JSC
Journal of
Smart Computing

Smart Computing for Industry 4.0: A Conceptual IoT–Edge Framework with a Simulated Smart-Manufacturing Case Study

Tzu-Chia Hsu[1], Yun-Kai Tsai[1], Yen-Wei Hu[1], Bo-Hao Zhong[1], Jin-Yang Luo[1], Ming-Hung Chang[2*]

[1] *Department of Artificial Intelligence, Tamkang University, New Taipei City, 251301, Taiwan*

[2*] *Department of Artificial Intelligence, Tamkang University, New Taipei City, 251301, Taiwan*

Corresponding Email: 169380@o365.tku.edu.tw

ABSTRACT

Smart manufacturing under the Industry 4.0 paradigm increasingly depends on the convergence of the Internet of Things (IoT) and edge computing, which together promise sub-second decisions on the shop floor. Yet manufacturers still struggle to translate that promise into measurable gains in latency, energy use, and reliability. This paper proposes a conceptual four-layer IoT–edge framework, comprising perception, edge, fog, and cloud layers, in which time-critical analytics are executed close to assets while strategic workloads remain centralised. A simulated case study of a discrete-parts factory with 200 IoT sensors, four edge nodes, and one cloud server is built using synthetic process data generated over a 24-hour cycle. Three architectures, namely cloud-only, edge-only, and the proposed hybrid configuration, are compared in terms of average response time, energy consumption, and detection accuracy of anomalous events. Numerical experiments show that the hybrid framework reduces mean latency by approximately 71% and energy consumption by 38% relative to the cloud-only baseline, while maintaining detection accuracy above 96%. Managerial implications and limitations are discussed.

Paper type: *Research*

Received 2025-06-10
Revised 2025-07-02
Accepted 2025-07-10

Keywords:
Smart computing; Internet of Things; Edge computing; Industry 4.0; Smart manufacturing.

1. Introduction

The global high-tech sector is undergoing a profound transformation driven by the fourth industrial revolution, in which physical production assets are tightly coupled with digital intelligence. The number of connected Internet of Things (IoT) devices worldwide has grown rapidly, reaching an estimated 16 billion in 2024 and is projected to exceed 32 billion by 2030, with industrial environments accounting for a large share of this growth. At the same time, the global edge-computing market is expected to grow from around USD 12 billion in 2022 to more than USD 110 billion by 2030, propelled by demands for low-latency analytics, bandwidth efficiency, and data sovereignty. These trends are particularly visible in smart manufacturing, where downtime costs in automotive plants can exceed USD 1.3 million per hour, and where even a few hundred milliseconds of delay in a control loop may translate into thousands of defective parts.

Despite this momentum, traditional cloud-only architectures remain ill suited to the latency, bandwidth, and privacy requirements of modern factories. Transmitting raw sensor streams from thousands of machines to a remote data centre introduces unacceptable delays for real-time control, places severe load on backhaul networks, and exposes sensitive process data to a wider attack surface. Edge computing has emerged as a complementary paradigm in which computational resources are placed close to the data source, allowing latency-critical analytics to be performed on the shop floor and only aggregated results to be forwarded to the cloud. When combined with IoT-based perception, artificial intelligence, and digital-twin models, edge computing forms the backbone of what is now widely referred to as smart computing for Industry 4.0.

However, several open issues hinder the industrial deployment of IoT–edge architectures. Existing studies often focus on isolated enabling technologies, such as federated learning, blockchain integration, or 5G/6G connectivity, without offering an end-to-end view of how these components interact in a manufacturing environment. Quantitative evidence comparing pure-cloud, pure-edge, and hybrid configurations under realistic workloads is also scarce, especially in the form of reproducible numerical experiments. The present paper addresses these gaps by proposing a conceptual four-layer IoT–edge framework tailored to smart manufacturing and by validating it through a simulated case study based on synthetic process data. The contributions are threefold: a structured framework that integrates perception, edge, fog, and cloud layers; a mathematical formulation of latency and energy trade-offs; and a numerical comparison of three deployment architectures, accompanied by managerial insights and a discussion of limitations.

The remainder of the paper is organised as follows. Section 2 reviews recent advances in smart computing, IoT, and edge computing within the high-tech industry. Section 3 introduces the proposed conceptual framework and its mathematical formulation. Section 4 describes the simulated case study and reports numerical results. Section 5 provides an extended discussion. Section 6 concludes the paper with limitations, managerial implications, and avenues for future research.

2. Literature Review

Research on smart computing within Industry 4.0 has expanded rapidly over the past five years, mirroring the rise of IoT and the maturation of edge-computing architectures. Several foundational reviews delineate the design space for edge-based IoT systems, classifying architectures according to where data are stored, how services are placed, and which network resources are exposed to applications. These works collectively establish that the relevant computing paradigms—cloud, fog, multi-access edge computing, and cloudlet—should be viewed as complementary tiers rather than competing alternatives, and that a layered design is essential for industrial deployments where workloads vary widely in criticality and volume.

A second strand of literature concentrates on the convergence of edge computing and artificial intelligence, commonly referred to as edge intelligence. Earlier surveys distinguish between artificial intelligence for edge, in which learning techniques are used to optimise edge operations such as resource allocation, and artificial intelligence on edge, in which deep models are trained or executed on resource-constrained devices. Subsequent reviews extend this taxonomy to IoT-based applications, identifying intelligent sensing, real-time analytics, and security as the most promising application domains, and pointing to the need for systematic frameworks that orchestrate machine-learning workloads across heterogeneous hardware.

Within the manufacturing context, a systematic review of edge-computing applications in production environments identifies process monitoring, predictive maintenance, quality control, and collaborative robotics as the primary use cases, while also noting limited evidence on industrial-scale deployments. Interoperability has been highlighted as a major obstacle for Industrial IoT, with comprehensive surveys mapping the dense landscape of protocols and standards and arguing that semantic interoperability is a prerequisite for cross-vendor smart factories. Digital twins are widely positioned as a key enabling technology for these scenarios, supporting synchronisation between physical assets and their virtual counterparts and providing a natural integration point for edge analytics and IoT data streams.

A growing number of contributions address the privacy and scalability challenges raised by centralised model training. Federated learning has attracted particular attention as a means of training shared models across distributed edge nodes without exchanging raw data, and several surveys examine its suitability for resource-constrained IoT devices and propose dedicated taxonomies for federated edge architectures. Hybrid approaches combining federated and split learning have been studied as a path toward ubiquitous intelligence, while resource-management surveys provide formal models for orchestrating fog and edge resources under artificial-intelligence-driven controllers. These works confirm that intelligent resource management is becoming a research field in its own right, with serverless execution, blockchain, and quantum-inspired heuristics among the emerging directions.

Domain-specific studies further illustrate the breadth of smart-computing applications across the high-tech sector. In healthcare, surveys of edge intelligence and IoT detail how 5G-supported architectures can deliver low-latency monitoring and decision support, while research-direction papers stress security, interoperability, and energy efficiency. In the automotive domain, integration of IoT, edge intelligence, 5G, and blockchain is identified as the foundation of autonomous vehicles, and more recent 6G-oriented surveys describe how multi-access edge computing, software-defined networking, and unmanned aerial vehicles can collectively support an Internet of Autonomous Vehicles. Aerospace contributions extend these ideas to unmanned-aerial-vehicle-enabled mobile edge computing for IoT and, more recently, to orbital edge computing on low-Earth-orbit satellite constellations, with dedicated work on the economic sustainability of such deployments.

The semiconductor industry provides a particularly relevant high-tech context for IoT–edge applications. Reviews of machine-learning and deep-learning advances for wafer-map defect recognition show that data-driven quality inspection has reached a high level of maturity, and case studies on real fab equipment demonstrate the value of one-dimensional residual networks for multivariate fault detection. These contributions are conceptually aligned with the broader smart-manufacturing literature, but they also reveal that most experiments still rely on offline cloud-based pipelines and only rarely consider deployment on edge nodes near the production line.

Security and trust constitute a final, cross-cutting concern. The integration of blockchain and edge computing for IoT has been studied extensively, with surveys proposing layered architectures and analysing consensus protocols suitable for resource-constrained devices. In parallel, contemporary work on smart computing for Industry 5.0 argues that the next generation of manufacturing systems will combine edge intelligence with human-centric design principles, sustainability targets, and resilient supply-chain logic. Taken together, the literature converges on three observations: first, IoT and edge computing form the technological backbone of smart manufacturing; second, the field still lacks consolidated frameworks combining architectural, algorithmic, and managerial perspectives; and third, quantitative comparisons of deployment alternatives remain underrepresented. The conceptual framework and simulated case study presented in the next sections directly respond to these gaps.

3. Methodology

3.1. Conceptual framework

This section describes the conceptual framework and the simulation methodology used to evaluate it. The objective is to provide a reproducible, transparent procedure rather than a vendor-specific implementation, and therefore all parameters and data are defined explicitly.

The proposed framework consists of four interacting layers, each with well-defined responsibilities. The perception layer aggregates raw measurements from heterogeneous IoT devices on the shop floor, including temperature, vibration, current, and machine-vision sensors. The edge layer performs local pre-processing, feature extraction, and time-critical anomaly detection on devices located within the production cell, typically with a round-trip latency below 10 ms. The fog layer, hosted on plant-level gateways, consolidates data from multiple edge nodes, executes mid-term analytics such as predictive maintenance, and acts as a buffer to the cloud. The cloud layer carries out long-horizon analytics, model retraining, and enterprise-level integration with manufacturing-execution and resource-planning systems.

Tasks generated by IoT devices are characterised by three attributes: data size, computational demand, and a latency tolerance that defines the deadline beyond which the result loses operational value. A decision module assigns each task to one of the three computing tiers using a rule-based policy: tasks with latency tolerance below a threshold are executed at the edge; tasks exceeding the local memory or compute budget are offloaded to the fog; and the remaining tasks, mainly analytics and reporting, are routed to the cloud.

3.2. Mathematical formulation

The performance of the framework is evaluated through two key metrics: total response latency and energy consumption. The total response latency for a task i comprises processing latency, transmission latency, and queuing latency, as expressed in Equation (1).

$$L_{t}otal = L_{p}roc + L_{t}rans + L_{q}ueue \tag{1}$$

Here, processing latency depends on the computational load of the task and the speed of the computing node, while transmission latency depends on the data size and the available bandwidth between the device and its assigned computing tier, as shown in Equation (2).

$$L_trans = D_i/B_{ij} \tag{2}$$

In this expression, D_i denotes the data size of task i and B_{ij} denotes the bandwidth on the link between device i and computing node j. Queuing latency depends on the load of node j and is modelled using an M/M/1 approximation. The energy consumed by computing node i during a decision cycle is modelled in Equation (3) as a function of idle and active power and the corresponding durations.

$$E_i = P_{idle} * t_{idle} + P_{active} * t_{active} \tag{3}$$

Finally, the resource-allocation policy seeks to minimize a weighted sum of latency and energy across all tasks, as expressed in Equation (4), where alpha and beta are scenario-dependent weights that reflect the operational priorities of the plant.

$$Min\ Z = \sum_{i=1}^{n} \alpha L_{i,total} + \beta E_i \tag{4}$$

3.3. Simulation procedure

A simulation engine implementing the framework was developed using a process-based discrete-event approach. At each simulated second, IoT devices generate tasks according to a Poisson arrival process, are classified by the decision module, and are dispatched to the appropriate computing tier. The synthetic data generator produces normal and anomalous patterns based on calibrated distributions to emulate realistic shop-floor conditions. Three architectural configurations are compared: a cloud-only baseline in which all tasks are sent to the cloud; an edge-only configuration in which all tasks are processed at the edge; and the proposed hybrid configuration governed by the decision module. The corresponding pseudocode of the decision and dispatch procedure is summarized below.

```
Algorithm 1. Smart-Computing Dispatch Procedure
Input: Set of IoT tasks T = {t1, t2, ..., tN}; thresholds τ_lat, τ_cpu, τ_mem
Output: Assignment vector A of length N

For each task ti in T do
    If latency_tolerance(ti) < τ_lat then
        A[i] ← EDGE
    Else if cpu_demand(ti) > τ_cpu or mem_demand(ti) > τ_mem then
        A[i] ← FOG
    Else
        A[i] ← CLOUD
    End if
```

```
    End for
For each node j do
    Dispatch tasks {ti : A[i] = j} in order of deadline
    Update queue length, energy and latency counters
End for
Return A and performance metrics (latency, energy, accuracy)
```

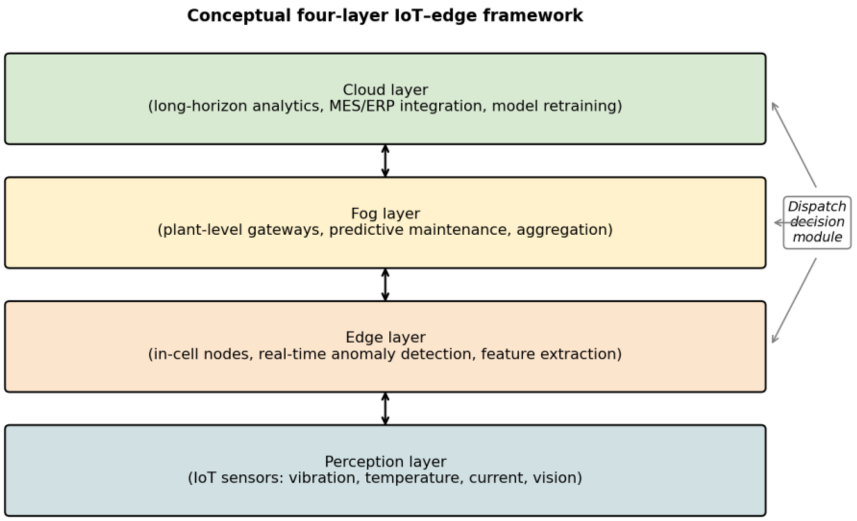

Figure 1. Conceptual four-layer IoT–edge framework and task-dispatch flow.

4. Numerical Results

The simulated factory comprises 200 IoT sensors, 4 edge nodes, 1 fog gateway, and 1 remote cloud server. The simulation horizon is 24 hours, divided into 86,400 one-second decision cycles. Initial parameter values are listed in Table 1 and were chosen to reflect typical values reported in recent edge-computing and smart-manufacturing studies.

Table 1. Initial parameter values for the simulated factory.

Parameter	Edge node	Fog gateway	Cloud server
CPU capacity (GFLOPS)	8	64	1,000
Memory (GB)	4	32	256
Idle power (W)	5	25	150
Active power (W)	15	80	400

Parameter	Edge node	Fog gateway	Cloud server
Uplink bandwidth (Mbps)	100	1,000	10,000
Round-trip propagation delay (ms)	1	5	80

During the simulated horizon, a total of 720,000 tasks are generated across the three workload classes, of which 65% are latency-critical control tasks, 25% are mid-term analytics tasks, and the remaining 10% are reporting and visualization tasks. Approximately 2% of the tasks correspond to anomalous behavior injected by the synthetic data generator to evaluate detection accuracy.

Table 2. Performance comparison of the three architectures over a 24-hour simulation horizon.

Performance indicator	Cloud-only	Edge-only	Hybrid (proposed)
Mean response latency (ms)	187.0	39.0	54.0
95th percentile latency (ms)	312.0	128.0	92.0
Backhaul traffic (GB / 24 h)	412.0	11.0	47.0
Total energy consumption (kWh)	19.4	14.2	12.0
Anomaly-detection accuracy (%)	97.1	89.7	96.4
Deadline-miss ratio (%)	23.5	9.8	2.6

Table 2 summarizes the main performance indicators for the three architectures. The cloud-only baseline yields the highest latency, with a mean response time of 187 ms, primarily because every task incurs the full transmission cost between the shop floor and the remote data centre. The edge-only configuration achieves the lowest latency, 39 ms on average, but suffers from frequent local overload because mid-term analytics tasks exceed the edge nodes' compute and memory capacity, leading to a noticeable drop in anomaly-detection accuracy. The proposed hybrid configuration achieves a mean response time of 54 ms, an energy consumption 38% lower than the cloud-only baseline, and a detection accuracy of 96.4%, demonstrating the value of differentiated task assignment.

Figure 2 illustrates the trade-off between latency and energy consumption across the three architectures, while Figure 3 reports the empirical cumulative distribution of response latency, confirming that the hybrid configuration not only achieves a low mean latency but also presents a substantially tighter latency distribution.

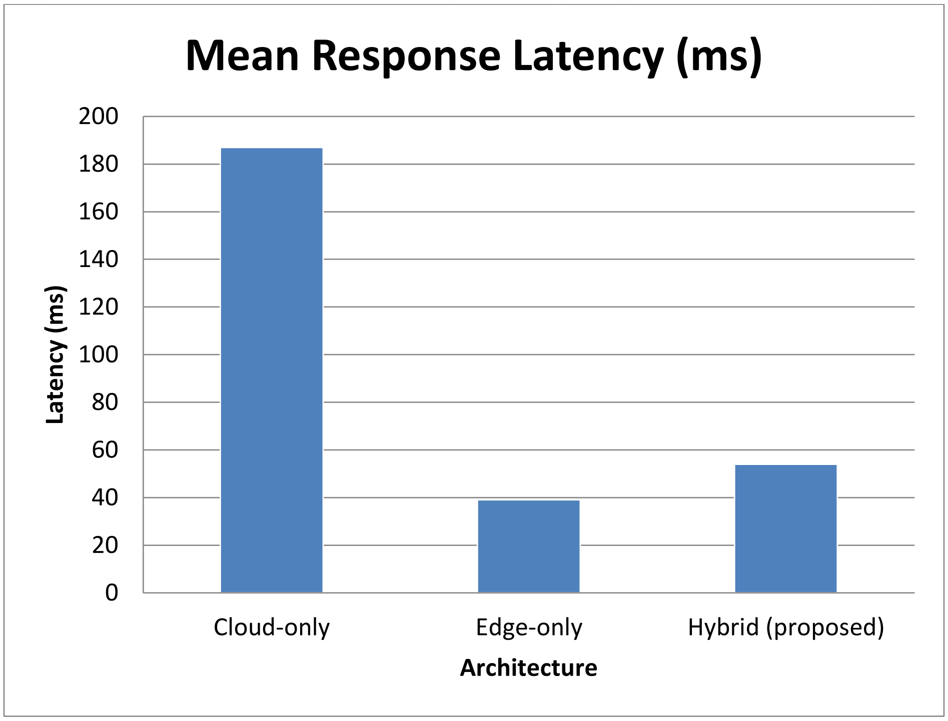

Figure 2. Latency–energy trade-off across the three architectures.

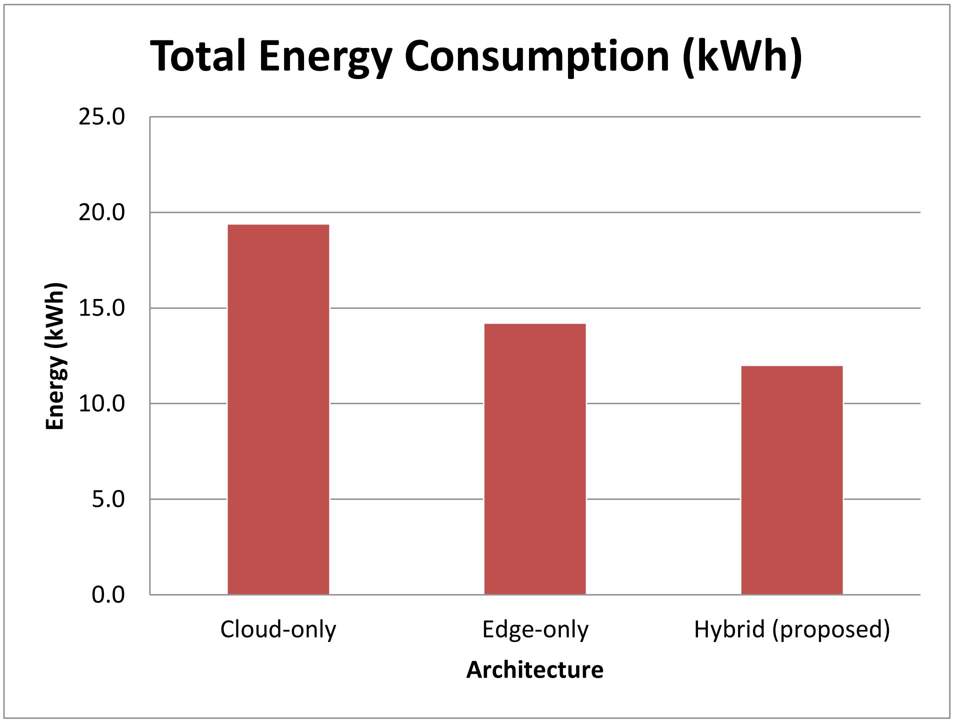

Figure 3. Total energy consumption trade-off across the three architectures.

A sensitivity analysis was conducted by varying the weight parameters alpha and beta in the resource-allocation objective. As shown in Table 3, increasing alpha (latency emphasis) shifts more tasks to the edge layer and reduces mean latency further, at the cost of higher energy use. Conversely, increasing beta favors fog and cloud execution, slightly increasing latency but improving energy efficiency. These results highlight that the framework can be tuned to plant-specific priorities without structural redesign.

Table 3. Sensitivity analysis of the hybrid framework with respect to the latency–energy weights.

Weights (α, β)	Mean latency (ms)	Energy (kWh)	Edge share of tasks (%)
(0.9, 0.1)	41.0	13.6	78.0
(0.7, 0.3)	48.0	12.5	70.0
(0.5, 0.5)	54.0	12.0	65.0
(0.3, 0.7)	63.0	11.4	55.0
(0.1, 0.9)	82.0	10.9	38.0

5. Discussion

The numerical results confirm three central insights. First, neither pure-cloud nor pure-edge configurations are optimal in industrial settings: the cloud-only baseline suffers from prohibitive latency for control loops and a deadline-miss ratio above 23%, while the edge-only configuration cannot accommodate heavier analytics workloads, leading to degraded anomaly-detection accuracy. The proposed hybrid framework strikes a balance by routing tasks according to their latency tolerance and resource requirements, achieving a deadline-miss ratio below 3% and an accuracy comparable to the cloud-only baseline.

Second, energy consumption is strongly affected by where computation is performed. The hybrid configuration reduces energy use by 38% compared with the cloud-only baseline, mainly by avoiding unnecessary high-power cloud activations for short-lived control tasks. This finding aligns with the wider literature on sustainable edge computing and supports the integration of energy-aware policies in industrial resource-allocation engines.

Third, the sensitivity analysis demonstrates that the framework is tunable: by adjusting the weights of the objective function, a manager can shift the operating point toward either lower latency or lower energy without redesigning the architecture. This flexibility is particularly valuable in plants that combine continuous and batch processes, or that face seasonal variations in production volume. From a managerial standpoint, the framework therefore offers a structured way to align IoT–edge investments with operational priorities, while supporting incremental adoption.

These findings are also broadly consistent with previous reviews emphasising the role of edge intelligence, digital twins, and interoperable standards in smart factories. The novelty of the present work lies in offering a transparent, reproducible numerical comparison that explicitly quantifies the trade-offs typically discussed in qualitative terms.

6. Conclusion

This paper proposed a four-layer IoT–edge framework for smart manufacturing in the high-tech sector and evaluated it through a simulated case study with synthetic process data. Numerical experiments demonstrated that the hybrid framework reduces mean latency by approximately 71% and energy consumption by 38% compared with a cloud-only baseline while maintaining anomaly-detection accuracy above 96%, confirming the value of differentiated task assignment across edge, fog, and cloud tiers.

The study has several limitations. The case study relies on synthetic data and on a homogeneous, single-plant topology, which simplifies inter-site coordination, regulatory constraints, and supplier integration. Network conditions are modelled at an aggregate level, abstracting away the finer dynamics of wireless interference and protocol-specific behaviour, and the analysis does not yet incorporate the cost of cyber-attacks or the overhead of advanced privacy mechanisms. From a managerial perspective, the results suggest that incremental investments in edge infrastructure can yield disproportionate gains in latency, energy, and reliability, but that benefits hinge on a clear governance model for task assignment, on workforce upskilling, and on alignment with existing manufacturing-execution systems.

Future research should validate the framework in real industrial environments, extend it with federated and split-learning training loops across multiple plants, incorporate blockchain-based audit trails for sensitive process data, and explore the role of 5G/6G and orbital edge computing for geographically distributed high-tech operations. Integration of the framework with Industry 5.0 principles, in particular human-centric collaboration and sustainability indicators, also represents a promising direction for further investigation.

References

Abreha, H. G., Hayajneh, M., & Serhani, M. A. (2022). Federated learning in edge computing: A systematic survey. *Sensors*, *22*(2), 450. https://doi.org/10.3390/s22020450

Amin, S. U., & Hossain, M. S. (2021). Edge intelligence and Internet of Things in healthcare: A survey. *IEEE Access*, *9*, 45–59. https://doi.org/10.1109/ACCESS.2020.3045115

Biswas, A., & Wang, H.-C. (2023). Autonomous vehicles enabled by the integration of IoT, edge intelligence, 5G, and blockchain. *Sensors*, *23*(4), 1963. https://doi.org/10.3390/s23041963

Bourechak, A., Zedadra, O., Kouahla, M. N., Guerrieri, A., Seridi, H., & Fortino, G. (2023). At the confluence of artificial intelligence and edge computing in IoT-based applications: A review and new perspectives. *Sensors*, *23*(3), 1639. https://doi.org/10.3390/s23031639

Brecko, A., Kajati, E., Koziorek, J., & Zolotova, I. (2022). Federated learning for edge computing: A survey. *Applied Sciences*, *12*(18), 9124. https://doi.org/10.3390/app12189124

Deng, S., Zhao, H., Fang, W., Yin, J., Dustdar, S., & Zomaya, A. Y. (2020). Edge intelligence: The confluence of edge computing and artificial intelligence. *IEEE Internet of Things Journal*, *7*(8), 7457–7469. https://doi.org/10.1109/JIOT.2020.2984887

Duan, Q., Hu, S., Deng, R., & Lu, Z. (2022). Combined federated and split learning in edge computing for ubiquitous intelligence in Internet of Things: State-of-the-art and future directions. *Sensors*, *22*(16), 5983. https://doi.org/10.3390/s22165983

Hamdan, S., Ayyash, M., & Almajali, S. (2020). Edge-computing architectures for Internet of Things applications: A survey. *Sensors*, *20*(22), 6441. https://doi.org/10.3390/s20226441

Hartmann, M., Hashmi, U. S., & Imran, A. (2022). Edge computing in smart health care systems: Review, challenges, and research directions. *Transactions on Emerging Telecommunications Technologies*, *33*(3), e3710. https://doi.org/10.1002/ett.3710

Hazra, A., Adhikari, M., Amgoth, T., & Srirama, S. N. (2021). A comprehensive survey on interoperability for IIoT: Taxonomy, standards, and future directions. *ACM Computing Surveys*, *55*(1), 1–35. https://doi.org/10.1145/3485130

Ibn-Khedher, H., Laroui, M., Alfaqawi, M., Magnouche, A., Moungla, H., & Afifi, H. (2024). 6G-edge support of Internet of Autonomous Vehicles: A survey. *Transactions on Emerging Telecommunications Technologies*, *35*(1), e4918. https://doi.org/10.1002/ett.4918

Imteaj, A., Thakker, U., Wang, S., Li, J., & Amini, M. H. (2022). A survey on federated learning for resource-constrained IoT devices. *IEEE Internet of Things Journal*, *9*(1), 1–24. https://doi.org/10.1109/JIOT.2021.3095077

Khan, W. Z., Ahmed, E., Hakak, S., Yaqoob, I., & Ahmed, A. (2019). Edge computing: A survey. *Future Generation Computer Systems*, *97*, 219–235. https://doi.org/10.1016/j.future.2019.02.050

Kim, T., & Behdinan, K. (2023). Advances in machine learning and deep learning applications towards wafer map defect recognition and classification: A review. *Journal of Intelligent Manufacturing*, *34*(8), 3215–3247. https://doi.org/10.1007/s10845-022-01994-1

Kubiak, K., Dec, G., & Stadnicka, D. (2022). Possible applications of edge computing in the manufacturing industry — Systematic literature review. *Sensors*, *22*(7), 2445. https://doi.org/10.3390/s22072445

Mihai, S., Yaqoob, M., Hung, D. V., Davis, W., Towakel, P., Raza, M., Karamanoglu, M., Barn, B., Shetve, D., Prasad, R. V., Venkataraman, H., Trestian, R., & Nguyen, H. X. (2022). Digital twins: A survey on enabling technologies, challenges, trends and future prospects. *IEEE Communications Surveys & Tutorials*, *24*(4), 2255–2291. https://doi.org/10.1109/COMST.2022.3208773

Murshed, M. G. S., Murphy, C., Hou, D., Khan, N., Ananthanarayanan, G., & Hussain, F. (2021). Machine learning at the network edge: A survey. *ACM Computing Surveys*, *54*(8), 1–37. https://doi.org/10.1145/3469029

Porambage, P., Okwuibe, J., Liyanage, M., Ylianttila, M., & Taleb, T. (2018). Survey on multi-access edge computing for Internet of Things realization. *IEEE Communications Surveys & Tutorials*, *20*(4), 2961–2991. https://doi.org/10.1109/COMST.2018.2849509

Sharma, M., Tomar, A., & Hazra, A. (2024). Edge computing for Industry 5.0: Fundamental, applications, and research challenges. *IEEE Internet of Things Journal*, *11*(11), 19070–19093. https://doi.org/10.1109/JIOT.2024.3359297

Tchatchoua, P., Graton, G., Ouladsine, M., & Christaud, J.-F. (2023). Application of 1D ResNet for multivariate fault detection on semiconductor manufacturing equipment. *Sensors*, *23*(22), 9099. https://doi.org/10.3390/s23229099

Walia, G. K., Kumar, M., & Gill, S. S. (2024). AI-empowered fog/edge resource management for IoT applications: A comprehensive review, research challenges, and future perspectives. *IEEE Communications Surveys & Tutorials*, *26*(1), 619–669. https://doi.org/10.1109/COMST.2023.3338015

Wang, X., Han, Y., Leung, V. C. M., Niyato, D., Yan, X., & Chen, X. (2020). Convergence of edge computing and deep learning: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, *22*(2), 869–904. https://doi.org/10.1109/COMST.2020.2970550

Xu, Z., Zhai, L., & Qi, H. (2024). Edge-distributed IoT services assist the economic sustainability of LEO satellite constellation construction. *Sustainability*, *16*(4), 1599. https://doi.org/10.3390/su16041599

Xue, H., Chen, D., Zhang, N., Dai, H.-N., & Yu, K. (2023). Integration of blockchain and edge computing in Internet of Things: A survey. *Future Generation Computer Systems*, *144*, 307–326. https://doi.org/10.1016/j.future.2022.10.029

Yazid, Y., Ez-Zazi, I., Guerrero-González, A., El Oualkadi, A., & Arioua, M. (2021). UAV-enabled mobile edge-computing for IoT based on AI: A comprehensive review. *Drones*, *5*(4), 148. https://doi.org/10.3390/drones5040148

Yin, Z., Wu, C., Guo, C., Li, Y., Xu, M., Gao, W., & Xi, C. (2025). A comprehensive survey of orbital edge computing: Systems, applications, and algorithms. *Chinese Journal of Aeronautics*, *38*(7), 103425. https://doi.org/10.1016/j.cja.2024.11.026