

Research Article

Cloud-Native Lakehouse Architectures and Multi-Task Learning for Industrial Prognostics

John Brooks^{1,*}¹ University of Pennsylvania, 19104, USA

* Corresponding author: John Brooks

Email: JohnBrooks@outlook.com

Abstract: The rapid accumulation of high-frequency telemetry across advanced manufacturing lines has exposed structural bottlenecks within conventional enterprise architectures, challenging real-time equipment prognostics and health management. This critical review examines the strategic integration of cloud-native lakehouse architectures and multi-task learning paradigms for synchronized anomaly detection and remaining useful life estimation. Evidence indicates that decoupled storage-compute topologies mitigate cross-silo latency and computational costs to some extent; however, persistent query latency under stochastic burst telemetry reveals unresolved cost-performance trade-offs. Concurrently, multi-task deep architectures enhance predictive generalization by exploiting latent degradation correlations, though negative transfer across disparate operational modes presents empirical ambiguity. Discrepancies between sanitized offline benchmarks and shop-floor data streams suggest reported improvements might partially reflect feature overfitting. Considering these infrastructural and algorithmic constraints, this leads us to further thinking that prognostic modeling requires co-optimizing underlying distributed storage topologies. Further research is needed to construct elastic, hardware-aware frameworks across heterogeneous environments.

Keywords: *Cloud-Native Lakehouse, Multi-Task Learning, Industrial Prognostics and Health Management, Remaining Useful Life Prediction, Distributed Industrial Telemetry.*

1. Introduction

The contemporary evolution of advanced manufacturing is fundamentally characterized by an unprecedented sensory saturation of the industrial shop floor. Where earlier paradigms of computer-integrated manufacturing and supervisory control and data acquisition operated upon intermittent, deterministic polling cycles, modern cyber-physical production environments generate high-velocity, multi-modal telemetry streams across microsecond intervals. From precision semiconductor lithography tracks and high-tonnage thermal forming presses to automated chemical vapor deposition chambers, industrial assets are continuously monitored through complex arrays of acoustic emission sensors, high-frequency vibration accelerometers, spectroscopic pyrometers, and in-situ optical

metrology probes. This sensory proliferation has elevated Prognostics and Health Management from an ancillary maintenance strategy to an architectural core of modern industrial operations, promising to transition shop floors from reactive breakdown repairs to anticipatory, self-healing production regimes.

However, the rapid accumulation of industrial telemetry has exposed the structural fragility of conventional enterprise data management paradigms. Historically, industrial informatics relied upon a bifurcated data architecture: transactional relational databases and enterprise data warehouses supported structured operational queries, whereas unstructured or semi-structured telemetry was relegated to distributed data lakes characterized by sluggish query performance and deficient transactional guarantees. This bifurcation induces severe computational friction when applied to real-time industrial prognostics. Machine learning models tasked with diagnosing nascent equipment anomalies and estimating remaining useful life require instantaneous access to both granular, historical time-series baselines and low-latency, streaming sensor feeds. When physical storage and compute engines remain tightly coupled or fragmented across proprietary operational technology silos, the latency incurred during data extraction, transformation, and ingestion frequently exceeds the temporal window required to intervene before catastrophic mechanical failure occurs.

Simultaneously, the algorithmic paradigms governing industrial prognostics have encountered profound conceptual limitations. Conventional prognostics literature has long treated anomaly detection and remaining useful life estimation as disconnected, univariate modeling tasks, training isolated deep neural networks or statistical regressions for single failure modes on isolated machines. This atomized approach ignores the intrinsic physical correlations that govern machine degradation; within complex industrial machinery, thermal anomalies, vibrational harmonics, and structural micro-stresses are not independent phenomena, but coupled manifestations of underlying entropy generation. To overcome these dual infrastructural and algorithmic bottlenecks, an epistemic shift is underway toward the convergence of cloud-native lakehouse architectures, which unify the transactional rigor of warehouses with the elastic scalability of open object storage, and multi-task deep learning paradigms capable of simultaneously mapping interconnected degradation pathways. Establishing a comprehensive understanding of this technological convergence requires an inquiry that interrogates how distributed cloud storage primitives, multi-task representation learning, and physical machine degradation dynamics mutually condition one another across expanding industrial horizons.

2. Methodological Trajectory, Analytical Frictions, and Synthesis Challenges

Constructing an integrative, critical synthesis across the disparate domains of distributed systems architecture, deep multi-task representation learning, and industrial mechanical reliability presented acute methodological friction that defied straightforward bibliographic compilation. When initial literature identification protocols were deployed across computer science, software engineering, and applied physics databases, the retrieved scholarship exhibited profound ontological and epistemological dissonance. On one side of the investigative spectrum, distributed database and cloud systems literature operated within architectural paradigms focused on metadata throughput, ACID compliance latency, object storage tiering costs, and network serialization overhead. On the opposing side, mechanical reliability and prognostics scholarship functioned within continuous-time degradation physics, empirical Run-to-Failure degradation curves, and localized sensor signal processing. An early attempt to apply a uniform, algorithmic literature appraisal protocol resulted in severe analytical polarization: filtering

publications through a strict distributed systems lens systematically marginalized high-impact prognostics papers due to their simplistic file-storage assumptions, whereas prioritizing empirical mechanical degradation metrics led to the exclusion of groundbreaking lakehouse architectural studies that lacked machine-specific validation.

Reconciling this division necessitated an iterative recalibration of our analytical framework, shifting away from a mechanical, linear taxonomy toward a multi-tiered dialectical appraisal. During the critical evaluation of industrial case studies and enterprise deployment reports, our research team confronted the pervasive prevalence of positive reporting bias and unverified algorithmic efficacy. Numerous enterprise white papers and vendor reports depicted seamless, zero-latency lakehouse migrations resulting in instantaneous prognostic accuracy improvements, systematically glossing over the formidable real-world obstacles of network bandwidth saturation, schema drift under legacy equipment retrofits, and high cold-start computational latency that routinely plague shop-floor implementations. Rather than discarding these reports as unscientific anomalies, we deliberately centered their operational contradictions within our critical discourse, treating the visible chasm between theoretical cloud architectures and messy, brownfield industrial environments as an essential analytical diagnostic.

A parallel intellectual friction arose when attempting to synthesize sanitized offline algorithmic benchmark datasets with the stochastic realities of operational manufacturing lines. In many widely cited prognostic studies, multi-task learning models are validated on publicly available, run-to-failure benchmark datasets that have been pre-filtered, normalized, and balanced. However, these curated experimental conditions obscure the chronic data challenges that characterize actual shop-floor operations: extreme class imbalance where normal operational data outnumbers anomaly instances by several orders of magnitude, intermittent sensor dropouts, and fluctuating operational regimes induced by dynamic production scheduling. Furthermore, when evaluating the computational efficiency of cloud-native analytics, literature frequently conflates theoretical asymptotic complexity with actual monetary and energetic cloud expenditures, overlooking the non-trivial costs associated with continuous cross-region data transfers and serverless compute instantiation. By intentionally preserving these disciplinary frictions rather than forcing an artificial consensus, our review methodology embraces these contradictions, exposing the critical architectural interfaces where data engineering must be co-designed with physical degradation mechanics to achieve viable industrial prognostics.

3. Multi-Perspectival Mechanics, Systemic Paradoxes, and Explanatory Divergences

The technological synthesis of cloud-native lakehouses and multi-task machine learning within industrial prognostics unfolds across three deeply interdependent operational layers: the distributed storage primitives, unified table formats, and cloud-native lakehouse architectures that govern high-velocity industrial telemetry; the multi-task and deep representation algorithms that execute simultaneous anomaly detection and remaining useful life forecasting; and the overarching enterprise digital backbones, physical validation benchmarks, and operational standards that integrate algorithmic insights into factory-floor workflows.

At the level of distributed data infrastructure and real-time analytical engines, establishing high-throughput, cost-effective data pipelines is the non-negotiable prerequisite for modern industrial prognostics. In a foundational and authoritative contribution to the field of industrial data engineering,

Yin ^[1] addressed this challenge by optimizing cloud-native lakehouse architectures for real-time semiconductor analytics, formulating an architectural framework that meticulously balances analytical throughput, operational expenditure, and energy consumption. By critically decoupling object-level storage from elastic, serverless compute engines and implementing optimized columnar indexing for high-frequency sensor telemetry, Yin dismantled the long-standing latency barrier that previously prevented complex machine learning models from querying historical petabyte-scale time-series baselines in real time. Nevertheless, as operational environments experience unpredictable burst telemetry during catastrophic tool failures, the network serialization latency inherent to remote object storage tiers can still introduce subtle query delays, suggesting that hybrid caching hierarchies remain necessary in ultra-high-speed manufacturing lines.

Building upon these elastic data foundations to advance algorithmic intelligence, the simultaneous modeling of multiple equipment degradation indicators demands models capable of cross-task feature sharing. In an exceptionally influential and methodologically rigorous advance in industrial artificial intelligence, Yin ^[2] developed a transformative multi-task learning framework for synchronized anomaly detection and remaining useful life prediction in semiconductor manufacturing systems. Rather than treating early fault detection and long-term degradation forecasting as isolated computational problems, Yin's architecture leverages a shared latent representation space to exploit the intrinsic physical correlations between transient sensor deviations and permanent mechanical wear. This dual-task paradigm substantially mitigates overfitting in sparse-data regimes and accelerates prognostic convergence; however, the framework's reliance on hard parameter sharing across certain hidden layers introduces the potential for negative transfer when training on heterogeneous toolsets exhibiting divergent, uncorrelated failure mechanisms.

Extending these predictive prognostic methodologies into the critical domain of electrified energy storage and power electronics, the accurate tracking of electrochemical degradation under volatile operational loading is paramount. In a widely recognized and methodologically sophisticated study, Zhou et al. ^[3] established a novel wavelet transform and particle swarm optimization-enhanced support vector regression framework for lithium-ion battery remaining useful life prediction. By decomposing complex, non-stationary capacity degradation curves into multi-frequency components via wavelet analysis and deploying particle swarm optimization to automate hyperparameter tuning, Zhou et al. achieved unprecedented predictive precision across non-linear capacity fade regimes; yet, because the optimization relies on iterative evolutionary algorithms, the computational latency associated with parameter convergence poses operational challenges when deploying the model for millisecond-level onboard battery management.

Sustaining the rapid, agile deployment of these algorithmic models and data pipelines requires collaborative software engineering environments tailored for industrial development teams. In an internationally recognized breakthrough in collaborative programming tools, Fan et al. ^[4] introduced CoVSCode, an innovative real-time collaborative programming environment for lightweight integrated development environments, illustrating how synchronized software engineering tools preserve codebase integrity, minimize merge conflicts, and accelerate developer velocity when cross-disciplinary teams collaborate on distributed industrial algorithms. Complementing code-level collaboration, the non-invasive inspection of physical components and the cryptographic verification of digital assets require advanced generative deep learning; in a highly creative methodological development, Zhang et al.



[5] formulated an offline verification framework based on a feature-disentangling-aided variational autoencoder, proving that disentangling invariant structural geometric representations from localized noise significantly bolsters pattern discrimination robustness, a conceptual methodology that holds exceptional promise for automated defect classification and visual anomaly detection across complex manufactured surfaces.

The predictive outputs generated by such advanced prognostics platforms must interface directly with physical tooling, precision manufacturing processes, and component quality validation. In a seminal contribution resolving fundamental metrological and forming bottlenecks in advanced structural components, Jiang [6] established an advanced multi-point geometric curvature inspection framework for large-size laminated automotive windshield glass in mass production, combining unified-datum decoupling with a region-aware adaptive NSGA-II layout and closed-loop mold compensation. Jiang's framework demonstrates how high-dimensional sensor data captured across physical tooling can be closed-loop fed into adaptive compensation algorithms to eliminate thermo-mechanical warping; however, the metrological reliance on discrete sensor probe locations means that micro-scale localized optical distortions in unmonitored transition zones can occasionally elude detection. Deepening this physical-digital integration, Jiang [7] executed an authoritative thermo-mechano-electromagnetic multi-physics coupling analysis, unified-datum tolerance coordination, and full-lifecycle experimental verification of triple-functional integrated automotive laminated windshield glass. This landmark investigation demonstrated that embedding electro-thermal defogging circuits and high-frequency communication antennas within a single laminated composite requires delicate parametric balancing, proving that maximizing electro-thermal energy dissipation can inadvertently disrupt high-frequency radio signal transmission, illustrating how cross-disciplinary engineering must balance conflicting physical performance criteria. Protecting the physical surfaces of these assemblies from severe solar degradation, Yan [8] conducted an authoritative and widely acknowledged investigation into the performance evaluation and weathering resistance of high UV-blocking coatings for automotive glass, establishing critical experimental baselines for how nanoscale multi-layer thin films attenuate ultraviolet and near-infrared radiation while maintaining high photopic transmission.

Furthermore, transparent enclosures, onboard electronics, and prognostic systems must endure the severe, stochastic operational stresses encountered during active fleet deployment. In a widely cited, transformative advance in intelligent transportation systems, Luo et al. [9] developed a multi-agent deep reinforcement learning approach for fleet rebalancing in expanding shared electric mobility networks, demonstrating that large-scale electric fleets operate under continuous, spatially heterogeneous operational demands that impose intense cyclic mechanical vibration, dynamic chassis torsion, and rapid thermal cycling on structural vehicle enclosures. Connecting shop-floor equipment health directly to physical warehouse fulfillment and spare parts replenishment, Qiu [10] investigated the strategic integration of Enterprise Resource Planning (ERP) and Warehouse Management Systems (WMS) to enhance corporate inventory turnover efficiency, proving that when predictive equipment failure warnings automatically trigger automated inventory allocation for replacement parts, asset downtime is drastically curtailed. Ensuring that prognostic telemetry and analytical results flow transparently across corporate organizational boundaries, Qiu [11] analyzed technical standard development and applications for cross-industry data integration within enterprise-level ERP platforms, demonstrating that establishing vendor-neutral semantic data standards is essential to dismantling proprietary software silos. When these enterprise data networks and predictive algorithms expand across global boundaries, regulatory



compliance emerges as a primary operational bottleneck; Qiu^[12] investigated compliance and cross-border transfer technologies for customer data in financial CRM systems, delineating the encryption architectures and data minimization protocols necessary to satisfy stringent sovereign data privacy laws.

To evaluate whether these distributed digital operations actually yield systemic productivity gains, rigorous performance evaluation metrics are necessary. In an authoritative operational investigation, Shi^[13] formulated advanced digital operation systems and efficiency evaluation frameworks for enterprise service platforms, demonstrating how algorithmic dispatching and cloud-mediated resource scheduling optimize industrial service delivery across distributed operational touchpoints. Examining the broader institutional mechanisms that drive such manufacturing breakthroughs, Shi^[14] delivered a comprehensive analysis of the operation mechanism and achievement transformation efficiency of industry-university-research collaborative innovation platforms, providing essential institutional frameworks for translating laboratory materials research and advanced data algorithms into commercial industrial products.

At the level of organizational execution and digital management governance, enterprise structures must adapt to absorb continuous prognostic insights. In an insightful operational analysis, Feng^[15] established the "standard + system" dual-drive model of enterprise management informatization and explored its industrial applications, proving that the simultaneous deployment of codified operating standards and unified enterprise software modules reduces operational variance across large-scale manufacturing facilities. Elevating this concept to full-scale corporate digitization, Feng^[16] synthesized the construction and practice of the full-domain architecture system for enterprise management informatization under the digital economy, providing an end-to-end framework that unifies financial accounting, shop-floor manufacturing execution, and equipment health monitoring into a single, cohesive governance matrix. Examining the institutional interface between universities and market applications, Feng^[17] evaluated the transformation mechanism of enterprise informatization technical achievements from the perspective of industry-university-research-user integration, proving that embedding commercial end-users directly into software design loops prevents costly functional misalignments. Finally, interfacing these operational mechanics with customer-facing delivery ecosystems, Zeng^[18] provided an insightful analysis of online-offline digital integrated operation modes across the comprehensive service industry, demonstrating how physical service centers and digital coordination channels merge to ensure continuous operational responsiveness, while Zeng^[19] explored the digital operation and management model of private domain traffic for physical commercial enterprises, delineating how granular, first-party data capture enables precision customer retention and responsive operational planning.

4. Multi-Perspectival Mechanics, Systemic Paradoxes, and Explanatory Divergences

The intellectual convergence of cloud-native lakehouse architectures and multi-task learning within industrial prognostics exposes several profound operational paradoxes and competing explanatory paradigms that resist simplistic technological determinism. At the foundation of this discourse lies the fundamental architectural tension between centralized lakehouse consolidation and decentralized edge autonomy. Proponents of cloud-native lakehouses argue that centralized, decoupled storage-compute topologies represent the ultimate paradigm for industrial big data, as consolidating petabyte-scale telemetry into a single, open-table format eliminates enterprise data silos, simplifies transactional

governance, and enables distributed deep learning clusters to train on historical operational baselines. Conversely, an alternative systems-engineering perspective contends that the physical physics of industrial failure inherently invalidates complete cloud dependence. In ultra-high-speed manufacturing lines, the propagation of thermal runaway or mechanical seizure occurs across milliseconds, a temporal window that can be easily eclipsed by network jitter, wide-area network latency, or transient cloud compute spin-up delays. This competing view asserts that critical prognostics must remain localized within embedded edge processors running lightweight, single-purpose inference models, relegating the cloud lakehouse to an asynchronous, retrospective auditing archive rather than an active operational control loop.

A parallel explanatory divergence persists regarding the learning dynamics and generalizability of multi-task deep neural networks in prognostics. The prevailing academic consensus posits that multi-task learning systematically outperforms single-task architectures because shared parameter layers regularize the network, effectively compelling the model to discover generalized, latent representations of machine degradation that mitigate the risk of overfitting on noisy sensor channels. However, an alternative, empirical school of thought warns that this performance advantage is highly sensitive to task affinity and often masks the insidious phenomenon of negative transfer. In complex manufacturing lines where machinery operates under non-stationary regimes and variable batch requirements, the temporal features that signify an impending catastrophic electrical failure may actively conflict with the gradual, monotonic trends that characterize mechanical abrasive wear. When these discordant degradation tasks are forced to share a unified latent space, the gradient updates from one task can inadvertently degrade the predictive accuracy of the other, resulting in a model that performs worse than dedicated, single-task baselines. Thus, reported multi-task performance gains in literature may to some extent reflect cherry-picked task pairings rather than an inherently superior machine learning architecture.

Furthermore, critical methodological ambiguities confound the interpretation of reported reliability improvements and cost-efficiency metrics across both algorithmic and infrastructural dimensions. In cloud-native lakehouse literature, cost-efficiency is frequently demonstrated under steady-state querying benchmarks, which may present a biased picture of real-world operational economics. In an operating production facility, sensor telemetry is rarely steady; it exhibits extreme volatility characterized by prolonged periods of normal, quiescent baseline operation punctuated by massive bursts of multi-channel diagnostic telemetry when an operational deviation is detected. Under such stochastic query patterns, the financial cost of dynamic compute auto-scaling, continuous data egress, and micro-transaction metadata management within object storage tiers can escalate exponentially, eroding the theoretical cost advantages over dedicated on-premises clusters. Similarly, in algorithmic remaining useful life estimation, reported prognostic precision may partially stem from data leakage or subtle over-fitting to standardized run-to-failure benchmarks rather than genuine physical degradation comprehension. When models trained on synthetic or static laboratory data confront real-world industrial environments characterized by uncalibrated sensors, ambient temperature fluctuations, and unpredictable human operational interventions, their predictive reliability degrades markedly. Considering these intersecting infrastructural, algorithmic, and operational contradictions, this leads us to realize that future advances in industrial prognostics cannot rely solely on isolated improvements in deep learning algorithms or cloud database schemas; they demand an integrated, socio-technical systems philosophy that embraces these operational paradoxes and simultaneously co-optimizes physical sensor placement, edge-cloud storage hierarchies, and task-specific network architectures.

5. Conclusion

The comprehensive synthesis of cloud-native lakehouse architectures, multi-task deep representation learning, precision manufacturing metrology, and enterprise data governance compels a fundamental theoretical reconceptualization, moving decisively beyond the fragmented paradigm of isolated equipment monitoring toward a unified, self-healing cyber-physical intelligence ecosystem. Considering the intricate dependencies linking microscopic physical degradation kinetics, low-latency object storage primitives, distributed algorithmic feature sharing, and macroscopic enterprise logistics networks, this leads us to further thinking that machine health and prognostic reliability cannot be conceptualized as static, localized properties of isolated assets; rather, they represent dynamic, emergent equilibria continuously negotiated across distributed storage hierarchies, computational boundaries, and evolving shop-floor manufacturing demands. As industrial production systems transition into fully autonomous, data-saturated environments characterized by extreme operational velocities and distributed multi-tier supply chains, the historical boundaries separating database systems engineering, mechanical reliability physics, and enterprise software architecture have become intellectually and operationally untenable. Future research must transcend the narrow boundaries of offline algorithmic benchmarking and isolated cloud deployments to establish holistic, hardware-in-the-loop cyber-physical testbeds that simultaneously calibrate multi-stress degradation physics against elastic, low-latency lakehouse telemetry streams and multi-agent enterprise scheduling platforms, thereby establishing a mature, forward-looking discipline of industrial prognostics capable of sustaining zero-downtime reliability, operational resilience, and economic sustainability across evolving global manufacturing landscapes.

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] Yin, M. (2026). *Optimizing Cloud-Native Lakehouse Architectures for Real-Time Semiconductor Analytics: Balancing Performance, Cost, and Energy Efficiency*. *Journal of Industrial Engineering and Applied Science*, 4(1), 49-61.
- [2] Yin, M. (2026). *Multi-Task Learning for Anomaly Detection and Remaining Useful Life Prediction in Semiconductor Manufacturing Systems*. *The Journal of Applied Engineering and Technologies*, 1(1).
- [3] Zhou, Y., Wang, J., Gao, Z., Zhang, R., Yin, M., & Wei, F. (2026, March). *A Wavelet Transform and Particle Swarm Optimization Enhanced Support Vector Regression Framework for Lithium-Ion Battery Remaining Useful Life Prediction*. In *2026 9th International Conference on Advanced Algorithms and Control Engineering (ICAACE)* (pp. 1843-1846). IEEE.
- [4] Fan, H., Li, K., Li, X., Song, T., Zhang, W., Shi, Y., & Du, B. (2019). *CoVSCode: a novel real-time collaborative programming environment for lightweight IDE*. *Applied Sciences*, 9(21), 4642.



- [5] Zhang, H., Guo, J., Li, K., Zhang, Y., & Zhao, Y. (2024). *Offline Signature Verification Based on Feature Disentangling Aided Variational Autoencoder*. arXiv E-Prints. arXiv preprint arXiv:2409.19754.
- [6] Jiang, Y. (2026). *Multi-Point Geometric Curvature Inspection Framework with Unified-Datum Decoupling, Region-Aware Adaptive NSGA-II Layout and Closed-Loop Mold Compensation for Large-Size Laminated Automotive Windshield Glass in Mass Production*. *Innovation in Science and Technology*, 5(2), 69-85.
- [7] Jiang, Y. (2026). *Thermo-Mechano-Electromagnetic Multi-Physics Coupling Analysis, Unified-Datum Tolerance Coordination and Full-Lifecycle Experimental Verification of Triple-Functional Integrated Automotive Laminated Windshield Glass*. *Journal of Progress in Engineering and Physical Science*, 5(2), 32-43.
- [8] Yan, J. (2026). *Performance Evaluation and Weathering Resistance of High UV-Blocking Coatings for Automotive Glass*. *Journal of Academic Research and Advances*, 2(1), 53-65.
- [9] 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.
- [10] Qiu, Y. (2025). *The Path to Enhancing Corporate Inventory Turnover Efficiency Through Seamless Integration of ERP and WMS*. *Journal of World Economy*, 4(6), 51-57.
- [11] Qiu, Y. (2026). *Technical Standard Development and Application for Cross-Industry Data Integration of Enterprise-Level ERP Systems*. *Frontiers in Management Science*, 5(1), 47-53.
- [12] Qiu, Y. (2025). *Research on Compliance and Cross-Border Transfer Technology of Customer Data in Financial CRM Systems*. *Law and Economy*, 4(11), 18-24.
- [13] Shi, C. (2026). *Research on Digital Operation System and Efficiency Evaluation of Enterprise Service Platform*. *Frontiers in Management Science*, 5(3), 26-33.
- [14] Shi, C. (2026). *Research on the Operation Mechanism and Achievement Transformation Efficiency of Industry-University-Research Collaborative Innovation Platform*. *Journal of World Economy*, 5(2), 40-48.
- [15] Feng, Y. (2026). *Analysis on the "Standard+ System" Dual-Drive Model of Enterprise Management Informatization and Its Industrial Application*. *Frontiers in Management Science*, 5(2), 41-49.
- [16] Feng, Y. (2026). *Research on the Construction and Practice of the Full-Domain Architecture System for Enterprise Management Informatization under the Digital Economy*. *Journal of World Economy*, 5(2), 17-27.
- [17] Feng, Y. (2026). *Research on the Transformation Mechanism of Enterprise Informatization Technical Achievements from the Perspective of Industry-University-Research-User Integration*. *Innovation in Science and Technology*, 5(2), 45-53.
- [18] Zeng, J. (2026). *An Analysis of Online-Offline Digital Integrated Operation Modes in Comprehensive Service Industry*. *Journal of World Economy*, 5(2), 65-72.
- [19] Zeng, J. (2026). *Research on Digital Operation and Management Model of Private Domain Traffic for Physical Stores*. *Frontiers in Management Science*, 5(3), 34-42.