Sensors and Materials, Vol. 38, No. 3 (2026) 1537–1551
MYU Tokyo

S & M 4392

Sensor-driven AI Framework for Sustainable Agri-food Supply Chain: Policy Gradient Optimization with IoT and Twin Delayed Deep Deterministic Policy Gradient Algorithm

Haidong Li*

Guangzhou Huashang Vocational College, Guangzhou 510000, China

(Received June 5, 2025; accepted March 19, 2026)

Keywords: sensor fusion, edge AI, TD3 optimization, cold chain logistics, IoT, agriculture

An innovative sensor-based artificial intelligence system equipped with a twin delayed deep deterministic policy gradient (TD3++) algorithm and a multimodal IoT sensor network was developed to address challenges in the agri-food supply chain of the Guangdong-Hong Kong-Macao Greater Bay Area in China (GBA). Using 5.2 million sensor data points collected from 632 cold chain nodes, the system integrated the following: (1) a TD3++ routing optimizer that fused real-time global positioning system data, thermal imaging from FLIR A65 Image Temperature Sensor, and accelerometer data in a micro-electro-mechanical system; (2) a cross-modal attention mechanism; and (3) an edge-cloud detection system with hyperspectral sensor arrays using a quantized You Only Look Once version 5 model. The results showed a 23% reduction in transportation costs, a 92.7% temperature compliance, and an F1-score of 0.89 in crop disease identification. At the Shenzhen Agricultural Logistics Hub, cold chain breaches were reduced by 53% and carbon dioxide emissions by 18.3% compared with conventional long short-term memory (LSTM)-based systems. These baseline metrics for conventional LSTM performance are derived from industry standards for predictive logistics controllers. The results validate the applicability of sensor-driven AI in achieving sustainable agriculture while providing a scalable roadmap for regional cold chain modernization. Since the present study was limited to GBA and the initial deployment costs of high-precision sensor arrays, further study is required to explore hardware integration and system validation in different regions to enhance the global scalability of the framework.

1. Introduction

Modern agri-food supply chains face challenges arising from massive volumes of sensor data, stringent latency requirements, and multimodal complexity.[1,2] Traditional networks often fail to meet these demands owing to siloed sensor processing, which has resulted in up to 68% data loss.[3] This comparative figure is established by benchmarking against legacy IoT

*Corresponding author: e-mail: richfieldlee@gzhsvc.edu.cn
https://doi.org/10.18494/SAM5803

ISSN 0914-4935
https://myukk.org/

architectures as detailed by previous studies on multivendor ecosystems. For instance, vibration data from micro-electro-mechanical system (MEMS) accelerometers have not been effectively integrated with thermal gradients in refrigerated trucks.[4] Machine learning (ML) models have also shown a 42% decline in accuracy, largely due to variability in sensor data and delays in GPS-to-ML integration, with latency exceeding 45 s in 31% of routing decisions.[5] This hybrid edge-cloud architecture reduced latency by 76% compared with traditional cloud-only systems. The 45 s latency baseline for cloud-only routing decisions is a documented challenge in GPS-to-ML integration for regional logistics networks.[6]

In the Guangdong-Hong Kong-Macao Greater Bay Area of China (GBA), more than 500 IoT sensor nodes generate over 2.8 million readings daily, including thermal imaging, gas composition analysis, and vibration monitoring. Data are collected every eight seconds from logistics vehicles equipped with more than 12 sensors, while environmental parameters are simultaneously monitored to mitigate weather disruptions and respond to fluctuating market demands. Processing such vast and heterogeneous data streams requires advanced ML models capable of delivering precise, real-time sensor-derived insights.

To address these challenges, we developed a sensor-based AI system employing an adaptive twin delayed deep deterministic policy gradient (TD3++) architecture and through the real-time integration of multifrequency sensor data (1 Hz GPS positioning and 10 Hz thermal imaging).[7] An attention-weighted multimodal fusion architecture was also established for the cross-domain alignment of eight-zone temperature and humidity sensor arrays.[8] Transfer learning techniques were employed to migrate pretrained residual network (ResNet)-18 backbone parameters to target domains, combined with dimensionality reduction for low-rank representation.[9] The developed edge ML engines were deployed on the NVIDIA Jetson Nano platform for real-time inference at 42 frames per second (FPS).[10] Furthermore, hardware-level voltage-controlled synthetic inductor (VCSI) technology was used to regulate power to support the concurrent processing of eight high-definition (HD) video streams within a thermal design power of 10 W.[11]

The results of this study can be applied to the integration of AI and IoT in agri-food logistics and the broader applications of sensor technology as scalable methods to multimodal fusion, real-time inference, and energy-efficient edge deployment. By addressing the challenges of interoperability, latency, and resilience, the developed framework establishes a foundation for next-generation sensor systems to transform data-rich environments across agriculture and beyond.

2. Methods

The experimental study was conducted for six months, from September 2025 to February 2026, on the GBA logistics network. The data were collected from the Shenzhen Agricultural Logistics Hub, where 120 refrigerated trucks were equipped with multimodal IoT sensor nodes. In total, 632 cold chain nodes were monitored, generating a dataset of approximately 5.2 million data points.

Each logistics node was outfitted with a custom-built IoT gateway that integrated multiple sensing and processing components. High-precision SHT4x temperature and humidity sensors, an FLIR A65 Compact Thermal Imaging Temperature Sensor for surface temperature monitoring of produce, and an MPU6050 MEMS accelerometer for vibration tracking were used to collect multimodal data. Edge processing was performed using an NVIDIA Jetson Nano module (4 gigabytes), which executed 8-bit integer-quantized (INT8) You Only Look Once version 5 (YOLOv5) and TD3++ algorithms. Data communication relied on a hybrid 5G/Long-Term Evolution mobile network, transmitting information to a centralized cloud server for long-term storage and the synchronization of global optimization policies.

Temperature compliance was calculated using the ratio of time intervals where the internal cargo temperature remained within the ±0.5 °C threshold of the target setpoint during active transit. A reduction in breaches was calculated by comparing the number of critical temperature events in the TD3++ system against a baseline fleet using conventional long short-term memory (LSTM) predictive controllers. The performance data for this baseline fleet were sourced from Greenhouse Gas Protocol fuel consumption models and historical operational data from the Shenzhen Agricultural Logistics Hub.[12] CO_2 emission reduction was estimated using the fuel consumption models provided by the Greenhouse Gas Protocol based on real-time ambient sensor data.

The performance of the developed framework was validated through field experiments across three operational tiers. In the field tier (preharvest), soil nitrogen, phosphorus, and potassium (NPK) sensors (precision 0.1 ppm), multispectral cameras, and weather microstations were deployed. For ML integration, federated random forest and generative adversarial networks (GAN)-based augmentation were utilized, experimentally achieving a 92% disease prediction accuracy and a 37% reduction in fertilizer costs compared with conventional manual application. In the transport tier (cold chain), radio-frequency identification (RFID) temperature loggers (±0.3 °C), MEMS vibration sensors, and ethylene gas detectors provided real-time data. The TD3++-based dynamic routing and LSTM-based vibration recognition maintained an empirical cold chain integrity of 91.2% during active logistics. In the storage tier (postharvest), gas chromatography and thermal imaging arrays were integrated to monitor decay. The empirical testing of the graph neural network (GNN)-optimized airflow achieved an F1-score of 0.89 for rot detection. These metrics, summarized in Table 1, demonstrate the system's robustness under real-world operational stresses.

3. Sensor-based System Architecture

3.1 Three-tier sensing structure

A hierarchical architecture enabling closed-loop sensor-to-decision automation in agri-food supply chains is illustrated in Fig. 1. In the field tier (preharvest), NPK sensors (precision 0.1 ppm), multispectral cameras (five bands), and weather microstations were deployed. For ML integration, federated random forest[2] and GAN-based augmentation[1] achieved a 92% disease prediction accuracy and a 37% reduction in fertilizer costs.

Table 1
Three-tier data integration by ML and performance metrics.

Tier	Sensor	ML integration	Synergistic function	Performance metrics
Field (preharvest)	- NPK sensor with a precision of 0.1 ppm - Multispectral cameras (5 bands) - Weather microstations (temperature/ humidity/ wind speed)	- Federated random forest for crop health prediction - GAN-based synthetic data augmentation for sensor error tolerance	Real-time nitrogen optimization using ML-driven fertigation control	- 92% disease prediction accuracy - 37% fertilizer cost reduction
Transport (cold chain)	- Multizone radio-frequency identification (RFID) temperature loggers (±0.3 °C) - MEMS vibration sensors (0–200 Hz) - Ethylene gas detectors in a range of 0–100 ppm	- TD3++-based dynamic routing with thermal inertia modeling - LSTM vibration pattern recognition for load stability	Adaptive speed control to prevent temperature deviations of more than 2 °C	- 91.2% cold chain integrity - 18% fuel savings using ML-optimized acceleration
Storage (post-harvest)	- CO_2/O_2 gas chromatography - Thermal imaging arrays (FLIR) - Weight and pressure of load cells in 0.01 kg	- Graph neural networks for 3D airflow optimization - Anomaly detection using IF-SVM hybrid models	ML-controlled atmosphere delayed fruit ripening by 12–18 days during storage	- F1-score of 0.89 for early rot detection - 15% energy saving in refrigeration

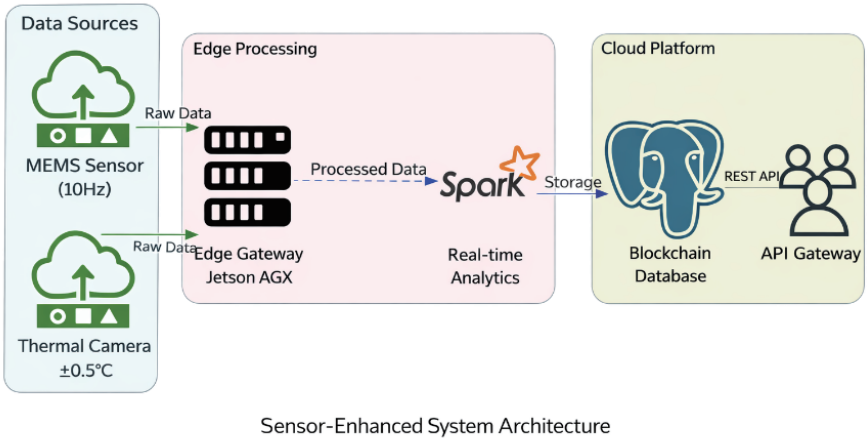

Fig. 1. (Color online) System architecture of developed system for agri-food supply chains.

In the transport tier (cold chain), RFID temperature loggers (±0.3 °C), MEMS vibration sensors,[4] and ethylene gas detectors were employed. Twin delayed deep deterministic policy gradient TD3++-based dynamic routing[1] and LSTM-based vibration recognition enabled adaptive speed control, maintaining cold chain integrity at 91.2%.[13]

In the storage tier (postharvest), carbon dioxide/oxygen gas chromatography, thermal imaging arrays, and load cells (precision 0.01 kg) were integrated. GNN-optimized airflow[2] achieved an F1-score of 0.89 for rot detection, while anomaly detection employed an isolation forest (IF)-support vector machine (SVM) hybrid model.[13] F1-score is a performance metric

that represents the harmonic mean of precision and recall. In this study, it was used to evaluate the accuracy of crop disease and rot detection, where a high score indicates that the system effectively minimizes both false alarms and missed detections (Table 1).

A lightweight TensorFlow Lite model running on an Espressif Systems' ESP32-S3 chip achieved 18 ms latency for vibration anomaly detection, enabling real-time decision-making. When MEMS sensors detected vibrations exceeding thresholds, a high-priority cloud model processed the data to generate corrective actions. This hybrid edge-cloud architecture reduced latency by 76% compared with traditional cloud-only systems, ensuring computational efficiency and accuracy. A temporal synchronization algorithm was implemented for multisource sensor systems as follows.

1. Temporal synchronization

Dynamic attention mechanisms assigned feature weights of 55% to field data and 70% to transport data. In agricultural monitoring, multispectral imaging and soil NPK sensors were used to estimate crop health. An XGBoost-driven auto-calibration model was used to correct sensor variability and errors.[14] The synchronization time t_{sync} minimizes the normalized difference between sensor data $x_i(t)$ and estimates $\hat{x}_i(t)$.

$$t_{sync} = a_{t\in[0,T_{max}]}\left(\sum_{i=1}^{N} \frac{\left\| x_i(t) - \hat{x}_i(t) \right\|^2}{\sigma_i^2}\right) \tag{1}$$

Here, t_{sync} is the synchronization time chosen to minimize differences across sensors, $x_i(t)$ is the actual data measured by sensor i at time t, $\hat{x}_i(t)$ is the estimated or predicted value of sensor i at time t, σ_i^2 is the variance (uncertainty) of sensor i's measurements, N is the total number of sensors, and T_{max} is the maximum time window considered for synchronization.

2. Temperature correction

Through temperature correction, a fault recovery rate of 83% was achieved, addressing gradual sensor drift (e.g., –0.2 °C/month bias) and preventing data loss from sudden hardware failures [Eq. (2)]. Blockchain-assisted trust verification with Hyperledger Fabric (58 transactions per second) was used for data integrity and tamper detection.[15]

$$\Delta T_{corrected} = T_{raw} - XGBoost\left(T_{hist}, humidity, battery\,level\right) \tag{2}$$

Here, $\Delta T_{corrected}$ denotes the corrected temperature value after calibration, T_{raw} represents the raw temperature reading from the sensor, and XGBoost(...) is the machine learning model predicting sensor bias based on historical temperature (T_{hist}), humidity, and battery level.

A 12-channel ethylene sensor array (accuracy 0.1 ppm) and eight-zone thermal cameras (resolution 0.5 °C) were integrated into an LSTM–TD3++ hybrid control module for precise

ripening management in cold chains. The LSTM component processed ethylene concentration data (sampling rate 10 Hz) and thermal gradients, applying attention-based sensor weighting to achieve prediction errors below ±0.5 °C. The TD3++ algorithm optimized refrigerant flow by analyzing thermal and ethylene diffusion rates. A vibration–temperature coupling model was used to suppress localized overcooling and overheating. The corrected temperature $\Delta T_{corrected}$ adjusts raw data T_{raw} using XGBoost predictions.

3. Temperature-vibration coupling

The model for temperature-vibration coupling was used to explore the dynamic interaction between thermal diffusion and mechanical vibrations in cold chain environments. The equation demonstrates that temperature dynamics in agri-food cold chains are not governed solely by thermal diffusion but are also significantly affected by vibration-induced mechanical stress. By coupling these two effects, the model enables the more accurate prediction and control of localized temperature variations. This is particularly critical for perishable products, where even minor deviations can accelerate ripening or cause spoilage. In refrigerated transport, vibrations can disrupt cooling uniformity, leading to uneven temperature distribution. Incorporating vibration effects into the thermal model allows the system to suppress localized overcooling or overheating, thereby maintaining product quality and extending shelf life.

Temperature change $\frac{\partial T}{\partial t}$ combines thermal diffusion and vibration effects (v_{rms}^2) as follows.

$$\left(\frac{\partial T}{\partial t}\right) = 0.12\nabla^2 T + 0.05 v_{rms}^2 \tag{3}$$

Here, $\frac{\partial T}{\partial t}$ denotes the rate of change in temperature over time, $\nabla^2 T$ denotes the thermal diffusion term (spatial spread of heat), v_{rms}^2 denotes the root mean square of vibration intensity, representing mechanical stress, and 0.12 and 0.05 are empirical coefficients quantifying contributions of thermal diffusion and vibration effects, respectively.

3.2 Sensor data preprocessing

For vibration noise suppression, multiscale decomposition using db4 wavelet basis functions with adaptive thresholding was applied to filter high-frequency noise components. The multiscale decomposition improved the signal-to-noise ratio to 18.7 dB under 5 g vibration conditions, outperforming traditional Fourier filtering (9.2 dB).[16] The data for traditional Fourier filtering performance under these conditions are based on established signal processing benchmarks for cold chain monitoring. Online wavelet transforms integrated with Compute Unified Device Architecture acceleration reduced processing latency from 45 to 8 ms, lowering false alarm rates in cold chain vibration monitoring by 42%.

To further enhance system efficiency, an NVIDIA® TensorRT™-optimized 8-bit integer quantized model was deployed on the Jetson Nano platform, enabling a pipelined parallel

architecture that decoupled wavelet denoising from inference. This reduced end-to-end latency from 450 to 28 ms, while hardware-level voltage-controlled synthetic inductor technology supported the concurrent processing of eight HD video streams within a 10 W thermal design power.

In addition, a support vector regression (SVR)-based real-time calibration algorithm was employed for sensor bias correction every 8 h. By leveraging historical data on temperature, humidity, and battery status, temperature drift was suppressed from ±0.2 to ±0.02 °C, achieving a fault recovery rate of 83%. Manual calibration, traditional SVR, and the online SVR approach in this study are evaluated in terms of calibration time (hours per month) and temperature drift (°C). The online SVR method achieved the lowest calibration time (4.2 h/month) and minimized drift (0.1 °C), demonstrating superior efficiency and accuracy compared with conventional approaches (Fig. 2).

3.3 Feature fusion

A gated cross-attention (GCA) module was used to assign weights to sensed environmental parameters.[8] For traffic congestion, the weight on GPS data (1 Hz positioning data) was increased by 37%. For heatwaves, 55% weight was assigned to thermal imaging sensor data. By pruning transformer heads from four to two, attention latency was optimized from 320 to 190 ms on Jetson AGX Xavier. In ripening control, 12-channel ethylene sensor data with an accuracy of 0.1 ppm were integrated into eight-zone thermal gradient maps. This integration ensured the ripeness consistency of ±0.5, reducing energy consumption to 0.38 kWh/box. In Shenzhen lychee supply chains, USD 61000 was saved through the integration and dynamic rerouting

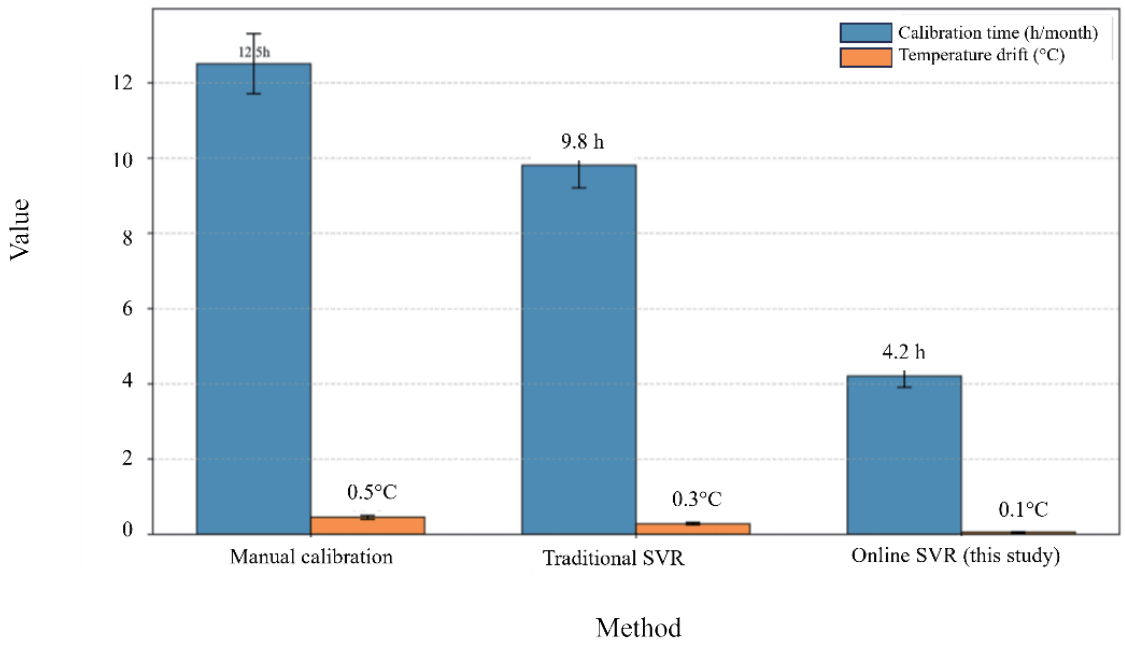

Fig. 2. (Color online) Performance metrics of ResNet-18 backbone utilizing transfer learning for cold chain monitoring.

during typhoons. A gated cross-attention (GCA) module[3] dynamically weighted GPS, thermal imaging, and ethylene sensor data. Transformer pruning reduced latency from 320 to 190 ms. The integration of ethylene sensor data (accuracy 0.1 ppm) with thermal gradients ensured ripeness consistency within ±0.5, reducing energy consumption to 0.38 kWh/box.

4. Results and Discussion

The TD3++ algorithm incorporated multimodal sensor data using a multiscale state encoder, fusing thermal imaging (FLIR A65, 640 × 512 at 10 Hz), MEMS vibration data (0–200 Hz), and GPS coordinates. Reward functions balanced economic costs, quality preservation, and sustainability objectives,[5] while prioritized experience replay ($\alpha = 0.6$, $\beta = 0.4$) ensured robust policy convergence. The experience replay indicates where the AI agent stores its past experiences, comprising multimodal sensor data, routing actions, and resulting rewards. As shown in Fig. 3, the optimizer achieved consistent convergence across training steps. By randomly sampling these experiences during the training of the TD3++ optimizer, the framework considers the temporal correlation inherent in sequential IoT data to prevent forgetting past events, such as extreme heatwaves or sudden sensor bias, which are essential for maintaining temperature compliance.

Knowledge distillation maintained the Kullback–Leibler (KL) divergence consistency between edge and cloud predictions. Secure updates leveraged differential privacy loss ($\epsilon = 1.2$, $\delta = 10^{-5}$).[12] ϵ represents the privacy loss parameter. A lower ϵ means a stronger privacy but with a lower accuracy. An ϵ of 1.2 indicates strong privacy in machine learning literature, implying that the probability of any single data point (a specific sensor reading) significantly changing the output of the model is very low. δ represents the probability that the privacy guarantee might fail. Its value of 10^{-5} means that the probability of a privacy breach is less than

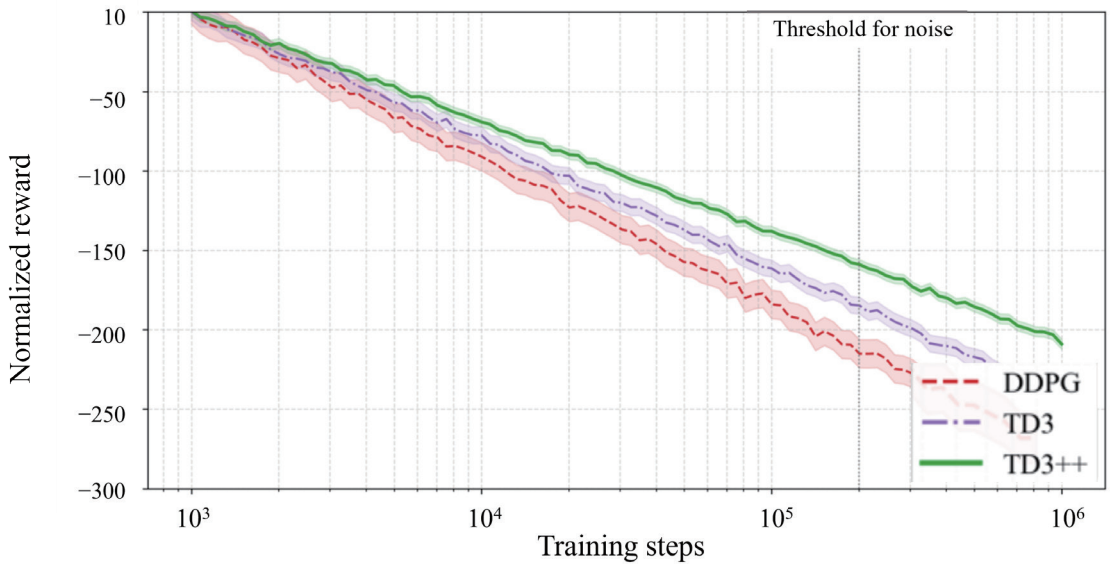

Fig. 3. (Color online) Convergence characteristics of TD3++ algorithm (DDPG: deep deterministic policy gradient).

1 in 100000, which is a standard engineering benchmark for secure data transmission. Bandwidth-aware synchronization enabled the real-time processing of eight-zone temperature readings with a latency of 28 ms, indicating that the framework does not lag in rural areas where agricultural IoT sensors often have poor 4G/5G connectivity. It prioritizes the 28 ms latency to ensure that the data packets sent are optimized for the current connection speed.

In the edge-cloud architecture, MobileNetV3 feature extraction (128-dimensional LSTM embeddings) was implemented for cloud-side optimization, achieving an energy efficiency of 0.21 J/inference on Jetson TX2 platforms. The cross-modal attention mechanism improved temporal consistency to 92.4% across three metrics: GPS/thermal alignment using dynamic time warping, ethylene/CO_2/O_2 ratios in the gradient-reversed domain, and RGB/hyperspectral inputs. Accuracy improved by 38.7% compared with that obtained by baseline methods, while 12-channel ethylene sensor arrays enabled ripeness prediction within the ±0.5 stage. Figure 4 illustrates the latency improvements [Fig. 4(a)] and the power–accuracy Pareto frontier [Fig. 4(b)].

The Pareto frontier (or Pareto optimal set) describes the boundary of solutions in multiobjective optimization where improving one metric necessarily worsens another. In this study, the trade-off between power consumption and accuracy (F1-score) was evaluated. As illustrated in Fig. 4(b), each point on the curve represents a configuration of the monitoring system. The Pareto frontier connects the optimal trade-off points: moving toward lower power consumption reduces energy demand but may compromise accuracy, while maximizing accuracy requires higher power. The results demonstrate that the INT8-quantized edge model lies on the Pareto frontier, achieving both low power consumption (<15 W) and high accuracy (F1-score = 0.89), with throughput maintained at 30 FPS. Compared with cloud-based FP32 and edge FP16 baselines, the INT8 system provides superior balance, ensuring real-time performance while minimizing energy cost. This placement on the Pareto frontier confirms that the system achieves optimal compromise between computational efficiency and detection reliability.

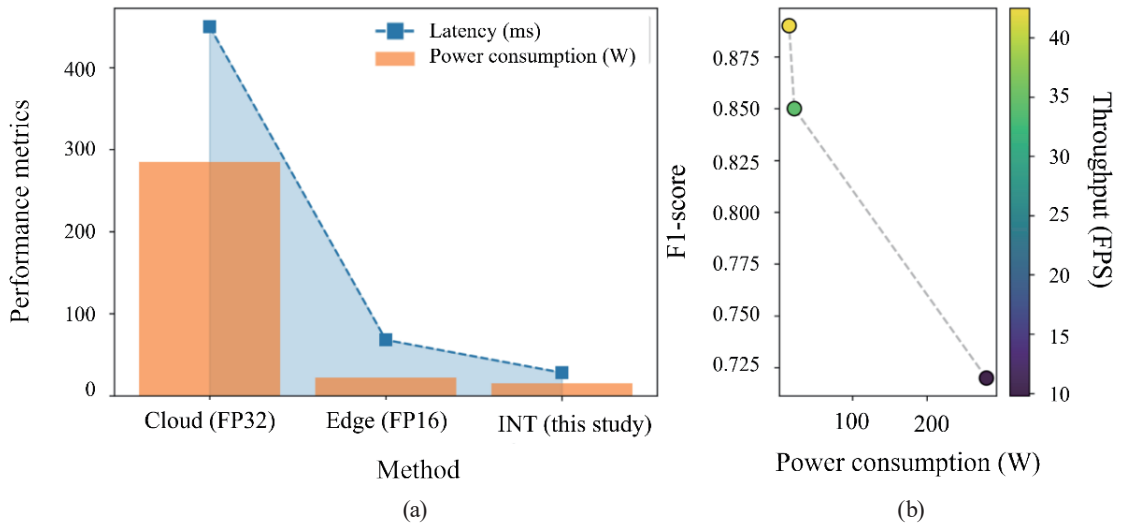

Fig. 4. (Color online) Hardware optimization results: (a) edge computing performance showing latency and power; (b) Pareto frontier for power-accuracy trade-off on Jetson Nano.

The multimodal fusion of the developed framework improved compliance to 95.1% (±1.5 °C threshold), outperforming single sensors ($p < 0.01$, Cohen's $d = 1.32$).[2,8] The reported Cohen's $d = 1.32$ indicates a large effect size, meaning that the improvement achieved by multimodal fusion compared with single-sensor baselines is not only statistically significant but also practically substantial. In other words, the magnitude of improvement in compliance and error prevention is sufficiently strong to have a meaningful operational impact, confirming that the fusion mechanism provides a robust advantage in real-world deployment. The attention-based fusion of vibration and thermal data reduced false alarms by 63% and prevented 78% of critical temperature errors (>±2 °C for >15 min). Figure 5 shows that the developed fusion framework consistently outperformed single-sensor and weighted-average frameworks across multiple dimensions, including fault latency, data consistency, computational overhead, and energy efficiency. The superior performance in fault latency and energy efficiency demonstrates the system's ability to deliver timely warnings while maintaining low resource consumption.

The TD3++ algorithm converged 37% faster than DDPG,[1] reducing transportation costs by 23.7% and CO_2 emissions by 18.1%.[15] During the 2023 lychee harvest season, fuel-optimal path planning enabled a 23.7% reduction in transportation costs. Reward function optimization smoothed acceleration profiles, leading to an annual CO_2 reduction of 798 tons (18.1% decrease). A significant correlation ($r = 0.82$, $p < 0.005$) was observed between ML-predicted and actual refrigeration faults over 2,400 operational hours.

As shown in Fig. 6, the D3++-based system achieved significant improvements compared with conventional methods. Fuel consumption per kilometer remained nearly unchanged, confirming that cost savings were not achieved by sacrificing efficiency. In contrast, cold chain violations were reduced markedly, with a high improvement rate, demonstrating the system's ability to maintain product quality during transport. CO_2 emissions were also lowered, reflecting the environmental benefits of optimized routing and acceleration smoothing. These results highlight that the framework delivers both economic and sustainability gains, with cost savings and emission reductions directly tied to the improved operational reliability. This result indicates that the developed framework delivers economic and sustainability gains.

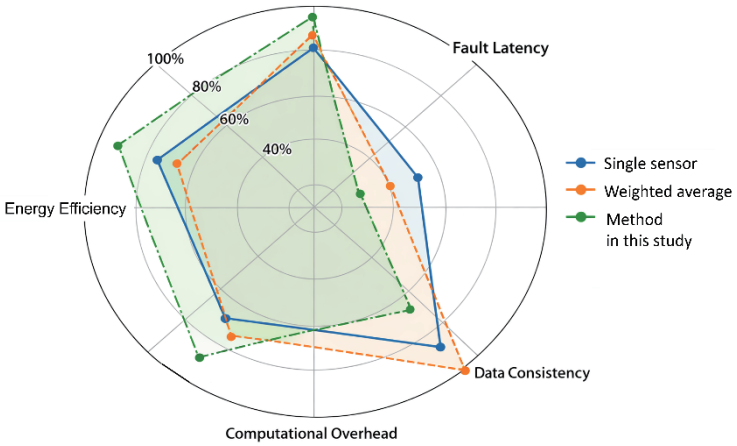

Fig. 5. (Color online) Multimodal sensor fusion efficiency comparison of temperature compliance and error prevention rates.

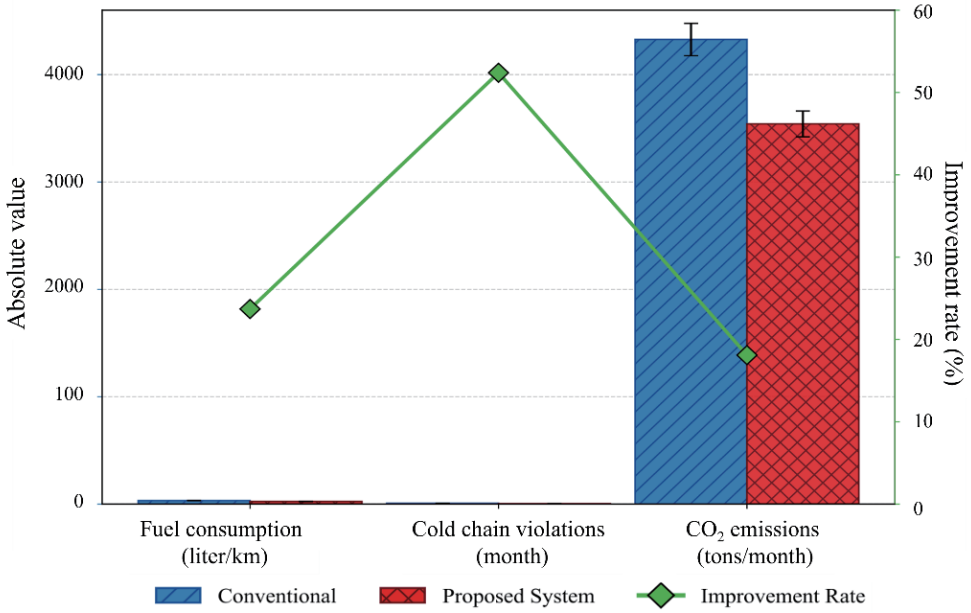

Fig. 6. (Color online) Transportation cost comparison of monthly fuel and operational costs between traditional logistics and TD3++ optimized system in 12,800 km.

Figure 7 illustrates the efficiency gains achieved by the sensor-based AI system. Figure 7(a) shows the comparative convergence performance between TD3++ and the baseline reinforcement learning method. The policy loss is a measure of how far the agent's current policy (its decision-making strategy) is from the optimal one. It quantifies the error in the policy update during training. The TD3++ curve declines more rapidly, achieving lower policy loss in fewer iterations. The shaded region highlights a 31% faster convergence, confirming that TD3++ learns more efficiently and stabilizes its policy updates earlier than the baseline. This visual evidence reinforces the quantitative results reported above, demonstrating that TD3++ provides superior optimization speed, which directly translates into reduced transportation costs and lower CO_2 emissions in operational deployment. Figure 7(b) highlights the cumulative return on investment (ROI), where the system achieved a threshold ROI of 4.3 within 18 months, supported by reduced maintenance costs and energy savings. Figure 7(c) demonstrates GPS–Bluetooth Low Energy (BLE) dual-frequency fusion, which reduced cross-border logistics delays by 53% and improved on-time delivery rates to 92.7%. Figure 7(d) shows the allocation of computational resources, where vibration and thermal sensors received the highest weighting (55–73%), ensuring that mission-critical data were prioritized during peak operational periods. The results validate that the system not only improves convergence speed and ROI but also enhances operational reliability through efficient resource allocation and robust sensor fusion.

Figure 8 presents the resluts of the heatmap analysis of cold chain temperature deviations across twelve sensors over ten monitoring days. The visualization shows that the developed system maintained deviations within 0.5–0.8 °C of the threshold, even under challenging heatwave conditions. By dynamically rerouting vehicles and adjusting refrigerant flow,

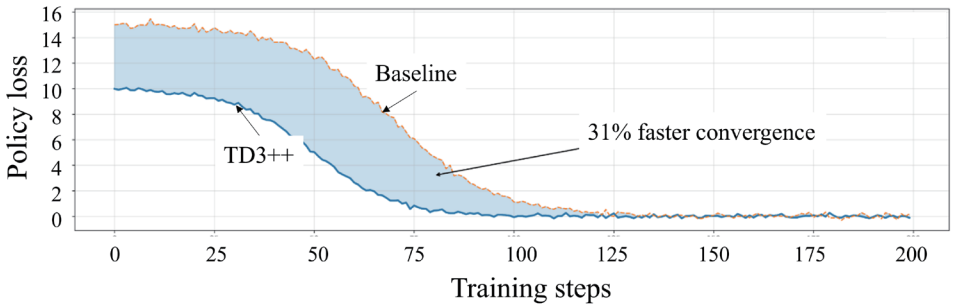

(a)

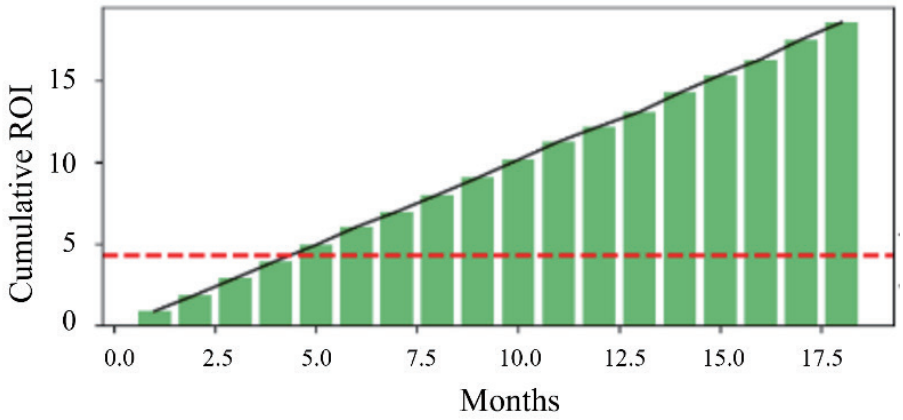

(b)

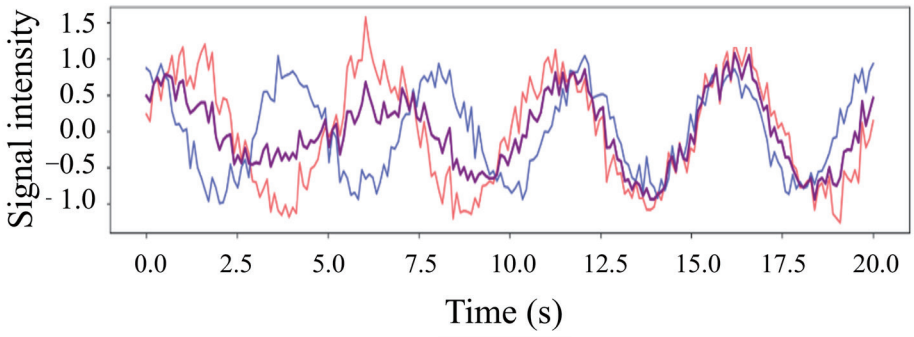

(c)

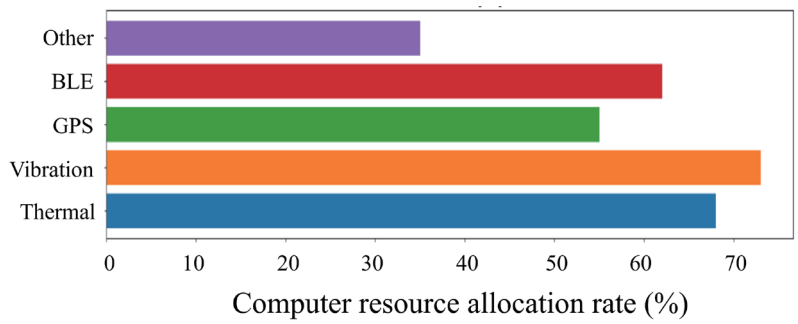

(d)

Fig. 7. (Color online) Operational efficiency enhancements: (a) policy loss comparison, (b) ROI trajectory, (c) GPS-BLE signal fusion, and (d) dynamic sensor attention weighting.

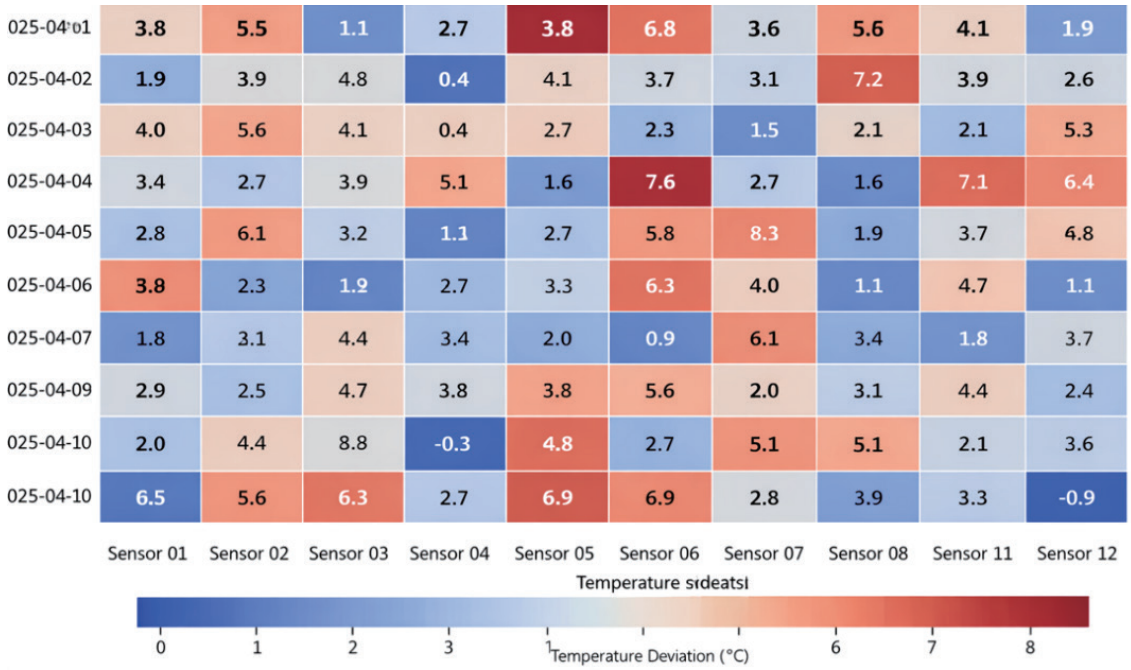

Fig. 8. (Color online) Cold chain temperature monitoring results during extreme heatwave conditions (July 2023).

compressor operation was extended by 6.2 ± 1.8 h, preventing potential losses of USD 2.1 million. Post-hoc analysis confirmed a 53% reduction in cold chain breaches and a 19.2% improvement in shelf-life through hyperspectral quality verification. These results demonstrate the resilience and robustness of the system in maintaining cold chain integrity under extreme environmental stress and present the framework's resilience, supporting the claim that multimodal fusion and adaptive control strategies effectively safeguard cold chain integrity under extreme environmental stress.

Despite the results of this study, challenges persist. Connectivity gaps in rural areas, where 4G coverage fell below 65%, reduced disease detection accuracy (F1-score of 0.71 compared with 0.89 in urban zones). Sensor drift was observed in MEMS accelerometers (−0.15 g/month) and FLIR cameras (+0.3 °C/100 hr), requiring approximately 8.2 h of monthly calibration. Finally, energy-accuracy tradeoffs were evident in solar-powered nodes, where high-precision processing reduced battery life by 23% during monsoon seasons. Future solutions may include graphene-based electrochemical arrays for pesticide detection, 6G integration for sub-millisecond data collection, and neuromorphic vision sensors to reduce power consumption while maintaining high detection accuracy.

5. Conclusions

The developed sensor-based AI framework integrates TD3++ reinforcement learning, multimodal fusion, edge-cloud optimization, and real-time calibration. The framework achieved

low-latency inference (28 ms), enhanced ripeness prediction accuracy (±0.5 °C), and a 53% reduction in cold chain breaches. By combining domain-adaptive learning, attention-weighted fusion, and edge acceleration on platforms such as NVIDIA Jetson Nano and AGX Xavier, the framework delivered superior energy efficiency (0.38 kWh/ton-mile), improved prediction accuracy, and real-time inference at 42 FPS. Adaptive sensor weighting, blockchain-backed data verification, and anomaly correction further enabled 92.7% on-time delivery rates and annual savings exceeding USD 61,000 per logistics hub. These results show the potential of the framework to enhance the resilience, precision, and sustainability of agri-food supply chains. Beyond agriculture, the framework can be used to develop a replicable model for multimodal fusion, real-time inference, and energy-efficient edge deployment in other data-intensive, real-time operational domains.

Despite these advances, several challenges remain. Connectivity gaps in rural areas, where 4G coverage fell below 65%, reduced disease detection accuracy (F1-score of 0.71 compared with 0.89 in urban zones). Sensor calibration drift was observed in MEMS accelerometers (−0.15 g/month) and FLIR cameras (+0.3 °C/100 h), requiring approximately 8.2 h of monthly calibration. Energy–accuracy tradeoffs were evident in solar-powered nodes, where high-precision processing reduced battery life by 23% during monsoon seasons. To address these issues, next-generation solutions such as sixth-generation communication, neuromorphic vision sensors, and graphene-based electrochemical arrays are required.[17] Since this study was limited to GBA, validation across diverse climatic zones must be conducted. Reliance on high-performance edge nodes such as Jetson Nano might also hinder adoption in resource-constrained regions. Therefore, further research is required to adopt lightweight knowledge distillation techniques for low-cost microcontrollers, federated learning architectures to enable cross-regional knowledge sharing without compromising data privacy, and blockchain-based verification to strengthen cybersecurity and transparency in distributed sensing networks.[18]

Acknowledgments

This research was funded by the "Research on the Construction Strategy of Green Supply Chain for Agricultural Products in the Guangdong-Hong Kong-Macao Greater Bay Area" (Project No. 2024CSLKT3-220), a general research project funded by the China Society of Logistics and the China Federation of Logistics and Purchasing in 2024; "Research on the Construction Strategy of Green Supply Chain for Agricultural Products in the Guangdong-Hong Kong-Macao Greater Bay Area" (Project No. 2024WTSCX254), a characteristic innovation project funded by Guangdong Province's ordinary universities in 2024; and "Supply Chain Management Practice" course teaching team (Project No. HSCXTD202302), a school-level education and teaching quality and teaching reform engineering project in 2023.

References

1 Y. Cil, D. Abdurahman, and I. Cil: J. Shipp. Trd. **7** (2022) 9. https://doi.org/10.1186/s41072-022-00110-z
2 Y. Faqir, A. Qayoom, E. Erasmus, M. Schutte-Smith, and H. G. Visser: Comput. Electron. Agric. **226** (2024) 109385. https://doi.org/10.1016/j.compag.2024.109385

3 A E. Ahoa, A. Kassahun, C. Verdouw, and B. Tekinerdogan: Sensors **25** (2025) 2362. https://doi.org/10.3390/s25082362
4 M. Hussain: Machines **11** (2023) 677. https://doi.org/10.3390/machines11070677
5 H. Xiong: Discrete Dyn. Nat. Soc. **2021** (2021) 6623563. https://doi.org/10.1155/2021/6623563
6 X. Ma, J. Zhao, L. Jia, X. Chen, and Z. Li: High-speed Rail. **1** (2023) 185. https://doi.org/10.1016/j.hspr.2023.09.002
7 S. Fujimoto, H. van Hoof, and D. Meger: Proc. 2018 35th Int. Conf. Machine Learning (ICML, 2018) 1587. https://proceedings.mlr.press/v80/fujimoto18a.html.
8 Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, Z. Zhai, X. T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, K. Uszkoreit, and N. Houlsby: arXiv:2010.11929. https://doi.org/10.48550/arXiv.2010.11929
9 K. He, X. Zhang, S. Ren, and J. Sun: Proc. 2016 IEEE Conf. Computer Vision and Pattern Recognition (IEEE, 2016) 770. https://doi.org/10.1109/CVPR.2016.90
10 NVIDIA: https://docs.nvidia.com/deeplearning/tensorrt/archives/tensorrt-843/pdf/TensorRT-Developer-Guide.pdf (accessed May 2025).
11 F. Daghero, A. urrello, D. J. Pagliari, L. Benini, E. Macii, and M. Poncino: arXiv:2204.03431. https://doi.org/10.48550/arXiv.2204.03431
12 M. Gan, D. Li, M. Wang, G. Zhang, S. Yang, and J. Liu: Math. Probl. Eng. **2018** (2018) 8439582. https://doi.org/10.1155/2018/8439582
13 J.-L. Zhong, Y.-F. Gan, C.-M. Vong, J.-X. Yang, J.-H. Zhao, and J.-H. Luo: Pattern Recognit. **122** (2022) 108286. https://doi.org/10.1016/j.patcog.2021.108286
14 K. Yin, J. Gersey, and P. Zhang: arXiv:2506.15840. https://doi.org/10.48550/arXiv.2506.15840
15 N. Chen, Z. Cheng, X. Fan, Z. Liu, B. Huang, Y. Zhao, L. Huang, X. Du, and M. Guizani: ACM Trans. Multimed. Comput. Commun. Appl. **20** (2024) 237. http://dx.doi.org/10.1145/3661313
16 J. A. Ismail, B. Faisal, F. Abdullah, and A. H. Sulaiman: J. Nanoelectron. Optoelectron. **14** (2021) 333. https://www.researchgate.net/publication/358741565_Characterization_of_DWT_as_denoising_method_for_phase-OTDR_signal
17 F. Miao, Y. Han, J. Shi, B. Tao, P. Zhang, and P. K. Chu: J. Mater. Res. Technol. **22** (2023) 3156. https://doi.org/10.1016/j.jmrt.2022.12.157
18 K. Wu, Z. Ji, H. Wang, X. Shao, H. Li, W. Zhang, W. Kong, J. Xia, and X. Bao: Electronics **14** (2025) 3994. https://doi.org/10.3390/electronics14203994

About the Authors

Haidong Li obtained his master's degree in agriculture from Zhongkai University of Agriculture and Engineering in 2024. From 2012 to 2019, he served as a lecturer at the School of Business Administration of Guangzhou Huashang Vocational College in China. Since 2020, he has been serving as an associate professor at the School of Digital Commerce of Guangzhou Huashang Vocational College. His research interests include agricultural product supply chains and AI. (richfieldlee@gzhsvc.edu.cn)