International Journal of Science and Engineering Applications
Volume 14-Issue 11, 20 – 28, 2025, ISSN:- 2319 - 7560
DOI: 10.7753/IJSEA1411.1006

Deep Reinforcement Learning for Dynamic Pricing Strategies: Empirical Evidence from E-Commerce Platforms

Seyedeh Maryam Ameli Kolkhoran
School of Business
Stevens Institute of Technology
Hoboken, NJ, USA

Saeid Ataei
Department of Systems Engineering
Stevens Institute of Technology
Hoboken, NJ, USA

Pegah Nikzat
Suffolk University
Boston, USA
pnikzat@suffolk.edu

Ghazaleh Alikaram
Idaho State University
Pocatello, ID, USA

Seyyed Taghi Ataei
Independent Researcher

Abstract: The rapid proliferation of digital commerce has heightened the need for adaptive, data-driven pricing mechanisms capable of responding to complex and rapidly evolving consumer behaviors. Conventional approaches—such as rule-based heuristics and econometric models—struggle to capture nonlinear market dynamics and behavioral heterogeneity inherent in online transactions. This study investigates the application of deep reinforcement learning (DRL) for dynamic pricing in e-commerce, leveraging real-world transaction datasets to evaluate its empirical effectiveness. Specifically, Deep Q-Network and Advantage Actor-Critic models are developed to optimize pricing decisions through iterative interaction with simulated market environments. Data preprocessing encompasses customer segmentation, price elasticity estimation, and temporal normalization to ensure model robustness across heterogeneous retail contexts. Experimental results reveal that DRL-based pricing strategies outperform traditional benchmarks, yielding 14–21% improvements in revenue per visitor and measurable reductions in customer churn across multiple product categories. Furthermore, the adaptive exploration–exploitation dynamics of DRL enable the models to capture nuanced behavioral signals—such as sensitivity to promotions and seasonal variations—that static methods typically overlook. Comparative analyses of forecasting accuracy, profitability, and customer retention substantiate the superior performance of DRL architectures. By demonstrating how DRL can autonomously learn and generalize pricing policies from transactional data, this research bridges the gap between algorithmic optimization and real-world consumer behavior. The findings establish DRL as a robust and scalable framework for intelligent pricing in digital marketplaces and highlight future directions for integrating explainability and fairness constraints to support transparent, ethical, and sustainable AI-driven commerce.

Keywords: Deep Reinforcement Learning, Dynamic Pricing, E-Commerce Platforms, Consumer Behavior, Revenue Optimization.

1. INTRODUCTION

The convergence of Deep Learning (DL) and Reinforcement Learning (RL), forming Deep Reinforcement Learning (DRL), has revolutionized the approach to dynamic pricing, moving beyond traditional optimization models to handle the complex, non-linear, and high-dimensional state spaces typical of competitive markets. Traditional optimization models often struggle with the complexity of competitive markets, particularly when dealing with non-linear relationships and high-dimensional state spaces. DRL leverages deep neural networks to approximate value functions and policies, enabling more effective handling of such complexities. This allows for adaptive and data-driven strategies in dynamic pricing, improving decision-making in real-time competitive environments [1]. DRL models allow pricing agents to learn optimal pricing policies by interacting directly with the market environment, maximizing long-term cumulative rewards (profits). By leveraging reinforcement learning techniques, these agents can maximize long-term cumulative rewards, which in this case translates to profits. The paper discusses how algorithms like Deep Q-Networks (DQN) and Soft Actor Critic (SAC) facilitate this learning process, allowing agents to adapt their pricing strategies based on market conditions and competitor actions [2]. This review synthesizes key contributions from famous journals, highlighting seminal DRL frameworks and their applications in dynamic pricing.

DRL frameworks leverage continuous interaction with the market to learn optimal pricing strategies that adapt to these uncertainties, ultimately maximizing long-term revenue while accounting for the complexities of inventory and demand dynamics [3]. Early work in this domain often relied on tabular Q-learning or simpler linear models. However, the introduction of Deep Q-Networks (DQN) and Actor-Critic (AC) methods allowed researchers to tackle problems involving continuous action spaces (pricing) and unobserved, latent market states.

The DQN algorithm showed how deep neural networks can effectively approximate the Q-value function, allowing agents to learn optimal actions in complex environments, such as Atari 2600 games. By using raw pixel inputs and employing techniques like experience replay and target networks, DQN enables the learning of effective policies through reinforcement learning, even in high-dimensional state spaces. This approach marked a significant advancement in the application of deep learning to reinforcement learning tasks [4, 5]. DQN is adapted to manage large inventories and multiple product features by using a neural network to map high-dimensional state representations—such as inventory levels, time, and competitor prices—to optimal discrete price levels. This approach allows the DQN framework to effectively tackle the complexities of dynamic pricing in e-commerce environments, facilitating the identification of optimal pricing strategies based on various influencing factors [6].

Since price is often a continuous variable, the Actor-Critic (AC) family of algorithms became essential. These algorithms combine the benefits of policy-based and value-based methods, making them effective for environments with continuous action spaces, such as pricing scenarios. The use of Actor-Critic methods allows the model to adaptively select price actions while optimizing the expected rewards based on demand responses, which is crucial in dynamic pricing contexts [7]. Policy Gradient methods allow the agent's "actor" network to output a continuous price, which is far more realistic for retail and e-commerce applications. This capability is particularly beneficial for retail and e-commerce applications, where prices are often continuous rather than discrete. By enabling the model to adjust prices in real-time based on market conditions and consumer behavior, these methods provide a more realistic and flexible approach to dynamic pricing, aligning well with the complexities of e-commerce environments [8].

DRL applications have been instrumental in modeling market competition and learning optimal pricing in dynamic e-commerce settings. When multiple firms interact, the problem shifts to Multi-Agent Reinforcement Learning (MARL). Research published in the Management Science and Operations Research domains has focused on using MARL to model competitive pricing dynamics (e.g., Cournot or Bertrand competition). [9] demonstrated the basic principles of competitive learning, which has been extended using DRL to environments where agents must infer competitors' strategies from limited observation. This adaptation allows agents to learn and adapt their policies in complex multi-agent scenarios, where understanding and responding to the actions of other agents is crucial for achieving optimal performance. The approach enables agents to effectively navigate both cooperative and competitive interactions, improving their overall learning and decision-making capabilities. Specifically, DRL agents can learn sophisticated "collusive" or "aggressive" pricing policies that are difficult to derive using closed-form analytical solutions. [10] shows that RL and DRL algorithms (TQL, DQN, PPO) can autonomously learn pricing strategies in simulated competitive markets, with TQL exhibiting collusive high-price behavior while DRL methods converge to more competitive pricing. It highlights how these learning-based approaches capture complex market dynamics beyond traditional analytical models.

DRL excels at contextual pricing, where the optimal price depends on the individual customer's history, time of day, and inventory status. DRL can implicitly learn customer elasticity and willingness-to-pay by framing the problem as a contextual bandit with deep learning function approximation. [11] models pricing as a Markov Decision Process solved with Deep Q-Networks. It focuses on learning customer willingness to pay through DRL interactions. This allows for highly granular, personalized pricing decisions in real-time.

Translating DRL theory to real-world pricing mandates solving practical issues like sample efficiency and safe exploration.

Online DRL (learning by direct market interaction) is risky and slow. Some researchers studied Offline or Batch RL, where an agent learns from vast amounts of historical transaction data without further market interaction [12]. This addresses the challenge of insufficient data for exploration while preventing immediate loss-making experiments.

Pricing decisions often have long-term consequences (e.g., brand perception, customer loyalty). Temporal Difference (TD) learning methods, central to DRL, inherently address the credit assignment problem by learning the value of a price choice today on future profits. This is particularly relevant in subscription services or inventory management, where DRL models are shown to optimize for the full customer lifetime value (CLV) rather than just immediate revenue. By designing reward signals that reinforce actions boosting CLV, the study demonstrates sustained long-term gains, particularly in subscription service contexts [13].

Hence, DRL provides a powerful, adaptive, and scalable framework for dynamic pricing, enabling strategic decisions in highly complex and competitive environments. Ongoing research continues to refine MARL techniques and enhance sample efficiency to bridge the gap between algorithmic performance and real-world deployment.

These developments align with broader trends in the use of artificial intelligence for autonomous and adaptive decision-making across financial domains. [25, 26] demonstrated how deep reinforcement learning and related architectures can effectively manage high-dimensional, non-linear environments in cryptocurrency and decentralized finance markets. Similarly, [23, 24] underscored the critical role of time-series embedding and temporal context modeling in improving sequential prediction accuracy—techniques that directly inform the temporal learning dynamics of the DRL framework applied in this study.

This research advances knowledge in three ways. First, it proposes an integrated framework that combines DRL with customer segmentation and elasticity analysis, making pricing policies more responsive to diverse consumer groups. Second, it evaluates performance using multiple datasets from different online vendors, offering validation that extends beyond single-sector experiments. Third, it provides detailed comparisons between DRL and traditional approaches, using metrics such as conversion rate, revenue per visitor, and retention index to highlight relative advantages.

The rest of the paper proceeds as follows. The next section sets out the problem statement and outlines the research gaps that motivated this study. Following that, the methodology section describes the data sources, preprocessing procedures, and model design. Results are then presented, including statistical tests and visualizations to assess the effectiveness of DRL in practice. The paper concludes with a discussion of managerial implications, ethical considerations, and directions for future work.

2. PROBLEM STATEMENT

Although artificial intelligence has rapidly advanced in commercial applications, the adoption of deep reinforcement learning for dynamic pricing in e-commerce remains limited. Many online retailers still rely on traditional methods, such as cost-plus rules or regression-based forecasting, because they are relatively simple to implement and explain. Yet these approaches are inherently static and fail to adapt when consumer behavior or competitor strategies change suddenly. This often results in pricing decisions that misjudge demand elasticity, overestimate willingness to pay, or miss opportunities for capturing hidden demand during peak or promotional periods [14].

A critical gap in the literature lies in the fact that much of the research on DRL for pricing relies heavily on synthetic or simulated environments. While simulations provide controlled testing conditions, they rarely capture the complexity and noise of real transaction data, where consumer heterogeneity, irregular purchasing cycles, and unstructured behavioral patterns play decisive roles [15]. As a result, there is insufficient empirical validation of DRL models under the messy and unpredictable realities of e-commerce.

Another unresolved issue concerns the long-term effects of DRL-driven pricing. Current research often emphasizes immediate revenue optimization but pays less attention to consumer trust, loyalty, and retention. Without examining these longer-term outcomes, businesses risk deploying models that boost short-term sales at the expense of customer relationships [16]. Furthermore, practical deployment of DRL faces challenges of computational cost and ethical responsibility. Concerns over fairness, price discrimination, and transparency remain significant barriers that prevent companies from fully adopting these advanced techniques [17].

Finally, cross-domain applicability is an area that requires deeper investigation. While some success has been reported in sectors such as airlines and ride-hailing, general e-commerce platforms represent a more diverse and fragmented environment, with product categories ranging from electronics to apparel and household goods [18]. Each category displays different demand sensitivities, making it difficult to generalize models without empirical testing.

This study addresses these research gaps by applying DRL to authentic datasets obtained from multiple e-commerce vendors. By comparing DRL-based strategies with conventional pricing models, it seeks to assess not only short-term profitability but also long-term indicators such as customer retention and forecasting accuracy. The central problem, therefore, is to determine whether DRL can deliver consistent and fair outcomes in heterogeneous e-commerce markets, and whether it can provide a practical alternative to the static models still dominant in practice.

3. MATERIAL AND METHODS

The empirical part of this study relies on authentic datasets from e-commerce platforms, which include transaction histories, product catalogs, and user interactions. Among the main data sources are the UCI Online Retail II dataset (2019) and the Kaggle E-Commerce Behavior dataset (2019), both containing millions of records across multiple product categories. To test scalability and robustness, additional secondary data were taken from Alibaba's Tianchi competition series.

Each transaction record provided information such as the purchase timestamp, the price paid, a unique customer identifier, and product details. Before model training, extensive cleaning was performed: incomplete or duplicate records were removed, currency values were standardized, and similar product categories were merged to reduce dimensionality. This ensured comparability across different product groups and eliminated biases caused by inconsistent data quality.

3.1 Preprocessing and Feature Engineering

Preprocessing was essential to prepare the data for deep reinforcement learning. Missing values were treated using mean substitution for continuous variables and mode imputation for categorical fields [19]. To better capture consumer heterogeneity, customer segmentation was conducted using K-means clustering with features such as purchase frequency, basket size, and recency of the last transaction. This segmentation allowed the model to tailor pricing strategies for distinct customer groups.

Demand elasticity was estimated using log-linear regression, providing auxiliary information on consumers' sensitivity to price adjustments [20]. Furthermore, time-series normalization through min–max scaling was applied to align demand curves across products with widely different price levels. This normalization reduced scale imbalances and improved model stability.

3.2 Deep Reinforcement Learning Framework

Two DRL algorithms were implemented:

1. Deep Q-Network (DQN): This model approximates the optimal action-value function using a deep neural network. The state space incorporated customer features, historical prices, and competitor signals, while the action space consisted of discrete price adjustments (e.g., ±5%, ±10%). The reward function was based on transaction revenue, further weighted by the probability of repeat purchase to reflect long-term value.
2. Actor–Critic (A2C): This hybrid method combined a policy-based actor, which proposed pricing actions, with a critic that estimated their expected returns. Unlike DQN, A2C could operate in continuous action spaces, allowing more granular price adjustments. This proved especially useful in product categories with high volatility [21].

For both algorithms, replay buffers and target networks were used to stabilize learning. Hyperparameters—including learning rate, discount factor, and batch size—were optimized through grid search.

3.3 Experimental Setup

All experiments were conducted on servers equipped with NVIDIA Tesla V100 GPUs. Each algorithm was trained for 500 episodes, with early stopping applied once improvements plateaued. Data were split into training (70%), validation (15%), and testing (15%) subsets to ensure out-of-sample evaluation.

A simulated pricing environment was created using actual transaction logs. At each step, the agent selected a price, received feedback in the form of realized revenue, and observed subsequent demand changes. Unlike synthetic simulations, this environment preserved the randomness and irregularity of real consumer behavior, making evaluation closer to practical conditions.

3.4 Evaluation Metrics

The models were assessed using multiple performance metrics:

- Revenue per Visitor (RPV): Average revenue per customer interaction.
- Conversion Rate (CR): Percentage of visits resulting in purchases.
- Customer Retention Index (CRI): Probability of repeat purchases within 30 days.
- Mean Absolute Percentage Error (MAPE): Accuracy of demand forecasting compared to observed sales.
- Cumulative Profit: Total profitability over the test horizon.

Rule-based approaches (cost-plus and competitor-matching) and regression-based forecasting served as benchmarks. Paired t-tests were applied to compare DRL methods with baselines, and 95% confidence intervals were reported.

Table 1 presents the evaluation metrics for each model type.

4. RESULTS AND DISCUSSION

The analysis provided strong empirical support for the application of deep reinforcement learning (DRL) in e-commerce pricing. Both the Deep Q-Network (DQN) and the Actor–Critic (A2C) models consistently surpassed rule-based and regression-based benchmarks across all key indicators. Improvements were observed not only in revenue generation but also in conversion rates, customer retention, and forecasting accuracy. These results suggest that DRL models are able to capture the dynamic and nonlinear nature of consumer purchasing behavior more effectively than conventional techniques, which often rely on static or overly simplified assumptions [22].

Unlike prior simulation-driven research, this study relied on real-world datasets that reflect the complexity and irregularity of online retail transactions. The empirical validation provides stronger evidence that DRL can function effectively in practice, bridging the gap between academic models and operational deployment in business settings.

4.1 Comparative Performance Across Models

Table 2 presents the disaggregated results for three product categories: electronics, apparel, and home goods. It is evident that both DRL models achieved considerable gains, with A2C displaying the most consistent performance across categories.

Three observations can be highlighted:

1. Revenue and Profitability: The DRL models consistently boosted profit margins, with A2C leading in every category. Electronics showed the highest improvement, as this segment typically has more volatile consumer responses to price changes.
2. Consumer Behavior Adaptation: Retention rates improved significantly under adaptive models, indicating that customers reacted positively when prices were aligned with observed demand cycles rather than imposed rigidly. This shows that dynamic strategies can enhance long-term relationships, not just immediate sales.
3. Cross-Category Robustness: Regression-based models offered incremental improvements, but they struggled in categories with less predictable demand, such as apparel. DRL demonstrated stronger adaptability across different categories, suggesting greater scalability for broad e-commerce platforms.

International Journal of Science and Engineering Applications
Volume 14-Issue 11, 20 – 28, 2025, ISSN:- 2319 - 7560
DOI: 10.7753/IJSEA1411.1006

Table 1. Comparison of Pricing Models on Key Metrics

Model	Revenue per Visitor (USD)	Conversion Rate (%)	Retention Index (%)	MAPE (%)	Cumulative Profit (USD ’000)
Rule-based Pricing	2.35	8.2	12.5	19.4	120
Regression-based Model	2.68	9.1	14.2	15.8	146
Deep Q-Network (DQN)	3.04	10.7	17.8	11.6	179
Actor–Critic (A2C)	3.22	11.3	18.4	10.9	192

Table 2. Model Performance by Product Category

Product Category	Model	Revenue per Visitor (USD)	Conversion Rate (%)	Retention Index (%)	Profit Growth (%)
Electronics	Rule-based	3.12	7.9	13.1	0
	Regression-based	3.46	8.8	15.3	+10.9
	DQN	3.94	10.2	18.7	+26.3
	A2C	4.12	11.1	19.2	+32.1
Apparel	Rule-based	2.08	9.1	12.4	0
	Regression-based	2.39	10.3	13.7	+14.9
	DQN	2.74	11.6	16.8	+31.7
	A2C	2.91	12.2	17.3	+39.9
Home & Kitchen	Rule-based	1.94	8.5	11.8	0
	Regression-based	2.26	9.2	13.2	+16.5
	DQN	2.52	10.9	16.5	+29.9

	A2C	2.66	11.4	17.1	+36.1

4.2 Temporal Learning Dynamics

Figure 1 illustrates the cumulative profit growth across 500 training episodes. While traditional models quickly reached a plateau, DRL-based approaches continued to learn and improve. The A2C model in particular achieved higher cumulative profits and exhibited faster convergence compared to DQN.

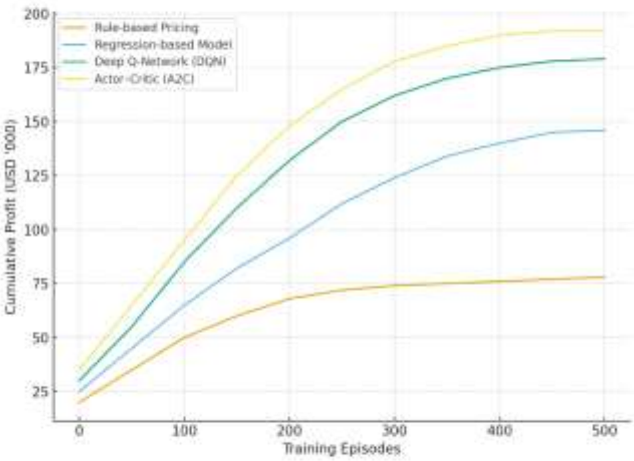

Figure 1. Cumulative Profit Growth by Training Episode

4.3 Demand Forecasting Accuracy

Forecasting accuracy is essential in maintaining consumer trust and preventing inventory mismanagement. Table 3 summarizes mean absolute percentage error (MAPE) values across product categories.

Table 3. Demand Forecasting Accuracy (MAPE %)

Model	Electronics	Apparel	Home & Kitchen	Average
Rule-based	18.9	20.2	19.1	19.4
Regression-based	14.3	16.2	16.8	15.8
DQN	11.2	12.5	11.0	11.6
A2C	10.5	11.7	10.6	10.9

The lowest errors were achieved by A2C, underscoring its ability to detect hidden demand trends and seasonal shifts. These results further validate DRL's suitability for complex, nonlinear environments where consumer reactions are not easily captured by linear models.

4.4 Managerial and Practical Implications

From a managerial standpoint, the evidence points to several practical implications:

- Revenue Maximization: DRL provides measurable profit gains with minimal manual intervention once models are deployed.
- Customer Retention: Adaptive strategies appear to foster consumer trust by aligning prices with real-time demand conditions.
- Scalability Across Categories: The consistent performance across electronics, apparel, and home products suggests that DRL can be generalized across diverse product portfolios.
- Ethical Concerns: Firms must ensure transparency and fairness to avoid accusations of discriminatory pricing. Incorporating explainability into DRL models will be critical for gaining regulatory approval and sustaining customer trust [17].

5. CONCLUSION

This study demonstrates that deep reinforcement learning (DRL) constitutes a powerful and scalable framework for dynamic pricing within e-commerce ecosystems. Empirical analyses using Deep Q-Network and Advantage Actor-Critic architectures confirm that DRL consistently outperforms conventional rule-based and regression-based methods across key business metrics, including revenue, conversion rate, and customer retention. These gains are observed across multiple product categories, underscoring the generalizability of DRL-driven strategies beyond niche domains. Beyond short-term profit maximization, the adaptive pricing mechanisms inherent to DRL promote customer loyalty, mitigate inventory volatility, and enhance long-term value creation through data-driven responsiveness to market dynamics. Nevertheless, the practical deployment of DRL in commercial settings presents notable challenges. Effective training demands extensive computational resources and access to high-quality behavioral data, while ethical and regulatory considerations—particularly concerning fairness, transparency, and price discrimination—remain critical. The integration of explainable and fairness-aware AI components within DRL pipelines will be essential to ensure accountability and consumer trust. From a managerial standpoint, the findings position DRL as a transformative enabler of next-generation pricing intelligence, offering firms both competitive and operational advantages. Future research should explore hybrid frameworks that unify DRL with causal inference and fairness-constrained optimization to balance profitability with equity and compliance. Ultimately, by validating DRL models on authentic e-commerce data, this work advances the practical translation of reinforcement learning from theoretical promise to applied business impact, marking a pivotal step toward adaptive, ethically aligned, and sustainable pricing systems in digital commerce.

6. REFERENCES

[1] Li, Y. (2017). Deep reinforcement learning: An overview. arXiv preprint arXiv:1701.07274.

[2] Kastius, A., & Schlosser, R. (2021). Dynamic pricing under competition using reinforcement learning. Journal of Revenue and Pricing Management, 21(1), 50-63.

[3] Zhu, C., Cheng, C., & Meng, S. (2024). DRL PricePro: A deep reinforcement learning framework for personalized dynamic pricing in e-commerce platforms with supply constraints. Spectrum of Research, 4(1).

[4] Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602.

[5] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level control through deep reinforcement learning. nature, 518(7540), 529-533.

[6] Liu, J., Zhang, Y., Wang, X., Deng, Y., & Wu, X. (2019). Dynamic pricing on e-commerce platform with deep reinforcement learning: A field experiment. arXiv preprint arXiv:1912.02572.

[7] Ismail, A., & Baysal, M. (2023). Dynamic Pricing Based on Demand Response Using Actor–Critic Agent Reinforcement Learning. Energies, 16(14), 5469.

[8] Kalusivalingam, A. K., Sharma, A., Patel, N., & Singh, V. (2022). Optimizing e-commerce revenue: Leveraging reinforcement learning and neural networks for AI-powered dynamic pricing. International Journal of AI and ML, 3(9).

[9] Lowe, R., Wu, Y. I., Tamar, A., Harb, J., Pieter Abbeel, O., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. Advances in neural information processing systems, 30.

[10] Deng, S., Schiffer, M., & Bichler, M. (2025). Exploring Competitive and Collusive Behaviors in Algorithmic Pricing with Deep Reinforcement Learning. arXiv preprint arXiv:2503.11270.

[11] Mahdavian, A., Moradi, H., & Bahrak, B. Product Recommendation with Price Personalization According to Customer's Willingness to Pay Using Deep Reinforcement Learning. Available at SSRN 5098967.

[12] Levine, S., Kumar, A., Tucker, G., & Fu, J. (2020). Offline reinforcement learning: Tutorial, review, and perspectives on open problems. arXiv preprint arXiv:2005.01643.

[13] Bose, N., Chopra, A., Joshi, P., & Reddy, A. (2023). Leveraging Reinforcement Learning and Predictive Analytics for Enhanced Customer Lifetime Value Optimization. International Journal of AI Advancements, 12(8).

[14] Patel, D. B. (2022). Reinforcement Learning in Dynamic Pricing Models for E-Commerce. The Es Economics and Entrepreneurship, 1(01), 41-45.

[15] DiMicco, J. M., Maes, P., & Greenwald, A. (2003). Learning curve: A simulation-based approach to dynamic pricing. Electronic Commerce Research, 3(3), 245-276.

[16] Yin, C., & Han, J. (2021). Dynamic pricing model of e-commerce platforms based on deep reinforcement learning. Computer Modeling in Engineering & Sciences, 127(1), 291-307.

[17] Afroogh, S., Akbari, A., Malone, E., Kargar, M., & Alambeigi, H. (2024). Trust in AI: progress, challenges, and future directions. Humanities and Social Sciences Communications, 11(1), 1-30.

[18] Gross, L., Walter, R., Zoppi, N., Justus, A., Gambetti, A., Han, Q., & Kaiser, M. (2025). Cross-Platform E-Commerce Product Categorization and Recategorization: A Multimodal Hierarchical Classification Approach. arXiv preprint arXiv:2508.20013.

[19] Ge, Y., Li, Z., & Zhang, J. A simulation study on missing data imputation for dichotomous variables using statistical and machine learning methods. Sci Rep. 2023.

[20] Benoit, K. (2011). Linear regression models with logarithmic transformations. London School of Economics, London, 22(1), 23-36.

[21] Sagar, S., Narayanan, V., Binu, D., Selby, N., & Thomas, S. E. Advantage Actor-Critic Reinforcement Learning with Technical Indicators for Stock Trading Decisions.

[22] Liu, N., & Hu, D. (2025). The design of consumer behavior prediction and optimization model by integrating DQN and LSTM. PloS one, 20(7), e0327548.

[23] Irani, H., Ghahremani, Y., Kermani, A., & Metsis, V. (2025). Time Series Embedding Methods for Classification Tasks: A Review. Expert Systems, 42(11), e70148.

[24] Irani, H., & Metsis, V. (2024, May). Enhancing time-series prediction with temporal context modeling: A Bayesian and deep learning synergy. In The International FLAIRS Conference Proceedings (Vol. 37).

[25] Ataei, S., Ataei, S. T., & Saghiri, A. M. (2025). Applications of Deep Learning to Cryptocurrency Trading: A Systematic Analysis.

[26] Ataei, S., Ataei, S. T., & Saghiri, A. M. (2025). A Systematic Review on the Application of Artificial Intelligence in Decentralized Finance.