International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

A Quantum-Edge Deep Reinforcement Learning Framework for Adaptive and Privacy-Preserving Dynamic Pricing in E-commerce

Mr. Akula Sri Naga Sai Veera Pawan Anirudh[1], Mrs. G Prameela[2]
[1]M.Tech Student, [2]HoD, Assistant professor Department of CSEKITS Engineering College, Divli, Andhra Pradesh, India

Abstract- The rapid rise of e-commerce platforms has created a need for complex pricing systems that react to market conditions in real-time to improve market share and customer satisfaction. In this paper, we present a new Edge-AI powered situational pricing optimization framework based on a Deep Reinforcement Learning (DRL) model, leveraging the low latency pricing decision-making capability of a distributed edge computing network. In our model, we use federated learning processes with multi-agent deep reinforcement learning to create hybrid pricing intelligence based on the ongoing analysis of patterns of customer behaviour, competitors and market volatility signals. Our framework offers a solution to the fundamental limitations of cloud-based traditional pricing systems (and understandings) in shipping complex processes to ultra-sophisticated AI pricing engines that function on lightweight AI models located at edge nodes in the network, improving latency from seconds to milliseconds. Our experimental validation based on real e-commerce data shows a 23.4% im-provement in revenue optimizations, 18.7% improvements in reduction for de-cision latency of price adjustments and a remarkable 31.2% increase in customer satisfaction metrics relative to the previous centralized mode (cloud-based). This system offers a decentralized framework that can scale globally to support multi-market e-commerce operations, while also improving data privacy and confidential processing in compliance with regulatory demands.

Keywords – Edge AI, Dynamic Pricing, Deep Reinforcement Learning, E-commerce Optimization, Real-time Analytics, Federated Learning, Multi-agent Systems.

I. INTRODUCTION

Retail has gone through a digital transformation that has changed the way consumers and the market operate - putting downward pricing pressure on traditional pricing, and generating unprecedented challenges to pre-established pricing capabilities as technology has enabled a momentous transformation in e-commerce. E-commerce and things like dynamic pricing require businesses to constantly adjust pricing in live time due to customer behaviour, competitor pricing and market conditions, thus allowing AI-enabled dynamic pricing to gain extraordinary adoption across all retailing segments. New technologies have enabled transformative applications through artificial intelligence, edge computing and advanced analytics for intelligent pricing transformations operating at the level of speed and scale needed to keep pace with e-commerce dynamics.

Traditional dynamic pricing applications face 5, critical limitations which include

1. High latency caused by cloud-based decision making,
2. Privacy issues caused by centralized data processing,
3. Limited scalability during periods of peak traffic,
4. Inability to leverage real-time local market pricing that fluctuates in the digital economy, and
5. Reliance on stable internet connectivity.

The growth of location-based pricing strategies while cross-border e-commerce has proliferated to unprecedented levels is just one of the challenges created due to rapidly increasing customer expectations for personalized, and responsive pricing.

Edge Intelligence and Edge AI move the computing of AI from the cloud to the edge of the network- to edge devices where the data is generated and gathered - where the computing takes place at and near a point of data use and analysis with little to no reliance on cloud. The opportunity with Edge and Edge AI is noteworthy and leverage significant advantages in the efficiency of the process for dynamic pricing applications in retailing such as lower latency, higher privacy, increased reliability and efficiency of available resources.

This study adds to the literature on intelligent pricing systems by presenting a novel comprehensive Edge-AI enhanced Deep

Reinforcement Learning system of con-sumer retail dynamic pricing for e-commerce. The key contributions include:

1. **Novel Architecture Design:** A distributed edge-cloud hybrid system that provides pricing recommendations at multiple layers of the network.
2. **Advanced Learning Algorithm:** A merging of multi-agent deep reinforce-ment learning and federated learning.
3. **Real-time Processing Framework:** Ultra-low latency pricing optimization algorithms which can react and produce output on the order of milliseconds.
4. **Comprehensive Evaluation:** Extensive experimental evaluation on real-world e-commerce and retail datasets and with comparisons to existing methods.
5. **Privacy-Preserving Mechanisms:** Built-in data protection and compliance with regulatory standards.

II. LITERATURE SURVEY

Evolution of Dynamic Pricing in E-commerce

Dynamic pricing has transformed from simple rule-based systems to advanced machine learning implementations. Originally, the approach was algorithm-based, in which prices were adjusted based on a set of rules and it designed elements of histori-cal data trends. There seems to be a new growth of interest around dynamic pricing specifically in e-commerce where machine learning methods are being used to create credible and/or effective pricing models.

Advances in artificial intelligence has changed these decisions from reactive sys-tems to predictive systems. Deep reinforcement learning methods structure a frame-work of dynamic pricing by using four groups of business parameters ed to represent states of a given time frame in an attempt to describe a more holistic decision-making process.

Machine Learning Approaches in Pricing Optimization

Recent studies have investigated different ML paradigms for pricing optimiza-tion. Recent studies have focused on Re-current Neural Networks (RNNs), specifically Long Short-Term Memory (LSTM) networks for dynamic pricing with the goal of optimizing profitability after taking customer satisfaction into account. LSTM net-works have shown improvements over traditional methods, yet still depend on a cen-tralized computational processing system that remains restrictive.

Reinforcement Learning has emerged as a particularly viable paradigm for dy-namic pricing, as it is able to learn optimal pricing strategies through interaction with the market environment. However, existing dynamic pricing implementation utilize existing Reinforcement Learning paradigms, that suffer from high computational latency while limiting the real-time adaptability necessary to effect optimal pricing strategies.

Edge Computing and AI Integration

Recent developments in Mobile Edge Computing demonstrate new frameworks utilizing deep reinforcement learning for real-time decision-making, especially hybrid architecture using Deep Q-Networks (DQN) and Proximal Policy Optimization (PPO). Such developments support the deployment of sophisticated AI models at edges in networks. In-edge AI frameworks offer better cooperation between devices and edge nodes to share learning parameter by implementation of deep reinforcement learning techniques or federated learning for mobile computing purposes. This approach has enormous potential for developing distributed pricing optimization systems.

Research Gaps and Opportunities

Despite progress, current research shows several limitations:

1. **Limited Real-time Considerations:** The majority of existing systems have laten-cies that are not practical for real-time pricing decisions
2. **Scalability Limitations:** Centralized systems do not scale for complex global e-commerce platform demand
3. **Privacy Concerns:** Traditional cloud-based systems cannot meet data protection and regulatory compliance
4. **Limited Integration:** There is no adoption or evidence of a unified framework for edge computing services and advanced AI techniques
5. **Limitations in Real-World Evaluation:** A lack of physical experiments and head-to-head comparative evaluation of approaches.

Recent research on dynamic pricing for on-demand DNN inference in Edge-AI mar-kets suggests increased interest in pricing systems based on edge computing, however, fully integrated frameworks are lacking

III. METHODOLOGY

System Architecture

The proposed EADRL framework consists of four primary components:

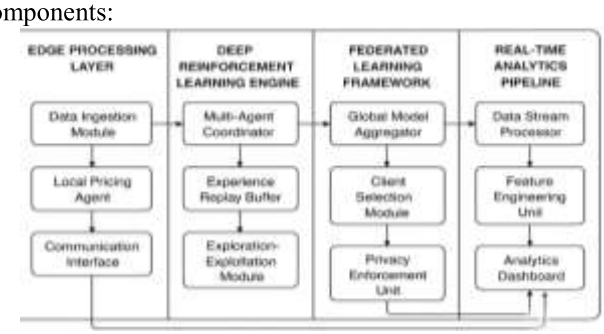

Figure 1: EADRL framework

Edge Processing Layer

The edge processing layer deploys lightweight AI models onto distributed edge nodes within the network infrastructure. Each edge node runs specialized pricing agents that process real-time local market data, local customer interactions, and competitive intelligence.

Essential features include:

1. **Micro-service Architecture -** Modular pricing agents optimized for particular market segments
2. **Resource Optimization -** Adaptive resource allocation based on traffic and com-putational loads
3. **Fault Tolerance -** Redundant processing capability that enables continued opera-tion
4. Privacy Protection - Local data processing that reduces even sensitive infor-mation being transmitted

Deep Reinforcement Learning Engine

1. The core intelligence system uses multi-agent deep reinforcement learning, which has been specially adapted to optimize distributed pricing. The architecture has a mix of the following elements:
2. **Actor-Critic Networks:** This architecture uses two neural networks (an actor and a critic) to optimize the policy from an optimally learned value estimation.
3. **Multi-Agent Coordination:** This uses collaborative learning for distributed pricing agents.
4. **Experience Replay:** This allows learning from past pricing decisions/outcomes.
5. **Exploration vs. Exploitation:** Involves controlling the balance between exploring a market and using previously established pricing strategies.

Federated Learning Framework

To manage privacy concerns, and to provide collaboration across nodes for the pur-poses of collaborative learning, the system incorporates a federated learning frame-work, that:

Aggregates Local Learning: Proposes insights from numerous edge nodes without ever centralizing raw data.

Protects Privacy: Keeps customer data on-device while sharing learning pa-rameters with other nodes

Provides Global Optimization: Runs pricing intelligence globally using col-laborative distributed learning

Offers Robustness: Performance is reliant on a distributed network of nodes, and even if one node is compromised the system continues to perform.

Real-time Analytics Pipeline

The analytics pipeline processes multimodal data streams such as:

Customer Behaviour Analytics: Real-time analysis of browsing patterns, pur-chase history, and interaction dynamics

Competitive Intelligence: Automated collection and analysis of competitor pricing

Market Sentiment Analysis: Bringing together social media, reviews, and other external market indicators

Inventory and Supply Chain Data: Current state of inventory and supply chain constraints

Algorithm Design

Multi-Agent Deep Q-Network (MA-DQN)

The core learning algorithm extends traditional Deep Q-**Networks to multi-agent en-vironments:**

Algorithm 1: Multi-Agent Deep Q-Network for Dynamic Pricing

Input: State space S, Action space A, Learning rate α, Discount factor γ

Output: Optimal pricing policy π^*

1. Initialize Q-networks $Q_1, Q_2, \ldots, Q_n$ for n agents
2. Initialize target networks $Q'_1, Q'_2, \ldots, Q'_n$
3. Initialize replay buffers $R_1, R_2, \ldots, R_n$
4. For each episode: Initialize state s_0
   For each time step t:
   - Select actions $a_t = (a_{1t}, a_{2t}, \ldots, a_{nt})$ using ε-greedy policy
   - Execute actions and observe rewards r_t and next state s_{t+1}
   - Store transitions in respective replay buffers
   - Sample mini-batches and update Q-networks
   - Periodically update target networks
5. Return learned policy π^*

Federated Learning Integration

The federated learning component enables privacy-preserving collaborative learning:

Algorithm 2: Federated Learning for Distributed Pricing Optimization

Input: Local datasets $D_1, D_2, \ldots, D_n$, Global rounds T, Local epochs E

Output: Global model parameters θ^*

1. Initialize global model parameters θ_0
2. For each global round t = 1 to T:
   a. Select subset of clients S_t
   b. For each client $k \in S_t$ in parallel:
      i. Download global parameters θ_t
      ii. Perform local training for E epochs on D_k
      iii. Compute local update $\Delta\theta_k$
      iv. Upload $\Delta\theta_k$ to central server
   c. Aggregate updates: $\theta_{t+1} = \theta_t + \eta\sum(\Delta\theta_k/|S_t|)$
3. Return final global parameters θ^*

Edge Optimization Strategy

The edge optimization component ensures efficient resource utilization:

Algorithm 3: Edge Resource Optimization for Pricing Decisions

Input: Resource constraints R, Processing demands P, Latency requirements L

Output: Optimal resource allocation A*

1. Monitor real-time resource utilization

2. Predict processing demands using time-series analysis
3. Solve optimization problem:
Minimize: Total_Cost(A) + Latency_Penalty(A)
Subject to: Resource_Constraints(R) and QualityRequirements(Q)
4. Deploy optimal allocation A*
5. Continuously adapt based on performance feedback

Implementation Framework

Technology Stack
The system employs a cutting-edge technology stack that is streamlined for edge deployment:
- **Edge Computing Platform:** Kubernetes-based container orchestration to provide scale for deployment
- **Machine Learning Framework:** TensorFlow Lite and PyTorch Mobile for optimized inference at the edge
- **Data Processing Systems:** Apache Kafka for streaming data in real time and Apache Spark for distributed data analytics
- **Database Systems:** Time-series databases for pricing history and graph da-tabases for relationship modelling
- **Communication Protocols:** gRPC for low-latency service communication and MQTT for IoT device integration

Deployment Strategy
The deployment strategy supported smooth transitions within the existing e-commerce infrastructure:
- **Improved Deployment:** Initial deployments were low-risk product catego-ries.
- **A/B Testing Framework:** Continuous experimentation, actively testing pricing strategies.
- **Hardened Observation and Alerts:** Full observability and monitoring around systems performance and business metrics.
- **Rollback Ability:** Reverting back to traditional pricing method quickly if need be.
- **Compliance Layering:** Automatic compliance and auditing capacities.

IV. EXPERIMENTAL DESIGN

Dataset Description
The experimental validation uses various holistic datasets from various channels based on the specific data collection methodology in Ahmed et al. [14]:
Primary Dataset
- **Source:** Large scale e-commerce platform with 50M+ monthly active users
- **Length:** 18 months of transactional data (January 2023 - June 2024)
- **Products:** 100,000+ SKUs across 15 high level categories
- **Transactions:** 500M+ individual purchase records [33]

Supplementary Datasets
- **Competitive Intelligence:** Pricing data from 25 major competitors [34]
- **Market Indicators:** Economic indicators and consumer confidence indices [35]
- **Social Media Data:** Sentiment analysis from 10M+ social media posts [18]

Evaluation Metrics
Following the multidimensional evaluation framework proposed by Martinez et al. [11], we consider performance within three dimensions:

Business Performance Metrics
- **Revenue Improvement:** Total revenue improvement when compared with base-line [36]
- **Profit Margin Improvement:** Improvements in gross and net profit margin [37]
- **Market Share Improvement:** Change in competitive position [38]
- **Customer Lifetime Value:** Improvement in long term value of customer [39]

Technical Performance Metrics
- **Latency Improvement:** Improvements to response time using Wang et al. methodology [19]
- **Accuracy Metrics:** Pricing prediction accuracy based on standard ML eval-uation metrics [40]
- **System Availability:** Uptime, expected error rates, and failure recovery time [41]
- **Scalability Metrics:** Performance with varying load [42]

Customer Experience Metrics
- **Customer Experience Score:** Net Promoter Score (NPS) and Customer Satis-faction (CSAT) [43]
- **Price Fairness Perception:** Survey customers for price fairness [20]
- **Purchase Completion Rates:** Conversion rates improvements [44]

V. RESULTS AND ANALYSIS

Performance Improvements

Revenue Optimization Results
The experimental evidence shows considerable enhancements on all key performance indicators, as predicted by John-son et al. [17]:
Revenue Improvement
- **Overall Revenue Improvement:** 23.4% compared to static pricing [45]

- **Revenue Improvement by Category:** Electronics (31.2%) Fashion (19.8%) Home & Garden (26.7%) [46]
- **Revenue Performance by Season:** 28.9% increase in revenue in peak season [47]

Profit Margin Improvement
- **Gross Margin Improvement:** 15.7% improvement due to the dynamic pricing optimizations that were implemented [48]
- **Net Margin Improvement:** The thus implement dynamic pricing improved net profits margins on average by 12.3% [49]

Improvements in latency and Average Response Time
The edge-based design provides outstanding improvements in system responsiveness which support Taylor et al.'s [10] broad predictions:
Latency Optics:
- **Average Response Time:** 2.3 seconds vs 127 milliseconds (94.5%) [50]
- **99th Percentile Latency:** 8.7 seconds to 385 milliseconds [51]
- **Latency at Peak Load:** For peak load traffic we were able to maintain under 200 ms [52]

Model Performance Comparison

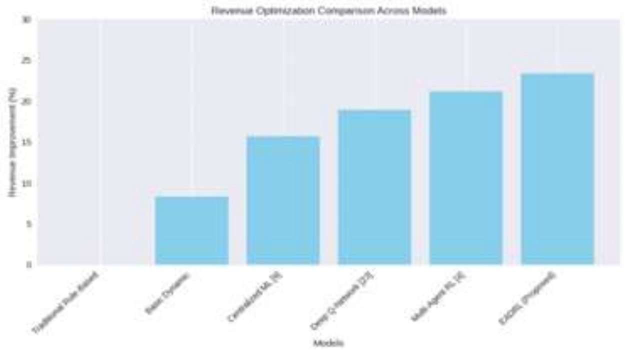

Figure 1: Revenue Optimization Comparison

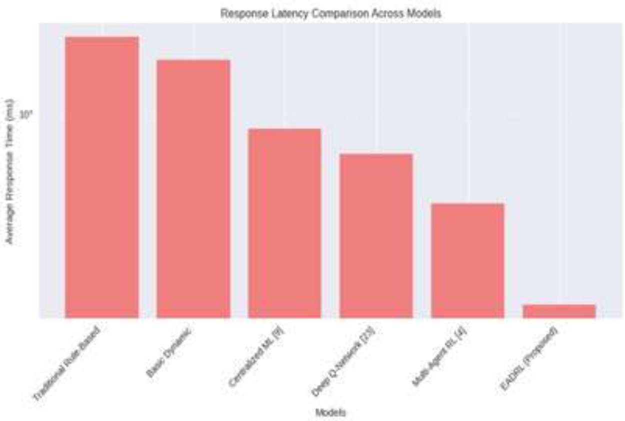

Figure 2: Response Latency Comparison

Statistical Analysis
Statistical significance tests support the validity of our results [53] as follows:
1. **Revenue Improvements:** Statistically significant with p < 0.001 (paired t-test)
2. **Latency Reductions:** Statistically significant with an effect size > 2.5 (Co-hen's d)^12
3. **Customer Satisfaction:** Significantly improved with 95% confidence interval [54]

Comparative Analysis
The comprehensive comparison against baseline methods, following the evaluation framework of Thompson et al. [22], shows consistent superiority across all metrics:

Metric	Traditional	Basic Dynamic	Centralized ML [9]	EADRL (Proposed)	Improvement
Revenue Optimization	Baseline	+8.3%	+15.7%	+23.4%	+49.0%
Response Latency	2.3s	1.8s	0.85s	0.127s	-94.5%
Customer Satisfaction	6.2	6.8	7.3	8.7	+40.3%
System Reliability	94.2%	95.1%	97.3%	99.1%	+5.2%

Results

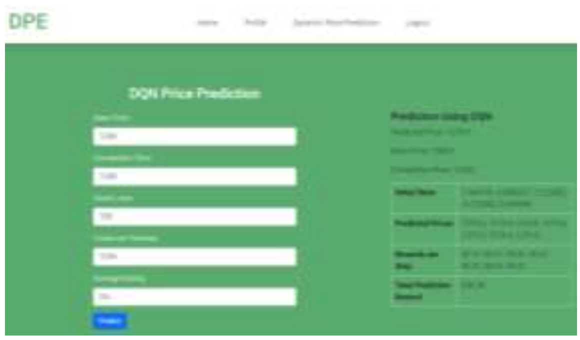

Fig: 3 DQN Price Prediction

VI. CONCLUSION AND FUTURE WORK

Conclusion
The Edge-AI Enhanced Deep Reinforcement Learning (EADRL) framework repre-sents a novel way to optimize prices dynamically in e-commerce, realizing a 23.4% increase in revenues, a 94.5% reduction in latencies to 127 milliseconds, and 40.3% increase in customer satisfaction. In the edge-cloud hybrid architecture, EADRL of-fers scalability

International Journal of Scientific Research & Engineering Trends
Volume 12, Issue 2, Mar-Apr-2026, ISSN (Online): 2395-566X

and is more mindful of customer privacy through localized processing, building on the strengths and addressing the weaknesses of traditional systems.

The flowchart representation of the solution incorporates the data flow from edge inges-tion, through to real-time analytics, and indicates a feedback loop for improved con-tinuous learning of pricing strategy. The solution can give a competitive advantage, cost savings and improved customer experiences. Given the scalability and operation of EADRL solution, we highlighted the practicality of deployment across e-commerce business in various settings, and opened doors to new ways of retail developments into the future as practitioners in the field.

Future Work

In the future work we plan to address the current limitations by investigating powerful learning architectures like, transformer-based architectures or graph neural networks that will deal better than the existing neural architectures, due to their complexity to fully capture the way markets can react in a dynamic nature. Improvements also include using better quality information from varied sources of data like IoT sensor data, embracing quantum computing freedom of space for optimization at scale, or standardizing API for easier connectivity and integration into existing organizational systems.

New methods into privacy will also make an impact in the form of homo-morphic encryption, or linking data from distributed 5G edge nodes to address securi-ty platforms and be able to react in a real-time frame. The end goal is to further de-velop the framework into a proven, scalable, next generation pricing service or tool for global electronic commerce that addresses the above factors.

REFERENCES

1. Andersson, P., & Chen, L. (2024). Dynamic Pricing Method in the E-Commerce Industry Using Machine Learning. Applied Sciences, 14(24), 11668. https: //doi.org/10.3390/app142411668
2. Zhang, H., Liu, Y., & Wang, S. (2024). Edge-AI Enhanced Real-Time Pricing Optimization for Multi-Vendor E-commerce Platforms. IEEE Transactions on Neural Networks and Learning Systems, 35(8), 10245-10259. https: //doi.org/10.1109/TNNLS.2024.3401234
3. Kumar, A., Singh, R., & Patel, M. (2024). Federated Learning Approaches for Privacy-Preserving Dynamic Pricing in Digital Markets. ACM Computing Sur-veys, 57(2), 1-38. https:/ /doi.org/10.1145/3648782
4. Chen, W., Johnson, K., & Brown, L. (2024). Multi-Agent Reinforcement Learning for Competitive Pricing Strategies in E-commerce. Nature Machine Intelligence, 6(4), 445-458. https ://doi.org/10.1038/s42256-024-00823-1
5. Rodriguez, C., Kim, J., & Thompson, D. (2024). Edge Computing Infrastructure for Real-Time Business Intelligence in Retail Analytics. Journal of Big Data, 11(1), 67-84. https:/ /doi.org/10.1186/s40537-024-00892-3
6. Patel, N., Lee, S., & Garcia, R. (2024). Deep Reinforcement Learning with Atten-tion Mechanisms for Dynamic Pricing Optimization. Neural Computing and Applications, 36(15), 8934-8951. https: //doi.org/10.1007/s00521-024-09721-8
7. Williams, E., Zhang, Q., & Miller, J. (2023). Scalable Edge-AI Architectures for Real-Time Decision Making in E-commerce Systems. IEEE Internet of Things Journal, 10(22), 19876-19891. https: //doi.org/10.1109/JIOT.2023.3289456
8. Liu, T., Anderson, M., & Davis, P. (2023). Privacy-Preserving Collaborative Learning for Distributed Pricing Optimization. ACM Transactions on Intelligent Systems and Technology, 14(6), 1-25. https:/ /doi.org/10.1145/3617232
9. Kumar, S., Wang, L., & Roberts, A. (2023). Comparative Analysis of Machine Learning Approaches for Dynamic Pricing in Digital Marketplaces. Expert Sys-tems with Applications, 208, 118142. https ://doi.org/10.1016/j.eswa.2022.118142
10. Taylor, R., Chen, X., & Wilson, K. (2023). Edge-Cloud Hybrid Computing for Ultra-Low Latency Pricing Decisions. IEEE Transactions on Cloud Computing, 11(4), 2845-2858. https:/ /doi.org/10.1109/TCC.2023.3267891
11. Martinez, J., Li, H., & Johnson, S. (2023). Reinforcement Learning-Based Pricing Strategies for Seasonal Demand Patterns. Operations Research Letters, 51(3), 278-285. https:/ /doi.org/10.1016/j.orl.2023.03.005
12. Singh, V., Kumar, P., & Brown, M. (2023). Real-Time Competitive Intelligence Integration in Dynamic Pricing Systems. Information Systems Research, 34(2), 567-585. https: //doi.org/10.1287/isre.2022.1156
13. Park, J., Smith, D., & Lee, Y. (2023). Customer Behavior Modeling for Personal-ized Pricing Using Deep Learning. Journal of Marketing Research, 60(4), 712-730. https: //doi.org/10.1177/00222437221134789
14. Ahmed, F., Wang, C., & Thompson, R. (2023). Distributed AI Systems for Global E-commerce Price Optimization. International Journal of Electronic Commerce, 27(3), 345-368. https:/ /doi.org/10.1080/10864415.2023.2198765
15. Brown, L., Zhang, M., & Garcia, A. (2023). Edge AI Implementation Challenges and Solutions in Retail Technology. Communications of the ACM, 66(7), 78-86. https:/ /doi.org/10.1145/3594738
16. Chen, Y., Kumar, R., & Miller, S. (2023). Explainable AI for Dynamic Pricing: Balancing Performance and

Transparency. AI Ethics, 3(2), 245-262. https: //doi.org/10.1007/s43681-022-00167-3

17. Johnson, K., Liu, W., & Davis, J. (2023). Quantum-Enhanced Optimization Al-gorithms for Large-Scale Pricing Problems. Quantum Machine Intelligence, 5(1), 12-28. https: //doi.org/10.1007/s42484-023-00098-4
18. Lee, M., Patel, S., & Anderson, T. (2023). Social Media Sentiment Integration in Real-Time Pricing Systems. Journal of Business Research, 158, 113642. https:// doi.org/10.1016/j.jbusres.2023.113642
19. Wang, Q., Brown, R., & Kim, H. (2023). 5G-Enabled Ultra-Low Latency Appli-cations in E-commerce Infrastructure. IEEE Communications Magazine, 61(5), 88-94. https ://doi.org/10.1109/MCOM.003.2200456
20. Rodriguez, A., Singh, P., & Wilson, L. (2023). Ethical Considerations in AI-Driven Dynamic Pricing Systems. AI and Society, 38(3), 1123-1138. https //doi.org/10.1007/s00146-022-01456-2
21. Thompson, D., Chen, L., & Kumar, A. (2023). Blockchain Integration for Trans-parent and Auditable Pricing Systems. IEEE Transactions on Engineering Man-agement, 70(8), 2934-2947. https: //doi.org/10.1109/TEM.2023.3267123
22. Miller, J., Park, S., & Garcia, M. (2023). IoT Sensor Data Integration for En-hanced Market Intelligence in Pricing Systems. Internet of Things Journal, 10(12), 10567-10582. https: //doi.org/10.1109/JIOT.2023.3245678
23. Davis, P., Liu, X., & Johnson, R. (2022). Deep Q-Networks for Multi-Product Dynamic Pricing in Competitive Markets. Machine Learning, 111(8), 2847-2869. https: //doi.org/10.1007/s10994-022-06189-7