling should be considered.

Candidate models may include logistic regression as a transparent baseline, random forest or gradient boosting for nonlinear tabular relationships, support vector machines for classification/regression, and recurrent or temporal neural networks when sufficient sequential data are available. The final model should be selected on out-of-sample performance, stability, interpretability and deployment cost rather than accuracy alone.

3.6 Evaluation Metrics

Table 2: Evaluation metrics for future empirical validation

Evaluation Domain	Recommended Metrics	Purpose
Fault/anomaly classification	Precision, recall, F1-score, ROC-AUC, false-alarm rate	Measures detection quality and operational alarm burden
Regression/forecasting	MAE, RMSE, MAPE where appropriate	Measures error in temperature, load, energy or degradation prediction
Maintenance performance	Unplanned downtime, MTBF, MTTR, maintenance cost, emergency work orders	Tests operational value of predictive maintenance
Energy performance	HVAC kWh, kWh/m ² , peak kW, COP/EER, energy cost	Tests energy and cost improvement
Comfort/constraint performance	Temperature/humidity violations, comfort-hours outside limits	Ensures energy savings do not compromise indoor conditions

3.7 Reliability, Validity and Research Ethics

Technical validity should be established through out-of-sample testing, sensitivity analysis and comparison with baseline maintenance/control strategies. Where labelled fault data are created by technicians, inter-rater agreement and a documented fault taxonomy should be used. Sensor calibration and data-quality rules are necessary to support measurement reliability. Operational datasets should be de-identified where they contain building, occupant or staff information, and cybersecurity controls should protect building-management-system interfaces.

4. ANALYSIS AND FINDINGS

4.1 Descriptive Profile of the Evidence

The ten-source evidence base is recent and strongly technology-oriented. It includes studies focused directly on HVAC optimisation and intelligent control, as well as broader predictive-maintenance applications that contribute methods transferable to HVAC. The synthesis shows a convergence around continuous sensing, data-driven condition assessment and proactive intervention.

Table 3: Profile of the ten-source evidence base

Source	Primary Focus	ML/AI Role	Outcome Emphases	Framework Contribution
Moveh et al. (2025)	HVAC thermodynamic optimisation	Advanced ML/dynamic modelling	Energy/cost and system performance	Links prediction with operating optimisation
Reshma et al. (2025)	Industrial predictive maintenance	Failure prediction	Downtime and maintenance	Supports proactive maintenance logic
Jin (2025)	AI/ML in maintenance and production	Predictive analytics	Resource allocation and scheduling	Supports decision-support integration
Jerald et al. (2024)	HVAC intelligent control	Intelligent control	Energy efficiency	Supports energy-control layer
Nwulu et al. (2023)	Predictive maintenance applications	ML-based maintenance	Efficiency and reliability	Transferable ML maintenance workflow
Köse et al. (2024)	Commercial-building demand management	Real-time optimisation	Energy/load management	Supports real-time optimisation
Nikolaevich (2025)	AI in technical-system maintenance	AI condition management	Reliability and maintenance	Supports state/condition monitoring
Kaur & Sharma (2025)	AI predictive maintenance	AI/ML analytics	Maintenance effectiveness	Supports predictive-maintenance capability
Kumar et al. (2025)	ML predictive-maintenance system	ML deployment/API	Operational implementation	Supports deployable analytics architecture

Almazan et al. (2025)	HVAC ML + model predictive control	ML/MPC integration	Sustainable energy efficiency	Strong support for integrated framework
-----------------------	------------------------------------	--------------------	-------------------------------	---

4.2 Theme 1: Real-Time Data Are the Foundation of Predictive Capability

The first finding is that predictive maintenance cannot be separated from data quality. Real-time equipment signals provide early evidence of abnormal behaviour, but raw sensor data require synchronisation, validation and contextualisation. The framework therefore treats data engineering as a distinct layer rather than assuming that sensor data are immediately suitable for ML. This finding directly answers RQ1.

4.3 Theme 2: ML Must Produce Actionable Condition Intelligence

The second finding is that the value of ML lies in converting sensor patterns into operationally meaningful outputs. Four analytical functions recur across the evidence base: anomaly detection, fault classification, performance/degradation forecasting and failure-risk estimation. A model that achieves high statistical accuracy but produces excessive false alarms or cannot be linked to a maintenance action has limited operational value. This supports the inclusion of false-alarm rate, precision/recall and maintenance outcomes alongside conventional prediction accuracy.

4.4 Theme 3: Predictive Maintenance and Energy Optimisation Are Mutually Reinforcing

The third finding is the central integrative result. Equipment degradation can increase energy use, while operating strategy can accelerate or reduce equipment stress. Predictive health information should therefore become a constraint or decision variable in energy optimisation. For example, an optimiser may avoid aggressively loading a component with a high failure-risk score, while maintenance scheduling can prioritise faults associated with large energy penalties. This answers RQ2 and RQ3.

4.5 Theme 4: Reported Benefits Are Material but Context-Dependent

The source manuscript cites a reported energy-cost reduction of approximately 21.75% in one ML-related HVAC optimisation study and energy-use reductions above 40% in an intelligent-control context. These figures demonstrate the potential scale of benefit but should not be interpreted as pooled effects or guaranteed outcomes. The reviewed sources differ in system boundaries, baseline controls, operating conditions and evaluation methods. Accordingly, the present study does not compute a meta-analytic average from these heterogeneous reports.

4.6 Integrated Framework Finding

The analysis yields a closed-loop framework: (1) acquire real-time equipment, environmental and energy data; (2) validate and engineer features; (3) estimate anomaly, fault and failure risk using ML; (4) convert predictions into maintenance priorities; (5) feed equipment-health information into an energy-optimisation controller; (6) implement maintenance/control actions; and (7) measure the resulting reliability, energy and comfort outcomes. Performance feedback is then used for model monitoring and periodic retraining.

Sense	Prepare	Predict	Prioritise	Optimise	Act	Learn
-------	---------	---------	------------	----------	-----	-------

Fig -1: Logical flow of the proposed intelligent HVAC predictive-maintenance and energy-optimisation framework.

4.7 Proposed Decision Logic

A practical implementation can combine a health-risk score with an energy-penalty score. Low-risk/low-energy-penalty conditions remain under normal control. Low-risk/high-energy-penalty conditions trigger control optimisation or efficiency inspection. High-risk/low-energy-penalty conditions trigger condition-based maintenance planning. High-risk/high-energy-penalty conditions receive the highest intervention priority because they threaten both reliability and energy performance.

Table 4: Proposed health-risk and energy-penalty decision logic

Condition	Interpretation	Recommended Action
Low failure risk + low energy penalty	Normal/efficient operation	Continue monitoring
Low failure risk + high energy penalty	Efficiency/control problem likely	Optimise set points; inspect controls/sensors
High failure risk + low energy penalty	Reliability risk	Plan condition-based maintenance
High failure risk + high energy penalty	Critical combined risk	Prioritise intervention and apply safe energy-control constraints

5. DISCUSSION

The revised findings support the argument that intelligent HVAC management should be evaluated as a cyber-physical decision system rather than as a standalone ML model. Prediction performance is necessary but insufficient. The model must be connected to maintenance workflows, energy controls and measurable operational outcomes. This changes the research question from 'Can a model predict a fault?' to 'Does the integrated prediction-and-control system reduce failures and energy use without compromising comfort or safety?'

The framework also clarifies why energy savings reported in individual studies vary. Savings depend on the initial inefficiency of the HVAC system, quality of controls, climate, occupancy pattern, equipment age, maintenance backlog, tariff structure and the degree of control authority granted to the optimisation layer. A well-maintained and already-optimised system may show smaller incremental savings than a poorly controlled baseline.

5.1 Theoretical Implications

The study extends the predictive-maintenance logic by explicitly incorporating energy performance as a co-optimised outcome. It proposes that equipment health is not merely an output of monitoring but an input to operational optimisation. This provides a basis for future models that jointly minimise expected failure cost and energy cost under comfort and safety constraints.

5.2 Practical Implications

Facilities managers should begin with a limited number of critical HVAC assets and establish reliable sensor, alarm and work-order data before deploying complex ML. A transparent baseline model is useful for benchmarking. Deployment should include alarm governance, technician feedback, model-drift monitoring and clear rules defining when automated control is permitted versus when human approval is required.

5.3 Implementation Roadmap

Table 5: Implementation roadmap for future deployment

Phase	Main Activities	Deliverable	Go/No-Go Criterion
1. Baseline	Asset criticality, metering review, fault taxonomy, baseline energy/downtime	Baseline performance profile	Sufficient data coverage and sensor reliability
2. Data pipeline	BMS/IoT integration, cleaning, feature engineering	Validated time-series dataset	Acceptable missingness and timestamp quality
3. ML pilot	Train/test candidate models, threshold selection	Validated condition model	Out-of-sample performance exceeds baseline
4. Decision support	Risk scoring, CMMS/work-order integration	Maintenance prioritisation workflow	Technician acceptance and manageable false alarms
5. Energy integration	Health-aware optimisation/MP C under comfort constraints	Closed-loop optimisation pilot	Energy benefit without unacceptable comfort/safety violations
6. Scale-up	Model monitoring, retraining, governance, portfolio rollout	Operational framework	Stable benefit and maintainable lifecycle cost

6. LIMITATIONS AND FUTURE RESEARCH

The most important limitation is that the uploaded source article contains no raw HVAC sensor dataset, maintenance log, experimental sample, respondent dataset or model output. It would therefore be methodologically inappropriate to invent regression coefficients, accuracy values, p-values, confidence intervals or energy savings and present them as observed findings. The present revision instead provides a transparent evidence synthesis and a complete empirical validation protocol.

Future research should validate the framework using longitudinal building data and compare at least three operating conditions: conventional/reactive or time-based maintenance, ML-enabled predictive maintenance, and integrated predictive maintenance plus energy optimisation. A quasi-experimental interrupted time-series or matched-period design would allow changes in energy use and downtime to be assessed while controlling for weather and load. Where multiple buildings are available, hierarchical or mixed-effects models can account for building-level differences.

7. CONCLUSION

This study develops a complete research framework for intelligent HVAC predictive maintenance and energy optimisation from the evidence and concepts contained in the original manuscript. The synthesis shows that four capabilities must operate together: high-quality real-time data, ML-based condition intelligence, maintenance decision support and health-aware energy optimisation. The proposed closed-loop architecture converts sensor data into maintenance and control actions and evaluates success through prediction, reliability,

energy and comfort metrics. The article therefore moves beyond a descriptive discussion toward a research-testable framework suitable for implementation and empirical validation. The next stage should use real operational HVAC data so that model accuracy, downtime reduction, energy savings and cost effectiveness can be quantified without relying on assumed results.

REFERENCES

- Almazam, K., Humaidan, O., Shannan, N. M., Bashir, F. M., Gammoudi, T., & Dodo, Y. (2025). Innovative energy efficiency in HVAC systems with an integrated machine learning and model predictive control technique: A prospective toward sustainable buildings.
- Jerald, F., Gopalan, A., Shanmugasundaram, V., Santhoshi, B., & Sundaram, K. M. (2024). Enhancing energy efficiency in HVAC systems through intelligent control strategies.
- Jin, J. (2025). Enhancing manufacturing performance with AI and machine learning: Applications in predictive maintenance and production optimization.
- Kaur, C., & Sharma, A. (2025). Enhancing predictive maintenance with AI: Applications and impact.
- Köse, A., Sukhanov, I., Belikov, J., Tepljakov, A., & Petlenkov, E. (2024). Demand side management for commercial buildings: Multiple heating systems based optimization in real-time.
- Kumar, C. P., Jagadeeswari, B., Naveen, G., Sai, K. P., & Babu, M. M. (2025). Predictive maintenance system using machine learning and FastAPI.
- Moveh, S., Merchán-Cruz, E. A., Ibrahim, A. O., Elhassan, Z., Abdelhai, N. M. R., & Abdelrazig, M. D. (2025). Thermodynamic optimization of building HVAC systems through dynamic modeling and advanced machine learning.
- Nikolaevich, E. M. (2025). Overview of the use of artificial intelligence in predictive maintenance and state management of technical systems.
- Nwulu, E. O., Elete, T. Y., Erhue, O. V., Akano, O. A., & Omomo, K. O. (2023). Machine learning applications in predictive maintenance: Enhancing efficiency across the oil and gas industry.
- Reshma, R., Divya, A., & Dharneeshwaran, D. (2025). Predictive maintenance for industrial equipment using machine learning.

APPENDIX A. RECOMMENDED DATA DICTIONARY FOR FUTURE FIELD VALIDATION

Table 6: Recommended data dictionary for future field validation

Variable Group	Examples	Sampling/Source	Analytical Use
Equipment condition	Vibration, motor current, refrigerant/loop pressure, valve/damper position	BMS/IoT; seconds to minutes	Anomaly/fault/degradation features
Thermal operation	Supply/return air or water temperature, zone temperature, humidity	BMS; 1–15 min	Load, comfort and fault context
Energy	kW, kWh, compressor/fan/pump power, COP/EER	Submeter/BMS; 1–15 min	Energy baseline and optimisation objective

External/context	Outdoor temperature/humidity, occupancy, schedules	Weather/BMS/access systems	Normalisation and control context
Maintenance labels	Fault code, work order, component, failure time, corrective action	CMMS/manual records	Supervised labels and outcome validation

APPENDIX B. SUGGESTED STATISTICAL ANALYSIS PLAN FOR A FUTURE EMPIRICAL DATASET

1. Descriptive analysis: missingness, distributions, operating ranges, event counts and baseline energy intensity.
2. Data-quality analysis: sensor completeness, drift checks, timestamp alignment and duplicate-event review.
3. Predictive modelling: chronological train/validation/test split; compare baseline and ML models.
4. Classification evaluation: confusion matrix, precision, recall, F1-score, ROC-AUC and false alarms per operating day.
5. Regression/forecast evaluation: MAE and RMSE; use MAPE only where denominators are stable and non-zero.
6. Operational impact: compare unplanned downtime, emergency work orders, MTBF, MTTR and maintenance cost before and after deployment.
7. Energy impact: weather/load-normalised comparison of HVAC kWh and peak demand; report absolute and percentage change with confidence intervals.
8. Robustness: sensitivity to model thresholds, seasons, equipment type and building operating schedule.
9. Economic evaluation: implementation cost, avoided maintenance cost, energy-cost savings, payback period and net benefit.