GSJ: Volume 5, Issue 5, May 2017, Online: ISSN 2320-9186

www.globalscientificjournal.com

IoT-based predictive maintenance in critical facilities: a literature review

Fabio Justo de Sant´Anna

Author Details

Fabio Justo de Sant'Anna is the founder and principal consultant at JP Management America LLC, working with project management and facilities optimization for critical infrastructure and smart buildings in Brazil and the United States. E-mail: fajusa74@gmail.com

KeyWords

Critical infrastructure, energy efficiency, facilities management, HVAC systems, Internet of Things (IoT), machine learning, predictive maintenance, smart buildings

ABSTRACT

Predictive maintenance has become a practical way to reduce downtime, energy waste and unexpected failures in building facilities, especially where HVAC and other critical systems operate continuously. Over the last decade, Internet of Things (IoT) technologies and machine-learning methods have enabled continuous monitoring of equipment and data-driven detection of anomalies before they turn into major problems. This paper presents a concise literature review on IoT-based predictive maintenance in building facilities, with emphasis on HVAC systems and critical infrastructure contexts. The review synthesizes how recent studies describe IoT architectures, data sources, and predictive models, and how they report the benefits and limitations of real deployments. Particular attention is given to three recurring challenges: heterogeneous and often incomplete data, building-specific operating conditions, and the difficulty of integrating predictive models into existing maintenance workflows. The findings suggest that predictive maintenance in facilities should be understood as a systems problem, combining sensing, data handling, analytics and process change, rather than as an isolated algorithmic task. The paper concludes by outlining practical implications for facility managers

and project teams, and by highlighting the need for more long-term, real-world evaluations of predictive maintenance strategies in different types of critical facilities.

Introduction

Operation and maintenance activities continue to represent a major challenge in both commercial and institutional buildings, especially where infrastructure is aging and service interruptions have consequences beyond simple repair costs. Building systems consume a substantial share of total energy use, and faults in HVAC and related installations can materially increase energy consumption while reducing comfort, safety, and reliability. For this reason, maintenance strategy is no longer a purely technical concern. It has become a core management issue tied to operational performance, sustainability, and business continuity.

Traditional approaches to building maintenance tend to fall into two categories: corrective maintenance, which reacts after failure, and preventive maintenance, which follows time-based or usage-based schedules. Although both approaches remain common, each has clear limitations. Corrective maintenance exposes facilities to unnecessary downtime and emergency intervention, while preventive maintenance often results in unnecessary inspections or replacement of components that are still functioning adequately. In mission-critical environments, these inefficiencies can be particularly costly because system failure may affect healthcare delivery, data operations, manufacturing continuity, or occupant safety.

The spread of IoT technologies has opened a more sophisticated path. Sensors, gateways, and connected control platforms can now collect real-time data from equipment and environmental conditions at relatively low cost, while machine-learning techniques make it possible to process that information into useful predictions about faults, anomalies, or deterioration trends. This combination has made predictive maintenance increasingly attractive in the built environment. Instead of relying only on schedules or human observation, facilities can use data patterns to identify when intervention is actually needed.

However, the literature also shows that predictive maintenance in buildings is different from predictive maintenance in more standardized industrial settings. Buildings are operationally diverse, and similar equipment may behave differently depending on occupancy patterns, architecture, climate exposure, control logic, and maintenance history. This means that predictive maintenance in facilities requires more than algorithm selection. It requires a structured understanding of building data sources, user workflows, and the specific realities of facilities management.

This article reviews literature that addresses IoT-enabled predictive maintenance in facilities, with emphasis on practical deployment rather than abstract modelling alone. The purpose is to synthesize the architectural, analytical, and managerial lessons emerging from the literature and to discuss how they apply to critical facilities that need stronger reliability and energy performance under real-world constraints.

Survey methodology

This article follows a narrative literature review approach oriented toward practical application. The review was anchored primarily in two peer-reviewed articles that together provide a strong foundation for the topic. The first is *Predictive Maintenance in Building Facilities: A Machine Learning-Based Approach*, published in *Sensors* in 2021, which proposes a structured predictive maintenance framework for buildings and demonstrates it through an HVAC case study in a sports facility. The second is *Recent Advances in Internet of Things (IoT) Infrastructures for Building Energy Systems: A Review*, also published in *Sensors* in 2021, which surveys IoT architectures, control layers, sensors, and applications for building energy systems, with specific attention to smart buildings and HVAC controls.

These two sources were selected because they complement each other well. One approaches the subject from the perspective of building energy IoT infrastructure, while the other approaches it from the perspective of maintenance implementation and machine-learning workflow in real facilities. Together, they offer enough conceptual and practical material to support a focused literature review for publication in a technical outlet.

The review process involved four steps. First, the selected studies were read in full, with attention to research objectives, system architectures, data sources, methods, case results, and implementation challenges. Second, the material was grouped into recurring themes: data collection and infrastructure, predictive modelling, operational integration, and barriers to adoption. Third, the findings were compared for convergence and divergence, especially regarding what is technically possible versus what is operationally scalable. Finally, these findings were interpreted from the viewpoint of critical facilities management, where uptime, energy performance, and lifecycle risk are central concerns.

This review does not claim to be exhaustive. Its purpose is to present a concise, technically grounded synthesis suitable for practitioners, project managers, and facility leaders seeking to understand the current literature and its practical implications.

IoT as the enabling layer for predictive maintenance

The reviewed literature shows that predictive maintenance in facilities depends first on the availability of a functional IoT infrastructure. Without reliable sensing, communication, storage, and control layers, there is no practical way to transform physical equipment behaviour into usable predictive intelligence. The article by Yaïci et al. frames this clearly by describing IoT systems as multi-layer structures consisting of hardware or sensor layers, software control layers, and application layers designed to support monitoring, forecasting, and operational response in buildings.

This architectural view is important because it places predictive maintenance within a larger digital ecosystem. In buildings, sensors do not operate in isolation. They are part of a network that may include occupancy sensors, temperature and humidity devices, energy meters, actuators, wireless communication protocols, cloud or local databases, and user-facing dashboards. The significance of this is practical: predictive maintenance is not merely a machine-learning exercise, but an infrastructure problem involving interoperability, data flow, and decision support.

The literature also makes clear that the same IoT infrastructure used for predictive maintenance can simultaneously support energy management and comfort optimization. This dual use is valuable in critical facilities because maintenance interventions often affect energy performance, and energy anomalies can be early indicators of asset degradation. An inefficient chiller, unstable air-handling unit, or poorly controlled HVAC sequence may represent both an energy issue and a maintenance issue.

A second relevant point is that building IoT solutions face architectural constraints not always present in industrial contexts. Buildings may contain legacy automation systems, fragmented communication standards, and mixed generations of equipment. In older facilities, data access may be limited to a building automation system that was not designed for advanced analytics. This means that predictive maintenance in facilities often begins with imperfect and incomplete data environments, which immediately affects the sophistication and reliability of the models that can be developed.

Data sources and the challenge of building-specific variability

One of the strongest contributions in the literature is the recognition that data in buildings are both abundant and deficient at the same time. Bouabdallaoui et al. identify multiple potential data sources for predictive maintenance in facilities, including building automation systems, IoT devices, computerized maintenance management systems, and building information modelling environments. In theory, these create a rich information base. In practice, however, data are often incomplete, poorly structured, weakly labelled, or simply not retained in ways that support machine learning.

This mismatch is one of the central tensions in predictive maintenance for facilities. Building automation systems can generate streams of temperatures, pressures, flow rates, alarms, and binary equipment states. IoT devices can supplement this with vibration, electrical, and environmental measurements. Yet the existence of data does not automatically mean the data are useful. Historical records may contain gaps, maintenance work orders may not clearly identify actual faults, and operational context may be missing.

The reviewed literature repeatedly emphasizes the importance of preprocessing, cleaning, and structuring data before meaningful prediction can occur. This is especially relevant in facilities because buildings are not uniform environments. The same air-handling unit model can behave differently in a hospital, a sports facility, and a commercial office, even under similar weather conditions, because occupancy, schedules, control sequences, and operational priorities differ. For that reason, building-specific variability becomes a defining constraint on predictive maintenance in the built environment.

This insight has major practical implications. It means that predictive maintenance models developed for one facility cannot always be transferred directly to another without adaptation. It also explains why many theoretically promising models fail to scale beyond pilot projects. The issue is not only algorithm performance; it is the difficulty of turning diverse, context-sensitive building data into stable and reusable predictive assets.

Machine-learning approaches and their practical meaning

The literature presents machine learning as a useful but not self-sufficient tool in predictive maintenance. Bouabdallaoui et al. propose a five-step framework consisting of data collection, data processing, model development, fault notification, and model improvement, and they test this structure using HVAC equipment in a sports facility. Their model uses an autoencoder architecture with LSTM layers, chosen largely because labelled fault data are scarce in buildings and time-series information dominates facility telemetry.

This choice is analytically sensible. In facilities management, faults are often not formally labelled in a machine-learning-ready manner, and anomaly detection methods therefore become more attractive than fully supervised classification systems. The use of reconstruction error as an anomaly score reflects a practical attempt to work with real constraints rather than ideal datasets. It also illustrates a broader lesson from the literature: successful predictive maintenance in buildings is often based on methods that are robust under incomplete information, not only on the most advanced algorithm available.

The case study described in the article is especially valuable because it shifts the discussion from theory to implementation. The authors report that true positive alerts were achieved and that some failures were anticipated in advance, but they also acknowledge false positives, undetected events, and the limitations created by a short observation period. This honesty is useful because it reflects how predictive maintenance behaves in real operations. Models can help, but they rarely produce perfect foresight, particularly when the dataset is small or operational conditions are atypical.

An important takeaway is that practical predictive maintenance requires a learning loop, not a one-time model deployment. Feedback from the maintenance team, threshold adjustments, and periodic retraining are all part of the framework. This gives the approach managerial significance: predictive maintenance should be integrated into facilities operations as an evolving process, not treated as a plug-and-play technology product.

Energy efficiency, reliability, and critical facilities

The review by Yaïci et al. broadens the discussion by showing how IoT infrastructures in buildings are closely linked to energy performance. Smart sensors, occupancy-based control, adaptive and predictive strategies, and connected HVAC systems can reduce waste and improve comfort when properly integrated into building operations. This matters for predictive maintenance because faults in HVAC and related systems are often energy-intensive before they become operationally disruptive.

In critical facilities, this connection between maintenance and energy is even more important. A hospital, data center, industrial plant, or command facility cannot afford unstable climate control, unplanned pump failure, or deteriorating electrical support systems. In such settings, predictive maintenance can be understood as part of resilience strategy. It helps facility teams move from delayed reaction to earlier intervention, reducing both operational risk and hidden energy losses.

The literature supports this integrated view. Building IoT systems are designed not only to record what is happening now, but also to support both predictive and adaptive control actions. This opens the possibility of combining asset diagnostics with broader building performance objectives. For example, the same data environment that identifies deteriorating HVAC behaviour can also support continuous commissioning, setpoint refinement, and more accurate scheduling of loads.

From a facilities management perspective, this is one of the strongest arguments for adoption. Predictive maintenance is not useful only because it reduces failures. It is useful because it aligns maintenance, energy, comfort, and lifecycle management within one digital operating logic.

Implementation barriers and why many pilots do not scale

Although the literature is optimistic about technical potential, it is equally clear about the barriers to implementation. One barrier is data scarcity in operationally meaningful form. Another is the uniqueness of buildings, which makes standardization difficult. A third is return on investment: predictive maintenance may require months of data collection, integration effort, and model tuning before delivering consistent visible value.

This time horizon matters. Facility managers are often under pressure to control budgets, respond to complaints, and maintain compliance with limited staff. A predictive maintenance initiative that requires substantial upfront work before producing reliable outcomes may be technically sound yet operationally difficult to justify. Bouabdallaoui et al. explicitly note this tension and point to it as one of the reasons why implementation remains challenging in buildings.

Cybersecurity and privacy concerns also deserve attention. Yaïci et al. stress that IoT systems in buildings involve the collection and transmission of behavioural and operational data, which can create risk if poorly designed or inadequately protected. In critical facilities, this concern is amplified because operational technology environments often support essential services and may be connected to regulated systems.

There is also a governance issue. Predictive maintenance changes how teams interpret alarms, assign priorities, and decide when to intervene. If the maintenance workflow remains disconnected from analytics outputs, then even a good predictive model may not create practical value. This is why the literature increasingly points toward integrated frameworks rather than isolated technical pilots.

Implications for practice in Brazil and the United States

The lessons in the reviewed literature are especially relevant for Brazil and the United States, although the institutional settings differ. In both countries, facility operators face pressure to improve uptime, energy performance, and cost control in the context of aging assets and uneven modernization across building portfolios. Predictive maintenance offers a route toward smarter operations, but only if the implementation strategy is realistic.

A practical starting point is to focus on critical systems rather than trying to digitize every asset at once. HVAC plants, pumps, boilers, power-support systems, and other equipment with clear operational and energy impact are usually the most suitable entry points. From there, teams can build a limited but high-value data environment, connect it to existing maintenance processes, and gradually mature the predictive model.

For Brazil, where many facilities still operate under budget constraints and with mixed automation maturity, the literature suggests that a phased approach is more realistic than a full digital overhaul. For the United States, where many facilities already possess more sophisticated building systems but remain fragmented across vendors and platforms, the priority may be interoperability and workflow integration rather than pure sensor deployment. In both contexts, the literature supports a strategy centered on pragmatic deployment, building-specific adaptation, and operational feedback loops.

Conclusion

The literature reviewed in this article shows that IoT-enabled predictive maintenance in facilities has moved from conceptual promise toward operational feasibility. Research now provides credible frameworks for collecting building data, processing time-series information, detecting anomalies, and integrating predictive alerts into maintenance decision-making. At the same time, the literature does not support technological optimism without qualification. Real implementation remains difficult because of fragmented data, context-specific equipment behaviour, integration complexity, and the organizational effort required to turn predictive insight into action.

Two major conclusions emerge. First, predictive maintenance in buildings should be treated as a systems problem, not just an analytics problem. Sensors, communication protocols, storage, automation logic, maintenance records, user feedback, and management processes all influence whether predictive models become useful operational tools. Second, in critical facilities, predictive maintenance offers value not only by reducing failure risk but also by supporting energy efficiency, resilience, and better lifecycle control.

For practitioners, the implication is clear: begin with critical assets, invest in usable data, and align predictive tools with actual maintenance workflow. For researchers, the most useful next step is to generate more long-duration, real-world studies showing how predictive maintenance performs across different building types and operational conditions. Bridging this gap between technical capability and practical facility deployment will define the next stage of maturity in the field.

Acknowledgment

The author acknowledges the contribution of facility management professionals, project teams, and technical colleagues whose day-to-day operational challenges continue to shape practical thinking about reliability, energy performance, and infrastructure resilience.

References

Bouabdallaoui, Y., Lafhaj, Z., Yim, P., Ducoulombier, L., & Bennadji, B. Predictive maintenance in building facilities: A machine learning-based approach. *Sensors*, 21(4), 1044, 2021. https://doi.org/10.3390/s21041044

Yaïci, W., Krishnamurthy, K., Entchev, E., & Longo, M. Recent advances in Internet of Things (IoT) infrastructures for building energy systems: A review. *Sensors*, 21(6), 2152, 2021. https://doi.org/10.3390/s21062152