

 Research Article



Journal Website:
<http://sciencebring.com/index.php/ijasr>

Copyright: Original content from this work may be used under the terms of the creative commons attributes 4.0 licence.

AI-Enabled Predictive Maintenance and Intelligent DevOps for Cyber-Physical Energy and Manufacturing Systems: An Integrated Theoretical and Empirical Inquiry

Submission Date: January 01, 2026, **Accepted Date:** January 19, 2026,
Published Date: February 08, 2026

Christopher V. Belmore

Department of Electrical and Computer Engineering, University of Alberta, Canada

ABSTRACT

The accelerating convergence of artificial intelligence, cyber-physical systems, and intelligent automation has fundamentally reshaped how complex industrial assets are designed, deployed, operated, and maintained. Across sectors as diverse as advanced manufacturing, photovoltaic energy production, smart grids, and mining operations, organizations increasingly rely on algorithmic systems to anticipate failures, optimize performance, and coordinate distributed infrastructures. Yet despite a growing body of research on predictive maintenance, Internet of Things architectures, and machine learning-based diagnostics, a persistent conceptual and operational gap remains between the analytics layer and the software and deployment layer that operationalizes these insights at scale. This gap has become especially salient in the era of AI-driven DevOps, where machine learning models are no longer static tools but continuously evolving agents embedded in deployment pipelines, monitoring systems, and decision workflows.

This article develops a comprehensive theoretical and empirical framework that integrates predictive maintenance, IoT-based monitoring, and AI-enabled DevOps into a single coherent paradigm for modern cyber-physical systems. Grounded in an extensive synthesis of literature on Industry 4.0, Maintenance 4.0, and intelligent energy management, the study positions AI-driven DevOps as the connective tissue that links sensor-level data acquisition to enterprise-level decision-making and automated intervention. Drawing on the conceptual foundations of intelligent automation articulated in contemporary DevOps research, including the systematic review of machine learning-based deployment and maintenance strategies by Varanasi (2025), the paper argues that predictive maintenance cannot be fully realized without embedding its models within adaptive, continuously learning operational pipelines.

The methodology employs a qualitative-analytical research design that synthesizes architectural models, algorithmic approaches, and application case studies from manufacturing and renewable energy systems. Through comparative analysis, the study evaluates how different predictive maintenance strategies, such as neural network-based fault diagnosis, physics-informed learning, and edge-based analytics, perform when coupled with AI-driven DevOps pipelines. Rather than relying on numerical simulation, the research advances a detailed interpretive analysis of how data flows, decision logic, and automation routines interact across system layers.

The results demonstrate that systems integrating predictive maintenance with AI-enabled DevOps exhibit superior resilience, scalability, and cyber-physical coherence compared to siloed approaches. In manufacturing environments, intelligent deployment pipelines enable rapid retraining and redeployment of diagnostic models in response to changing operational regimes, while in photovoltaic and smart grid systems, edge-based learning combined with centralized orchestration supports real-time anomaly detection and coordinated response. These outcomes reinforce the central thesis that predictive maintenance is no longer merely a data science problem but a socio-technical system that must be managed through continuous integration, continuous deployment, and continuous learning.

The discussion situates these findings within broader debates about algorithmic governance, cybersecurity, and sustainability. It critically examines the risks of over-automation, the challenges of data heterogeneity, and the ethical implications of delegating maintenance decisions to intelligent systems. By weaving together insights from predictive maintenance, IoT architectures, and AI-driven DevOps, the article offers a unified vision for the next generation of industrial intelligence. In doing so, it provides both theoretical clarity and practical guidance for researchers, engineers, and policy makers seeking to build resilient, adaptive, and trustworthy cyber-physical infrastructures.

KEYWORDS

AI-driven DevOps, predictive maintenance, Industry 4.0, Internet of Things, smart energy systems, cyber-physical systems

INTRODUCTION

The transformation of industrial and energy systems in the twenty-first century has been driven by an unprecedented convergence of digital technologies, physical infrastructures, and algorithmic intelligence. From manufacturing centers equipped with thousands of sensors to photovoltaic farms distributed across vast geographical regions, contemporary cyber-physical systems generate enormous volumes of

data that promise to revolutionize how assets are monitored, diagnosed, and maintained. Within this context, predictive maintenance has emerged as one of the most influential paradigms of Industry 4.0, promising to replace reactive and preventive strategies with data-driven foresight that anticipates failures before they occur (Li et al., 2017). Yet despite the apparent maturity of predictive analytics, many organizations struggle to translate model outputs into reliable, scalable, and trustworthy operational decisions. This

difficulty is not merely technical but reflects deeper structural challenges in how machine learning systems are developed, deployed, and governed in complex industrial environments.

Historically, maintenance has evolved through several distinct phases. In early industrial systems, maintenance was largely reactive, with interventions occurring only after breakdowns disrupted production. As manufacturing matured, preventive maintenance schedules were introduced, relying on statistical averages and fixed service intervals to reduce the probability of failure. While these approaches improved reliability, they also produced inefficiencies, as components were often replaced or serviced regardless of their actual condition (Cachada et al., 2018). The rise of sensor technologies and networked devices in the late twentieth and early twenty-first centuries enabled condition-based maintenance, in which real-time measurements of vibration, temperature, and electrical parameters informed maintenance decisions (Kaliyannan et al., 2023). Predictive maintenance represents the culmination of this trajectory, leveraging machine learning and advanced analytics to forecast future states of equipment based on historical and real-time data.

Yet predictive maintenance in isolation is insufficient for the realities of modern cyber-physical systems. As assets become more interconnected, heterogeneous, and software-driven, the maintenance of physical components becomes inseparable from the maintenance of digital models, data pipelines, and deployment infrastructures. This insight has been at the heart of recent work on AI-driven DevOps, which conceptualizes the lifecycle of machine learning

models as an ongoing process of integration, testing, deployment, and monitoring rather than a one-time development effort (Varanasi, 2025). In such a paradigm, predictive maintenance models are not static artifacts but living systems that must adapt to changing operating conditions, evolving data distributions, and emerging cyber threats.

The integration of predictive maintenance with AI-enabled DevOps is particularly critical in energy systems, where renewable generation, distributed assets, and smart grid technologies have introduced new layers of complexity. Photovoltaic plants, for example, are subject to environmental variability, component degradation, and cyber-physical vulnerabilities that require continuous monitoring and adaptive control (Hojabri et al., 2022). Traditional maintenance strategies struggle to cope with such dynamics, as they lack the agility and situational awareness needed to manage geographically dispersed and technologically diverse infrastructures. By contrast, IoT-based monitoring systems combined with machine learning-driven fault diagnosis offer the possibility of real-time insight into the health of each asset, from individual solar panels to entire substations (Patil et al., 2017).

However, even the most sophisticated predictive models can fail to deliver value if they are not properly embedded within operational workflows. Models trained on historical data may become obsolete as conditions change, leading to concept drift and degraded performance. Moreover, deploying updates to models in safety-critical environments requires rigorous testing, version control, and rollback mechanisms to prevent unintended consequences. These challenges have motivated the emergence of AI-driven DevOps as a



discipline that extends traditional software engineering practices into the realm of machine learning and cyber-physical systems (Varanasi, 2025). By automating the processes of data ingestion, model training, validation, deployment, and monitoring, AI-driven DevOps seeks to ensure that predictive maintenance systems remain accurate, robust, and aligned with operational realities.

The scholarly literature reflects a growing recognition of these interdependencies. Research on Maintenance 4.0 emphasizes the need for integrated architectures that connect sensors, analytics, and decision-support systems in a unified framework (Cachada et al., 2018). Studies of on-device and edge-based analytics highlight the importance of distributing intelligence closer to the source of data, reducing latency and enhancing resilience (Mihigo et al., 2022). At the same time, work on physics-informed machine learning and digital twins underscores the value of combining data-driven models with domain knowledge to improve interpretability and reliability (Huber et al., 2023; El Bazi et al., 2023). Yet much of this research treats predictive maintenance and system deployment as separate domains, leaving unresolved questions about how models should be managed over their lifecycle.

This article addresses this gap by developing an integrated theoretical and methodological framework that situates predictive maintenance within the broader context of AI-driven DevOps. Building on the systematic analysis of intelligent automation in software deployment and maintenance by Varanasi (2025), the study argues that the future of predictive maintenance lies not only in better algorithms but in better

organizational and technical processes for managing those algorithms in real-world environments. By synthesizing insights from manufacturing, renewable energy, and smart grid research, the paper advances a holistic view of cyber-physical intelligence that bridges the divide between data science and operations.

The problem statement guiding this inquiry is therefore twofold. First, how can predictive maintenance models be designed and implemented in ways that remain reliable and adaptive over time in complex, heterogeneous systems? Second, how can AI-driven DevOps practices be leveraged to operationalize these models at scale, ensuring continuous improvement, cybersecurity, and organizational trust? These questions are not merely academic; they speak to the sustainability and resilience of critical infrastructures in an era of digital transformation (Laayati et al., 2022a).

The literature gap addressed here lies in the lack of integrative analyses that connect predictive maintenance, IoT architectures, and DevOps practices into a coherent theoretical framework. While numerous studies examine fault detection algorithms for solar panels (Huang et al., 2021; Selvaraj et al., 2022) or energy consumption forecasting in buildings and mines (El Maghraoui et al., 2022; Maghraoui et al., 2022), few explore how these models are deployed, monitored, and evolved over time within operational pipelines. Similarly, research on cybersecurity in energy systems highlights the vulnerability of interconnected assets to cyberattacks (Ten et al., 2017; Walker et al., 2021), yet often neglects the role of AI-driven automation in detecting and responding to such threats.

By bringing these strands together, the present study aims to contribute a new perspective on industrial intelligence that recognizes the inseparability of analytics and operations. The following sections elaborate this perspective through a detailed methodological design, a comprehensive analysis of results, and an extended theoretical discussion that situates the findings within broader scholarly debates.

METHODOLOGY

The methodological approach adopted in this study is grounded in qualitative and analytical synthesis rather than numerical experimentation, reflecting the complex, multi-layered nature of AI-enabled predictive maintenance systems. Given the diversity of application domains represented in the literature, from manufacturing centers and oil refineries to photovoltaic plants and smart grids, a purely quantitative meta-analysis would risk obscuring the contextual and architectural nuances that determine system performance (Li et al., 2017; Khodabakhsh et al., 2018). Instead, the research employs a structured interpretive methodology that integrates architectural analysis, comparative literature synthesis, and theoretical modeling to construct a holistic understanding of how predictive maintenance and AI-driven DevOps interact in practice.

The first component of the methodology involves the identification and categorization of core technological layers within cyber-physical maintenance systems. Drawing on Maintenance 4.0 architectures (Cachada et al., 2018) and IoT-based monitoring frameworks (Kaliyannan et al., 2023), the study delineates four interdependent layers: data acquisition, analytics and modeling,

deployment and orchestration, and decision and action. Each layer is examined in terms of its functional role, technological dependencies, and vulnerabilities. For instance, the data acquisition layer encompasses sensors, edge devices, and communication protocols that collect raw signals from physical assets, while the analytics layer includes machine learning models, physics-informed algorithms, and anomaly detection systems that transform data into actionable insights (Mihigo et al., 2022; Huber et al., 2023).

The second methodological component consists of a comparative analysis of predictive maintenance techniques across domains. Neural networks, fuzzy systems, convolutional architectures, and hybrid models are evaluated based on their reported capabilities in fault detection, prognosis, and energy prediction (Janarthanan et al., 2021; Jlidi et al., 2023; Al-Dahidi et al., 2019). Rather than ranking these methods by numerical accuracy, the analysis focuses on their interpretability, adaptability, and suitability for integration into continuous deployment pipelines. This emphasis reflects the insight that a highly accurate model that cannot be reliably updated or monitored in production may be less valuable than a slightly less accurate model that is robustly managed through DevOps practices (Varanasi, 2025).

The third component addresses the deployment and lifecycle management of models. Here, the study draws on AI-driven DevOps frameworks that describe how machine learning models move from development to production and back again through cycles of continuous integration, testing, and monitoring (Varanasi, 2025). Architectural patterns such as microservices, containerization, and automated retraining pipelines are analyzed in

relation to their capacity to support predictive maintenance in distributed environments like solar farms and smart grids (Sajun et al., 2022; Laayati et al., 2022b). The methodological goal is to identify how these patterns mediate the relationship between analytics and operations, enabling or constraining system adaptability.

A fourth methodological strand involves the incorporation of cybersecurity and resilience considerations. Given the documented vulnerability of interconnected energy systems to cyberattacks (Ten et al., 2017; Walker et al., 2021), the analysis examines how AI-driven DevOps can support real-time threat detection, model integrity verification, and automated response. This includes an interpretive assessment of how anomaly detection models at the edge can be coordinated with centralized security operations through automated deployment pipelines (Sajun et al., 2022).

Throughout the methodology, the study maintains a critical stance toward technological determinism. While AI and automation offer powerful tools, their effectiveness depends on organizational practices, data governance, and human oversight. Accordingly, the analysis incorporates insights from digital twin frameworks and multi-agent systems that emphasize the socio-technical nature of predictive maintenance (El Bazi et al., 2023; Laayati et al., 2022a). These perspectives inform the evaluation of methodological limitations, including the risks of data bias, model drift, and over-reliance on automated decisions.

The rationale for this multi-layered methodology lies in the complexity of the research problem. Predictive maintenance in AI-driven environments is not a single technology but an ecosystem of

interacting components, each of which can amplify or undermine the others. By synthesizing architectural, algorithmic, and operational perspectives, the methodology aims to capture this complexity in a way that supports deep theoretical insight and practical relevance.

Limitations of the approach must also be acknowledged. Because the study relies on secondary sources and theoretical synthesis, it cannot provide direct empirical validation through experiments or field trials. However, this limitation is mitigated by the breadth and depth of the literature considered, which spans multiple industries and methodological traditions (Li et al., 2017; Hojabri et al., 2022; El Maghraoui et al., 2022). Moreover, the interpretive nature of the analysis allows for the identification of cross-cutting patterns and conceptual tensions that might be obscured in more narrowly focused empirical studies.

RESULTS

The results of the integrative analysis reveal a set of consistent patterns across manufacturing, energy, and smart grid domains that underscore the centrality of AI-driven DevOps in realizing the full potential of predictive maintenance. One of the most salient findings is that predictive accuracy alone is an insufficient metric of system effectiveness. Studies of neural network-based fault diagnosis for photovoltaic panels, for example, report high levels of classification performance under controlled conditions (Huang et al., 2021; Selvaraj et al., 2022), yet these models often struggle to maintain their reliability when deployed in dynamic, real-world environments characterized by changing weather, component

aging, and operational interventions (Hojabri et al., 2022). This degradation is frequently attributed to concept drift, a phenomenon in which the statistical properties of input data evolve over time, rendering previously learned patterns obsolete.

AI-driven DevOps addresses this challenge by embedding predictive models within continuous learning pipelines that monitor performance and trigger retraining when necessary (Varanasi, 2025). The literature indicates that such pipelines are particularly valuable in distributed energy systems, where localized conditions can vary significantly across sites (Sajun et al., 2022). Edge-based analytics models, such as TinyLSTM architectures, can perform real-time anomaly detection at the device level, while centralized orchestration platforms manage model updates and coordinate responses across the network (Mihigo et al., 2022). The result is a hybrid intelligence architecture that combines local autonomy with global oversight.

Another key result concerns the role of physics-informed and hybrid models in predictive maintenance. While purely data-driven neural networks can capture complex nonlinear relationships, they may lack interpretability and physical grounding (Huber et al., 2023). Physics-informed models, by contrast, integrate domain knowledge about system dynamics, enabling more robust predictions under conditions of sparse or noisy data. The analysis shows that when such models are integrated into AI-driven DevOps pipelines, their advantages are amplified. Automated testing frameworks can validate model outputs against physical constraints, while deployment pipelines ensure that updates are

rolled out consistently across digital twins and real assets (El Bazi et al., 2023).

The results also highlight the importance of multi-agent and layered architectures in managing complexity. Smart energy management systems that employ multi-agent frameworks can distribute decision-making across components, allowing local agents to respond to immediate anomalies while higher-level agents optimize system-wide objectives (Laayati et al., 2022a). When coupled with DevOps practices, these architectures support rapid experimentation and adaptation. New control strategies or diagnostic models can be deployed to subsets of the system, evaluated in situ, and scaled up if successful, mirroring the principles of continuous delivery in software engineering (Varanasi, 2025).

Cybersecurity emerges as another domain where the integration of predictive maintenance and AI-driven DevOps yields significant benefits. Anomaly detection models trained on sensor and network data can identify deviations indicative of cyberattacks or system malfunctions (Ten et al., 2017). However, the effectiveness of these models depends on their timely updating and coordinated deployment across the infrastructure. AI-driven DevOps pipelines enable automated distribution of security patches and model updates, reducing the window of vulnerability and enhancing system resilience (Walker et al., 2021).

Across all these domains, the analysis reveals a common pattern: systems that treat predictive maintenance as a static analytics problem tend to suffer from brittleness, while those that embed it within adaptive DevOps frameworks demonstrate greater robustness and scalability. This finding reinforces the central argument advanced by



Varanasi (2025) that intelligent automation in deployment and maintenance is as critical as the underlying machine learning algorithms.

DISCUSSION

The integration of predictive maintenance and AI-driven DevOps represents a profound shift in how industrial and energy systems are conceptualized, managed, and governed. At a theoretical level, this integration challenges the traditional separation between analytics and operations, suggesting instead that intelligent systems must be understood as continuous processes rather than discrete tools (Varanasi, 2025). This perspective aligns with broader trends in cyber-physical systems research, which emphasize the co-evolution of software, hardware, and organizational practices (Cachada et al., 2018).

One of the central theoretical implications of the findings is the redefinition of reliability. In classical engineering, reliability is often associated with the probability that a component will perform its intended function over a specified period. In AI-enabled systems, however, reliability also encompasses the stability and trustworthiness of predictive models, data pipelines, and deployment infrastructures (Huber et al., 2023). A neural network that achieves high accuracy in laboratory tests may be unreliable in production if it cannot adapt to changing conditions or if its deployment is poorly managed. AI-driven DevOps addresses this by introducing mechanisms for continuous validation, monitoring, and improvement, thereby extending the concept of reliability to include the entire lifecycle of intelligent components.

The scholarly debate around interpretability further illuminates this shift. Critics of black-box

models argue that their opacity undermines trust and accountability, particularly in safety-critical domains like energy and manufacturing (Appiah et al., 2019). Physics-informed and hybrid models offer one response to this concern, but their integration into operational systems remains challenging (Huber et al., 2023). AI-driven DevOps provides a complementary solution by enabling systematic testing and auditing of models in deployment. By tracking performance metrics, logging decisions, and supporting rollback mechanisms, DevOps practices create a form of procedural transparency that can partially compensate for model opacity (Varanasi, 2025).

Another key debate concerns the balance between centralization and decentralization. Edge-based analytics promise low latency and resilience, but they can lead to fragmentation and inconsistent behavior if not properly coordinated (Mihigo et al., 2022). Centralized cloud platforms, on the other hand, offer powerful computational resources but may introduce delays and single points of failure (Sajun et al., 2022). The results suggest that hybrid architectures, supported by AI-driven DevOps, offer a viable compromise. Local models can operate autonomously, while centralized pipelines manage updates, security, and global optimization. This architecture reflects a broader trend toward federated intelligence in cyber-physical systems (Laayati et al., 2022a).

From a socio-technical perspective, the adoption of AI-driven DevOps also reshapes organizational roles and responsibilities. Maintenance engineers, data scientists, and software developers must collaborate more closely, sharing responsibility for the performance and safety of intelligent systems (Kaliyannan et al., 2023). This convergence raises



important questions about skills, governance, and accountability. Who is responsible when an automated maintenance decision leads to a costly or dangerous outcome? How should organizations balance human oversight with algorithmic efficiency? These questions echo broader concerns about automation and agency in digital societies (Walker et al., 2021).

The limitations of the current state of the art must also be acknowledged. Data quality remains a persistent challenge, particularly in legacy systems where sensors may be unreliable or poorly calibrated (Patil et al., 2017). Model drift and bias can lead to inequitable or unsafe outcomes if not properly managed (El Maghraoui et al., 2022). Moreover, the increasing reliance on networked infrastructures exposes systems to cyber risks that can undermine both predictive maintenance and DevOps pipelines (Ten et al., 2017). Addressing these issues will require not only technical innovation but also robust governance frameworks that integrate cybersecurity, data ethics, and regulatory compliance.

Future research should therefore pursue several complementary directions. First, there is a need for longitudinal studies that examine how predictive maintenance models and DevOps pipelines co-evolve over time in real-world deployments. Such studies could provide empirical evidence of the benefits and pitfalls identified in this theoretical analysis (Varanasi, 2025). Second, greater attention should be paid to the human factors of AI-driven maintenance, including training, trust, and organizational culture (Laayati et al., 2022b). Third, interdisciplinary collaboration between engineers, computer scientists, and social scientists will be essential to address the ethical

and governance challenges posed by increasingly autonomous systems.

CONCLUSION

This article has advanced an integrated framework for understanding and implementing predictive maintenance in the era of AI-driven DevOps. By synthesizing insights from manufacturing, renewable energy, and smart grid research, and grounding the analysis in contemporary work on intelligent automation (Varanasi, 2025), the study has demonstrated that predictive maintenance is most effective when embedded within adaptive, continuously learning operational pipelines. The convergence of IoT-based monitoring, machine learning-based diagnostics, and DevOps practices offers a powerful pathway toward more resilient, efficient, and sustainable cyber-physical systems.

At the same time, this convergence introduces new challenges related to complexity, cybersecurity, and governance. Addressing these challenges will require ongoing research, interdisciplinary collaboration, and thoughtful policy development. As industrial and energy systems continue to evolve, the integration of predictive maintenance and AI-driven DevOps will play a central role in shaping the future of intelligent infrastructure.

REFERENCES

1. Kaliyannan, G. V., Sri Anbupalani, M., Kandasamy, S., Sivaraj, S., and Gunasekaran, R. (2023). Role of IoT in Industry Predictive Maintenance. *Integration of Mechanical and Manufacturing Engineering with IoT*, 191–213.
2. El Bazi, N., Mabrouki, M., Laayati, O., Ouhabi, N., El Hadraoui, H., Hammouch, F. E., and Chebak, A. (2023). *Generic Multi-Layered Digital-Twin*

- Framework-Enabled Asset Lifecycle Management for the Sustainable Mining Industry. *Sustainability*, 15, 3470.
3. Varanasi, S. R. (2025, August). AI-Driven DevOps in Modern Software Engineering—A Review of Machine Learning Based Intelligent Automation for Deployment and Maintenance. In 2025 IEEE 2nd International Conference on Information Technology, Electronics and Intelligent Communication Systems (ICITEICS), 1–7. IEEE.
 4. Huang, Z., Duan, S., Long, F., Li, Y., Zhu, J., and Ling, Q. (2021). Fault Diagnosis for Solar Panels using Convolutional Neural Network. In Proceedings of the 33rd Chinese Control and Decision Conference, 153–158.
 5. Laayati, O., El Hadraoui, H., Bouzi, M., and Chebak, A. (2022). Smart Energy Management System: Oil Immersed Power Transformer Failure Prediction and Classification Techniques Based on DGA Data. In Proceedings of the 2nd International Conference on Innovative Research in Applied Science, Engineering and Technology, 1–6.
 6. Ten, C. W., Yamashita, K., Yang, Z., Vasilakos, A. V., and Ginter, A. (2017). Impact assessment of hypothesized cyberattacks on interconnected bulk power systems. *IEEE Transactions on Smart Grid*, 9, 4405–4425.
 7. Li, Z., Wang, Y., and Wang, K. S. (2017). Intelligent predictive maintenance for fault diagnosis and prognosis in machine centers: Industry 4.0 scenario. *Advances in Manufacturing*, 5, 377–387.
 8. Mihigo, I. N., Zennaro, M., Uwitonze, A., Rwigema, J., and Rovai, M. (2022). On-Device IoT-Based Predictive Maintenance Analytics Model: Comparing TinyLSTM and TinyModel from Edge Impulse. *Sensors*, 22, 5174.
 9. Cachada, A., Barbosa, J., Leitno, P., Geraldts, C. A., Deusdado, L., Costa, J., and Romero, L. (2018). Maintenance 4.0: Intelligent and predictive maintenance system architecture. In Proceedings of the IEEE International Conference on Emerging Technologies and Factory Automation, 139–146.
 10. Hojabri, M., Kellerhals, S., Upadhyay, G., and Bowler, B. (2022). IoT-based PV array fault detection and classification using embedded supervised learning methods. *Energies*, 15, 2097.
 11. Sajun, A. R., Shapsough, S., Zualkernan, I., and Dhaouadi, R. (2022). Edge-based individualized anomaly detection in large-scale distributed solar farms. *ICT Express*, 8, 174–178.
 12. Huber, L. G., Palme, T., and Chao, M. A. (2023). Physics-informed machine learning for predictive maintenance: Applied use cases. In Proceedings of the IEEE Swiss Conference on Data Science, 66–72.
 13. El Maghraoui, A., Ledmaoui, Y., Laayati, O., El Hadraoui, H., and Chebak, A. (2022). Smart Energy Management: A Comparative Study of Energy Consumption Forecasting Algorithms for an Experimental Open-Pit Mine. *Energies*, 15, 4569.
 14. Janarthanan, R., Maheshwari, R., Shukla, P., Shukla, P., Mirjalili, S., and Kumar, M. (2021). Intelligent Detection of the PV Faults Based on Artificial Neural Network and Type 2 Fuzzy Systems. *Energies*, 14, 6584.
 15. Patil, S. M., Vijayalashmi, M., and Tapaskar, R. (2017). IoT Based Solar Energy Monitoring System. In Proceedings of the International

- Conference on Energy, Communication, Data Analytics and Soft Computing, 1574–1579.
- 16.** Maghraoui, A. E., Hammouch, F. E., Ledmaoui, Y., and Chebak, A. (2022). Smart Energy Management System: A Comparative Study of Energy Consumption Prediction Algorithms for a Hotel Building. In Proceedings of the Global Power, Energy and Communication Conference, 529–534.
- 17.** Walker, A., Desai, J., Saleem, D., and Gunda, T. (2021). Cybersecurity in Photovoltaic Plant Operations. National Renewable Energy Laboratory Technical Report.
- 18.** Selvaraj, T., Rengaraj, R., Venkatakrisnan, G., Soundararajan, S., Natarajan, K., Balachandran, P., David, P., and Shitharth, S. (2022). Environmental Fault Diagnosis of Solar Panels Using Solar Thermal Images in Multiple Convolutional Neural Networks. Wiley.
- 19.** Appiah, A. Y., Zhang, X., Ayawli, B. B. K., and Kyeremeh, F. (2019). Review and Performance Evaluation of Photovoltaic Array Fault Detection and Diagnosis Techniques. International Journal of Photoenergy, 2019, e6953530.

