Edge Intelligence, IoT and Cyber-Physical Systems for Next-Generation Automation

Adrian Mercer, Leila Rahman, Daniel Okafor

Article type: Editorial

Volume 1, Issue 1 (2026)

Read Current Issue →

Abstract

Intelligent automation increasingly depends on computation occurring close to sensors, machines, and physical processes rather than exclusively in centralized cloud infrastructures. Industrial Internet of Things deployments, embedded artificial intelligence, autonomous devices, digital twins, and cyber-physical systems generate continuous data streams whose operational value often depends on low latency, reliable connectivity, privacy preservation, energy efficiency, and immediate local response. This editorial argues that edge intelligence should therefore be understood as a fundamental architectural component of next-generation automation rather than merely a reduced form of cloud computing. The emerging landscape combines distributed inference, embedded machine learning, federated learning, real-time networking, digital twins, predictive maintenance, and edge-cloud coordination, but substantial challenges remain in constrained computation, heterogeneous hardware, interoperability, cybersecurity, synchronization, model lifecycle management, fault tolerance, and distributed-system validation. Future progress requires energy-efficient models, hardware-software co-design, resilient communication and control, secure edge architectures, standardized benchmarking, trustworthy federated learning, and validation under realistic hardware and network conditions. Dependable automation will ultimately require coordinated design across sensing, communication, computation, intelligence, and physical control.

Keywords: edge intelligence; industrial Internet of Things; cyber-physical systems; embedded artificial intelligence; smart manufacturing; edge computing; distributed automation

Citation

Mercer, A., Rahman, L., & Okafor, D. (2026). Edge Intelligence, IoT and Cyber-Physical Systems for Next-Generation Automation. Journal of AI, Automation & Intelligent Systems, 1(1). https://eldranacademic.com/jaais/articles/3/

Received	Revised	Accepted	Published
02 February 2026	21 April 2026	15 May 2026	04 June 2026

Copyright

© 2026 Adrian Mercer, Leila Rahman, Daniel Okafor. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY 4.0) ↗. Use, distribution or reproduction in other forums is permitted, provided the original authors and copyright owners are credited and the original publication in this journal is cited in accordance with accepted academic practice.

Disclaimer

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, the publisher, the editors or the reviewers. Any product evaluated in this article, or claim made by its manufacturer, is not guaranteed or endorsed by the publisher.

1.0 Introduction

1.1 Context and Problem Framing

The architecture of intelligent automation is shifting from a predominantly centralized computing model toward distributed intelligence embedded throughout physical environments. Networked sensors, industrial Internet of Things (IIoT) devices, intelligent controllers, robotic platforms, smart infrastructure, autonomous machines, and embedded artificial intelligence now generate and consume data at points increasingly distant from conventional data centres. The Internet of Things created the underlying model of pervasive sensing and connectivity, linking smart objects through interoperable communication and computing infrastructures (Al-Fuqaha et al., 2015). Industrial adoption subsequently extended these ideas into production environments in which connected sensing, control, communication, and enterprise systems interact continuously (Xu et al., 2014). Edge computing brought storage and computation closer to data sources, principally to reduce response time, bandwidth consumption, and dependence on remote infrastructure (Shi et al., 2016). Satyanarayanan (2017) similarly framed edge computing as an architectural response to physical limits associated with relying exclusively on distant cloud resources. In manufacturing, such distributed computation increasingly forms part of the cyber-physical structure connecting factory equipment with computational models and intelligent decision processes (Lee et al., 2015).

Centralized cloud infrastructure remains essential for large-scale storage, model development, fleet-level analytics, historical analysis, and computationally intensive workloads. It is nevertheless insufficient as the sole computational foundation for many automation applications. Transmitting continuous high-volume sensor streams to remote data centres can introduce latency, increase bandwidth requirements, expose sensitive operational information, and make critical system behaviour dependent on wide-area connectivity. Mobile-edge-computing research therefore emphasizes allocating computation according to latency, communication, mobility, and energy constraints rather than assuming remote execution is always optimal (Mach & Bečvář, 2017). Industrial IoT systems intensify these requirements because real-time performance, interoperability, reliability, energy efficiency, security, privacy, and coexistence among heterogeneous networks must often be achieved simultaneously (Sisinni et al., 2018). A cloud connection that is acceptable for production reporting may consequently be unsuitable for robotic coordination, machine-vision inspection, protective monitoring, or a tightly timed control loop.

1.2 Why the Issue Matters

The central scholarly question is therefore not whether computation should occur in the cloud or at the edge, but how intelligence should be distributed across devices, gateways, local servers, communication networks, and centralized platforms. Fog architectures similarly conceptualize computation, networking, storage, and control as a continuum extending from cloud infrastructure toward physical devices (Chiang & Zhang, 2016). Edge intelligence adds machine learning and adaptive inference to this continuum, creating new decisions about where models execute, what information is communicated, how distributed nodes coordinate, and how systems behave when networks or hardware degrade. The argument developed here is that edge intelligence should not be treated as "smaller cloud AI." It is an architectural approach to intelligent automation in which sensing, communication, computation, learning, and control must be jointly designed around the temporal, energetic, security, and physical requirements of the process being automated.

2.0 Current Landscape

2.1 Edge AI Architectures

Contemporary edge intelligence spans a wide computational continuum, from microcontroller-scale inference to industrial gateways and powerful on-premises servers coordinated with cloud platforms. Zhou et al. (2019) conceptualized edge intelligence as the convergence of artificial intelligence and edge computing, encompassing distributed inference, learning, resource allocation, and model deployment. Deep learning at the edge similarly involves decisions about partitioning inference and training among end devices, edge nodes, and cloud resources while balancing latency, bandwidth, privacy, and computational capacity (Chen & Ran, 2019). This distributed architecture is increasingly important because automation applications rarely have uniform latency, accuracy, energy, and reliability requirements.

Embedded machine learning represents the most resource-constrained end of this continuum. Deep neural networks require substantial computation and data movement, making energy-efficient processing a central systems problem rather than an algorithmic afterthought (Sze et al., 2017). Hardware accelerators such as Eyeriss demonstrate how memory organization and dataflow can significantly reduce the energy cost of neural-network inference (Chen et al., 2017). Hardware-software co-design has pushed these capabilities further. MCUNet jointly optimizes neural architecture and inference execution for microcontrollers with severe memory constraints (Lin et al., 2020), while once-for-all networks allow models to be specialized for different hardware and latency budgets without complete retraining for every target platform (Cai et al., 2020). TinyML research consequently illustrates how model architecture, compilers, memory management, accelerators, and device characteristics must be considered together (Shafique et al., 2021).

2.2 Intelligent IoT

Distributed learning extends the edge from inference toward decentralized model development. Federated learning allows multiple devices or organizations to contribute to training while keeping raw datasets decentralized, making it attractive for privacy-sensitive and geographically distributed automation. Yet federated deployment creates communication, resource-allocation, device-heterogeneity, statistical, privacy, and security challenges (Lim et al., 2020). Kairouz et al. (2021) accordingly position federated learning as a systems problem involving communication efficiency, robustness, heterogeneity, privacy, fairness, and operational management rather than simply a decentralized implementation of conventional training.

2.3 Cyber-Physical Integration

Industrial implementations increasingly connect edge intelligence with digital twins, predictive maintenance, inspection, and adaptive production. Digital twins link physical assets with continuously updated computational representations (Fuller et al., 2020), while industrial digital-twin research emphasizes integration between cyber and physical spaces across design, production, prognostics, and lifecycle management (Tao et al., 2019). Lu et al. (2020) further identify industrial communication, data processing, modelling, and architecture as essential components of digital-twin-driven smart manufacturing. Predictive-maintenance research similarly demonstrates how machine learning can convert distributed sensor data into equipment-health and failure information (Carvalho et al., 2019), while recent smart-manufacturing reviews position digital twins as platforms for monitoring, optimization, quality improvement, and predictive maintenance (Soori et al., 2023).

2.4 Distributed Automation

Communication technologies are simultaneously evolving to support distributed automation. Industrial networking increasingly seeks convergence between deterministic communication, interoperable Ethernet, wireless connectivity, and cyber-physical control (Wollschlaeger et al., 2017). Private 5G provides one pathway for supporting industrial mobility, low-latency communication, and local network control where wireless operation is required (Aijaz, 2020), while time-sensitive networking aims to provide predictable Ethernet communication for critical industrial traffic (Zhang et al., 2025). Edge intelligence therefore represents a convergence of machine learning, embedded computation, industrial networking, and control rather than a standalone computing technology.

3.0 Critical Challenges

3.1 Computational and Energy Constraints

Resource constraints remain a fundamental obstacle. Edge hardware ranges from microcontrollers with severe memory and power limitations to GPUs and specialized accelerators installed in industrial gateways. A model that executes efficiently on one platform may be infeasible on another, while computational demand varies with sensing rates, model complexity, thermal conditions, and application workload. Hardware specialization can improve energy efficiency but also increase dependence on particular toolchains, accelerators, instruction sets, or memory architectures. The resulting heterogeneity complicates portability, reproducibility, benchmarking, model updating, and long-term maintenance. Energy must likewise be treated as a system-level constraint: local processing may reduce communication demand while increasing device power consumption, thermal load, or battery depletion.

Interoperability and communication reliability constitute a second challenge. IIoT environments combine legacy controllers, proprietary fieldbuses, industrial Ethernet, wireless links, robotics platforms, gateways, databases, and cloud interfaces developed across different generations and vendors. Boyes et al. (2018) demonstrate that IIoT architectures are heterogeneous not merely in communication technology but in device function, location, capability, and security characteristics. Time-sensitive networking can improve deterministic Ethernet communication, yet scheduling, synchronization, configuration, wireless integration, and multi-vendor interoperability remain substantial engineering issues (Zhang et al., 2025). Distributed intelligence cannot therefore be assumed reliable merely because individual computing nodes operate correctly; system performance also depends on communication timing, packet loss, synchronization, topology, and recovery behaviour.

3.2 Latency and Reliability

Cybersecurity becomes more difficult as computational capability moves closer to physical assets. Edge deployments enlarge the attack surface by introducing gateways, wireless interfaces, remote update mechanisms, embedded operating systems, local storage, APIs, accelerators, and potentially physically accessible devices. Cyber-physical attacks can affect both information and physical processes, making conventional information-security reasoning insufficient when compromised computation can influence actuation or control (Humayed et al., 2017). Edge and fog architectures additionally introduce threats involving malicious nodes, compromised virtualization, insecure interfaces, distributed denial of service, and weak trust boundaries (Roman et al., 2018). Security architectures must therefore protect data, models, communications, devices, and control processes without imposing computational or latency overhead that defeats the purpose of edge deployment.

Federated and distributed learning does not eliminate security and privacy risks. Decentralized data can reduce some centralized exposure, but model updates can leak information or be manipulated through poisoning and other adversarial behaviours. Mothukuri et al. (2021) show that federated learning introduces distinctive

confidentiality, integrity, inference, and poisoning threats that interact with its distributed architecture. Device heterogeneity further complicates participation, aggregation, update frequency, and robustness. A trustworthy federated automation system must therefore determine not only how models are trained collaboratively but also how devices are authenticated, updates are validated, anomalous participants are detected, and degraded or compromised clients are isolated.

3.3 Cybersecurity and Privacy

Lifecycle management creates another unresolved problem. Edge models operate on devices that may remain deployed for years while sensors, environments, firmware, data distributions, hardware, and operational requirements change. Updating models across large fleets can introduce version inconsistency, rollback problems, synchronization errors, and unexpected interactions with existing control logic. Connectivity interruptions may leave different nodes running different software or models. Distributed digital twins create related synchronization questions because useful virtual representations depend on timely and credible correspondence with evolving physical assets.

3.4 Interoperability and Reproducibility

Validation is consequently more difficult than evaluating a centralized model in isolation. Edge automation must be tested under real computational loads, network contention, variable latency, packet loss, device failure, clock drift, sensor degradation, power limitations, cybersecurity events, and partial connectivity. Benchmark results obtained on powerful development hardware may reveal little about behaviour after quantization or deployment on constrained devices. MLPerf Tiny responds to part of this problem by evaluating accuracy, latency, and energy for ultra-low-power inference (Banbury et al., 2021), but automation requires broader benchmarks that include networking, control consequences, resilience, and lifecycle behaviour. Distributed-system validation must therefore evaluate the interaction among machine intelligence, hardware, communication, and physical processes.

4.0 Discussion and Priorities

4.1 Integrative Discussion

Edge intelligence changes intelligent automation because it alters where decisions are made and how rapidly computational outputs can influence physical systems. Sensing, communication, computation, learning, and control cannot be optimized independently. A local inference model may reduce latency but consume additional energy; offloading computation may conserve device resources but increase dependence on communication; cloud coordination may improve fleet-level optimization while creating bandwidth and privacy costs. The relevant engineering objective is therefore not to maximize local computation but to distribute intelligence according to the operational requirements of each function.

This perspective also clarifies why edge AI should not be understood as a miniature version of cloud AI. Cloud systems are typically designed around abundant centralized resources, comparatively stable infrastructure, and large-scale aggregation. Edge systems operate across heterogeneous hardware, intermittent networks, real-time processes, environmental exposure, and constrained power envelopes. Deep-learning efficiency at the edge therefore depends on co-design across models, memory, accelerators, scheduling, and communication (Sze et al., 2017). Similarly, industrial edge architectures must coexist with deterministic control requirements and legacy equipment that cannot be replaced simply because a new inference model becomes available.

Local autonomy and centralized coordination should instead be treated as complementary architectural capabilities. Safety-critical or time-sensitive decisions may need to remain local, whereas long-horizon optimization, model training, fleet analytics, and cross-site learning may be better suited to regional or cloud infrastructure. Hierarchical systems can distribute functions accordingly, but additional layers also create synchronization, security, configuration, and accountability problems. Resilience therefore depends on graceful degradation: when a network fails, an edge node should know which functions can continue locally, which should enter a safe state, and which require restored coordination.

The same interdependence applies to privacy and security. Processing data locally can reduce unnecessary transmission, but physical accessibility and large device populations create new attack surfaces. Federated learning can limit raw-data centralization while simultaneously exposing model-update channels to manipulation. Edge intelligence thus redistributes risk rather than eliminating it. Successful architectures must explicitly balance latency, bandwidth, privacy, cybersecurity, energy, scalability, and safety according to the physical consequences of each automated function.

4.2 Research and Practice Priorities

A first priority is energy-efficient intelligence designed jointly with hardware. Future work should report not only accuracy but memory footprint, energy per inference, latency, thermal behaviour, and computational utilization on the actual target platform. Hardware-software co-design should become routine for embedded automation, building on evidence that neural architectures and execution engines can be optimized jointly for constrained devices (Lin et al., 2020). Progress should be demonstrated through reproducible measurements on representative hardware rather than estimated complexity alone.

A second priority is interoperable edge architecture. Industrial systems require protocols, information models, interfaces, and orchestration mechanisms capable of operating across heterogeneous vendors and generations of equipment. Research should examine the integration of industrial Ethernet, TSN, wireless systems, gateways, digital twins, and legacy control networks under realistic timing constraints. Standardization should emphasize interoperability and predictable behaviour rather than merely connectivity.

Third, secure and resilient edge architectures must treat cybersecurity as part of control-system engineering. Authentication, secure boot, protected updates, trusted execution, network segmentation, anomaly detection, model integrity, and recovery mechanisms should be evaluated under resource and timing constraints. Edge-based anomaly detection is particularly important because local systems may identify abnormal machine or network behaviour even when upstream connectivity is unavailable. Security mechanisms should nevertheless be tested for their effects on latency and deterministic operation.

Fourth, distributed-control resilience requires explicit investigation of communication failure, device loss, clock desynchronization, degraded sensing, and partial network partition. Systems should specify which decisions remain local and what operational state is entered when coordination fails. Validation should include fault injection and network impairment rather than assuming ideal connectivity.

Fifth, model lifecycle management at the edge requires better mechanisms for deployment, versioning, drift detection, update verification, rollback, and revalidation. Edge models should be treated as maintained operational components rather than static files flashed once onto hardware. Digital twins may provide useful infrastructure for tracking the relationship among physical assets, operational data, and model versions, but lifecycle evidence should show that synchronization remains reliable over time.

Sixth, trustworthy federated and distributed learning should be evaluated against realistic heterogeneity and attack conditions. Communication efficiency alone is insufficient. Research should demonstrate robustness to non-independent data, intermittent participation, compromised clients, privacy attacks, hardware variation, and

delayed updates. Federated systems intended for industrial deployment need explicit security and governance assumptions.

Finally, standardized end-to-end benchmarking should connect AI metrics with systems and control metrics. Benchmarks such as MLPerf Tiny provide an important foundation for comparable measurements of constrained inference (Banbury et al., 2021), but next-generation automation requires evaluation that combines accuracy, latency, energy, communication reliability, fault recovery, security overhead, and physical-process consequences. Testing should occur on real devices and networks wherever feasible. Simulation remains useful, but deployment claims should ultimately be supported by evidence obtained under realistic hardware, workload, connectivity, and environmental constraints.

4.3 Implications for the Journal's Scope

For the Journal of AI, Automation & Intelligent Systems (JAAIS), edge intelligence represents a natural intersection of several research communities that have too often been evaluated separately. Progress in next-generation automation increasingly depends on connecting machine learning with embedded systems, industrial IoT, communications, robotics, intelligent control, cybersecurity, and cyber-physical engineering. Research contributions should therefore move beyond demonstrating that an algorithm can execute at the edge and establish why its placement, timing, resource consumption, resilience, and interaction with physical processes improve the complete automation system.

JAAIS should particularly encourage interdisciplinary studies of embedded and edge AI, intelligent sensors, distributed robotic systems, edge-cloud control architectures, digital twins, federated industrial learning, low-power inference, industrial networking, predictive maintenance, cyber-physical security, and resilient autonomous infrastructure. Strong contributions should report realistic hardware and communication conditions, describe system architecture clearly, make benchmarking reproducible, and analyse failure behaviour rather than only nominal performance. Work that integrates algorithms with communications and control is especially important because edge intelligence derives its operational significance from precisely these interactions.

The journal can consequently provide a forum for research in which intelligence is evaluated as part of an engineered physical system. Such work can help bridge the persistent gap between advances in machine-learning capability and the dependable distributed infrastructure required to use those capabilities in real-time automation.

4.4 Conclusion

The next generation of intelligent automation will not be built exclusively in remote data centres. Increasingly capable sensors, embedded processors, accelerators, industrial networks, digital twins, and distributed learning systems are relocating computational intelligence toward the machines and physical processes where decisions must take effect. This shift can reduce latency, conserve bandwidth, improve local responsiveness, and support greater data autonomy, but it simultaneously exposes severe constraints in energy, hardware heterogeneity, interoperability, cybersecurity, synchronization, maintenance, and system validation.

The appropriate response is not an unconditional preference for edge over cloud computation. Intelligent automation requires architectures that assign computation to the locations where it can satisfy operational requirements most effectively. Some decisions require immediate local execution; others benefit from edge aggregation, plant-level coordination, or centralized cloud analytics. Dependability emerges from coordinating these layers rather than treating any one of them as universally sufficient.

Edge intelligence should therefore mature from a collection of deployment techniques into a systems discipline integrating machine learning, embedded computation, communication, security, and control. The decisive question for future automation will not simply be whether intelligence can be moved closer to the physical world,

but whether distributed intelligence can operate predictably, securely, efficiently, and resiliently within it.

5.0 Declarations

5.1 Author Contributions

Prof. Adrian Mercer, Dr. Leila Rahman, and Prof. Daniel Okafor contributed to the conceptual development, literature synthesis, critical interpretation, drafting, and revision of the editorial.

5.2 Funding

No external funding was received for the preparation of this editorial.

5.3 Competing Interests

The authors declare that they have no conflicts of interest or competing interests relevant to this editorial.

5.4 Data Availability

No original datasets were generated or analysed for this editorial.

5.5 AI Use Disclosure

Generative AI was used to support literature organization and language drafting during preparation of this editorial. The authors remain responsible for independently verifying the accuracy of the arguments, citations, references, and final submitted text in accordance with the publisher's policy.

6.0 References

Aijaz, A. (2020). Private 5G: The future of industrial wireless. IEEE Industrial Electronics Magazine, 14(4), 136–145. https://doi.org/10.1109/MIE.2020.3004975

Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of Things: A survey on enabling technologies, protocols, and applications. IEEE Communications Surveys & Tutorials, 17(4), 2347–2376. https://doi.org/10.1109/COMST.2015.2444095

Banbury, C., Reddi, V. J., Torelli, P., Holleman, J., Jeffries, N., Kiraly, C., Montino, P., Kanter, D., Ahmed, S., Pau, D., Thakker, U., Torrini, A., Warden, P., Cordaro, J., Di Guglielmo, G., Duarte, J., Gibellini, S., Parekh, V., Tran, H., . . . Xuesong, X. (2021). MLPerf Tiny benchmark. Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks, 1.

Boyes, H., Hallaq, B., Cunningham, J., & Watson, T. (2018). The industrial internet of things (IIoT): An analysis framework. Computers in Industry, 101, 1–12. https://doi.org/10.1016/j.compind.2018.04.015

Cai, H., Gan, C., Wang, T., Zhang, Z., & Han, S. (2020). Once for all: Train one network and specialize it for efficient deployment. International Conference on Learning Representations.

Carvalho, T. P., Soares, F. A. A. M. N., Vita, R., Francisco, R. P., Basto, J. P., & Alcalá, S. G. S. (2019). A systematic literature review of machine learning methods applied to predictive maintenance. Computers & Industrial Engineering, 137, 106024. https://doi.org/10.1016/j.cie.2019.106024

Chen, J., & Ran, X. (2019). Deep learning with edge computing: A review. Proceedings of the IEEE, 107(8), 1655–1674. https://doi.org/10.1109/JPROC.2019.2921977

Chen, Y.-H., Krishna, T., Emer, J. S., & Sze, V. (2017). Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks. IEEE Journal of Solid-State Circuits, 52(1), 127–138. https://doi.org/10.1109/JSSC.2016.2616357

Chiang, M., & Zhang, T. (2016). Fog and IoT: An overview of research opportunities. IEEE Internet of Things Journal, 3(6), 854–864. https://doi.org/10.1109/JIOT.2016.2584538

Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. IEEE Access, 8, 108952–108971. https://doi.org/10.1109/ACCESS.2020.2998358

Humayed, A., Lin, J., Li, F., & Luo, B. (2017). Cyber-physical systems security—A survey. IEEE Internet of Things Journal, 4(6), 1802–1831. https://doi.org/10.1109/JIOT.2017.2703172

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., . . . Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210. https://doi.org/10.1561/2200000083

Lee, J., Bagheri, B., & Kao, H.-A. (2015). A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. Manufacturing Letters, 3, 18–23. https://doi.org/10.1016/j.mfglet.2014.12.001

Lim, W. Y. B., Luong, N. C., Hoang, D. T., Jiao, Y., Liang, Y.-C., Yang, Q., Niyato, D., & Miao, C. (2020). Federated learning in mobile edge networks: A comprehensive survey. IEEE Communications Surveys & Tutorials, 22(3), 2031–2063. https://doi.org/10.1109/COMST.2020.2986024

Lin, J., Chen, W.-M., Lin, Y., Cohn, J., Gan, C., & Han, S. (2020). MCUNet: Tiny deep learning on IoT devices. Advances in Neural Information Processing Systems, 33, 11711–11722.

Lu, Y., Liu, C., Wang, K. I.-K., Huang, H., & Xu, X. (2020). Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. Robotics and Computer-Integrated Manufacturing, 61, 101837. https://doi.org/10.1016/j.rcim.2019.101837

Mach, P., & Bečvář, Z. (2017). Mobile edge computing: A survey on architecture and computation offloading. IEEE Communications Surveys & Tutorials, 19(3), 1628–1656. https://doi.org/10.1109/COMST.2017.2682318

Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619–640. https://doi.org/10.1016/j.future.2020.10.007

Roman, R., Lopez, J., & Mambo, M. (2018). Mobile edge computing, fog et al.: A survey and analysis of security threats and challenges. Future Generation Computer Systems, 78, 680–698. https://doi.org/10.1016/j.future.2016.11.009

Satyanarayanan, M. (2017). The emergence of edge computing. Computer, 50(1), 30–39. https://doi.org/10.1109/MC.2017.9

Shafique, M., Theocharides, T., Reddi, V. J., & Murmann, B. (2021). TinyML: Current progress, research challenges, and future roadmap. 2021 58th ACM/IEEE Design Automation Conference (DAC), 1303–1306. https://doi.org/10.1109/DAC18074.2021.9586232

Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637–646. https://doi.org/10.1109/JIOT.2016.2579198

Sisinni, E., Saifullah, A., Han, S., Jennehag, U., & Gidlund, M. (2018). Industrial Internet of Things: Challenges, opportunities, and directions. IEEE Transactions on Industrial Informatics, 14(11), 4724–4734. https://doi.org/10.1109/TII.2018.2852491

Soori, M., Arezoo, B., & Dastres, R. (2023). Digital twin for smart manufacturing, a review. Sustainable Manufacturing and Service Economics, 2, 100017. https://doi.org/10.1016/j.smse.2023.100017

Sze, V., Chen, Y.-H., Yang, T.-J., & Emer, J. S. (2017). Efficient processing of deep neural networks: A tutorial and survey. Proceedings of the IEEE, 105(12), 2295–2329. https://doi.org/10.1109/JPROC.2017.2761740

Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405–2415. https://doi.org/10.1109/TII.2018.2873186

Wollschlaeger, M., Sauter, T., & Jasperneite, J. (2017). The future of industrial communication: Automation networks in the era of the Internet of Things and Industry 4.0. IEEE Industrial Electronics Magazine, 11(1), 17–27. https://doi.org/10.1109/MIE.2017.2649104

Xu, L. D., He, W., & Li, S. (2014). Internet of Things in industries: A survey. IEEE Transactions on Industrial Informatics, 10(4), 2233–2243. https://doi.org/10.1109/TII.2014.2300753

Zhang, T., Wang, G., Xue, C., Wang, J., Nixon, M., & Han, S. (2025). Time-sensitive networking (TSN) for industrial automation: Current advances and future directions. ACM Computing Surveys, 57(2), Article 30, 1–38. https://doi.org/10.1145/3695248

Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., & Zhang, J. (2019). Edge intelligence: Paving the last mile of artificial intelligence with edge computing. Proceedings of the IEEE, 107(8), 1738–1762. https://doi.org/10.1109/JPROC.2019.2918951