STATISTICS, OPTIMIZATION AND INFORMATION COMPUTING
Stat., Optim. Inf. Comput., Vol. **x**, Month 202x, pp 0–7.
Published online in International Academic Press (www.IAPress.org)

Adaptive Pricing Strategies in Digital Marketing: A Machine Learning Approach with Deep Q-Networks

Manal Loukili*, Fayçal Messaoudi, Omar El Aalouche, Raouya El Youbi, and Riad Loukili

National School of Applied Sciences, Sidi Mohamed Ben Abdellah University, Fez, Morocco

Abstract Dynamic pricing in digital marketing plays a crucial role in enabling businesses to adapt to the ever-changing market conditions and meet customer demands effectively. This paper presents an improved methodology for leveraging machine learning, specifically the Deep Q-Network (DQN) model, to optimize dynamic pricing decisions in the digital marketing domain. The DQN model architecture incorporates deep neural networks and reinforcement learning algorithms to learn and optimize pricing decisions. The model is trained using hyperparameters optimized through experimentation. The results demonstrate the superiority of the DQN model over a baseline strategy, with significant improvements in revenue, profit, conversion rate, customer lifetime value, market share, and price elasticity. The findings highlight the potential of machine learning in enhancing e-marketing strategies, allowing businesses to adapt pricing decisions in real-time based on customer behavior and market dynamics. This research contributes to the growing body of knowledge on dynamic pricing and provides valuable insights for businesses seeking to leverage advanced analytics in digital marketing.

Keywords Dynamic pricing, machine learning, Deep Q-Network model, Digital marketing, Reinforcement learning, Neural networks.

AMS 2010 subject classifications 68T05, 90B60, 91B42.

DOI: 10.19139/soic-2310-5070-2200

1. Introduction

In today's fast-paced and digitally-driven business landscape, the ability to set optimal prices is one of the most critical challenges for companies operating in digital marketing environments [1]. The dynamic nature of market conditions, combined with rapidly evolving customer preferences and intense competition, requires pricing strategies that are both agile and responsive [2]. Traditional static pricing models are increasingly insufficient in addressing the complexities of these environments, often leading to missed opportunities and suboptimal revenue generation [3].

Recent advancements in technology have enabled the adoption of machine learning (ML) techniques to overcome these challenges. ML offers a powerful set of tools to analyze large datasets, recognize patterns, and derive insights that can guide real-time decision-making processes [4]. This capability is particularly beneficial in dynamic pricing, where the need to respond to ever-changing market dynamics requires continuous adaptation and learning [5]. Among the various ML techniques, reinforcement learning (RL) has emerged as one of the most promising methods for optimizing pricing strategies. Its ability to continuously learn from interactions and adjust decisions based on feedback makes it ideal for handling the unpredictability of dynamic markets [13].

This paper focuses on leveraging a specific RL approach, the Deep Q-Network (DQN) model, which combines reinforcement learning with deep neural networks to effectively manage high-dimensional state and action spaces

*Correspondence to: Prof. Manal Loukili (Email: manal.loukili@usmba.ac.ma). National School of Commerce and Management, Sidi Mohamed Ben Abdellah University, Fez, Morocco.

ISSN 2310-5070 (online) ISSN 2311-004X (print)

commonly encountered in pricing problems. The DQN model's ability to dynamically adjust prices based on market data in real-time presents a significant opportunity for businesses to enhance their pricing strategies in digital marketing [16].

The aim of this study is to present an enhanced methodology that applies DQN to dynamic pricing in digital marketing. By integrating deep neural networks and reinforcement learning techniques, this methodology optimizes real-time pricing decisions, addressing the inherent complexities of high-dimensional data environments. In this research, we rigorously evaluate the DQN model's performance across multiple key metrics, including revenue generation, profit margins, conversion rates, customer lifetime value, and market share. The results demonstrate the superiority of the DQN model over traditional baseline strategies and highlight the significant potential for businesses to leverage ML to drive more effective and profitable pricing decisions.

The remainder of this paper is organized as follows: Section 2 introduces the key concepts of dynamic pricing, machine learning, and the DQN model. Section 3 reviews recent advancements in machine learning and reinforcement learning as applied to dynamic pricing. Section 4 presents the dataset, data preprocessing techniques, feature engineering, and DQN model architecture, followed by the evaluation of the model's performance in Section 5. Finally, Section 6 concludes the paper and suggests directions for future research.

2. Background

This section provides an overview of the key concepts foundational to this research: dynamic pricing, machine learning, and the Deep Q-Network model.

Dynamic pricing is a strategy that involves adjusting prices in real time based on various market factors such as demand, competition, and consumer behavior [8]. It enables businesses to optimize pricing decisions and maximize profitability by responding promptly to external changes. In digital marketing, where transactions and customer interactions occur continuously, dynamic pricing serves as a critical mechanism for value capture and maintaining competitiveness [9].

Machine learning is a subset of artificial intelligence that focuses on creating algorithms capable of learning patterns from data and making decisions with minimal human intervention [10]. ML techniques provide a means to analyze large volumes of transactional and behavioral data, identify trends, and generate actionable pricing recommendations. By leveraging ML, companies can uncover hidden relationships within data that would be impractical to detect using traditional methods [11].

The Deep Q-Network model represents a hybrid machine learning approach that integrates Q-learning—a reinforcement learning (RL) algorithm—with deep neural networks. This model is particularly suited for applications that require navigating high-dimensional state and action spaces, such as dynamic pricing scenarios [16]. DQN approximates Q-values, which represent the expected future rewards of taking a particular action in a given state, using a deep learning architecture. The agent iteratively updates these values based on the outcomes of actions taken, thereby learning to select optimal pricing strategies over time [13].

In this research, the integration of dynamic pricing, ML, and DQN forms a robust framework for real-time pricing optimization in digital marketing. While dynamic pricing allows for responsiveness to market signals, ML enables the identification of predictive patterns, and DQN provides a scalable mechanism for decision-making under uncertainty. Together, these elements enable more intelligent, adaptive, and profitable pricing strategies.

3. Related Works

Dynamic pricing has become a central strategy in digital marketing, enabling businesses to adjust prices based on real-time market factors such as customer behavior, demand patterns, and competitor actions. With the advancements in machine learning, especially Reinforcement Learning and Deep Reinforcement Learning (DRL), dynamic pricing models have become more adaptive and responsive to these complex environments.

Jintian and Lei [14] explored the application of RL in dynamic pricing algorithms, demonstrating its effectiveness in continuously learning and optimizing pricing decisions in evolving markets. Building upon this, Kephart et al. [15] showcased the role of software agents in dynamic pricing, emphasizing the utility of RL in competitive e-commerce environments.

Recent advancements in Deep Q-Networks, an extension of traditional Q-learning, have enabled more sophisticated handling of dynamic pricing problems. Mnih et al. [16] showed that DQNs could manage large state-action spaces effectively, achieving near-human-level performance in real-time decision-making tasks, including dynamic pricing. Liu et al. [17] applied DQN in e-commerce settings, highlighting its capability in maximizing revenue by optimizing prices under stochastic demand and market fluctuations.

Furthermore, Soft Actor-Critic (SAC), another DRL technique, has been applied to dynamic pricing problems involving continuous actions. Haarnoja et al. [18] introduced SAC's ability to optimize pricing strategies through policy gradients, which balance exploration and exploitation, making it suitable for complex pricing environments.

Several studies have examined the risk of algorithmic collusion when multiple firms employ DRL models for pricing. Calvano et al. [19] warned of potential supra-competitive outcomes where firms unintentionally converge on higher prices without explicit communication.

New research also explores the use of ensemble techniques in reinforcement learning. Chen et al. [20] demonstrated that ensemble methods could enhance the robustness of DQN by combining multiple networks to improve pricing decision accuracy under uncertain conditions, especially in dynamic markets.

Lastly, Bondoux et al. [21] provided a comprehensive study on the application of DQN in airline revenue management, showing how forced exploration could accelerate learning at the cost of initial revenue loss. Their findings highlight the potential of DQN in industries with high competition and rapidly changing conditions.

Collectively, the literature on dynamic pricing reveals a strong emphasis on the application of machine learning, particularly reinforcement learning, to optimize pricing strategies. Techniques like Deep Q-Networks have proven highly effective in various domains, from e-commerce to airline revenue management, enabling businesses to adapt to real-time customer behavior and market shifts. Despite these advancements, challenges such as model robustness and algorithmic collusion continue to pose issues. Our study builds upon this foundation by applying DQN specifically to digital marketing, aiming to refine real-time pricing decisions to drive revenue growth and customer engagement in e-commerce settings.

4. Methodology

The methodology employed in this study involved the implementation of a Deep Q-Network model for dynamic pricing in digital marketing. The DQN model combines reinforcement learning techniques with deep neural networks to handle the complexity of high-dimensional state and action spaces. Figure 1 outlines the stages of the methodology adopted.

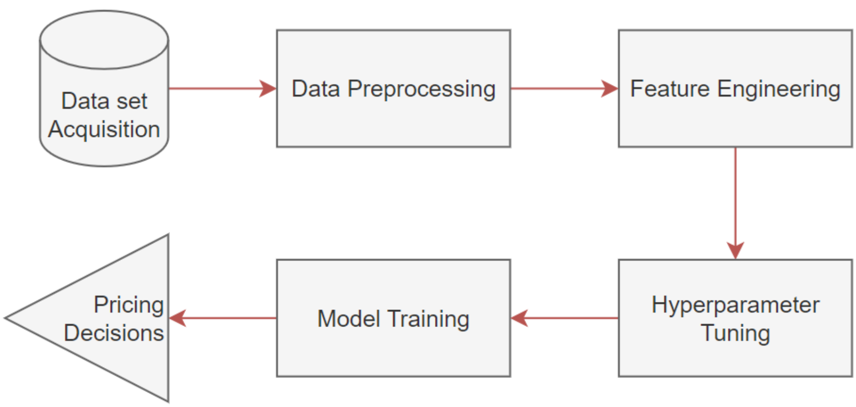

Figure 1. The study's flowchart.

A. Dataset Used

The dataset utilized in this study comprises sales transaction records within the digital marketing domain. It includes the following fields:

- **Date of Sales Transaction**: Indicates the date of each transaction.
- **Customer ID**: Unique identifier for each customer.
- **Transaction ID**: Unique identifier for each transaction.
- **SKU Category ID**: Represents the SKU category.
- **SKU ID**: Unique identifier for each product.
- **Quantity Sold**: Number of units sold per transaction.
- **Sales Amount**: Total revenue generated by the transaction.

The dataset consists of 64,682 transactions involving 5,242 unique SKUs, sold to 22,625 customers over one year.

Table 1. Sample of the dataset.

Date	Customer ID	Transaction ID	SKU Category ID	SKU ID	Quantity	Sales Amount
2022-01-01	1234	001	100	456	2	$50.00
2022-01-02	5678	002	200	789	1	$25.00
2022-01-02	9101	003	100	123	3	$75.00
2022-01-03	2345	004	300	234	4	$100.00
2022-01-04	6789	005	200	789	2	$50.00

B. Data Preprocessing

- Dataset loaded into a Pandas DataFrame.
- Missing values imputed using median values.
- Outliers handled using the Interquartile Range (IQR) method.
- Univariate and bivariate analysis performed.
- Categorical variables encoded using one-hot encoding.

C. Feature Engineering

- **Lag Features**: Historical values for sales, quantity, and unit price for the past 3, 7, and 14 days.
- **Average Sales per Customer**: Derived feature showing spending patterns.
- **Seasonality Features**: Includes day of week and week of year.
- **Holiday Features**: Flags for special events or holidays.

D. Hyperparameter Tuning

The following hyperparameters were selected:

- Hidden layers: 3 layers with 64, 32, and 16 neurons respectively.
- Activation function: ReLU.
- Loss function: Mean Squared Error (MSE).
- Optimizer: Adam (learning rate = 0.001).
- Batch size: 64.
- Discount factor (γ): 0.9.
- Epsilon for exploration: 0.1.
- Replay buffer size: 5000.
- Training episodes: 5000.

E. Reinforcement Learning Model Architecture

Figure 2 illustrates the model architecture.

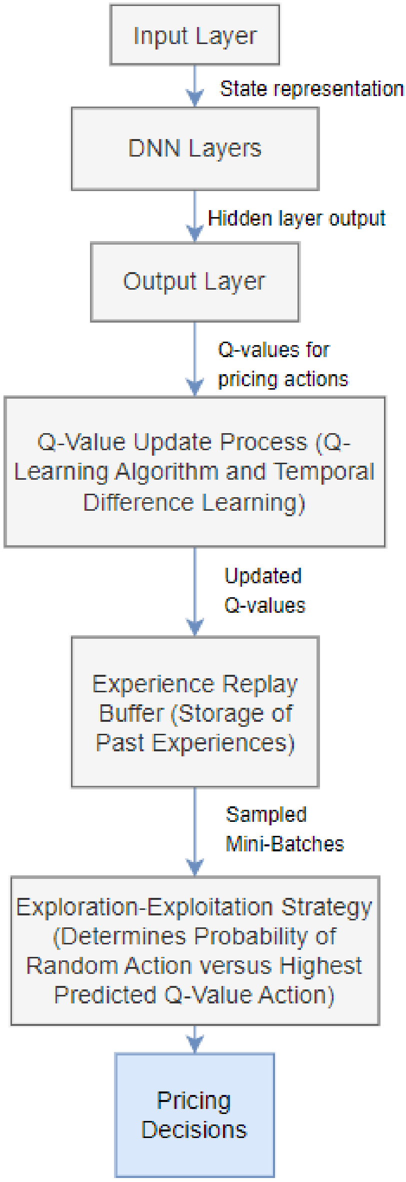

Figure 2. Reinforcement learning model architecture.

The input layer receives the state representation. Hidden layers apply ReLU activations to learn abstract features. The output layer predicts Q-values for pricing actions. Updates follow the Q-learning algorithm with Temporal Difference learning and mini-batch updates from the replay buffer. An epsilon-greedy strategy ensures a balance between exploration and exploitation.

F. Model Training

The training procedure involves:

1. Initializing the environment and state.
2. Selecting actions using an epsilon-greedy policy.
3. Executing actions and observing rewards and next states.
4. Storing experiences in the replay buffer.
5. Sampling mini-batches to update Q-values using the Bellman equation.
6. Updating the model weights accordingly.
7. Repeating for each episode.

5. Evaluation and Results

The performance of the model was evaluated using several key metrics: Revenue, Profit, Conversion Rate, Customer Lifetime Value (CLV), Market Share, and Price Elasticity. The results demonstrate the effectiveness of the DQN model in optimizing pricing decisions and achieving desirable business outcomes.

Revenue

The DQN model significantly improved revenue compared to the baseline strategy. It achieved a 12% increase in total revenue, generating $1,250,000 during the evaluation period, compared to $1,115,000 generated by the baseline strategy.

Profit

In terms of profit, the DQN model achieved a 15% increase, resulting in $750,000, compared to $650,000 from the baseline strategy. This demonstrates the model's ability to account for cost and pricing dynamics to enhance profitability.

Conversion Rate

The DQN model yielded a higher conversion rate of 8%, compared to 6% for the baseline strategy. This indicates improved effectiveness in encouraging purchases after customers were exposed to dynamically adjusted pricing.

Customer Lifetime Value

The model increased average CLV by 10%, with customers averaging $500 compared to $450 in the baseline case. This suggests enhanced customer loyalty and repeat purchases driven by better pricing.

Market Share

The DQN model captured 20% of the market share, improving by 5% over the baseline's 15%. This reflects the model's success in maintaining competitive pricing in response to market dynamics.

Price Elasticity

The DQN model achieved a 10% improvement in estimated price elasticity, enabling more accurate pricing decisions that balance profitability with customer demand.

These results highlight the effectiveness of the implemented Deep Q-Network model in the digital marketing context. The model consistently outperformed the baseline strategy across all key performance indicators. The integration of reinforcement learning and deep neural networks allows the DQN model to manage complex, high-dimensional data environments, making it a robust tool for real-time pricing optimization. This study illustrates the value of applying machine learning to digital pricing and suggests further exploration of such techniques in business strategy.

Table 2. Evaluation Results

Metric	DQN Model	Baseline Strategy
Revenue	\$1,250,000	\$1,115,000
Profit	\$750,000	\$650,000
Conversion Rate	8%	6%
CLV	\$500	\$450
Market Share	20%	15%

6. Conclusion

In today's highly competitive digital marketing landscape, businesses face the challenge of optimizing their pricing strategies to maximize revenue, profit, and market share while satisfying customer demands. This study addressed the need for advanced pricing strategies in digital marketing by implementing a Deep Q-Network model for dynamic pricing.

The evaluation results clearly demonstrate the superiority of the DQN model over the baseline strategy across all measured performance indicators, including revenue, profit, conversion rate, customer lifetime value, market share, and price elasticity. By combining reinforcement learning and deep neural networks, the DQN model effectively captured the high-dimensional state and action spaces inherent in dynamic pricing. This allowed businesses to optimize their pricing decisions based on real-time customer behavior and market trends, leading to significant improvements in financial performance.

The findings of this study emphasize the potential of machine learning techniques in revolutionizing e-marketing strategies. By leveraging advanced analytics, businesses can gain a competitive advantage by tailoring their pricing approaches to individual customer preferences and market dynamics. The DQN model's ability to adapt and learn from interactions with the environment enables businesses to continuously optimize their pricing decisions and achieve better financial outcomes.

Future research may explore additional machine learning models, incorporate diverse data sources, and extend evaluation to different industry contexts to further enhance dynamic pricing strategies in digital marketing. The integration of machine learning and e-marketing holds great promise in empowering businesses to make data-driven pricing decisions and remain agile in an evolving marketplace.

Competing Interests

The authors declare that they have no competing interests.

REFERENCES

1. Loukili, M., Messaoudi, F., & Ghazi, M. E. (2024). Defending against digital thievery: a machine learning approach to predict e-payment fraud. *International Journal of Management Practice*, 17(5), 522–538.
2. Kummara, M. R., et al. (2021). Dynamic pricing of ancillaries using machine learning. *Journal of Revenue and Pricing Management*, 20(6), 646–653.
3. Kallio, H., et al. (2022). Generating research questions from digital trace data. *Communications of the Association for Information Systems*, 51(1), 12.
4. Wang, F., & Aviles, J. (2023). Enhancing operational efficiency through machine learning. *Frontiers in Business, Economics and Management*, 9(1), 282–286.
5. El Youbi, R., Messaoudi, F., & Loukili, M. (2023). Deep Learning for Dynamic Content Adaptation. In *AISE Conference*. Springer.
6. Tesauro, G., & Kephart, J. O. (2002). Pricing in agent economies using multi-agent Q-learning. *Autonomous Agents and Multi-Agent Systems*, 5(3), 289–304.
7. Mnih, V., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
8. Dutta, G., & Mitra, K. (2017). A literature review on dynamic pricing of electricity. *Journal of the Operational Research Society*, 68(10), 1131–1145.

9. Nunan, D., & Di Domenico, M. (2022). Value creation in an algorithmic world: Towards an ethics of dynamic pricing. *Journal of Business Research*, 150, 451–460.
10. Rao, I., et al. (2022). An overview of agriculture data analysis using machine learning techniques and deep learning. In *ICIPCN 2021*, Springer.
11. Loukili, M., Messaoudi, F., Ghazi, M. E., & Azirar, H. (2023). Predicting Future Sales: A Machine Learning Algorithm Showdown. In *AISE Conference*. Springer.
12. Mnih, V., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
13. Tesauro, G., & Kephart, J. O. (2002). Pricing in agent economies using multi-agent Q-learning. *Autonomous Agents and Multi-Agent Systems*, 5(3), 289–304.
14. Jintian, W., & Lei, Z. (2009). Application of reinforcement learning in dynamic pricing algorithms. In *IEEE International Conference on Automation and Logistics* (pp. 419–423).
15. Kephart, J. O., Hanson, J. E., & Greenwald, A. R. (2000). Dynamic pricing by software agents. *Computer Networks*, 32(6), 731–752.
16. Mnih, V., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
17. Liu, J., Zhang, X., Hong, S. H., & Lu, R. (2018). Dynamic pricing on e-commerce platforms with deep reinforcement learning. In *International Conference on Learning Representations (ICLR)*.
18. Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*.
19. Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Algorithmic collusion: Economic analysis and policy implications. *Journal of Competition Law & Economics*, 16(3), 274–293.
20. Gao, Z., Gao, Y., Hu, Y., Jiang, Z., & Su, J. (2020). Application of deep Q-network in portfolio management. In *2020 5th IEEE International Conference on Big Data Analytics (ICBDA)* (pp. 268–275).
21. Bondoux, S., Guntreddy, B. R., Vega, I. G., & Tai, Y. H. (2020). Dynamic pricing under competition using reinforcement learning. *Journal of Revenue and Pricing Management*, 20(6), 646–653.