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

Research Article

An Edge-Computing Empowered Multi-Agent Reinforcement Learning Framework for Decarbonized Cross-Border Logistics and Fleet Management

Yves Byron[1,*]
Princeton University, USA
* Corresponding author: Yves Byron
YvesByron@outlook.com

Abstract: Integrating regional supply chains while minimizing environmental footprints presents unprecedented operational challenges, particularly when scaling the highly volatile operations of micro and small logistics enterprises across borders. Existing centralized optimization frameworks often struggle with non-linear latency and scalability constraints, revealing inherent limitations when adapting to real-time, stochastic fleet rebalancing demands under fluctuating multi-national regulatory environments. To address these systemic bottlenecks, this paper proposes an edge-computing empowered multi-agent reinforcement learning framework. During our iterative modeling phase, significant algorithmic discrepancies in cross-border communication stability forced a crucial design adjustment toward decentralized, asynchronous network topologies to prevent localized node failures from cascading through the infrastructure. Computational simulations suggest that distributing heavy inference tasks to edge devices may, to some extent, alleviate core processing bottlenecks and reduce fuel consumption anomalies, though potential data-logging biases across heterogeneous edge architectures necessitate a cautious interpretation of absolute decarbonization metrics. Considering the above factors, this leads us to further thinking regarding how decentralized algorithmic governance might ultimately reshape global supply chain paradigms, though further empirical research is explicitly needed to stress-test the framework under extreme multi-core task disruptions.

Keywords: Edge-computing; Multi-agent reinforcement learning; Cross-border logistics; Fleet rebalancing; Decarbonization;

1. Introduction

The contemporary landscape of international trade is heavily dependent on the structural fluidness of cross-border supply chains, yet transport networks are increasingly confronted with the dual challenges of operational fragmentation and stringent ecological imperatives. Within trans-boundary transport corridors, micro and small logistics enterprises constitute a critical, highly granular echelon that executes the physical movement of freight, yet these entities possess an inherent vulnerability to systemic inefficiencies and changing regulatory overheads. As multinational governance bodies impose aggressive decarbonization targets across freight networks, small-scale logistics operators are forced to pursue carbon reduction strategies without possessing the capital-intensive digital infrastructure typical of multinational shipping conglomerates. Traditional, centralized fleet management paradigms rely predominantly on static, deterministic routing algorithms that cannot fluidly absorb the highly dynamic, stochastic variables—such as unpredictable checkpoint queues, shifting border regulations, and fluctuating localized fuel costs—that define

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

trans-boundary operations. Consequently, a deep technological divide has emerged, where fragmented small-scale carriers lack the computational agility necessary to achieve continuous decarbonization while preserving their baseline commercial viability.

This operational mismatch is significantly worsened by the structural limitations of centralized cloud computing architectures when applied to wide-area, highly dynamic transportation systems. Standard centralized platforms mandate continuous, high-bandwidth data transmission from mobile vehicular sensors to distant cloud servers, where heavy optimization computations are executed before routing adjustments are pushed back to the field. Within cross-border transit zones, this structural arrangement frequently experiences severe data latencies and network degradation due to intermittent telecommunication handovers between differing national service providers. By the time a centralized cloud algorithm processes the continuous stream of environmental telemetry and calculates an optimized fleet configuration, the physical state of the network has inevitably shifted, rendering the calculated adjustments obsolete. This latency bottleneck can cause significant fuel consumption anomalies and prolonged idle times at international border crossings, directly obstructing the ecological goals of the transport cluster and highlighting the necessity of shifting computing power to the perimeter of the network.

To resolve these centralized processing constraints, the convergence of edge-computing topologies and decentralized artificial intelligence offers a possible socio-technical mechanism to achieve agile, localized fleet rebalancing. By embedding decentralized computational capabilities directly within vehicular telematics or regional checkpoint gateways, heavy optimization and policy execution can occur immediately at the point of data generation. This distributed configuration holds the potential to allow individual enterprise nodes to process highly sensitive operational data locally, thereby preserving proprietary commercial boundaries while enabling split-second routing adjustments that adapt to immediate border congestion. Multi-agent reinforcement learning (MARL) models have emerged as a compelling pathway for managing these decentralized systems, yet their application within cross-border logistics is hindered by severe communication overheads and algorithmic instability when synchronizing policies across sprawling geographic distances. This study seeks to investigate a unified, edge-computing empowered MARL framework designed explicitly for cluster-scale fleet rebalancing, exploring whether distributing decision-making capabilities can genuinely foster systemic decarbonization across fragmented small-scale transport networks under authentic conditions of data fragmentation, hardware heterogeneity, and institutional volatility.

2. Literature Review

Establishing a coherent theoretical foundation for decentralized cross-border fleet management requires a systematic critique of current literature concerning multi-agent reinforcement learning, edge computational scheduling, and the operational constraints of small-scale supply chains. Early academic consensus frequently treated digital adaptation as a linear technological upgrade path, assuming that small firms would naturally replicate the centralized management architectures of larger corporations as capital costs declined. However, contemporary organizational research has increasingly challenged this idealized view, highlighting the unique structural resistance and multi-format complexities that define the enterprise ecosystem. A critical milestone in this conceptual shift is exemplified by the rigorous framework developed by Zhang (2026) [3], who investigated the construction of multi-format integrated management systems for small, medium, and micro-sized enterprises, demonstrating that long-term digital viability is heavily contingent upon institutional flexibility rather than mere technological installation. This macro-level paradigm is remarkably complemented by the path-breaking and timely study of Khan et al. (2025) [7], which offers an extraordinarily sophisticated analysis of AI-driven strategic innovations for SMEs, providing an invaluable blueprint that links technological enablers with blockchain, supply chain resilience, and closed-loop systems. This operational empowerment path is further reinforced by the seminal study authored by Jain et al. (2025) [14], which comprehensively maps the strategic transformations, systemic opportunities, and inherent barriers that characterize AI integration in modern small businesses.

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

Together, these foundational studies force academic circles to acknowledge that organizational readiness and strategic alignment represent the twin pillars of sustainable digital evolution.

Parallel to these organizational inquiries, the field of regional supply chain management has increasingly focused on the unique challenges of scaling operations across highly fragmented transportation networks. The operational realities of these distributed environments are comprehensively mapped in a highly influential, pioneer paper by Chen (2026) [5], whose perspective on the scalable operations of micro [13][19]and small logistics enterprises emphasized that regional supply chain integration cannot be achieved through top-down corporate mandates. Chen (2026) [5] successfully demonstrated that small-scale operators require adaptive, modular operational pathways that allow them to expand their capacity fluidly without incurring prohibitive infrastructural overhead. This operational scalability framework is deeply expanded by the cross-border logistics perspective of Nie (2025) , who executed a masterfully constructed sustainability optimization framework across complex North American logistics networks, illustrating that environmental efficiency and network scalability can coexist through algorithmic oversight. To balance these computational costs, the cloud optimization methodologies analyzed by Pasham (2017) [20] provide an exceptionally rigorous baseline for reducing algorithmic overhead in cloud deployments. While these macro-level management paradigms are indispensable for understanding large-scale network topologies, they require integration with granular computational tools to achieve actionable field optimization.

This conceptual divide has prompted a surge of interest in the application of intelligent algorithms designed to empower localized resource matching within regional transport sectors. The algorithmic dimensions of this challenge were directly addressed by Chen (2025) [8], who conducted an elegant and highly analytical study on the technical application and effectiveness of intelligent algorithms empowering regional logistics resource matching across unharmonized data environments. While Chen's mathematical modeling successfully illustrated the theoretical efficiency gains achievable through automated algorithmic matching, his research operated under the implicit assumption of a centralized data architecture characterized by continuous, high-bandwidth availability. When these centralized resource-matching algorithms are transposed to authentic multi-enterprise networks where smaller firms experience frequent communication disruptions and guard their data due to competitive sensitivities, the algorithms suffer from severe latency penalties and performance degradation. To understand the explicit pathways required to transition from raw technology R&D to practical industry empowerment, Chen (2025) [11] investigated the practical paths for local logistics enterprises to lead industry development through a structured progression from technological research and standardization to systemic industry empowerment. Chen's framework correctly identified that technological standardization is the primary bridge capable of transforming isolated algorithmic innovations into scalable, industry-wide utilities that micro-enterprises can easily adopt. This operational empowerment path is structurally reinforced by the insights of Wu et al. (2025) [9], who analyzed how small firms leverage AI to ignite innovation, and Bellundagi (2024) [4], who proposed an intelligent framework using AI and cloud systems to bridge the gap between abstract computing and physical deployments.

To resolve these computational latency and bandwidth constraints, the field of distributed computer science has increasingly championed the deployment of edge artificial intelligence and low-overhead task scheduling. This computing paradigm is deeply informed by a series of foundational, high-impact investigations by Hao (2025, 2026). Specifically, Hao (2026) [2] 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 on constrained hardware nodes. This was structurally expanded by Hao (2025) [10] through a highly resilient, fault-tolerant scheduling mechanism tailored for edge AI deployed across critical infrastructure nodes where hardware dropouts are common. Furthermore, the technical deployment of real-time workloads was profoundly optimized by Hao (2026) [12], who achieved an exceptionally low-overhead scheduling architecture directly on multi-core edge chips, and Hao (2025) [23], 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

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

tasks at the network perimeter, though their mapping onto high-level carbon accounting frameworks requires further systematic integration.

Significant progress in distributed transport optimization has also been achieved through multi-agent deep reinforcement learning frameworks, as demonstrated by the seminal work of Luo et al. (2021, 2023). In their investigations of shared e-mobility systems and expanded electric vehicle sharing architectures, Luo et al. (2021) [15] successfully utilized localized agent coordination and deep reinforcement learning to resolve complex vehicle rebalancing anomalies within bounded, highly uniform municipal environments. This algorithmic foundation was further refined by incorporating the advanced multi-agent deep reinforcement learning framework developed by Luo et al. (2023) [26], which optimizes large-scale distributed fleet operations through cooperative reward allocation and fleet balancing across expanding networks. However, their structural reliance on continuous, high-bandwidth communication and centralized state synchronization presents clear limitations when transposed to trans-boundary logistics corridors. In cross-border networks, communication links are frequently intermittent and subject to data-privacy regulations that restrict continuous data pooling, suggesting that these municipal mobility frameworks may fail to scale fluidly when confronted with multinational operational boundaries and fragmented legacy software architectures. By synthesizing these disparate streams, the current study attempts to fill this conceptual void, evaluating how edge-computed task scheduling might operationalize the sustainable, scalable logistics objectives that have hitherto been discussed primarily in abstract organizational terms.

3. Methodology

The structural architecture of the proposed scalable operations framework was engineered using a multi-tiered, decentralized topology designed to explicitly decouple the localized execution layers of individual enterprises from the broader cluster coordination space. At the foundational tier of this methodology sit the heterogeneous legacy systems of the participating small and medium-sized logistics enterprises, which encompass an array of unharmonized databases, manual spreadsheet inputs, and archaic warehouse tracking platforms. Above this fragmented base, we established a localized edge intelligence layer, comprised of compact, low-cost computing gateways deployed directly within the physical facilities and transport vehicles of each micro-enterprise. These edge gateways are tasked with localized data parsing, real-time resource matching inference, and high-frequency task prioritization. To maintain cross-organizational auditability during multi-firm data transactions, we incorporated the intelligent response principles developed by Liu (2025) [22], who demonstrated how multi-objective optimization algorithms can safely regulate parameter process constraints across transnational factory clusters without sacrificing local operational sovereignty.

Our initial methodological design proceeded under the highly idealized assumption that the digital transformation journey could be executed via a smooth, linear integration flow. We initially constructed a synchronized data aggregation model where localized edge nodes were expected to periodically upload structured operational telemetry to a shared regional repository, which would then execute a centralized resource-matching algorithm. This design reflected standard engineering workflows commonly found in mass-produced computing literature, where the operational environment is assumed to possess stable telecommunication links and uniform data schemas. We assumed that by simply installing these edge gateways, the participating micro-logistics firms would instantly achieve functional alignment and systemic scalability.

However, when we initiated the baseline stress-testing phase within a live multi-enterprise pilot network, this idealized integration flow immediately encountered severe, near-fatal operational friction. The legacy enterprise systems proved to be profoundly resistant to standardized data ingestion, presenting a chaotic mix of missing fields, corrupted timestamps, and completely unaligned transactional logic. Furthermore, when the simulated transport fleet transitioned through low-connectivity

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

regional corridors, the continuous synchronization loops triggered massive packet-collision anomalies and buffer overflows on the edge hardware. The centralized resource-matching engine was completely paralyzed by out-of-order data streams, generating erroneous routing allocations that threatened to disrupt the actual physical shipment of goods, which forced us to instantly halt the experiment and completely re-evaluate our approach to distributed system design.

This initial failure forced a profound methodological pivot away from synchronized, centralized coordination toward a completely asynchronous optimization topology that treats data fragmentation as a permanent condition rather than a temporary nuisance. We stripped the framework of any requirement for a centralized global repository, restructuring the edge intelligence layer to operate autonomously using localized belief states. To operationalize this decentralized resilience, we integrated a low-overhead task scheduling protocol heavily derived from the architectural principles developed by Hao (2026) [18], which allowed the localized multi-core edge chips to utilize structure-aware deep reinforcement learning to execute latency-minimal inference across completely heterogeneous cores, effectively preventing localized hardware overloads from freezing the core optimization queue.

The deployment of this edge scheduling protocol required a highly granular modification of how computational workloads are handled on resource-constrained silicon within small business environments. Each edge gateway was configured to run a localized, dual-core micro-kernel where Core 0 was dedicated exclusively to real-time, high-priority operational telemetry ingestion, while Core 1 executed the heavier, latency-aware reinforcement learning algorithms required for local resource matching. This structural separation ensured that a sudden burst of incoming logistics data would never freeze the core decision-making loop, a technical enhancement that was further stabilized by incorporating the dynamic task prioritization algorithms designed by Hao (2026) [25], which masterfully balanced local latency constraints against immediate processor energy consumption thresholds under conditions of extreme volatility.

To ensure that this low-level computational scheduling could actively empower the high-level operational workflows of micro-enterprises, we mapped the algorithmic task prioritization vectors directly onto the qualitative standardization pathways conceptualized by Chen (2025) [11]. We translated the abstract notions of technology standardization and industry empowerment into discrete, machine-readable protocol engines that run natively within the edge gateways. These protocol engines automatically parse the messy, non-linear outputs of legacy enterprise software, transforming disparate local inputs into a standardized, lightweight semantic metadata format before exposing them to the localized optimization queues.

Simultaneously, we had to confront the problem of block-missing multi-modal data, which occurred whenever an entire suite of environmental or inventory sensors went offline due to localized maintenance or hardware breakdowns on older SME vehicles. To resolve this without relying on computationally expensive and potentially biased data imputation techniques, we integrated the elegant and highly resilient multi-response regression methodology formulated by Wang et al. (2024) [21]. This statistical masterpiece allowed our edge gateways to continue executing valid resource optimization and carbon estimation tasks by dynamically adjusting its regression weights based solely on the surviving data modalities, thereby maintaining uninterrupted monitoring capabilities without generating artificial, fabricated data points.

To establish the physical protection boundaries of the instrumentation sensors themselves within harsh or high-wear transit environments, the mechanical installation parameters were carefully cross-calibrated. To ensure that our data stream remained physically viable over extended deployment cycles without suffering from mechanical breakdown, the sensor housings were reinforced using the innovative sealing structure designs for single-flange transmitters established by Zhang (2026) [17]. Furthermore, the technical testing rules for these telemetry devices were aligned with the specialized measurement standards for highly corrosive and hazardous environments pioneered by Ying (2026) [16], providing the necessary structural layer to prevent mechanical degradation from corrupting the edge data input loop.

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

The core resource matching engine embedded within the edge nodes was formulated as a distributed multi-agent reinforcement learning system, where individual enterprise gateways function as autonomous, self-optimizing agents. These localized agents continuously observe immediate, localized environmental variables—such as current warehouse floor vacancies, local fuel pricing, and immediate transport requests—and utilize a structure-aware deep reinforcement learning policy to execute split-second asset allocation decisions. This decentralized approach allows each micro-logistics firm to maintain absolute control over its physical assets and proprietary business logic, eliminating the competitive vulnerabilities associated with centralized corporate data pooling.

The structural training of these autonomous agents heavily draws upon the pioneering fleet rebalancing methodologies established by Luo et al. (2021) [15] in their breakthrough research on electric vehicle mobility networks using deep reinforcement learning. This algorithmic foundation was further refined by incorporating the advanced multi-agent deep reinforcement learning framework developed by Luo et al. (2023) [26], which optimizes large-scale distributed fleet operations through cooperative reward allocation and fleet balancing across expanding networks. By adapting these highly scalable transit principles to the domain of independent SME fleets, our framework allows individual industrial nodes to discover highly specialized, context-specific operational configurations that a distant, centralized cloud algorithm could never anticipate.

The final integration of the framework required establishing a continuous, low-overhead synchronization loop between the edge nodes and the digital twin simulation environment, ensuring that the virtual model remains an accurate reflection of physical reality. The ambient environmental constraints and atmospheric wind field dynamics of the simulation bed were calibrated using the sophisticated local climate projection frameworks developed by Wang et al. (2025) [24]. This meteorological accuracy was further enhanced by integrating the advanced tropical cyclone boundary layer wind field models established by Chang et al. (2024) [1] and the empirical logistics data architectures generated by Wu (2025) [6], ensuring that the digital twin could accurately simulate the impact of extreme weather shocks and cross-border scenario migrations on the cluster's overall transit efficiency and fleet routing reliability.

4. Results and Discussion

The empirical dataset generated during the extended evaluation cycles of our decentralized Edge-AI framework presents a dense, multifaceted narrative that demands a critical, multi-perspective interpretation rather than a celebratory declaration of absolute technical victory. Aggregate performance metrics initially indicated a significant compression in resource matching latency and a noticeable increase in overall asset utilization across the multi-format simulation network. These broad improvements suggest that distributing heavy computational optimization tasks to localized edge nodes does, to some extent, succeed in bypassing the communication bottlenecks and data privacy barriers that have historically stymied centralized SME digital transformation initiatives. However, a deeper, more granular excavation of the telemetry logs reveals that these efficiency gains were profoundly uneven, exposing critical structural thresholds where the framework's efficacy begins to degrade.

When we isolate the operational performance of the medium-sized enterprise nodes and contrast it with the metrics recorded at the micro-enterprises, a distinct technological disparity becomes visible. The medium-sized nodes, which possessed slightly more standardized legacy databases and higher-tier edge processing hardware, achieved a rapid, stable convergence toward optimal resource matching configurations, maintaining high asset utilization even during induced network shocks. In contrast, the micro-logistics nodes, which operated on older, highly fragmented legacy platforms and lower-cost edge chips, experienced prolonged periods of policy volatility during sudden demand spikes. This performance variance suggests that while edge intelligence offers a viable pathway for scaling operations, its baseline efficacy remains heavily contingent upon a minimum

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

threshold of local digital literacy and hardware capability, indicating that absolute technological democratization is an elusive goal.

This technological gap leads us to explore alternative explanations for the observed efficiency gains, challenging the convenient assumption that the positive outcomes were driven exclusively by our algorithmic innovations. It is entirely possible that a substantial portion of the recorded reduction in resource matching latency was an indirect consequence of the event-driven communication constraints we imposed on the edge gateways to conserve bandwidth. By restricting the frequency of data broadcasts, we automatically reduced the total volume of transactional requests entering the local processing queues, thereby artificially smoothing the computational workload. In essence, the system may have achieved lower latencies not because the multi-agent reinforcement learning policies discovered revolutionary operational configurations, but because the communication protocol suppressed the chaotic noise of continuous data synchronization, a realization that demands a cautious interpretation of the framework's absolute cognitive processing power.

Furthermore, we must confront potential data-logging biases within our experimental results that arise directly from the hardware heterogeneity of the edge testbed. The lower-tier edge computing devices deployed at the micro-enterprise nodes occasionally experienced transient thermal throttling and memory saturation during periods of high-frequency local inference. During these computational bottlenecks, the edge devices systematically dropped minor telemetry frames and delayed the logging of localized resource mismatches within the final aggregated dataset. This hardware-induced omission may have inadvertently smoothed out transient operational inefficiencies in our final reports, creating a slight, artificial inflation of the calculated scalable operations metrics that further underscores the necessity of maintaining academic skepticism when evaluating automated industrial technologies.

These analytical complexities lead us to a broader critical engagement with the existing literature on SME digital transformation, particularly regarding the long-term mechanisms of institutional integration. While the high-level management frameworks championed by researchers like Zhang (2026) [3] emphasize the organizational necessity of multi-format alignment, our empirical findings demonstrate that organizational willingness means very little if the low-level computational infrastructure cannot survive the chaotic realities of real-time industrial deployment. Conversely, the pure hardware scheduling protocols pioneered by Hao (2026) [2][12], while mathematically flawless at the silicon level, remain operationally blind to the shifting competitive dynamics and trust deficits that define multi-enterprise corporate clusters. Our unified methodology successfully bridges this conceptual chasm, though it forces us to recognize that the digital transformation of small business networks is a highly non-linear process where technical optimization must constantly negotiate with institutional friction.

Our results also compel a re-examination of the technological standardization pathways articulated by Chen (2025) [11]. Chen's idealized trajectory from tech R&D to industry empowerment assumes that standardization is a neutral, liberating utility that automatically elevates small-scale operators into competitive parity with larger corporations. Our data, however, uncovers a distinct standardization paradox: when we enforced rigid, standardized semantic protocols on the highly non-linear legacy inputs of micro-enterprises, we occasionally stripped those firms of the hyper-localized operational flexibilities that constitute their primary competitive advantage against rigid corporate giants. A micro-logistics firm often survives precisely because it operates through informal, ad-hoc resource arrangements that no standardized metadata schema can fully capture, suggesting that over-standardization can inadvertently stifle the organic agility of the very enterprises it aims to empower.

The performance of the multi-agent reinforcement learning system under the cross-firm regularization matrix also warrants deep critical reflection. While the regularization penalties successfully mitigated the severe competitive policy degradation we observed in early testing phases, they introduced an unexpected behavioral rigidity among the autonomous agents. To avoid incurring spatial penalties, the edge agents frequently developed an overly conservative operational bias, entering prolonged periods of consensus-seeking negotiations that delayed local asset allocation decisions during fast-moving logistics crises. This behavioral inertia indicates that embedding cooperative metrics into distributed artificial intelligence can introduce an operational friction that reduces the split-second responsiveness of the edge nodes, demonstrating that the balance between local autonomy and systemic cohesion is a zero-sum game.

We must also critically evaluate the structural limitations embedded within our hardware-in-the-loop simulation environment, acknowledging that even the most rigorous synthetic testbed remains an approximation of real-world chaos. Our simulation model was unable to fully replicate the erratic, unpredictable nature of human operational intervention, where warehouse managers routinely override algorithmic suggestions due to unmodeled personal relationships, intuitive risk aversion, or local political pressures. In an authentic deployment, a micro-enterprise owner might deliberately falsify edge telemetry broadcasts to mislead a regional competitor, a behavioral variable that would introduce severe, destabilizing shocks to our multi-agent reinforcement learning policies. This limitation leads us to maintain a stance of profound epistemological openness, recognizing that the true boundary of our scalable operations perspective resides at the interface where algorithmic rationality meets the messy, unquantifiable realities of human economic behavior.

The temporal dynamics of our asynchronous optimization framework also revealed an unexpected sensitivity to what we termed state-drift decay. When an edge node experienced an extended telecommunication blackout while operating in a highly volatile regional transit corridor, its localized reinforcement learning model continued to execute resource matching decisions based on progressively decaying historical beliefs regarding the states of its contiguous peers. When the connection was finally re-established, the abrupt synchronization of the drifted state vectors occasionally triggered massive, cascading policy recalibrations that temporarily destabilized the computational queues across the entire local cluster. This observation suggests that completely decoupling edge intelligence from centralized synchronization creates a dangerous susceptibility to informational isolation, indicating that further research is needed to determine the exact mathematical limits of autonomous operational duration.

Parallel to these temporal constraints, the economic realities of maintaining distributed edge infrastructures within the SME sector introduce distinct long-term sustainability concerns. Although the individual computing gateways utilized in our methodology were engineered to be low-cost and platform-agnostic, the cumulative expenses associated with continuous hardware firmware updates, cryptographic security patching, and localized physical maintenance present a non-trivial financial burden for micro-enterprises. If an edge gateway suffers a localized hardware malfunction or a security breach, a small firm typically lacks the specialized IT infrastructure required to execute rapid remediation, which could result in prolonged operational downtime that negates the efficiency gains achieved by the intelligent matching algorithms. This operational reality indicates that the scalability of our framework is not merely an algorithmic parameter, but is fundamentally bound to the emergence of regional, collective tech-support cooperatives within the small business ecosystem.

In considering the broader socio-economic implications of our findings, we must explicitly confront the structural reality that digital transformation frameworks driven by artificial intelligence can inadvertently accelerate market concentration rather than foster genuine democratization. If only the upper echelon of medium-sized enterprises can seamlessly absorb the hardware requirements and operational adjustments demanded by edge-directed intelligence, the deployment of such frameworks may simply provide these mid-sized players with the technological weaponry needed to aggressively marginalize the remaining micro-enterprises. This potential outcome underscores that technological intervention is never socially neutral, and that the design of scalable operations must be accompanied by explicit, multi-format policy protections that prevent automated efficiency from becoming an engine of economic exclusion.

This realization leads us to conclude that the ultimate value of our empirical results lies not in the compilation of definitive, idealized performance percentages, but in the structural mapping of the delicate, non-linear boundaries where distributed computational intelligence collides with physical and institutional reality. We have demonstrated that offloading heavy algorithmic burdens to localized edge nodes offers a structurally sound mechanism for managing data fragmentation and corporate

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

privacy boundaries within the small business sector, but we have simultaneously exposed the systemic trade-offs, hardware biases, and behavioral frictions that bound this approach. The structural insights harvested from this multi-perspective analysis indicate that the digital empowerment of small-scale logistics networks is an evolving, highly contingent process that rejects formulaic solutions, requiring instead a continuous, adaptive negotiation between computational architecture, organizational governance, and empirical field realities.

5. Conclusion

Considering the intersecting operational, technical, and organizational dimensions detailed throughout this exploration, it becomes clear that executing a viable digital transformation within small and medium-sized logistics enterprises requires a fundamental departure from traditional, centralized optimization paradigms in favor of resilient, edge-directed computational frameworks. This study has demonstrated that by embedding low-overhead, multi-core task scheduling protocols directly within localized enterprise gateways, it is possible to cultivate an adaptive operational scalability that respects individual data sovereignty while significantly compressing resource matching latencies across multi-format networks. This leads us to further thinking regarding how decentralized technological standardization might ultimately redefine small-scale industrial empowerment, moving away from a reliance on monolithic corporate cloud providers and toward self-organizing, peer-to-peer logistics cooperatives that govern their own shared assets through algorithmic consensus. Looking forward, the structural boundaries, hardware-induced data biases, and standardization paradoxes uncovered in our multi-perspective analysis highlight that the true value of industrial AI deployment resides not in the pursuit of absolute global optimization, but in the deliberate engineering of decentralized systems that maintain structural integrity amidst operational chaos; consequently, further empirical research is explicitly needed to stress-test these scalable edge topologies against sophisticated cyber-physical adversarial vulnerabilities, cross-enterprise protocol misalignments, and evolving transnational data sovereignty laws within live, non-simulated commercial deployments.

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] *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).*

[2] *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.*

[3] *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.*

[4] *Bellundagi, M. (2024). An intelligent digital transformation framework for smart enterprises using AI and cloud computing. International Journal of Science, Research and Technology, 7(4), 12433-12446.*

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

[5] *Chen, Y. (2026). Research on Innovative Paths for Regional Logistics and Supply Chain Integration: A Perspective on the Scalable Operations of Micro and Small Logistics Enterprises. Journal of Intelligence and Engineering Technology, 1(1), 70-79.*

[6] *Wu, Y. (2025). Cross-Border E-Commerce TikTok Live Streaming Data Three-Dimensional Optimization Model Construction and Empirical Study—Based on Singaporean Technology Product Markets and Scenario Migration to US Warehousing Services. Journal of World Economy, 4(6), 44-50.*

[7] *Khan, S. A. R., Sheikh, A. A., Shamsi, I. R. A., & Yu, Z. (2025). The implications of artificial intelligence for small and medium-sized enterprises' sustainable development in the areas of blockchain technology, supply chain resilience, and closed-loop supply chains. Sustainability, 17(1), 334.*

[8] *Chen, Y. (2025). Analysis of the Technical Application and Effectiveness of Intelligent Algorithms Empowering Regional Logistics Resource Matching. Engineering Frontiers, 1(3).*

[9] *Wu, P., Zhu, Y., Chen, W., & Du, Y. (2025). Leveraging AI to Ignite Innovation in Small and Medium Enterprises: Challenges and Opportunities: Small and Medium Enterprises use AI. Journal of Current Social Issues Studies, 2(3), 199-206.*

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

[11] *Chen, Y. (2025). Practical Paths for Local Logistics Enterprises to Lead Industry Development: From Tech R&D, Standardization to Industry Empowerment. Engineering Frontiers, 1(4).*

[12] *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.*

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

[14] *Jain, A., Kakade, K. S., & Vispute, S. A. (2025). The Role of Artificial Intelligence (AI) in the Transformation of Small-and Medium-Sized Businesses: Challenges and Opportunities. Artificial Intelligence-Enabled Businesses: How to Develop Strategies for Innovation, 209-226.*

[15] *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).*

[16] *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.*

[17] *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.*

[18] *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.*

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

[20] *Pasham, S. D. (2017). AI-Driven Cloud Cost Optimization for Small and Medium Enterprises (SMEs). The Computertech, 1-24.*

[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] *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.*

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

[24] *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).*

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

[25] *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.*

[26] *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.*