ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

Research Article

# A Digital Twin and Edge-AI Framework for Carbon Emission Accounting and Resource Optimization in Supply Chain Clusters

**Virgil Pepys**[1,*]
University of California, Los Angeles, USA
* Corresponding author: Virgil Pepys
VirgilPepys@outlook.com

***Abstract:*** Quantifying carbon footprints while dynamically optimizing resource allocation within highly integrated supply chain clusters remains a complex computational task, largely due to data silo vulnerabilities and telemetry granular irregularities. Traditional centralized tracking methodologies frequently present severe limitations when adapting to real-time industrial fluctuations, exposing a significant operational gap in carbon accounting across multi-enterprise ecosystems. To address these systemic bottlenecks, this paper introduces a coupled digital twin and edge-directed artificial intelligence framework tailored for cluster-scale optimization. During our empirical deployment phase, unanticipated calibration misalignments between physical sensors and virtual representations initially caused substantial predictive skewing, necessitating an immediate structural correction through noise-robust pre-training and statistical multi-response regression techniques. Analytical simulations suggest that distributing accounting protocols to localized edge nodes may, to some extent, eliminate core telemetry latencies and uncover hidden processing inefficiencies, though underlying reporting discrepancies across heterogeneous hardware architectures could partially distort absolute sustainability metrics. Considering the above factors, this leads us to further thinking regarding how decentralized carbon ledger validation might ultimately reshape industrial ecological compliance, though further empirical research is explicitly needed to evaluate framework resilience under severe multi-modal data disruptions.

***Keywords:*** Digital twin; Edge artificial intelligence; Carbon accounting; Supply chain clusters; Resource optimization;

## 1. Introduction

The imperative for ecological transformation within contemporary industrial networks has forced a radical re-evaluation of how carbon emissions are tracked, verified, and mitigated across multi-enterprise ecosystems. In high-density supply chain clusters, where the production pipelines of disparate chemical, manufacturing, and raw material entities are intrinsically interwoven, traditional boundaries of corporate environmental responsibility become fundamentally blurred. These clusters produce massive amounts of greenhouse gases while requiring highly coordinated resource distribution networks to sustain continuous manufacturing operations. Because the carbon footprint of an individual product is heavily contingent upon the immediate, fluctuating operational efficiency of upstream suppliers, localized resource anomalies can ripple through the entire cluster, causing significant emission spikes. Consequently, modern industrial governance requires a holistic approach that can simultaneously manage real-time operational optimization and precise carbon accounting.

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

However, the practical realization of such integrated governance is severely obstructed by deep-seated technological and structural bottlenecks within existing industrial computing infrastructures. Multi-enterprise supply chain clusters are characterized by severe data silo vulnerabilities, where individual firms hesitate to share granular operational telemetry due to competitive sensitivities and intellectual property concerns. Even when data-sharing protocols are established, the resulting datasets frequently exhibit severe granular irregularities, characterized by mismatched sampling rates from disparate Internet-of-Things devices and legacy enterprise resource planning platforms. This lack of data harmonization prevents centralized management platforms from obtaining a cohesive, high-fidelity understanding of the cluster's collective thermodynamic and operational state. Without a unified, real-time data baseline, carbon emission accounting models remain inherently retrospective, relying on historic averages rather than capturing transient, high-emission operational anomalies.

Traditional centralized tracking methodologies, which depend heavily on periodic reporting and cloud-based data pooling, demonstrate manifest limitations when adapting to the rapid, non-linear fluctuations of real-time industrial environments. When continuous streams of multi-modal environmental telemetry from thousands of localized factory sensors are transmitted to a distant, centralized cloud infrastructure, the resulting bandwidth strain and processing latencies create severe systemic friction. By the time a centralized cloud algorithm identifies a resource misallocation or an emission anomaly, the physical operational state on the factory floor has inevitably shifted, rendering the calculated corrective interventions obsolete. Furthermore, centralized systems are highly vulnerable to localized network disruptions, where a single communication failure can completely sever the feedback loop between the monitoring platform and physical industrial assets. This operational gap underscores the necessity of pioneering alternative computational paradigms that can distribute processing power directly to the perimeter of industrial operations.

In response to these systemic limitations, the convergence of digital twin technologies and edge-directed artificial intelligence offers a possible socio-technical pathway toward reconciling real-time resource optimization with continuous carbon verification. A digital twin framework allows for the creation of a dynamic, virtual mirroring space that replicates the intricate physical, operational, and thermodynamic behaviors of the entire supply chain cluster in real time. By embedding edge-directed artificial intelligence directly within localized computing nodes on the factory floor, heavy computational and inference tasks can be executed immediately at the point of data generation. This decentralized arrangement could allow individual enterprise nodes to process highly sensitive sensor data locally, thereby preserving proprietary boundaries while exposing only aggregated, anonymized carbon metrics to the wider cluster digital twin. Such a configuration holds the potential to dramatically compress localized data latencies, enabling split-second resource adjustments that simultaneously maximize manufacturing throughput and minimize environmental degradation.

This study seeks to establish and critically evaluate a unified digital twin and edge-AI framework specifically engineered for cluster-scale resource and emission co-optimization. By exploring the complex interplay between decentralized computational intelligence and virtual systemic replication, we aim to uncover the underlying mechanisms that govern real-time ecological compliance in fragmented industrial ecosystems. Our structural inquiry moves beyond traditional, idealized simulations by explicitly incorporating the stochastic communication failures, hardware heterogeneities, and sensor anomalies that define actual industrial operations. Considering the profound structural complexities inherent in multi-enterprise coordination, this research does not pretend to offer an absolute panacea for industrial decarbonization; rather, it aims to clarify the structural trade-offs and computational boundaries that define the limits of digital environmental oversight. The subsequent sections will detail our non-linear methodological developments, analyze the multi-perspective empirical results, and deliberate on the theoretical transformations required for future industrial ecological compliance.

## 2. Literature Review

A rigorous critique of current literature concerning green supply chains reveals a historical reliance on macro-level, static conceptualizations of carbon accounting that struggle to keep pace with the digitalization of heavy industry. For many years, academic inquiry focused heavily on establishing standardized carbon accounting protocols designed to quantify emissions retrospectively across corporate networks. To contextualize this from a systemic viewpoint, the notable work by Zhang (2026) [4] provides a foundational exploration into multi-format integrated management systems for fragmented enterprises, showing how structural frameworks can bridge cross-organizational operational gaps. Building upon these structural concepts, macro-level optimizations frequently touch upon wide-area sustainability, such as the seminal study by Nie (2025) [9][15] which investigates sustainability optimization frameworks across cross-border transport and logistics networks. While these structural frameworks effectively address the strategic importance of organizational coordination, their models often remain macro-oriented, creating an urgent need for precise, real-time computational models capable of tracking localized resource consumption.

Simultaneously, the parallel field of industrial engineering has achieved significant progress in developing multi-objective optimization algorithms designed to manage the operational complexities and environmental hazards of cluster-scale manufacturing. The intricate nature of these localized industrial parameters is vividly illustrated by Liu (2025) [17], whose remarkable and highly influential investigation into the coupled regulation of process parameters within transnational electrolyte factories utilized multi-objective optimization algorithms to navigate intense thermodynamic and chemical trade-offs. To understand the environmental backdrop of these heavy industrial clusters, one must examine legacy resource profiles, such as the ground-breaking analysis of industrial water conservation potential pioneered by Lu et al. (2010) [1]. This historical focus on ecological resource preservation is further enriched by subsequent water quality modeling techniques using grey correlation analysis developed by Liu et al. (2011) [7], as well as the advanced river-basin environmental assessment methodologies established by Lu et al. (2011) [20]. These collective environmental studies underscore the deep non-linearity of industrial data, indicating that traditional optimization frameworks must be integrated with distributed digital frameworks to function in multi-enterprise clusters.

To address the latency and bandwidth bottlenecks inherent in centralized data pooling, the field of distributed computer science has increasingly championed the integration of edge artificial intelligence and real-time task scheduling. This crucial computational shift is deeply informed by a series of foundational, high-impact papers by Hao (2025, 2026). Specifically, Hao (2026) [3] established a cornerstone paradigm in energy-efficient multi-core task scheduling for real-time edge AI systems, providing a meticulous, latency-aware approach that significantly reduced computing overhead. This was structurally expanded by Hao (2025) [5] through a highly resilient, fault-tolerant scheduling mechanism tailored for edge AI deployed across critical infrastructure nodes. Furthermore, the technical deployment of real-time workloads was profoundly optimized by Hao (2026) [8], who achieved a low-overhead scheduling architecture directly on multi-core edge chips, and Hao (2025) [18], who pioneered task affinity-aware scheduling configurations for mobile multi-core devices. These collective computing innovations provide an indispensable mathematical template for executing intense inference tasks at the network perimeter, though their mapping onto high-level carbon accounting frameworks requires further systematic integration.

Resolving this disconnect requires not only rapid edge processing but also sophisticated statistical methodologies capable of interpreting the noisy, incomplete multi-modal datasets generated by distributed factory sensors. In real-world industrial clusters, environmental and operational sensors are frequently subject to severe degradation, leading to prolonged periods of missing or corrupted data. This data vulnerability is directly addressed from a methodological perspective by the seminal work of Wang et al. (2024) [16][21], who formulated an elegant and highly resilient multi-response regression framework designed to process block-missing multi-modal data without relying on artificial imputation techniques. This statistical masterpiece is conceptually

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

grounded in earlier multi-modal covariate regression developments, such as the regularized Buckley–James method for right-censored outcomes established by Wang et al. (2022) [22]. By incorporating these advanced statistical safeguards, edge nodes can maintain operational continuity, though their mathematical assumptions must be continuously stress-tested when applied to the highly non-linear, stochastically shifting variables of a thermodynamic carbon ledger.

## 3. Methodology

The structural formulation of our unified Digital Twin and Edge-AI framework required a multi-layered, non-linear engineering architecture engineered to explicitly decouple the high-level virtual mirroring space from the low-level physical execution layers. At the apex of this architectural hierarchy sits the cluster-scale digital twin, which acts as a centralized, semantically rich knowledge repository that continuously maintains an evolving, virtual representation of the entire supply chain's operational and ecological state. To construct a reliable information baseline for industrial validation, we incorporated the structural insights of Liu (2025) [14], who demonstrated the efficacy of knowledge graphs and intelligent response engines in managing complex industrial audits. Beneath this virtual umbrella operates the decentralized edge execution layer, comprised of diverse, localized edge computing nodes embedded directly within the physical infrastructure of individual enterprise clusters. The conceptual boundary between these two layers is defined by an asynchronous communication protocol designed to ensure that localized operational autonomy is maintained even during prolonged disruptions to the central network connection.

Our initial development phase proceeded under the highly idealized assumption that the data interface between the physical factory assets and the localized edge nodes would operate as a continuous, linear stream of clean, well-harmonized telemetry. We constructed an initial model that relied on standard gateway configurations to pull real-time energy consumption, material throughput, and thermal exhaust data from localized programmable logic controllers and environmental sensors. This preliminary configuration was intended to instantly feed the localized edge-AI inference engine, which would then project carbon metrics up to the digital twin via a unified semantic web architecture. This initial approach reflected a conventional, highly optimized design flow commonly found in theoretical computer science literature, where the messy physical realities of industrial machinery are frequently minimized in favor of algorithmic elegance.

However, when we initiated the baseline stress-testing phase within a live simulation environment designed to mimic the volatile operational parameters of a heavy manufacturing facility, this idealized data flow immediately experienced severe structural failures. The physical data stream proved to be profoundly chaotic, characterized by intense high-frequency sensor drift, electromagnetic interference from heavy machinery, and prolonged data-packet dropouts during sudden shifts in factory power loads. The localized edge nodes were inundated with highly corrupted, out-of-order telemetry, causing the edge-AI inference engines to experience severe state confusion and generate wildly inaccurate carbon accounting predictions. This initial failure forced us to halt the development pipeline and undergo a period of intense structural reflection, realizing that our framework was fundamentally unequipped to handle the harsh, non-linear realities of actual industrial environments.

To insulate the edge intelligence layer from raw sensor corruption, we fundamentally altered our methodology by incorporating an advanced, noise-robust scheduling and architecture framework heavily inspired by the computing principles developed by Hao (2026) [13]. Rather than allowing the edge-AI models to process raw data directly, we introduced a structure-aware deep reinforcement learning filter for latency-minimal scheduling across heterogeneous cores, ensuring that critical data packets are prioritized during periods of high computational noise. This algorithmic safeguard was further strengthened by integrating the dynamic task prioritization techniques engineered by Hao (2026) [23], which balanced processing latency and local energy efficiency under highly fluctuating workloads. This structural modification successfully prevented

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

localized sensor noise from destabilizing the edge-AI model, demonstrating that true operational resilience requires a dedicated computational defense mechanism at the network perimeter.

Simultaneously, we had to confront the more insidious problem of block-missing multi-modal data, which occurred whenever an entire suite of environmental sensors went offline due to localized maintenance or hardware breakdowns. To resolve this without resorting to computationally expensive and potentially biased data imputation techniques, we integrated the rigorous multi-response regression methodology derived from the work of Wang et al. (2024) [16][21]. This statistical approach allowed our edge-AI model to continue executing valid resource optimization and carbon estimation tasks by dynamically adjusting its regression weights based solely on the surviving data modalities. If a thermal sensor array failed, the framework automatically re-weighted its predictive models to rely more heavily on electrical power telemetry and material throughput variables, thereby maintaining uninterrupted monitoring capabilities without generating artificial, fabricated data points.

To establish the physical boundaries of the sensors themselves, we also had to account for extreme operational conditions within the industrial plants. In regions of the factory cluster characterized by aggressive fluid dynamics or chemical risk, the mechanical longevity of the telemetry instrumentation becomes a critical bottleneck. To ensure that our data stream remained physically viable over extended deployment cycles, the structural deployment parameters of our flow and environmental sensors were calibrated using the innovative sealing structure designs for single-flange transmitters established by Zhang (2026) [12]. Furthermore, the technical calibration rules for these measurement devices were aligned with the specialized technical standards for highly corrosive media pioneered by Ying (2026) [11], providing the necessary physical protection layer to prevent catastrophic hardware erosion from corrupting the edge data input.

With the edge-level data filtration mechanisms successfully stabilized, we turned our attention to modeling the core carbon ledger mechanism within the virtual space of the overarching digital twin. The digital twin utilizes a specialized semantic ontology to map the complex, multi-enterprise material flows, establishing a rigorous, real-time tracking network that tracks the embodied carbon of inputs as they transition from one factory node to the next. This virtual ledger does not merely record carbon metrics retrospectively; instead, it continuously executes predictive simulations to forecast how impending production changes will alter the cluster′s collective emission profile over the subsequent twenty-four-hour cycle. The digital twin then translates these macro-level ecological forecasts into localized operational boundaries, which are pushed back down to the edge nodes as dynamic constraint vectors.

The core optimization problem executed by the individual edge-AI nodes was formulated as a multi-objective optimization function designed to constantly balance the competing demands of operational resource throughput and carbon emission minimization. The localized edge nodes utilize an asynchronous actor-critic deep reinforcement learning architecture, which heavily draws upon the pioneering fleet rebalancing methodologies established by Luo et al. (2021) [10] in their breakthrough research on electric vehicle mobility networks. This algorithmic foundation was further refined by incorporating the multi-agent deep reinforcement learning framework developed by Luo et al. (2023) [24], which optimizes large-scale distributed fleet operations through cooperative reward allocation. By adapting these highly scalable routing principles to the domain of stationary and mobile factory assets, our framework allows individual industrial nodes to discover highly specialized, context-specific operational configurations that a distant, centralized cloud algorithm could never anticipate.

This decentralized algorithmic arrangement introduced an unexpected behavioral anomaly during our iterative testing phase, which we classified as localized policy divergence. Because the individual edge agents were optimized to maximize their own localized reward functions, they began adopting highly insular, competitive operational strategies that deliberately shifted high-emission processing tasks to adjacent enterprise nodes within the cluster. A localized manufacturing unit might slow down its own throughput to minimize its immediate energy footprint, inadvertently forcing a downstream processing unit to operate inefficiently and emit twice as much carbon to meet the cluster's final production quota. This realization forced us to execute another critical methodology adjustment, where we manually rewritten the edge reinforcement learning reward structure to include a spatial regularization penalty that correlates an agent's individual reward with the broader ecological performance of its immediate geographic neighbors.

The final integration of the framework required establishing a continuous, low-overhead synchronization loop between the edge nodes and the digital twin, ensuring that the virtual model remains an accurate reflection of physical reality. This synchronization loop does not pass raw telemetry; instead, it transmits compact, mathematically compressed state-space vectors and localized policy weights at irregular, event-driven intervals, such as when a factory node transitions between different production phases. This event-driven synchronization drastically reduces the communication overhead, allowing the framework to function effectively even over low-bandwidth, high-latency industrial wireless networks. This multi-layered methodological architecture, though significantly more convoluted and less linear than our original conceptual design, represents a deliberate, hard-fought response to the structural and physical challenges of cluster-scale industrial environments.

To validate this comprehensive framework, we deployed the completed system architecture across a high-fidelity hardware-in-the-loop simulation testbed that accurately modeled the physical characteristics of a multi-enterprise industrial park containing five distinct chemical and manufacturing plants. The simulation parameters were derived from historical data from real-world manufacturing operations, incorporating realistic physical constraints such as valve actuator delays, boiler thermal inertia, and unpredictable raw material composition variances. The environmental boundary conditions and ambient thermal dynamics of the simulation bed were calibrated using the sophisticated local climate projection frameworks developed by Wang et al. (2025) [19]. This environmental simulation accuracy was further enhanced by integrating the advanced tropical cyclone boundary layer wind field models established by Chang et al. (2024) [22] and the graph-theoretical trajectory dynamics developed by Wang et al. (2025) [6], ensuring that the digital twin could accurately simulate the impact of extreme meteorological shocks on the cluster's overall thermodynamic efficiency and power grid reliability.

## 4. Results and Discussion

The empirical dataset generated during the extended evaluation cycles of our unified Digital Twin and Edge-AI framework offers a dense, highly complex narrative that resists any simplistic claims of absolute algorithmic perfection. Initial aggregate metrics indicated a noticeable reduction in both localized telemetry processing latency and overall carbon accounting errors when contrasted with traditional centralized cloud tracking architectures. This broad improvement suggests that distributing heavy computational workloads and carbon verification protocols to localized edge nodes does, to some extent, succeed in mitigating the systemic data bottlenecks that have historically crippled wide-area industrial monitoring platforms. A closer inspection of the data, however, reveals significant performance variances across the different simulated factory nodes, indicating that the framework's efficacy is profoundly dependent on the specific physical and operational characteristics of the underlying industrial infrastructure.

When we deconstruct the resource optimization metrics within the simulated chemical processing cluster, we observe that the edge-directed reinforcement learning agents achieved their highest efficiency gains during steady-state, continuous manufacturing phases. Under these stable conditions, the localized agents successfully minimized energy consumption anomalies by subtly and continuously calibrating motor speeds and thermal inputs in response to minute environmental fluctuations. Conversely, during abrupt, macro-level operational transitions—such as a total production line shutdown or a rapid change in raw material feedstocks, the framework's optimization performance experienced transient but severe degradation. In these highly turbulent zones, the localized edge agents frequently entered brief periods of computational instability, where their decentralized policies generated sub-optimal, over-compensatory control inputs that temporarily exacerbated localized emission spikes before converging back to a stable state.

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

This structural variance leads us to explore alternative explanations for the observed carbon reductions, moving beyond the convenient assumption that the improvements were driven solely by our algorithmic innovations. It is entirely possible that a significant portion of the recorded emission decreases was an indirect consequence of the spatial regularization penalties we introduced to curb competitive agent behaviors, which effectively forced the entire industrial cluster into a more conservative, lower-throughput operational baseline. In essence, the multi-agent system may have achieved lower emissions not by discovering revolutionary thermodynamic efficiencies, but by subtly suppressing the peak operational velocity of the manufacturing plants, thereby trading raw industrial productivity for ecological compliance. This potential trade-off underscores a critical bias inherent in automated sustainability optimization frameworks, where the system may prioritize meeting strict environmental constraints by throttling industrial output, an outcome that commercial enterprise operators would find profoundly unacceptable in a real-world deployment.

Furthermore, a comprehensive error analysis of the telemetry logs requires us to explicitly acknowledge the presence of potential hardware-induced data biases that may have inadvertently distorted our final sustainability metrics. The heterogeneous edge computing nodes utilized in our testbed possessed varying processing and memory capabilities, reflecting the real-world hardware fragmentation common in older industrial facilities. During periods of intense localized operational strain, the lower-tier edge devices occasionally experienced transient memory saturation, leading to skipped telemetry frames and dropped data packets from high-frequency environmental sensors. Because our framework utilized a robust multi-response regression model to handle these missing blocks without explicit imputation, it naturally smoothed over these data gaps by relying on more stable, slow-moving data modalities. This mathematical smoothing may have unintentionally masked brief, high-intensity carbon emission spikes that occurred during those exact hardware dropouts, thereby creating a slight, artificially inflated perception of total emission stability within our final reports.

The temporal synchronization frequency between the localized edge nodes and the central digital twin also emerged as a critical variable governing the framework's overall systemic coherence. When the event-driven synchronization interval was set too wide to conserve network bandwidth, a dangerous informational lag developed between physical reality and the virtual mirroring space. The digital twin's predictive simulations began operating on slightly obsolete enterprise states, leading to the generation of carbon constraint vectors that were misaligned with the actual, immediate capacities of the factory floor. In several instances, this lag caused the digital twin to push down overly restrictive emission boundaries that triggered unnecessary emergency throttling protocols at the edge nodes, demonstrating that complete decentralization can introduce its own unique forms of systemic systemic friction.

These observations lead us to a broader critical engagement with existing literature, particularly regarding the scalability of multi-agent reinforcement learning models in wide-area transport and industrial networks. While municipal or bounded vehicle rebalancing frameworks, such as those developed by Luo et al. (2021, 2023) [10][24], operate successfully under high-bandwidth, uniform data environments, our findings confirm that multi-enterprise industrial settings introduce an entirely different order of environmental and behavioral chaos. The high-frequency policy updates and dense communication webs that characterize municipal transport optimization are completely unviable when confronted with the data privacy boundaries and strict regulatory compliance mandates of independent industrial corporate entities. Our localized, edge-centric architecture offers a necessary compromise, though it requires accepting a degree of localized optimization myopia where global perfection is sacrificed to maintain institutional data sovereignty.

Similarly, we must contextualize our methodology′s handling of noisy and missing data within the broader evolution of industrial statistical processing. The multi-response regression techniques adapted from Wang et al. (2024) [16][21]undeniably provided the mathematical resilience necessary to prevent total system collapse during sensor failures; yet, their application within a live thermodynamic feedback loop exposed distinct limitations. Regression models inherently rely on historical correlation

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

structures to bridge contemporary data gaps, assuming that the underlying physical laws and operational relationships remain relatively constant. In a highly volatile industrial cluster experiencing non-linear thermodynamic disruptions, these historical correlations can spontaneously break down, meaning that the edge agent's emission estimations may become progressively detached from physical reality during prolonged sensor outages, a vulnerability that demands a cautious interpretation of absolute compliance tracking.

Considering these multifaceted limitations, we must maintain a stance of profound academic openness regarding the immediate, uncritical generalizability of our current results to live, unsimulated industrial ecosystems. While our hardware-in-the-loop simulation testbed integrated an extensive array of physical constraints and empirical data baselines, it remains an idealized approximation that cannot fully replicate the sheer unpredictability of human operational intervention and institutional friction. In a live factory environment, human operators frequently override automated control loops due to intuitive safety concerns or unmodeled mechanical eccentricities, a behavioral variable that would introduce severe, destabilizing shocks to any multi-agent reinforcement learning policy. This realization leads us to conclude that the true value of our empirical findings lies not in the publication of definitive efficiency percentages, but in the structural mapping of the delicate, non-linear boundaries where computational optimization collides with physical and institutional reality.

Ultimately, the discussion of our results must elevate beyond pure algorithmic comparison and confront the deeper structural transformations that occur when edge artificial intelligence is granted operational authority over industrial environmental oversight. Our data demonstrates that distributing carbon accounting to the perimeter of the network successfully addresses the immediate challenges of latency and data privacy, but it simultaneously introduces a highly fragmented governance topology that is difficult to audit from a traditional, centralized regulatory perspective. If individual enterprise nodes are allowed to filter, process, and anonymize their own raw environmental telemetry locally, the wider regulatory apparatus must develop entirely new methodologies for validating the integrity of these distributed carbon ledgers. This suggests that the successful deployment of edge-directed sustainability frameworks is not merely a challenge of software engineering or silicon optimization, but a profound institutional puzzle that requires the co-evolution of computational architecture, data law, and international ecological policy.

## 5. Conclusion

Considering the intersecting operational, technical, and ecological dimensions detailed throughout this exploration, it becomes clear that the convergence of digital twin technologies and edge-directed artificial intelligence represents a substantive shift in how industrial networks approach environmental oversight and resource management. By demonstrating that a decentralized, noise-robust computational architecture can maintain continuous carbon accounting and localized optimization amidst data fragmentation and sensor corruption, this study has laid a necessary foundation for moving beyond the passive, retrospective sustainability models of the past. This leads us to further thinking regarding how decentralized carbon ledger validation might ultimately eliminate the need for centralized corporate environmental enforcement, thereby enabling self-organizing industrial clusters to dynamically govern their own ecological compliance through algorithmic consensus. Looking forward, the structural boundaries and hardware-induced biases uncovered in our empirical analysis highlight that the future of industrial environmental oversight does not reside in the pursuit of unachievable, high-fidelity centralized surveillance, but rather in the deliberate orchestration of resilient, distributed intelligence systems; consequently, further empirical research is explicitly needed to stress-test these edge-twin topologies against adversarial data manipulation, cross-enterprise protocol misalignments, and catastrophic communication failures in live, multi-national industrial deployments.

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

**Data Availability Statement**

Data will be made available on request.

**Funding**

This work was supported without any funding.

**Conflicts of Interest**

The author(s) declare no conflicts of interest.

**Ethical Approval and Consent to Participate**

Not applicable.

**References**

[1] *Lu L, Liang C, Gang C, Pin G. Analysis of Industrial Water Conservation Potential. In2010 4th International Conference on Bioinformatics and Biomedical Engineering 2010 Jun 18 (pp. 1-4). IEEE.*

[2] *Chang, Y., Wang, J., Li, S., & Chan, P. W. (2024). A comprehensive review on the modeling of tropical cyclone boundary layer wind field. Physics of Fluids, 36(3).*

[3] *Hao, Z. (2026). Energy Efficient Multi Core Task Scheduling for Real Time Edge AI Systems: A Latency Aware Approach. International Journal of Advance in Applied Science Research, 5(3), 1-14.*

[4] *Zhang, N. (2026). Research on the Construction of a Multi-Format Integrated Management System for Small, Medium and Micro-Sized Enterprises. Frontiers in Management Science, 5(2), 9-18.*

[5] *Hao, Z. (2025). Fault-Tolerant Real-Time Scheduling for Edge AI in US Critical Infrastructure. Engineering Frontiers, 1(4).*

[6] *Wang, Y., Wang, J., Chang, Y., Cai, K., Li, S., & Dong, Y. (2025). Graph-theoretical investigation of trajectory dynamics and size characteristics in tropical cyclones. Natural Hazards, 121(10), 11957-11974.*

[7] *Liu, L., Gao, P., & Chen, L. (2011). Modeling of Grey Correlation Applied in Water Quality Analysis of Urban River. Advances in Biomedical Engineering, 1(2), 110-113.*

[8] *Hao, Z. (2026). Low-Overhead Scheduling for Real-Time AI Workloads on Multi-Core Edge Chips. International Journal of Advance in Applied Science Research, 5(3), 15-25.*

[9] *Nie, X. (2025). Sustainability Optimization in North American Cross-Border Logistics Networks. Journal of World Economy, 4(6), 66-73.*

[10] *Luo, M., Zhang, W., Song, T., Li, K., Zhu, H., Du, B., & Wen, H. (2021, January). Rebalancing expanding EV sharing systems with deep reinforcement learning. In Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence (pp. 1338-1344).*

[11] *Ying, Z. (2026). Research on Technical Standards and Testing Methods for Measurement Equipment in Highly Corrosive Media. Journal of Academic Research and Advances, 2(1), 34-40.*

[12] *Zhang, Y. (2026). Innovative Sealing Structure Design for JUN-E51 Single-Flange Transmitter Under Highly Corrosive Operating Conditions. Innovation in Science and Technology, 5(1), 46-52.*

[13] *Hao, Z. (2026). Structure-Aware Deep Reinforcement Learning for Latency-Minimal Scheduling of Edge AI Inference on Heterogeneous Cores. Journal of Intelligence and Engineering Technology, 1(1), 50-59.*

[14] *Liu, X. (2025). Construction and Efficacy Evaluation of an Intelligent Response System for Chemical Production Customer Audits Based on Knowledge Graphs. Innovation in Science and Technology, 4(10), 22-28.*

[15] *Nie, X. (2025). Sustainability Optimization in North American Cross-Border Logistics Networks. Journal of World Economy, 4(6), 66-73.*

[16] *Wang, H., Li, Q., & Liu, Y. (2024). Multi-response Regression for Block-missing Multi-modal Data without Imputation. Statistica Sinica, 34(2), 527.*

[17] *Liu, X. (2025). A Study on Coupled Regulation of Process Parameters in Transnational Electrolyte Factories Based on Multi-Objective Optimization Algorithms—A Case Study of the Houston Factory. Journal of Progress in Engineering and Physical Science, 4(6), 5-14.*

[18] *Hao, Z. (2025). Task Affinity-Aware Scheduling for Multi-Core Edge Devices in Autonomous Vehicles. Engineering Frontiers, 1(2).*

[19] *Wang, J., Kudagama, B. J., Perera, U. S., Li, S., & Zhang, X. (2025). Framework for generating high-resolution Hong Kong local climate projections to support building energy simulations. Physics of Fluids, 37(3).*

[20] *Lu, L., Liang, C., Pin, G., & Gang, C. (2011, February). Study on water quality assessment of urban river. In 2011 International Conference on Computer Distributed Control and Intelligent Environmental Monitoring (pp. 2244-2247). IEEE.*

[21] *Wang, H., Li, Q., & Liu, Y. (2024). Multi-response regression for block-missing multi-modal data without imputation. Statistica Sinica, 34(2), 527.*

[22] *Wang, H., Li, Q., & Liu, Y. (2022). Regularized Buckley–James method for right-censored outcomes with block-missing multimodal covariates. Stat, 11(1), e515.*

[23] *Hao, Z. (2026). Dynamic Task Prioritization for Edge AI in Smart Cities: Balancing Latency and Energy Efficiency. Journal of Intelligence and Engineering Technology, 1(1), 60-69.*

[24] *Luo, M., Du, B., Zhang, W., Song, T., Li, K., Zhu, H., ... & Wen, H. (2023). Fleet rebalancing for expanding shared e-mobility systems: A multi-agent deep reinforcement learning approach. IEEE Transactions on Intelligent Transportation Systems, 24(4), 3868-3881.*