Integrating Real-Time Visual Return Grading with Deep Reinforcement Learning for Sustainable Reverse Logistics

Rahul Raj
Bentonville, USA

ABSTRACT
The exponential growth of e-commerce has intensified reverse logistics challenges, with product returns generating substantial carbon emissions through inefficient routing and processing. This research proposes an integrated architecture combining edge-deployed convolutional neural networks (CNNs) for real-time return quality assessment with Deep Reinforcement Learning (DRL) for carbon-aware dynamic routing optimization. The edge vision module classifies returned items into disposition categories with sub-second latency; these assessments feed directly into a DRL optimizer formulated as a Markov Decision Process (MDP). Simulation experiments on benchmark VRPSDP instances demonstrate 18–23% carbon emission reduction, 94.2% classification accuracy, and 67% lower decision latency compared to cloud-based alternatives. A six-month pilot with two retail partners validating 78,000 returns confirms operational viability with 99.2% system uptime. This is the first end-to-end framework integrating edge AI vision with DRL-based carbon-optimized routing for retail reverse logistics.

General Terms
Computer Vision; Deep Reinforcement Learning; Reverse Logistics; Sustainable Supply Chain; Carbon Emission Optimization; Edge Computing; Vehicle Routing.

Keywords
Edge computing; convolutional neural networks; deep reinforcement learning; carbon emission reduction; reverse logistics; vehicle routing problem; sustainable supply chain; retail returns management.

1. INTRODUCTION
The rapid expansion of e-commerce has produced return rates of 20–30% in online fashion retail and 15–20% across general categories. Transportation emissions from return processing contribute an estimated 5–7 million metric tons of CO_2 annually in the United States alone [3]. Traditional reverse logistics systems rely on fixed routes, manual inspection, and static schedules that disregard real-time traffic conditions, vehicle emission profiles, and return urgency.

AI-driven route optimization reduces logistics CO_2 by 15–25%, and automated visual inspection improves return processing efficiency by 40–60% [5]. These advances have, however, been developed independently. No existing framework couples real-time visual return assessment with carbon-optimized routing—a gap that constrains progress toward sustainable circular supply chains.

Edge-based computer vision combined with Deep Reinforcement Learning (DRL) addresses this gap. Edge computing enables item classification at collection points without cloud latency, while DRL learns adaptive routing policies balancing emission reduction, cost, and service quality. Proximal Policy Optimization (PPO) has demonstrated superior performance over traditional metaheuristics in dynamic routing scenarios [6].

This paper addresses three critical gaps: (1) existing vision systems focus on centralized warehouse inspection rather than distributed edge processing; (2) carbon-aware routing algorithms lack real-time return characteristic data; and (3) no end-to-end framework demonstrates edge vision integrated with DRL routing in operational reverse logistics. Sections 2–6 review literature, present the methodology, experimental results, discussion, and conclusions respectively.

2. LITERATURE REVIEW

2.1 Edge-Based Computer Vision for Retail Returns
Edge-based video analytics enables on-device defect detection and classification before items enter processing queues. Arvind et al. demonstrated an intelligent warehousing framework integrating edge analytics with machine learning and IoT, achieving 25% carbon reduction and 72% efficiency improvement [1]. However, this work addressed inventory management rather than return-specific grading and did not integrate vision outputs with dynamic routing. Giri et al. proposed an AI-enabled circular supply chain framework employing computer vision within its decision logic [8], but without explicit carbon minimization objectives or edge-deployment specifications.

2.2 Carbon-Aware Routing Optimization
Speed-based emission modelling underpins carbon-aware routing. Hou et al. proposed VRM-RCCE, predicting vehicle speed from real-time road conditions to calculate carbon emissions, validated using Gurobi with ant colony optimization [2]. Santos et al. explored VRPDDP in green reverse logistics, demonstrating that splitting customer visits reduces CO_2 using an augmented ε-constraint method [3]. Cheng et al. developed an INNC genetic algorithm producing a 4.69% decrease in carbon emission costs [4]. These static formulations limit applicability to dynamic reverse logistics environments.

2.3 Deep Reinforcement Learning in Logistics
DRL has demonstrated superior adaptability in dynamic routing. Zou et al. applied DRL to optimize reverse logistics transport in closed-loop supply chains [6]. Van der Linden et al. demonstrated reinforcement learning for efficient returns management, highlighting sample-efficiency advantages over classical optimization [12]. Juliet deployed IoT-integrated AI routing for smart logistics [10]. A key gap remains none of these frameworks incorporate real-time visual return characteristics into state representations or reward functions.

2.4 Research Gaps and Motivation
Three unaddressed gaps exist in the literature: (1) no lightweight edge CNN has been validated for real-time return grading at

collection points; (2) no DRL routing framework encodes visual assessments in its state-action formulation; and (3) no end-to-end validation exists for a combined vision-routing system in operational reverse logistics. This research closes all three gaps through an integrated architecture validated by simulation and six-month field deployment.

3. METHODOLOGY

3.1 System Architecture Overview

The proposed system comprises three integrated subsystems: (1) an edge vision module deployed at return collection points; (2) a carbon emission model providing real-time CO_2 cost estimates for candidate routes; and (3) a DRL routing optimizer generating adaptive vehicle assignments. An IoT integration layer connects all components via MQTT publish-subscribe messaging, enabling continuous data flow from edge devices to the central optimizer.

The architecture operates in a closed-loop cycle. The edge module classifies incoming returns and transmits disposition records; the emission model converts route proposals into CO_2 estimates; and the DRL agent updates routing decisions on a 15-minute rolling horizon. Returns with priority score > 8 trigger immediate route reconfiguration outside the regular cycle.

3.2 Edge Vision Module

The vision module employs a lightweight multi-task CNN derived from MobileNetV3, trained to simultaneously perform: (1) disposition classification (restock, refurbish, recycle, dispose); (2) ordinal condition grading on a five-point scale; (3) defect localization; and (4) product category recognition. INT8 quantization and structured pruning achieve sub-200 ms inference on NVIDIA Jetson Nano devices. Training used 180,000 annotated return images across 120 product categories, augmented with synthetic defect injection, lighting variation, and viewpoint perturbation.

3.3 Carbon Emission Model

Carbon emissions are modelled using a speed-dependent function calibrated from fleet telemetry. For a route segment at vehicle speed v (km/h), the emission rate E(v) in kg CO_2/km follows $E(v) = \alpha/v + \beta\cdot v^2$, where α and β are vehicle-class parameters. Emissions are integrated over each segment using traffic-API-predicted speeds, producing per-route CO_2 estimates that feed directly into the DRL reward function.

3.4 DRL Routing Framework

The routing problem is formulated as a Markov Decision Process. The state vector s_t = [V_t, R_t, T_t, E_t, C_t] encodes fleet positions and loads (V_t), pending returns with visual assessments (R_t), traffic conditions (T_t), cumulative emission tracking (E_t), and operational constraints (C_t), yielding 500–2,000 features for networks with 50–200 service nodes.

The action space decomposes into vehicle-to-return assignment (discrete) and visit-order sequencing (continuous, via multi-head attention). The reward function is $R(s_t, a_t, s_{t+1}) = -w_1\Delta E - w_2\Delta C - w_3\Delta T + w_4\Delta S$, where ΔE is incremental carbon emissions, ΔC is cost change, ΔT is travel time change, and ΔS is service quality improvement; weights are $w_1 = 0.45$, $w_2 = 0.25$, $w_3 = 0.20$, $w_4 = 0.10$.

The agent employs PPO with clipped surrogate objective. The policy network uses three fully-connected shared layers (512–256–128 units, ReLU, 0.2 dropout), a softmax assignment head, a linear value head, and a four-head attention mechanism for sequencing. Training employed Generalized Advantage Estimation ($\lambda = 0.95$), batch size 256, and a learning rate schedule from 3×10^{-4} to 1×10^{-5} over 500,000 steps.

3.5 System Integration and Data Flow

Vehicle telemetry (GPS, speed, load) is streamed at 10-second intervals. Traffic data is ingested at 5-minute intervals and fused with historical patterns to predict segment speeds. Fault tolerance is maintained through default-priority fallback for failed edge devices, last-received-route execution during connectivity loss, and redundant cloud infrastructure with automatic failover, yielding 99.2% operational uptime during field deployment.

4. RESULTS AND FINDINGS

4.1 Vision Module Performance

The edge vision module was evaluated on a held-out test set of 25,000 return images spanning 50 product categories from three retail partners over six months, encompassing diverse item conditions, lighting environments, and return reasons. Table 1 summarizes all classification and latency metrics against target thresholds.

Table 1. Vision Module Classification and Latency Performance

Performance Metric	Achieved Value	Target	Status
Disposition Classification Accuracy	94.2 %	≥ 90.0 %	Met
Condition Grade — Exact Match	91.7 %	—	—
Condition Grade — Within ±1 Grade	98.3 %	—	—
Defect Detection mAP@0.5	88.3 %	—	—
Defect Detection — Severe Defects	92.7–95.1 %	—	—
Defect Detection — Subtle Defects	81.4 %	—	—
Category Recognition — Top-1	96.8 %	—	—
Inference Latency — Jetson Nano	198 ms	< 200 ms	Met
Inference Latency — INT8 Quantized	142 ms	—	—

Inference Latency — Jetson Xavier	87 ms	—	—
End-to-End Latency — 4G Cellular	245 ms	< 500 ms	Met
End-to-End Latency — Local WIFI	216 ms	< 500 ms	Met

Disposition classification accuracy of 94.2% exceeded the 90% operational threshold for routing integration. Condition grading achieved 98.3% within-one-grade accuracy, demonstrating ordinal consistency suitable for priority-tiered routing. Defect detection performance varied by severity: severe defects reached 92.7–95.1% mAP@0.5, while subtle defects reached 81.4%, identifying low-contrast anomalies as the primary area for further model improvement. Table 2 compares the proposed edge approach against alternative classification methods.

Table 2. Comparison of Return Classification Approaches

Classification Approach	Accuracy	Decision Latency	Deployment Type
Proposed Edge Vision CNN	**94.2 %**	**198 ms**	**Edge node**
Manual Inspection	96.5 %	45 – 60 s	On-site personnel
Cloud-Based Vision	94.7 %	1.2 – 1.8 s	Remote cloud server
Rule-Based Metadata Only	73.2 %	< 5 ms	Central database

The proposed edge CNN achieves accuracy within 2.3 percentage points of manual inspection at more than 200 times lower latency. Cloud-based vision incurs 6–9 times higher latency, insufficient for real-time routing integration. Rule-based metadata classification falls 21 percentage points below edge vision accuracy, confirming the necessity of visual inspection for reliable disposition decisions. Error analysis identified four primary failure modes: ambiguous condition boundaries (42% of errors), lighting variation (18%), novel product types absent from training data (15%), and packaging obscuring item condition (12%). Vision-informed routing reduced average processing time by 34% (4.2 to 2.8 days) and carbon emissions by 11% compared to random disposition assignment.

4.2 Routing Optimization Results

The DRL routing optimizer was evaluated through simulation on modified Solomon benchmark instances [R1, C1, RC1] adapted for simultaneous delivery and pickup with dynamic return arrivals following Poisson processes (5–20 returns/hour). Networks ranged from 50 to 200 service nodes with 10–30 vehicles. Table 3 compares DRL-PPO against four baseline algorithms on a 100-node, 20-vehicle network.

Table 3. Routing Algorithm Performance Comparison (100-Node, 20-Vehicle Network)

Algorithm	Emission Reduction	Cost Increase	Computation Time	Adaptability (/ 10)
DRL — PPO (Proposed)	**21.3 %**	**2.1 %**	**0.8 s**	**9.2**
Genetic Algorithm (GA) [4]	18.7 %	1.8 %	12.4 s	4.3
Ant Colony Optim. [2]	17.2 %	2.3 %	8.7 s	5.1
Greedy Nearest-Neighbor	8.4 %	0.5 %	0.1 s	2.1
Static Routing (Baseline)	0 % (ref.)	0 % (ref.)	N/A	1.0

DRL-PPO outperformed Genetic Algorithm by 12.2 percentage points and Ant Colony Optimization by 19.2 percentage points in emission reduction. The adaptability score—measuring response to traffic disruptions, urgent returns, and vehicle breakdowns—strongly favored DRL (9.2/10) over metaheuristics (4.3–5.1/10), confirming the advantage of learned policies in dynamic operational environments. Computation time of 0.8 s fits comfortably within the 15-minute decision cycle. Table 4 presents emission reduction disaggregated by benchmark instance type.

Table 4. DRL Emission Reduction by Benchmark Instance Type (Modified Solomon Instances)

Benchmark Instance	DRL–PPO (Proposed)	GA Baseline	ACO Baseline	DRL Advantage vs. Best Baseline
R1 — Random Customers	21.3 % ± 3.2	9.1 %	8.7 %	+12.2 pp
C1 — Clustered Customers	19.7 % ± 2.8	7.8 %	8.2 %	+11.5 pp

RC1 — Mixed Distribution	22.1 % ± 3.5	9.8 %	9.3 %	+12.3 pp
Overall Average	**21.0 %**	8.9 %	8.7 %	**+12.3 pp**

DRL superiority is consistent across all geographic distribution patterns, with the greatest advantage on mixed RC1 instances (+12.3 pp) where adaptive routing benefits most from complex spatial-temporal patterns. Standard deviations below 3.5% confirm solution reliability. Service quality remained strong: average processing time increased by only 2.1% compared to cost-only optimization, while high-priority returns experienced 15% faster processing (2.1 vs. 2.5 days).

4.3 Carbon Emission Reduction Analysis

Table 5 presents a detailed breakdown of carbon emission sources and the DRL-achieved reduction per category, providing mechanistic insight into where environmental gains are realized.

Table 5. Carbon Emission Source Breakdown and DRL-Achieved Reduction per Category

Emission Source	Share of Total	DRL Reduction	Primary Mechanism	Vehicle Coverage
Line-haul (inter-stop travel)	62 %	15 %	Speed optimization (45–55 km/h)	All types
Stop-and-go (accel./decel.)	23 %	28 %	Consolidated stop batching	All types
Idling during item loading	11 %	19 %	Efficient loading sequences	All types
Congestion-related delays	4 %	35 %	Traffic-aware re-routing	All types

The greatest relative reduction (35%) occurs in congestion-related emissions, where traffic-aware re-routing bypasses high-emission idle periods. Stop-and-go emissions achieve 28% reduction through vision-informed consolidation of low-priority returns. Speed profile analysis shows the DRL agent maintains optimal efficiency speeds (45–55 km/h) for 73% of route distance, compared to 58% for baseline routing. Temporal analysis shows emission reduction peaks at 26–31% during peak traffic periods versus 15–18% off-peak, confirming greatest adaptive advantage when traffic variability is highest. For 500 returns per day across a 100-node urban network, the system reduces annual emissions by approximately 180 metric tons CO_2. Table 6 disaggregates reduction by vehicle type.

Table 6. Emission Reduction Disaggregated by Vehicle Type and Routing Strategy

Vehicle Type	Emission Reduction	Preferred Route Type	Key Optimization Strategy
Diesel Vans	23.4 %	Highway-dominant	Cruise efficiency at 45–55 km/h
Gasoline Trucks	19.8 %	Mixed urban/highway	Speed–emission curve optimization
Hybrid Vehicles	17.2 %	Balanced urban/highway	Optimal speed band maintenance
Electric Vehicles	12.1 % (equiv. CO_2)	Urban, frequent stops	Regenerative braking maximization

Vehicle-specific routing strategies emerged from DRL training: diesel vans are directed to highway segments where efficiency peaks at 45–55 km/h, while electric vehicles are assigned to urban routes to maximize regenerative braking benefit. Diesel vans achieved the greatest reduction (23.4%); electric vehicles achieved the smallest (12.1% equivalent CO_2), reflecting their already-efficient baseline profile.

4.4 System Integration and Scalability

End-to-end latency across the full pipeline totals 407 ms: vision processing (198 ms) + communication (47 ms) + policy evaluation (124 ms) + instruction transmission (38 ms). This represents a 100–200 times improvement over traditional systems combining manual inspection with batch routing. Table 7 presents six-month pilot deployment results.

Table 7. Six-Month Pilot Deployment Results — Two Retail Partners

Retail Partner	Returns Processed	Emission Reduction	Vision Accuracy	System Uptime	Duration
Fashion Retailer — 25 stores	47,000	21.7 %	92.8 %	99.2 %	6 months
Electronics Retailer — 18 stores	31,000	19.3 %	89.4 %	99.2 %	6 months
Combined Total	**78,000**	**20.7 % avg.**	**91.5 % avg.**	**99.2 %**	—

Field validation confirmed simulation performance. The fashion retailer achieved 21.7% emission reduction across 47,000 returns; the electronics retailer achieved 19.3% across 31,000 returns. Vision accuracy of 89.4–92.8%, validated against warehouse inspection records, confirms operational reliability. System uptime of 99.2% demonstrates production readiness. Operational challenges addressed during the pilot included: lighting variability at six collection points (LED installation); cellular gaps at three rural stores (signal boosters); and driver adaptation requiring two to three weeks of training. Monthly model updates addressed novel product types responsible for 15% of misclassifications. Scalability to networks of 500–1,000 nodes via hierarchical zone decomposition maintained 18–21% emission reduction with 4.2-second computation time—manageable with standard cloud infrastructure. Table 8 presents the economic cost-benefit analysis.

Table 8. Economic Cost-Benefit Analysis — 500 Returns/Day Reference Operation

Item	Value Range (USD)	Basis / Notes
COSTS		
Edge vision hardware — per collection point	\$800 – \$1,200	One-time capital expenditure
Cloud computing infrastructure	\$500 – \$1,000/month	Ongoing operational cost
System integration	\$50,000 – \$100,000	One-time capital expenditure
Maintenance and model updates	\$1,500 – \$2,500/month	Ongoing operational cost
ANNUAL BENEFITS (500 returns/day)		
Carbon offset cost savings	\$5,400 – \$9,000	180 MT CO_2/yr @ \$30–\$50 per MT
Inventory turnover improvement	\$15,000 – \$25,000	Faster high-value item processing
Labor savings (manual inspection)	\$30,000 – \$50,000	500 returns/day operation
ROI Breakeven Point	**12 – 18 months**	Operations > 300 returns/day

Economic analysis confirms positive ROI within 12–18 months for operations processing more than 300 returns per day. Labor savings from automated inspection (\$30,000–\$50,000/year) constitute the largest benefit, followed by inventory turnover improvements (\$15,000–\$25,000/year) and carbon offset value (\$5,400–\$9,000/year).

5. DISCUSSION

5.1 Theoretical Contributions

Three primary theoretical contributions are established. First, a formal framework is provided for integrating real-time visual perception with dynamic routing optimization, demonstrating that vision-derived item characteristics can be incorporated into MDP state representations and reward functions for DRL-based routing—addressing a gap where perception and routing have been treated as separate problems [1, 6, 8].

Second, lightweight CNNs with model compression are shown capable of achieving 94.2% classification accuracy with sub-200 ms latency on resource-constrained edge devices, validating distributed perception for time-critical logistics. This finding extends to quality control, inventory management, and automated handling contexts.

Third, multi-objective DRL is demonstrated to learn routing policies balancing carbon emissions, operational costs, and service quality without explicit Pareto frontier enumeration. The reward formulation—prioritizing emission reduction ($w_1 = 0.45$) while constraining cost and service degradation—provides a practical template for sustainable logistics optimization.

5.2 Practical Implications

The framework provides actionable deployment pathways for retail logistics operators. Operations processing more than 300 returns per day achieve positive ROI within 12–18 months. The six-month pilot with 78,000 processed returns demonstrates feasibility at realistic operational scale. The 180 metric tone annual CO_2 reduction for a mid-size operation supports corporate sustainability reporting and regulatory compliance. Pre-deployment surveys should assess lighting adequacy, cellular coverage, and driver training requirements.

5.3 Limitations

Four primary limitations are acknowledged: (1) a 6% vision misclassification rate requiring ongoing model maintenance for novel product types; (2) field validation covering only two retail partners over six months, limiting generalizability; (3) routing performance dependence on the accuracy of external traffic prediction and emission model calibration; and (4) partially validated scalability beyond 1,000 nodes and to multi-modal transportation networks.

6. CONCLUSION

This paper presents the first end-to-end framework integrating edge AI vision with DRL-based carbon-optimized routing for retail reverse logistics. The system achieves 18–23% carbon emission reduction, 94.2% return classification accuracy, and sub-500 ms decision latency. Pilot deployment with 78,000 returns across two retail partners confirms 99.2% system uptime and operational viability. Economic analysis demonstrates positive ROI within 12–18 months for operations exceeding 300 returns per day.

Key innovations include: (1) a lightweight edge CNN achieving real-time return grading at 198 ms on resource-constrained devices; (2) a DRL routing framework encoding visual assessments and speed-based emission models into state representation and reward functions; and (3) a comprehensive integration architecture validated through simulation and field deployment. Future research should address multi-modal sensing, blockchain-based traceability, federated learning for privacy-preserving model improvement, and extended field deployments across diverse retail contexts.

7. ACKNOWLEDGMENTS

The author thanks the retail partners who provided access to operational data during the pilot deployment phase, and the anonymous reviewers whose constructive feedback strengthened this manuscript.

8. REFERENCES

[1] Arvind, S. Srinivasan, K. Ramesh, and P. Kumar, "Intelligent Warehousing: A Machine Learning and IoT Framework for Precision Inventory Optimisation," IEEE Access, vol. 13, pp. 1–15, 2025. DOI: 10.1109/access.2025.3614679

[2] Y. Hou, L. Zhang, and X. Wang, "Vehicle Routing Model Based on Real-Time Road Conditions Considering Carbon Emissions," Proc. 2024 Int. Conf. NTCI, 2024, pp. 1–6. DOI: 10.1109/ntci64025.2024.10776332

[3] F. A. Santos, G. R. Mateus, and A. S. da Cunha, "Green Reverse Logistics: Exploring the VRPDDP," Omega — Int. J. Management Sci., vol. 118, art. 102864, 2023. DOI: 10.1016/j.omega.2023.102864

[4] Y. Cheng, X. Liu, and Z. Wang, "Low-Carbon Logistics Distribution Vehicle Routing Optimization Based on INNC-GA," Applied Sci., vol. 14, no. 7, art. 3061, 2024. DOI: 10.3390/app14073061

[5] S. Mandal, R. Kumar, and A. Sharma, "Implementation of AI Transportation Routing in Reverse Logistics to Reduce CO2 Footprint," Int. J. Supply Chain Management, vol. 13, no. 2, pp. 45–58, 2024. DOI: 10.47604/ijscm.3079

[6] W. Zou, Y. Chen, and H. Li, "Deep Reinforcement Learning-Based optimization of Reverse Logistics Transport in Closed-Loop Supply Chain," Proc. ICE – Transport, pp. 1–12, 2025. DOI: 10.1680/jtran.24.00163

[7] M. Younesse, A. Benali, and K. El Moutaouakil, "New Wave Logistics Framework: Enabling Sustainable Practices in Road Transportation," Proc. 2025 Int. Conf. ECAI, 2025. DOI: 10.1109/ecai65401.2025.11095484

[8] C. Giri, S. Jain, and M. K. Tiwari, "AI-Enabled Circular Supply Chain Framework for Sustainable Returns Management," 2025. [Preprint]

[9] S. Ullah, M. Ahmad, and F. Khan, "AI-Driven Intelligent Transportation for Sustainable Business Ecosystems," 2025. [Preprint]

[10] D. S. Juliet, "AI-Driven Smart Logistics Route Optimisation Using IoT and Machine Learning," Proc. 2025 Int. Conf. ICCSP, 2025. DOI: 10.1109/iccsp64183.2025.11089166

[11] Y. Qian, "Research on E-Commerce Smart Logistics Distribution Path optimization Based on AI," Proc. 2025 5th Int. Conf. ISCTT, 2025. DOI: 10.1109/isctt66403.2025.11137854

[12] E. van der Linden, N. Kempers, W. van Jaarsveld, and R. Dijkman, "Reinforcement Learning for Efficient Returns Management," arXiv:2501.14394, 2025. DOI: 10.48550/arxiv.2501.14394

[13] R. Sasikala, M. Priya, and S. Kumar, "Reducing Carbon Footprint with AI-Driven Logistics Planning," Advances in Computational Intelligence and Robotics, IGI Global, 2025, ch. 7. DOI: 10.4018/979-8-3693-9856-2.ch007

[14] O. O. Ilufoye, A. B. Adeyemi, and T. O. Ogunleye, "A Carbon Analytics Framework for Reducing Environmental Footprint in Retail Supply Chains," Eng. Tech. J., vol. 10, no. 8, 2025. DOI: 10.47191/etj/v10i08.19

[15] Y. Li, X. Zhang, and Q. Wang, "An Intelligent Distribution System for Green Logistics in the Blockchain Environment," Int. J. Logistics Res. Appl., pp. 1–18, 2023. DOI: 10.1080/13675567.2023.2289463

[16] S. Krishnamurthy, R. Patel, and A. Gupta, "Leveraging AI and ML to Optimize Retail Operations," DIRA Int. J., vol. 12, no. 3, 2024. DOI: 10.36676/dira.v12.i3.140

[17] J. C. Ferreira, A. L. Martins, and P. D. Gaspar, "Enhancing Sustainable Last-Mile Delivery: Impact of EVs and AI Optimisation," World Electric Vehicle J., vol. 16, no. 5, art. 242, 2025. DOI: 10.3390/wevj16050242

[18] G. Zahn, M. Spann, and B. Skiera, "The Smart Green Nudge: Reducing Product Returns through Enriched Digital Footprints and Causal Machine Learning," SSRN, 2022. DOI: 10.2139/ssrn.4262656

[19] A. Kulkarni, S. Mehta, and R. Sharma, "Utilising Gen AI and Computer Vision for Applications in the Retail Sector," J. AI and Cloud Computing, vol. 2, no. 2, 2023. DOI: 10.47363/jaicc/2023(2)181.

[20] F. Felder, M. Brossog, and J. Franke, "Smart Routing for Sustainable Supply Chain Networks: An AI and Knowledge Graph Driven Approach," Applied Sci., vol. 15, no. 14, art. 8001, 2025. DOI: 10.3390/app15148001