Global Academic Journal of Economics and Business

Available online at https://www.gajrc.com
DOI: https://doi.org/10.36348/gajeb.2026.v08i04.040

OPEN ACCESS

ISSN: 2706-9001 (P)
ISSN: 2707-2584 (O)

Review Article

AI-Driven Predictive Maintenance for MEP Assets in Large-Scale Commercial Facilities: A Systematic Literature Review

Muhammad Younas Khan[1*]
[1]Research Scholar

***Corresponding Author**
Muhammad Younas Khan
Research Scholar

Article History
Received: 17.06.2026
Accepted: 11.08.2026
Published: 12.08.2026

Abstract: Large-scale commercial facilities depend on mechanical, electrical, and plumbing (MEP) assets whose failures can disrupt tenant services, compromise safety, increase energy use, and generate emergency work. This systematic literature review examines how artificial intelligence supports predictive maintenance of MEP assets by transforming operational data into early warnings, fault diagnoses, health classifications, and maintenance decisions. A structured review protocol, aligned with PRISMA 2020 reporting principles, synthesized 24 studies published between 2020 and 2026 together with the reporting guideline. Evidence was organized around maintenance strategies, artificial intelligence techniques, asset failure modes, data requirements, implementation settings, benefits, and barriers. The review shows that supervised learning, deep neural networks, autoencoders, transfer learning, hybrid rule-based models, and digital twins are increasingly used for chillers, air-handling units, sensors, actuators, pumps, and building energy systems. However, most evidence remains concentrated on HVAC applications, while electrical distribution, plumbing, fire protection, and vertical transportation receive limited attention. Recurring implementation constraints include scarce labeled fault data, inconsistent sensor quality, weak interoperability between building automation and maintenance systems, limited model explainability, cybersecurity exposure, and insufficient feedback from technicians. The paper proposes a framework linking asset criticality, data governance, model validation, work-order integration, and continuous learning for reliable deployment in commercial facilities.

Keywords: Artificial Intelligence, Predictive Maintenance, MEP Assets, Commercial Facilities, HVAC, Fault Detection and Diagnosis; Digital Twin, Facilities Management.

1. INTRODUCTION

Commercial facilities such as shopping malls, mixed-use complexes, offices, hospitals, and transport hubs operate as interconnected service ecosystems. Their MEP assets include chillers, air-handling units, pumps, cooling towers, switchgear, transformers, generators, water systems, fire protection equipment, elevators, and control networks. These assets must sustain comfort, safety, indoor environmental quality, tenant operations, and business continuity under fluctuating loads. Conventional reactive maintenance intervenes after failure, whereas calendar-based preventive maintenance replaces or services components at predetermined intervals. Both approaches remain necessary, but neither consistently reflects actual asset condition. Reactive work creates disruption and secondary damage; excessive preventive work can waste labor, spare parts, and useful component life. Predictive maintenance addresses this gap by estimating degradation or failure risk before functional loss occurs and by directing intervention toward assets that need attention.

The growth of building automation systems, internet-connected sensors, computerized maintenance management systems, and cloud analytics has made data-driven maintenance increasingly feasible. Bouabdallaoui *et al.*, (2021) demonstrated a building-focused framework combining IoT and automation data with machine learning, while Wu *et al.*, (2021) showed that maintenance scheduling can be optimized jointly with HVAC operation. Digital-twin studies further connect live equipment data, virtual models, fault detection, and maintenance planning (Hosamo *et al.*, 2022; Wang, 2026). These developments reposition maintenance from a routine engineering function to a decision system informed by asset criticality, probability of failure, operational consequences, energy penalties, and service priorities.

Nevertheless, implementation in large commercial facilities remains difficult. Building datasets are heterogeneous, fault labels are scarce, occupancy and weather introduce nonstationary behavior, and different vendors use incompatible naming conventions and communication protocols. Research also concentrates heavily on HVAC equipment, leaving other MEP domains comparatively underrepresented. Deep-learning models may achieve strong diagnostic performance, yet their complexity, limited interpretability, and dependence on representative training data can reduce operator trust and transferability (Taheri *et al.*, 2021; Gao *et al.*, 2023). Facility managers therefore need evidence not only about algorithm accuracy, but also about data readiness, system integration, governance, maintenance workflow, and organizational capability.

This review asks: which AI techniques are most relevant to predictive maintenance of commercial MEP assets; what failure modes and data sources do they address; what benefits and constraints are reported; and how can the evidence be translated into a practical implementation framework? The study synthesizes recent literature, compares maintenance paradigms, identifies gaps beyond HVAC, and proposes a staged model for deployment in large-scale commercial facilities.

Its contribution is a facilities-management perspective that connects technical prediction with work-order execution, contractor coordination, lifecycle cost, compliance, resilience, and performance review rather than treating model development as an analytics exercise.

2. MEP Asset Maintenance in Commercial Facilities

MEP maintenance in commercial facilities is shaped by system interdependence. A chiller fault may reduce cooling capacity, increase electricity demand, overload standby units, affect tenant comfort, and trigger complaints before a complete shutdown occurs. A failing pump, valve, sensor, relay, or controller can create symptoms elsewhere in the system, making root-cause diagnosis difficult. Electrical disturbances can interrupt multiple services simultaneously, while plumbing leakage may damage finishes, electrical equipment, or occupied areas. Fire protection and vertical transportation assets carry additional life-safety and accessibility consequences. Maintenance priorities must therefore reflect criticality, redundancy, consequence of failure, regulatory obligations, and the availability of operational alternatives.

Reactive maintenance is appropriate for low-criticality, inexpensive, non-safety assets where failure has limited consequences. Preventive maintenance remains important for statutory inspections, manufacturer requirements, consumables, lubrication, calibration, and known wear mechanisms. Condition-based maintenance uses thresholds, inspections, vibration, temperature, pressure, current, or quality indicators to initiate work when measured condition deteriorates. Predictive maintenance extends condition-based practice by applying statistical or AI models to identify abnormal patterns, classify faults, estimate future health, or calculate remaining useful life. Prescriptive maintenance adds optimization by recommending the timing and type of intervention under cost, resource, and operational constraints.

Commercial portfolios also depend on structured asset information. CAFM or CMMS platforms contain asset registers, preventive schedules, failure histories, labor hours, spare parts, warranties, and work-order outcomes. Building automation systems provide high-frequency time-series data, while BIM can supply location, system relationships, specifications, and component attributes. Condotta and Scanagatta (2023) showed that BIM-CMMS interoperability can improve operation and maintenance information flows, although Durdyev *et al.*, (2022) identified cost and capability barriers to BIM adoption in facilities management. Marzouk and Hanafy (2022) similarly demonstrated the value of linking BIM with business intelligence for maintainability assessment.

The central management challenge is not selecting one maintenance strategy for every asset. It is creating a risk-based portfolio in which reactive, preventive, condition-based, predictive, and statutory tasks coexist. AI should be applied where failure consequences, data availability, repeatability, and economic value justify model development. Table 1 summarizes these complementary strategies and their most appropriate commercial-facility applications.

This portfolio approach prevents technology enthusiasm from displacing basic maintenance discipline and ensures that prediction supports, rather than replaces, competent engineering judgment, safe isolation procedures, inspections, and verified corrective action on site.

Table 1: Comparison of maintenance strategies for commercial MEP assets

Strategy	Primary trigger	Best suited applications	Main strength	Key limitation
Reactive	Functional failure or complaint	Low-criticality, inexpensive, redundant assets	Simple and low planning effort	Unplanned downtime and secondary damage
Preventive	Calendar, run hours, statutory or manufacturer interval	Regulated inspections, consumables, known wear mechanisms	Predictable resources and compliance	May over-maintain healthy assets
Condition-based	Measured threshold or inspection finding	Motors, pumps, switchgear, filters, water systems	Intervention reflects observed condition	Thresholds may miss complex degradation
Predictive	AI estimate of anomaly, fault, health, or failure risk	Critical, data-rich assets with repeatable failure patterns	Earlier targeted intervention and warning	Requires quality data, labels, validation, and monitoring
Prescriptive	Optimized recommendation under cost and operational constraints	Central plants and portfolios with scheduling flexibility	Links prediction to timing and resource decisions	Greater complexity and governance requirement

3. SYSTEMATIC REVIEW METHODOLOGY

The review used a structured evidence-synthesis protocol informed by PRISMA 2020, which emphasizes transparent identification, selection, appraisal, and reporting of sources (Page *et al.*, 2021). The review question combined four elements: the facility context, MEP assets, AI-enabled predictive maintenance, and operational outcomes. Searches covered literature published from January 2020 to July 2026, reflecting the rapid development of building analytics, digital twins, and deep learning. Scholarly search interfaces and publisher databases were queried across ScienceDirect, SpringerLink, IEEE Xplore, MDPI, and related journal repositories.

Search strings combined terms such as “predictive maintenance,” “fault detection and diagnosis,” “health prognosis,” “remaining useful life,” “machine learning,” “deep learning,” “autoencoder,” “digital twin,” “building,” “commercial facility,” “HVAC,” “chiller,” “air handling unit,” “MEP,” “BIM,” “IoT,” “CAFM,” and “CMMS.” Boolean combinations were adjusted for each interface. Backward citation checking was used to identify foundational or closely related studies cited by recent reviews. The search was restricted to English-language peer-reviewed journal articles, research articles, and systematic reviews. The PRISMA statement was retained as the reporting guideline.

Studies were eligible when they addressed building or facility assets and examined AI, data-driven fault detection, predictive maintenance, prognostics, maintenance scheduling, or enabling information integration. Empirical studies had to report a method, dataset, framework, architecture, or case application. Reviews were included when they synthesized current AI or building fault-diagnosis evidence. Studies focused solely on industrial manufacturing, vehicle systems, or energy control without a maintenance or fault-management connection were excluded. Duplicate records, editorials, promotional materials, and sources lacking sufficient methodological detail were also excluded.

Screening proceeded through title relevance, abstract assessment, and full-text or detailed metadata review. The final evidence base comprised 24 publications and one reporting guideline. Each technical source was coded against publication year, asset type, facility context, data source, AI method, maintenance function, validation approach, reported benefit, and implementation limitation. The synthesis used thematic comparison rather than meta-analysis because datasets, fault definitions, metrics, time horizons, and equipment configurations were too heterogeneous for defensible statistical pooling.

Quality appraisal considered clarity of the problem definition, transparency of data preparation, appropriateness of validation, realism of the operating context, comparison with baselines, and discussion of limitations. Greater weight was assigned to real-building cases and multi-source data studies, while simulation-only findings were interpreted cautiously. Figure 1 presents the review workflow. This method supports reproducibility while recognizing that the field remains fragmented and that publication bias may favor successful AI applications.

For consistency, diagnostic accuracy was not treated as the sole indicator of value. Evidence was also examined for warning horizon, false-alarm burden, computational demand, interpretability, integration with maintenance records, and whether predictions resulted in actionable work decisions. Findings were then mapped to the stages of a facilities-management framework and checked for consistency across empirical studies and reviews.

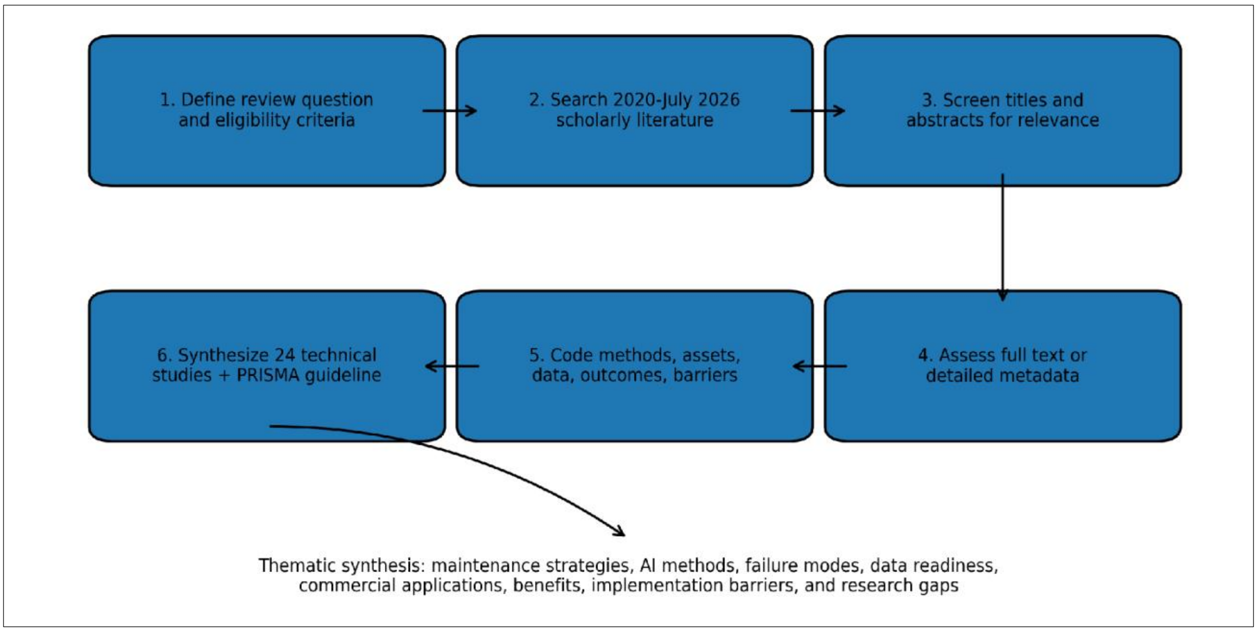

Figure 1: Structured literature identification, screening, coding, and synthesis workflow

4. Artificial Intelligence Techniques for Predictive Maintenance

Artificial intelligence techniques used for predictive maintenance can be grouped by the maintenance question they answer. Supervised classification models determine whether an asset is normal or faulty and may identify a fault class. Regression models estimate continuous variables such as efficiency loss, degradation level, failure probability, or remaining useful life. Common algorithms include logistic regression, decision trees, random forests, support vector machines, k-nearest neighbors, gradient boosting, and artificial neural networks. Their effectiveness depends on representative labeled data and meaningful features. Nelson and Culp (2022) note that machine-learning AFDD can shorten faulty operation, while Chen *et al*., (2023) show that data-driven HVAC diagnostics span classification, unsupervised learning, statistical analysis, and knowledge-based methods.

Deep learning is valuable when relationships are nonlinear, temporal, and complex. Recurrent neural networks and long short-term memory networks model sequences, making them suitable for temperature, pressure, flow, power, vibration, and control-signal histories. Convolutional neural networks can learn local patterns from one-dimensional time series or transformed images. Taheri *et al.,* (2021) applied deep recurrent architectures to HVAC fault diagnosis, and Zhang, Saeed, and Sadeghian (2023) found that convolutional approaches dominate deep-learning research. However, high accuracy in controlled datasets does not guarantee reliable performance under seasonal variation, equipment aging, maintenance changes, or different building configurations.

Unsupervised and semi-supervised learning address the scarcity of fault labels. Autoencoders learn a compressed representation of normal operation; large reconstruction errors can indicate anomalies. Bouabdallaoui *et al*., (2021) used autoencoder-based learning within a building predictive-maintenance framework. Tian, Gomez-Rosero, and Capretz (2023) combined autoencoder feature enrichment with health classification using power and weather data, demonstrating a pathway for assets without extensive internal sensing. Tra, Amayri, and Bouguila (2022) proposed a deep autoencoding Gaussian mixture method to manage outliers and limited labels in chiller diagnosis. Generative adversarial networks can exploit unlabeled data or generate additional fault examples; Li *et al*., (2021) reported improved HVAC diagnosis through a modified semi-supervised GAN.

Transfer learning aims to reuse knowledge from data-rich equipment in a new asset with limited data. Fan *et al*., (2022) transformed operational data into images and transferred learned representations across HVAC domains. This approach is attractive for commercial portfolios containing repeated equipment types, but differences in control logic, sensor calibration, climate, load, and maintenance history can cause negative transfer. Domain adaptation, recalibration, and validation remain essential. Explainable AI also matters because technicians must understand why a model generated an alarm. Gao, Miyata, and Akashi (2023) combined model pruning with layer-wise relevance analysis to reduce deployment cost and reveal influential variables. Such methods can support trust, root-cause review, and safer decisions.

Hybrid models combine data-driven learning with physical rules, fault trees, or expert knowledge. Bezyan and Zmeureanu (2022) showed that machine learning and rule-based reasoning can jointly diagnose dependent faults that create misleading symptoms. Hybridization is useful when safety constraints or engineering relationships are known but complete first-principles models are impractical. Model-based residuals can also become features for machine learning, increasing sensitivity while preserving interpretability.

Digital twins provide an integration layer rather than a single algorithm. A twin links physical assets, BIM or system models, IoT data, analytics, and maintenance services. Hosamo *et al*., (2022) combined automatic fault detection, machine-learning condition prediction, and maintenance planning for air-handling units. Hu *et al*., (2023) used a digital-twin context with parallel LSTM-autoencoders for indoor-climate failure prediction. Wang (2026) integrated digital twins, ensemble learning, and maintenance optimization in a commercial complex. These studies suggest that the greatest value emerges when predictions are connected to asset identity, system relationships, operational consequences, and work execution.

Finally, optimization converts prediction into maintenance action. Wu *et al*., (2021) used mixed-integer programming to coordinate HVAC operation and maintenance schedules, illustrating that the best intervention time depends on degradation, energy use, resources, and operating demand. Reinforcement learning and prescriptive analytics may further adapt decisions, but they require defined reward functions and safety limits. Table 2 compares the principal AI families, data needs, suitable MEP applications, and expected outputs.

Table 2: AI techniques, data requirements, MEP applications, and decision outputs

Technique family	**Typical data and MEP applications**	**Primary maintenance output**	**Main caution**
Supervised ML	Labeled sensor and work-order data; chillers, pumps, motors, sensors	Fault class, probability, degradation estimate	Needs representative labels and balanced classes
Deep sequence/CNN models	High-frequency multivariate time series; AHUs, chillers, control networks	Temporal fault detection and future-state prediction	Computational demand, explainability, and drift
Autoencoders and anomaly detection	Mostly normal-operation data; assets with few confirmed faults	Anomaly score and health state	Anomaly may not identify root cause
Semi-supervised/GAN methods	Small labeled set plus large unlabeled dataset	Improved classification where labels are scarce	Synthetic patterns may not reflect real faults
Transfer learning	Data from repeated equipment across buildings or climates	Reusable representation for a new asset	Negative transfer without local recalibration
Hybrid, digital twin, and optimization	Sensor data, physics, BIM, CMMS, system topology, cost and schedules	Explainable diagnosis, integrated planning, recommended action	Integration complexity and strong governance needs

Source: *Author's synthesis of the reviewed literature.*

5. MEP Failure Modes and Data Requirements

Predictive maintenance begins with credible failure definitions. HVAC assets dominate the literature because building automation systems already record temperatures, pressures, flows, valve positions, compressor states, fan speeds, currents, and setpoints. Common chiller faults include refrigerant leakage, condenser fouling, evaporator flow problems, sensor bias, compressor degradation, and abnormal pressure conditions. Air-handling units experience stuck dampers, leaking valves, fouled filters, fan faults, coil degradation, and sensor or actuator failures. Papadopoulos *et al*., (2022) emphasized the challenge of isolating sensor and actuator faults in multi-zone air-handling systems. Jiang *et al*., (2023) developed a time-series framework for predicting chiller faults from temperatures and flow rates, while Song *et al*., (2023) combined signal processing and machine learning for air-conditioning diagnosis.

Electrical MEP assets require different indicators. Switchgear, transformers, generators, motors, variable-frequency drives, and distribution boards may exhibit overheating, insulation deterioration, harmonic distortion, contact wear, bearing degradation, current imbalance, partial discharge, or repeated protective trips. Useful data include phase current, voltage, power factor, harmonics, thermal images, vibration, oil or insulation test results, event logs, loading, ambient temperature, and operating hours. Compared with HVAC, commercial-building evidence on integrated AI prediction for electrical distribution remains limited. Consequently, facilities often rely on condition monitoring, thermography, protection histories, and vendor diagnostics without a unified predictive model.

Plumbing and water systems present leakage, pressure instability, pump cavitation, valve failure, tank-level anomalies, water-quality deterioration, and blockage. Data may come from flow meters, pressure sensors, pump current, vibration, acoustic devices, moisture sensors, and water-quality instruments. Fire protection systems require special caution because prediction cannot replace statutory inspection, testing, and code compliance. AI may support anomaly detection for pumps, valves, tanks, alarm devices, and network pressure, but any maintenance recommendation must preserve fail-safe operation and verification. Vertical transportation assets similarly generate motor current, vibration, door-cycle, brake, fault-code, travel-time, and load data, yet proprietary controllers can restrict access.

Across all MEP domains, data quality is more important than data volume. Time stamps must be synchronized; sensor units, ranges, and sampling rates must be consistent; missing values and outliers must be handled; and maintenance events must be linked to operational data. Matetić *et al*., (2023) identify system complexity and access to reliable sensory measurements as recurring FDD barriers. Cross-domain variability also matters: the same equipment model may behave differently because of climate, occupancy, control sequences, commissioning quality, and maintenance condition.

Work-order data provide labels and context, but free-text descriptions are often incomplete. Failure codes, cause codes, action codes, parts used, downtime, and technician confirmation should be standardized. A practical dataset therefore combines sensor streams, alarms, control commands, asset metadata, maintenance history, inspection findings, energy data, weather, occupancy proxies, and system topology. The resulting information architecture should support both algorithm development and auditable engineering decisions.

Label design also matters. A model trained on the date a work order was closed may miss the degradation period, while alarms created after a failure provide limited warning value. Teams should define horizons, healthy reference periods, fault onset, failure, and recovery. Class imbalance should be addressed without creating unrealistic synthetic patterns.

6. Applications in Large-Scale Commercial Facilities

Large-scale commercial facilities provide a strong business case for predictive maintenance because they contain repeated equipment, long operating hours, variable occupancy, high cooling demand, and strict service expectations. Shopping malls are especially complex: retail, leisure, food, parking, and back-of-house zones have different schedules and environmental requirements. A single central-plant fault can affect many occupants and tenants, while emergency repair may require temporary cooling and cross-team coordination.

HVAC applications are the most mature. Bouabdallaoui *et al.,* (2021) demonstrated a five-stage framework covering data collection, processing, model development, fault notification, and improvement. Its case showed that building predictive maintenance is technically feasible but depends on reliable feedback and accessible data. Hosamo *et al*., (2022) extended this logic through a digital twin for air-handling units, combining rule-based fault detection, machine-learning condition prediction, and maintenance planning. For mall operations, the same architecture can prioritize degraded filters, unstable valves, abnormal fan behavior, coil fouling, or sensor drift before comfort complaints become widespread.

Chillers and central plants are high-value targets because they are energy intensive, critical, and data rich. Jiang *et al*., (2023) used time-series learning to predict chiller faults over future intervals. Tra *et al*., (2022) addressed outliers and limited labels, both common in operational chiller datasets. Wu *et al*., (2021) showed that predictive maintenance scheduling can be optimized alongside plant operation, enabling facility teams to plan intervention during lower-load periods, preserve redundancy, and reduce the energy penalty of degraded equipment. Portfolio models can compare similar chillers and identify abnormal efficiency or failure patterns.

AI can support broader operational decisions. Power-consumption and weather data may classify equipment health where internal sensing is unavailable (Tian *et al*., 2023). Deep-learning temperature predictions can support digital-twin monitoring and distinguish equipment degradation from changing loads (Norouzi *et al.,* 2023). Predictive control research shows how learned models can forecast equipment and zone behavior (Ra, Kim, and Park, 2023). When integrated with maintenance analytics, these forecasts help separate control problems, capacity limitations, and physical faults.

Mixed-use facilities benefit from contextual information. Occupancy, trading hours, outdoor conditions, tenant fit-out changes, cleaning schedules, and events influence MEP behavior. A model that ignores context may create false alarms during unusual operating periods. Digital twins can combine BIM relationships, sensor data, control sequences, and work history, creating a shared view for engineers, contractors, and managers. Wang (2026) reported a digital-twin and machine-learning framework for a large commercial complex, illustrating the movement toward integrated decision support rather than isolated alarms.

Implementation should begin with a bounded use case. A mall might select one chiller plant, several air-handling units, or critical pumps with consistent sensors and documented failures. The team should define the operational decision before choosing the algorithm: detect a fault, predict failure within a time window, rank risk, or recommend maintenance. Alerts should create or enrich CMMS work orders, display contributing variables, state confidence, and request technician confirmation. Confirmed outcomes then become training data.

Portfolio scaling requires standard asset naming, reusable data templates, common failure codes, and governance across vendors. Models should be monitored for drift after commissioning changes, sensor replacement, refurbishment, or seasonal transitions. Predictive maintenance creates value when it reduces avoidable downtime, energy waste, emergency cost, and service disruption without overwhelming staff with alarms. For commercial facilities, success is therefore measured through operational outcomes, not model accuracy alone under realistic commercial operating conditions.

7. Benefits, Limitations, and Implementation Challenges

The principal benefit of AI-driven predictive maintenance is earlier, more selective intervention. By detecting degradation before functional failure, facilities can schedule work during low-demand periods, protect redundancy, order parts, coordinate contractors, and avoid secondary damage. Energy performance may improve because fouled heat exchangers, leaking valves, sensor bias, excessive cycling, and control faults often increase consumption before causing shutdown. Predictive insight can also support lifecycle planning by distinguishing assets that require adjustment, repair, overhaul, or replacement.

Operational value depends on decision quality. False negatives allow failures to progress, while false positives consume labor and weaken trust. Accuracy, precision, recall, warning time, and false-alarm rate should therefore be assessed together. Deep models may perform well on historical datasets yet fail when equipment, controls, occupancy, or climate conditions change. Zhang, Saeed, and Sadeghian (2023) found that much HVAC deep-learning evidence relies on laboratory or simulation data, limiting direct generalization. Transfer learning and hybrid models help, but require site validation.

Data governance is another barrier. Missing sensors, drifting calibration, inconsistent point names, disconnected work orders, and incomplete failure codes reduce model reliability. Cybersecurity exposure increases when operational technology connects to cloud platforms or external analytics. Access controls, network segmentation, encryption, logging, backup, vendor responsibility, and incident procedures must be defined. Commercial agreements should also clarify data ownership, retention, model intellectual property, and responsibilities for incorrect recommendations.

Organizational capability can be more limiting than algorithms. Facilities teams need engineers who understand systems, analysts who understand time-series data, technicians who verify faults, and managers who connect predictions to budgets and service priorities. Durdyev *et al*., (2022) identified cost and lack of expertise as major BIM-FM barriers; similar constraints affect predictive maintenance. Explainability is essential where alarms influence safety, tenant service, or expensive intervention. Gao *et al.,* (2023) demonstrate that pruning and relevance methods can reduce model burden and reveal influential variables.

Economic assessment should include sensors, connectivity, data engineering, software, integration, training, cybersecurity, model monitoring, and change management, not only avoided repair cost. Benefits should be tracked through downtime, emergency work, energy variance, mean time between failures, maintenance cost, repeat faults, tenant complaints, and asset life. A phased pilot with clear baselines is usually more defensible than immediate portfolio-wide deployment.

Human authorization should remain mandatory for safety-critical isolation, statutory systems, and capital decisions. Predictive outputs are advisory evidence, not proof of failure; technicians must confirm conditions, record findings, and feed outcomes back into the governance process.

8. Research Gaps and Proposed Implementation Framework

The evidence reveals four gaps. First, research is concentrated on HVAC, especially chillers and air-handling units; commercial electrical, plumbing, fire protection, and vertical transportation assets need comparable datasets and validated methods. Second, many studies evaluate diagnostic accuracy but do not report warning horizon, maintenance action, avoided downtime, cost, or technician acceptance. Third, cross-building transfer remains weak because equipment configurations, control sequences, climates, and data standards differ. Fourth, limited explainability and poor CMMS integration prevent analytics from becoming routine maintenance practice.

Figure 2 proposes a six-stage implementation framework. Stage one establishes governance, business objectives, safety boundaries, and accountable roles. Stage two ranks assets by criticality, failure consequence, data availability, and economic opportunity. Stage three builds the data foundation by standardizing asset identity, sensor points, work-order codes, timestamps, and system relationships. Stage four develops and validates models using representative operating conditions, baseline comparisons, uncertainty estimates, and engineer review. Stage five integrates approved outputs with dashboards and CMMS workflows so alarms create prioritized, explainable actions rather than separate analytics reports. Stage six measures outcomes, captures technician feedback, monitors drift, and retrains or retires models.

The framework treats predictive maintenance as a controlled facilities-management process. A model should advance only when it demonstrates technical validity, operational usefulness, cybersecurity compliance, and measurable value. Future research should prioritize multi-site commercial datasets, non-HVAC MEP assets, interoperable data models, human-centered explanations, and prospective trials that compare AI-supported decisions with established maintenance practice. Such trials should document failure consequences, resource impacts, and changes in asset reliability and service continuity.

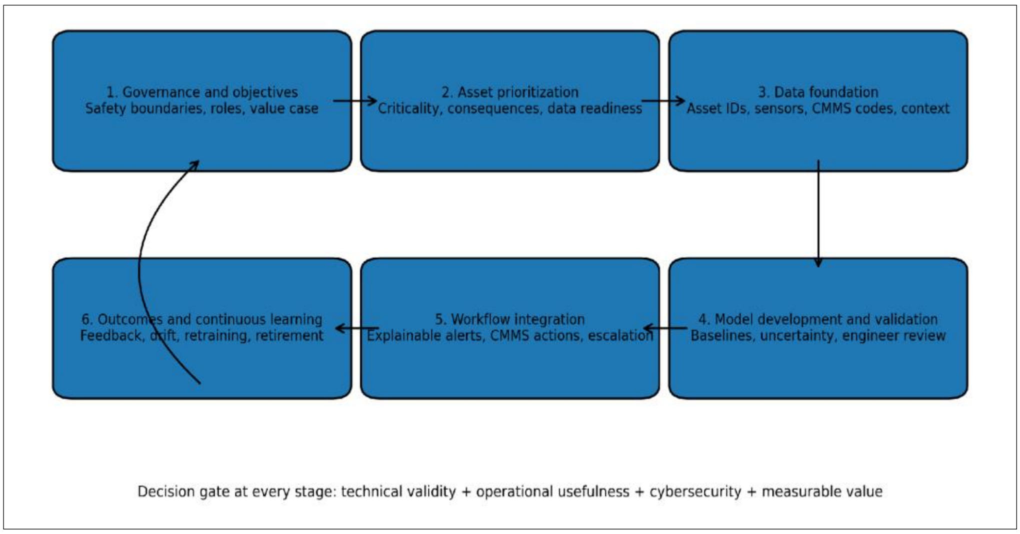

Figure 2: Six-stage framework for AI-driven predictive maintenance of commercial MEP assets

9. CONCLUSION

AI-driven predictive maintenance can strengthen the reliability, efficiency, and continuity of MEP services in large-scale commercial facilities. Evidence demonstrates applications of supervised learning, deep neural networks, autoencoders, transfer learning, hybrid diagnostics, optimization, and digital twins, particularly for HVAC assets. These methods can identify abnormal behavior, classify faults, estimate health, and improve maintenance timing. Their value, however, is not determined by algorithm accuracy alone.

Successful deployment requires a reliable asset register, synchronized sensor data, standardized maintenance records, validation, explainable outputs, cybersecurity controls, and integration with work management. Engineers and technicians remain essential for defining failure modes, confirming alarms, authorizing interventions, and converting outcomes into training data. The review also highlights an imbalance in the literature: chillers and air-handling systems are well represented, while electrical distribution, plumbing, fire protection, and vertical transportation need attention.

The proposed six-stage framework connects governance, asset criticality, data preparation, model validation, CMMS integration, and continuous improvement. It supports adoption through bounded pilots and measured outcomes. For commercial facilities, the appropriate goal is not autonomous maintenance, but evidence-based decision support that reduces preventable failures without compromising safety, compliance, or engineering judgment. Future multi-site studies should test transferability, lifecycle value, human acceptance, and service resilience across complete MEP portfolios.

REFERENCES

- Bezyan, B., & Zmeureanu, R. (2022). Detection and diagnosis of dependent faults that trigger false symptoms of heating and mechanical ventilation systems using combined machine learning and rule-based techniques. Energies, 15(5), 1691. https://doi.org/10.3390/en15051691
- Bouabdallaoui, Y., Lafhaj, Z., Yim, P., Ducoulombier, L., & Bennadji, B. (2021). Predictive maintenance in building facilities: A machine learning-based approach. Sensors, 21(4), 1044. https://doi.org/10.3390/s21041044
- Chen, Z., O'Neill, Z., Wen, J., et al. (2023). A review of data-driven fault detection and diagnostics for building HVAC systems. Applied Energy, 339, 121030. https://doi.org/10.1016/j.apenergy.2023.121030
- Condotta, M., & Scanagatta, C. (2023). BIM-based method to inform operation and maintenance phases through a simplified procedure. Journal of Building Engineering, 65, 105730. https://doi.org/10.1016/j.jobe.2022.105730
- Durdyev, S., Ashour, M., Connelly, S., & Mahdiyar, A. (2022). Barriers to the implementation of Building Information Modelling (BIM) for facility management. Journal of Building Engineering, 46, 103736. https://doi.org/10.1016/j.jobe.2021.103736
- Fan, C., He, W., Liu, Y., Xue, P., & Zhao, Y. (2022). A novel image-based transfer learning framework for cross-domain HVAC fault diagnosis: From multi-source data integration to knowledge sharing strategies. Energy and Buildings, 262, 111995. https://doi.org/10.1016/j.enbuild.2022.111995
- Gao, Y., Miyata, S., & Akashi, Y. (2023). How to improve the application potential of deep learning model in HVAC fault diagnosis: Based on pruning and interpretable deep learning method. Applied Energy, 348, 121591. https://doi.org/10.1016/j.apenergy.2023.121591
- Hosamo, H. H., Svennevig, P. R., Svidt, K., Han, D., & Nielsen, H. K. (2022). A Digital Twin predictive maintenance framework of air handling units based on automatic fault detection and diagnostics. Energy and Buildings, 261, 111988. https://doi.org/10.1016/j.enbuild.2022.111988
- Hu, W., Wang, X., Tan, K., & Cai, Y. (2023). Digital twin-enhanced predictive maintenance for indoor climate: A parallel LSTM-autoencoder failure prediction approach. Energy and Buildings, 301, 113738. https://doi.org/10.1016/j.enbuild.2023.113738
- Jiang, Z., Risbeck, M. J., Samy, S. C. K., Zhang, C., Cyrus, S., & Lee, Y. M. (2023). A timeseries supervised learning framework for fault prediction in chiller systems. Energy and Buildings, 285, 112876. https://doi.org/10.1016/j.enbuild.2023.112876
- Li, B., Cheng, F., Cai, H., Zhang, X., & Cai, W. (2021). A semi-supervised approach to fault detection and diagnosis for building HVAC systems based on the modified generative adversarial network. Energy and Buildings, 246, 111044. https://doi.org/10.1016/j.enbuild.2021.111044
- Marzouk, M., & Hanafy, M. (2022). Modelling maintainability of healthcare facilities services systems using BIM and business intelligence. Journal of Building Engineering, 46, 103820. https://doi.org/10.1016/j.jobe.2021.103820
- Matetić, I., Štajduhar, I., Wolf, I., & Ljubić, S. (2023). A review of data-driven approaches and

techniques for fault detection and diagnosis in HVAC systems. Sensors, 23(1), 1. https://doi.org/10.3390/s23010001

- Nelson, W., & Culp, C. (2022). Machine learning methods for automated fault detection and diagnostics in building systems—A review. Energies, 15(15), 5534. https://doi.org/10.3390/en15155534
- Norouzi, P., Maalej, S., & Mora, R. (2023). Applicability of deep learning algorithms for predicting indoor temperatures: Towards the development of digital twin HVAC systems. Buildings, 13(6), 1542. https://doi.org/10.3390/buildings13061542
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71
- Papadopoulos, P. M., Lymperopoulos, G., Polycarpou, M. M., & Ioannou, P. (2022). Distributed diagnosis of sensor and actuator faults in air handling units in multi-zone buildings: A model-based approach. Energy and Buildings, 256, 111709. https://doi.org/10.1016/j.enbuild.2021.111709
- Ra, S. J., Kim, J.-H., & Park, C. S. (2023). Real-time model predictive cooling control for an HVAC system in a factory building. Energy and Buildings, 285, 112860. https://doi.org/10.1016/j.enbuild.2023.112860
- Song, Y., Ma, Q., Zhang, T., Li, F., & Yu, Y. (2023). Research on fault diagnosis strategy of air-conditioning systems based on DPCA and machine learning. Processes, 11(4), 1192. https://doi.org/10.3390/pr11041192
- Taheri, S., Ahmadi, A., Mohammadi-Ivatloo, B., & Asadi, S. (2021). Fault detection diagnostic for HVAC systems via deep learning algorithms. Energy and Buildings, 250, 111275. https://doi.org/10.1016/j.enbuild.2021.111275
- Tian, R., Gomez-Rosero, S., & Capretz, M. A. (2023). Health prognostics classification with autoencoders for predictive maintenance of HVAC systems. Energies, 16(20), 7094. https://doi.org/10.3390/en16207094
- Tra, V., Amayri, M., & Bouguila, N. (2022). Outlier detection via multiclass deep autoencoding Gaussian mixture model for building chiller diagnosis. Energy and Buildings, 259, 111893. https://doi.org/10.1016/j.enbuild.2022.111893
- Wang, R. (2026). A data-driven predictive maintenance framework for smart buildings: Integrating digital twin and machine learning in HVAC systems. Journal of Building Engineering, 120, 115416. https://doi.org/10.1016/j.jobe.2026.115416
- Wu, Y., Maravelias, C. T., Wenzel, M. J., ElBsat, M. N., & Turney, R. T. (2021). Predictive maintenance scheduling optimization of building heating, ventilation, and air conditioning systems. Energy and Buildings, 231, 110487. https://doi.org/10.1016/j.enbuild.2020.110487
- Zhang, F., Saeed, N., & Sadeghian, P. (2023). Deep learning in fault detection and diagnosis of building HVAC systems: A systematic review with meta analysis. Energy and AI, 12, 100235. https://doi.org/10.1016/j.egyai.2023.100235.