WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Optimizing Dynamic Pricing with Deep Reinforcement Learning: A Comprehensive Review

Nagashree J [1], Ravikumar M L[2]

*[*1,2] PG Scholar, MTech in Artificial Intelligence*
*[*1,2] REVA Academy for Corporate Excellence (RACE)*
*[*1,2] REVA University, Bengaluru, India*
Doi : https://doi.org/10.55248/gengpi.5.0924.2702

A B S T R A C T

This literature review explore various reinforcement learning (RL)-based approaches to dynamic pricing (DP) and demand response (DR) in complex market environments. Utilizing deep reinforcement learning (DRL), these studies propose models that frame DP problems as Markov Decision Processes (MDPs), leveraging continuous price sets and newly defined reward functions, such as the difference of revenue conversion rates (DRCR), to outperform traditional revenue-based methods. Key contributions include addressing the cold-start problem through pre-training and product clustering, enabling better handling of sparse data. Additionally, multi-agent systems are employed to optimize energy consumption in demand response scenarios, while adaptive mechanisms are introduced to maintain system constraints and consumer comfort. These frameworks demonstrate significant improvements in pricing strategies, leading to enhanced revenue generation and operational efficiency, as evidenced by real-world applications, such as field experiments on platforms like Alibaba and Tmall. The collective insights highlight the versatility and effectiveness of RL techniques in dynamic and unpredictable markets.

Keywords: Machine Learning Algorithms, Statistical Modelling, Deep Reinforcement Learning (DRL) Framework, Markov Decision Process (MDP), Q-learning, Difference of Revenue Conversion Rates (DRCR), Multi-Armed Bandit Epsilon-Greedy, UCB1, and Thompson Sampling.

Introduction

The dynamic pricing (DP) problem, a key strategy for optimizing revenue in competitive markets, has gained increasing attention in various industries. Traditional pricing models, such as auctions, price discrimination, and first-come-first-served mechanisms, are being enhanced by cutting-edge approaches like reinforcement learning (RL), especially deep reinforcement learning (DRL). In online environments, DP strategies aim to maximize profit while balancing customer satisfaction and competitive pricing. However, newly introduced products, which often suffer from sparse data, pose challenges related to initial price setting, predicting customer and competitor reactions, and managing price adjustments across different time periods. To address these complexities, researchers have integrated techniques like clustering, transfer learning, and RL to develop more robust pricing strategies.

In the context of energy demand management, DP plays a crucial role in demand response (DR) programs, particularly in smart grids. DR programs aim to reduce peak load by adjusting consumption patterns based on dynamic price signals, providing benefits to both consumers and power system operators. Challenges in DR programs include managing thermal loads, addressing privacy concerns, and adhering to market regulations. Reinforcement learning has been proposed as a solution to optimize dynamic pricing in DR, allowing DR aggregators to generate pricing policies that meet both consumer comfort preferences and grid requirements. Moreover, recent applications of RL in e-commerce, such as Amazon's automated pricing system, demonstrate its ability to dynamically adjust prices in real time, overcoming challenges posed by fluctuating demand and market uncertainty. By employing DRL, researchers have made significant strides in developing adaptive pricing systems capable of long-term revenue optimization, real-world market applicability, and handling both discrete and continuous price sets.

In the energy sector, dynamic pricing is crucial for demand response (DR) programs, where electricity prices are adjusted to manage consumption patterns and reduce peak demand. These programs rely on capturing the flexibility of residential loads, particularly thermostatic devices, to improve grid reliability and operational costs. DR Aggregators (DRAs) act as intermediaries between consumers and the grid, generating price policies to incentivize reduced energy usage during peak times without compromising consumer comfort. Reinforcement learning has been applied to optimize these pricing policies, particularly in complex environments where user preferences and grid constraints must be balanced. By incorporating these advanced techniques, both the retail and energy sectors demonstrate the potential of dynamic pricing and reinforcement learning to enhance operational efficiency and profitability in uncertain market conditions.

Literature Review:

Dynamic pricing (DP) has gained considerable attention in recent years due to its ability to optimize revenue in changing and unpredictable markets. Several studies have employed Deep Reinforcement Learning (DRL) and Reinforcement Learning (RL) approaches to model DP problems as Markov Decision Processes (MDPs), with notable innovations in handling state representation, reward function design, and price-setting strategies.

1. Dynamic Pricing on E-Commerce Platform With Deep Reinforcement Learning: A Field Experiment

The authors propose a comprehensive DRL-based framework for DP in e-commerce, extending previous work by incorporating continuous price sets and introducing the Difference of Revenue Conversion Rates (DRCR) as a reward function. This model, tested in both offline and online field experiments using Alibaba's dataset, addresses the cold-start problem through pre-training with historical data. The results demonstrate superior performance in continuous pricing compared to manual methods.

2. Dynamic Pricing Using Reinforcement Learning

This paper offers an overview of the dynamic pricing challenge, emphasizing the importance of learning models like DRL in uncertain market environments. These models help in adapting to shifting consumer behaviours and external factors, with the potential to enhance company revenues by leveraging historical data patterns. DRL's flexibility makes it suitable for various scenarios, including parking prices and traffic congestion management.

3. Deep reinforcement learning based dynamic pricing for demand response considering market and supply constraints

This applies RL to a price-based Demand Response (DR) program for the energy sector, showcasing the utility of RL in managing consumption patterns in residential sectors. The model includes a multi-agent system for load flexibility and demonstrates how a price generator function can be used to optimize power capacity and profits while maintaining customer comfort. Comparative studies show the effectiveness of this RL-driven approach.

4. Enhancing Sparse Data Performance in E-commerce Dynamic Pricing with Reinforcement Learning and Pre-Trained Learning framework

In this article is designed to address dynamic pricing in low-traffic product scenarios, with a focus on overcoming the cold-start problem. The authors highlight the importance of product clustering and pre-trained learning to enhance the performance of predictive models in sparse data conditions.

Overall, these studies underline the promise of DRL and RL in dynamic pricing across diverse sectors, improving both business outcomes and consumer engagement by effectively managing price volatility and system constraints.

5. Dynamic Pricing with Multi-Armed Bandits

In this article, the author explores four Multi-armed Bandit algorithms to evaluate their efficacy against a well-defined (though not straightforward) demand curve. These studies then dissect the primary strengths and limitations of each algorithm and delve into the key metrics that are instrumental in gauging their performance.

Methodology

Paper 1: **Dynamic Pricing on E-Commerce Platform with Deep Reinforcement Learning: A Field Experiment**

The methodology of the literature review focuses on dynamic pricing applications in e-commerce platforms, particularly markdown pricing and daily pricing. It models the pricing process as a Markov Decision Process (MDP), where prices are adjusted at discrete time steps. Products are priced individually, and their states are described by four feature groups: price, sales, customer traffic, and competitiveness. The action space, either discrete or continuous, defines price adjustments based on historical data.

Two reward functions are explored: revenue conversion rate and a new difference-based function (DRCR). The methodology demonstrates that using revenue conversion rate is effective for luxury products but less so for fast-moving consumer goods (FMCGs) due to unstable sales and inventory dynamics. For model training, two reinforcement learning approaches are employed: Q-learning with Deep Q-networks (DQN) for discrete pricing, and actor-critic algorithms (DDPG) for continuous pricing. Pre-training strategies are implemented using historical pricing decisions as demonstrations to address the cold start problem.

In summary, this approach emphasizes balancing between pricing adjustments and customer behaviour dynamics, while leveraging advanced RL techniques to optimize pricing policies for different product categories in e-commerce.

Discrete pricing action models

To solve the dynamic pricing MDP we defined above, we first use Q-learning ([30]) to find the optimal pricing policy. Q-learning is a value iteration method to compute the optimal policy. It starts with randomly initialled Q value and recursively iterates using the transitions $t = (s, a, r, s')$ to get the optimal $Q*$ as well as the optimal policy:

$Q_{t+1}(s, a) \leftarrow (1 - \alpha) \cdot Q_t(s, a) + \alpha \cdot [r + \gamma \cdot \max Q_t(s', a')]$

where $\alpha \in (0, 1]$ is the learning rate and γ is the discount factor.

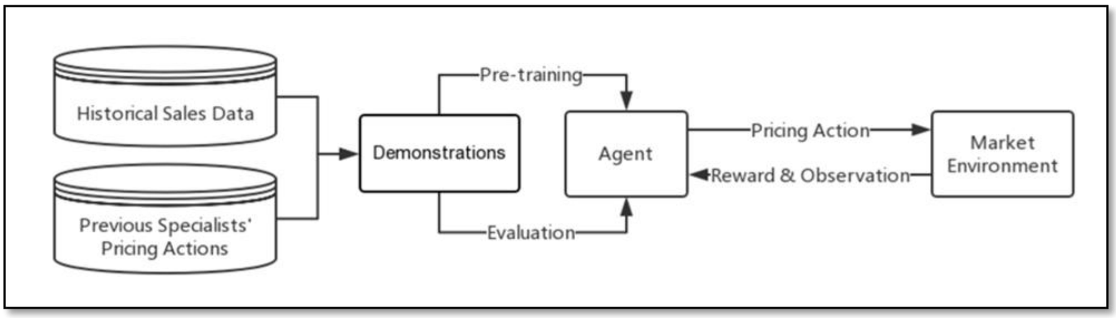

Fig 1: Dynamic pricing framework using DRL with demonstrations on E-commerce platform

Paper 2: **Dynamic Pricing Using Reinforcement Learning**

This study applies **Q-learning**, a reinforcement learning technique, to optimize dynamic pricing strategies aimed at maximizing long-term revenue. The methodology involves several key components:

1. **State Representation**: The system defines relevant states that encapsulate the current environment and market conditions, such as customer demand, purchase history, or competitor prices.
2. **Action Representation**: Each action corresponds to a specific pricing decision that can be adjusted over time based on feedback from the environment.
3. **Q-table Initialization**: A Q-table is initialized to store Q-values for each state-action pair. These Q-values represent the expected cumulative reward (revenue) associated with different pricing decisions in various states.
4. **Exploration-Exploitation Strategy**: The **epsilon-greedy strategy** is employed to manage the tradeoff between exploration and exploitation. With a probability of ε, the agent explores new pricing strategies, while with a probability of 1−ε, it exploits the pricing action with the highest Q-value for a given state.
5. **Reward Definition**: Rewards are defined in terms of revenue generation, where higher prices might lead to fewer purchases but larger profits per sale, while lower prices might result in more frequent purchases but lower per-unit revenue.
6. **Q-Value Update**: After each pricing decision, the Q-value for the chosen state-action pair is updated using the **Q-learning update rule**, adjusting the expected cumulative reward based on actual outcomes from the environment.

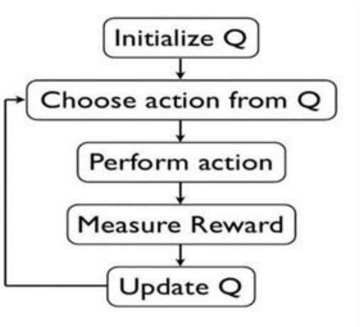

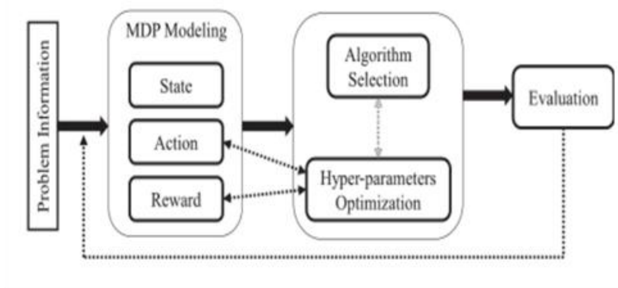

Fig 3: Q-Learning on Dynamic Pricing Architecture

Balancing **exploration** and **exploitation** are central to this methodology. Exploration allows the agent to discover new pricing actions that may lead to higher rewards, while exploitation focuses on selecting actions that have historically yielded positive results.

Additionally, a **Bayesian model-based approach** is integrated to further refine the dynamic pricing process by incorporating prior knowledge and continuously updating price strategies based on observed data. This combination of Q-learning and Bayesian methods ensures a robust, adaptive pricing system capable of responding to evolving market dynamics.

Paper 3: **Deep reinforcement learning based dynamic pricing for demand response considering market and supply constraints**

The dynamic pricing methodology outlined in the paper focuses on addressing demand response (DR) in residential power systems. To achieve this, the authors develop a multi-agent framework that utilizes reinforcement learning (RL) to optimize the pricing strategy while accounting for energy market constraints and capacity limitations.

The methodology begins by defining a **price generator function** for the Demand Response Aggregator (DRA), which sets price limits between a minimum (π_min) and maximum (π_max) value. This generator function, combined with a capacity factor M that reflects the power grid's physical limitations, enables the DRA to adjust pricing based on real-time demand. The DRA interacts with residential agents, which alter their consumption behaviour in response to dynamic price signals, helping reduce peak demand.

The pricing model is formulated as a **Stackelberg game**, where the DRA (leader) sets prices and residential agents (followers) react to minimize their energy costs. This dynamic interaction is supported by a **Proximal Decomposition Algorithm** to ensure convergence of the system. The RL method, specifically the **Proximal Policy Optimization (PPO) algorithm**, helps the DRA optimize the parameters of the price generator function by continuously learning from the agents' behaviour and adapting the price signals in a feedback loop.

This methodology is validated by simulations that illustrate how the RL-based framework enhances demand flexibility, reduces operational costs, and maximizes the DRA's profit, all while ensuring grid reliability. The **use of RL** allows the system to handle non-linearities and user deviations from consumption plans effectively, providing a robust solution to the dynamic pricing challenge in energy markets.

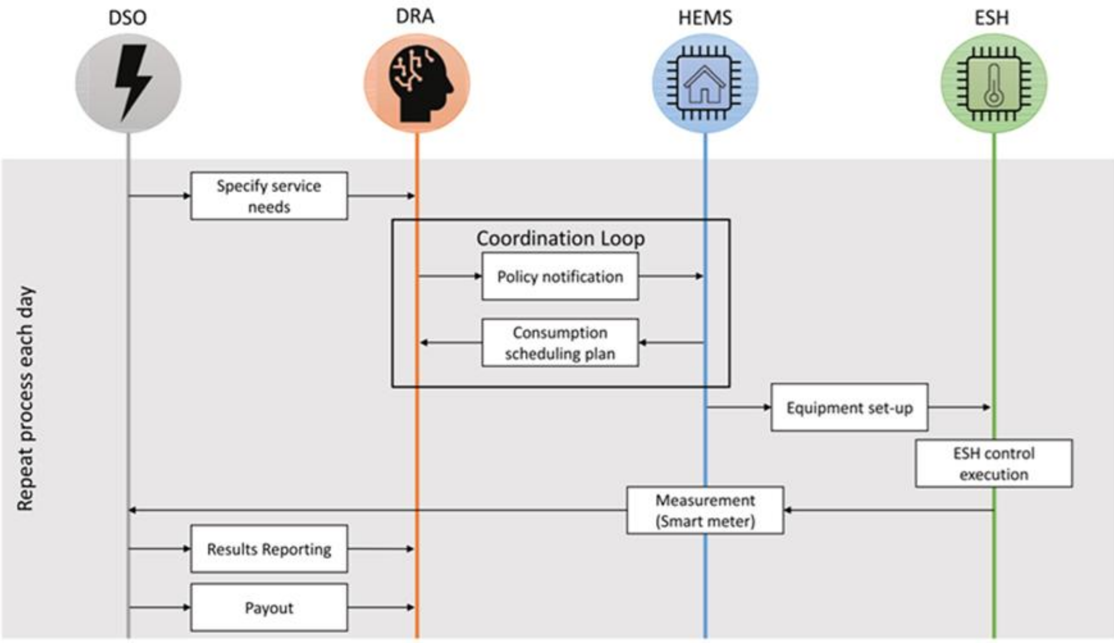

Fig 2: Automatic price-based DR sequence.

Author proposes the following price generator function,

$$\pi_k(y_k) = \pi_{min} + \frac{\pi_{max} - \pi_{min}}{1 + exp(\frac{-y_k+M}{\alpha})}$$

where y_k represents the aggregate consumption at time stamp $k \in \{1, \dots, N\}$. This value corresponds to the sum of individual household energy consumption, i.e. , where H represents the number of houses, and $u^i{}_k$ is the energy consumption of the i^{th} house at the time stamp k.

$$y_k = \sum_{i=1}^{H} u_k^i$$

Lastly, α is a positive parameter that controls the rate of price change. To properly determine this value, exploration must be conducted by the DRA agent due to the lack of existing information linked to the relationship between the users' elasticity and flexibility. The proposed price generator function, $\pi_k(y_k)$, has some particular properties that make it suitable for reducing aggregate load peaks of the aggregated demand profile.

Paper 4: **Enhancing Sparse Data Performance in E-commerce Dynamic Pricing with Reinforcement Learning and Pre-Trained Learning framework**

The literature review focuses on addressing the challenges of pricing optimization for newly introduced products in e-commerce, which often lack historical data. To mitigate this, clustering methods, specifically K-means, are explored as a strategy to reduce exploration by grouping similar products, which can then serve as benchmarks for reinforcement learning models.

1. **Clustering:**

- K-means is selected due to its simplicity and efficiency. It divides products into clusters based on features like brand, size, material, and color.
- This clustering method is applied to a dataset of quilt sets from Amazon, chosen for their uniform features. Feature engineering is performed using one-hot encoding and normalization, followed by K-means clustering with K=5 to group the products.
- The clustered groups allow the reinforcement learning model to identify appropriate pricing strategies for new products based on similar existing ones.

2. **Reinforcement Learning:**

- The pricing problem is modelled as a Markov Decision Process (MDP) consisting of state, action, and reward components.
- The Sarsa algorithm is used, which follows an on-policy temporal difference (TD) control method. This approach allows the system to learn optimal pricing strategies by updating the Q-values based on the ε-greedy policy.
- Price elasticity of demand is integrated to further refine the pricing model, helping capture the relationship between price changes and sales volume.

3. **Transfer Learning:**

- Transfer learning is employed to expedite learning by transferring knowledge from similar domains, which helps overcome the problem of sparse rewards and insufficient data.

By combining clustering and reinforcement learning, the methodology addresses data scarcity in dynamic pricing and demonstrates how these techniques can optimize pricing for newly introduced products in e-commerce.

Paper 5: **Dynamic Pricing with Multi-Armed Bandits**

This methodology examines the application of **Multi-armed Bandit (MAB) algorithms** and demand curve modelling in the context of dynamic pricing, especially for digital goods. The key challenge is determining the price that maximizes revenue, not just sales volume. The review highlights how different MAB strategies can help balance exploration and exploitation in identifying optimal pricing strategies.

1. **Modelling the Demand Curve:**

- Traditional demand curves depict the relationship between price and quantity; however, in dynamic pricing, the focus is shifted to modelling price against the probability of purchase.
- The demand curve is represented using a **logistic function**, which captures the nonlinear relationship between price and probability.
- The model calculates expected revenue by multiplying price by the probability of purchase. The **optimal price** is derived from this formula using calculus, allowing for the identification of the price that maximizes revenue.

2. **Multi-Armed Bandit (MAB) Algorithms:** Several MAB strategies are explored to dynamically adjust pricing:

- **Greedy Algorithm:** Selects the price that initially provides the highest reward, but suffers from insufficient exploration of other prices.
- **ϵ-greedy Algorithm:** Introduces a small probability of random price selection to promote exploration while primarily exploiting the best-known price.
- **Upper Confidence Bound (UCB1):** Balances exploitation and exploration by considering both the average reward and an "uncertainty bonus" for prices not frequently tested.
- **Thompson Sampling:** Uses Bayesian inference to probabilistically choose prices based on their posterior reward distributions. A modified version incorporates price into the sampling process, focusing on expected revenue rather than purchase probability.

By leveraging these MAB strategies, the methodology seeks to determine the price that maximizes revenue rather than solely focusing on maximizing purchases. This is achieved through a blend of exploration (testing untried prices) and exploitation (utilizing known rewards). Each algorithm has distinct strengths and weaknesses in balancing this trade-off.

$$\frac{a}{1 + e^{b \cdot price}}$$

4. Result and Discussion:

This high-level diagram illustrates the workflow of dynamic pricing using **Deep Reinforcement Learning (DRL)**. It breaks down the process into different components, starting from collecting market data to generating optimal pricing strategies and adapting to market changes.

1. Market Environment (Input Layer)

This layer represents the real-time data collection process from multiple sources. The dynamic pricing model requires detailed and continuous data to understand the current state of the market, consumer behaviour, and external factors like competitor prices. This data acts as the "state" for the reinforcement learning model to make informed pricing decisions.

2. Reinforcement Learning Model (Core Framework)

The reinforcement learning model is the engine of dynamic pricing, where data from the market environment is used to make pricing decisions. The **Markov Decision Process** framework defines the possible states, actions, and rewards, allowing the model to learn the most profitable pricing strategies. This process involves trial-and-error learning, where the model explores different price points and receives feedback (reward) to improve over time.

3. Hyperparameter Optimization

Hyperparameter optimization is crucial to improving model performance. This step ensures that the model is learning efficiently and making optimal pricing decisions in real-time. Techniques like **Bayesian optimization** help fine-tune parameters like learning rate, while **transfer learning** allows knowledge gained from one set of products to be applied to another, accelerating the learning process for new items.

4. Pricing Policy (Output Layer)

Once the model has learned an optimal pricing strategy, it produces a **dynamic pricing policy** that adapts to real-time market changes. The output could be either continuous or discrete price points, depending on the use case. For example, markdown pricing during seasonal sales might use discrete price changes, while fast-moving consumer goods (FMCGs) might need continuous adjustments to stay competitive.

5. Evaluation & Feedback Loop

The model continuously evaluates the performance of its pricing decisions. The **feedback loop** plays a critical role, as the rewards (e.g., higher revenue or improved customer conversion rates) are fed back into the model to refine its future pricing actions. This creates a **self-learning system** that evolves and improves over time, ensuring that the pricing strategy stays relevant in dynamic market conditions.

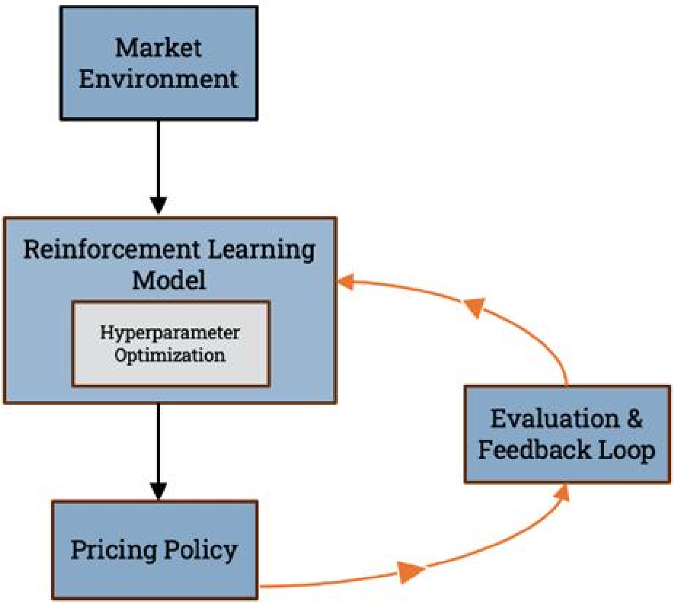

Fig 4: Dynamic Pricing Framework with DRL

Study	Approach	Techniques/Methods	Key Contributions	Advantages	Challenges	Results

Dynamic Pricing Using Reinforcement Learning	Deep Reinforcement Learning (DRL)	DRL pipeline, Markov Decision Process (MDP), Bayesian-genetic hyperparameter optimization	Automates dynamic pricing decisions using real-time data	Flexible and adaptable pricing strategies, improves revenue and competitiveness in uncertain markets	Uncertainty in demand prediction	Successfully optimizes real-time pricing in unpredictable markets
Enhancing Sparse Data Performance in E-Commerce Dynamic Pricing with Reinforcement Learning and Pre-Trained Learning	DRL for dynamic pricing	MDP, continuous price sets, DRCR (Difference of Revenue Conversion Rates) as reward function	Introduces DRCR reward function, solves cold-start problem via pre-training	Real-time adaptive pricing, improves over manual pricing	Requires extensive pre-training, computation-heavy	DRCR-based model outperforms traditional manual pricing, yielding increased revenue
E-commerce Tech Trends: Reinforcement Learning for Dynamic Pricing	Reinforcement Learning (RL)	Q-learning, sales simulation model (factors: seasonality, competitor pricing)	RL balances short-term and long-term rewards, addresses short-term optimization issues	Maximizes long-term cumulative revenue, adjusts pricing based on market changes	Difficulty in real-world data collection and simulation	RL improves cumulative revenue compared to short-term optimization methods
Dynamic Pricing, Reinforcement Learning and Multi-Armed Bandit	Multi-Armed Bandit (MAB) algorithms	Greedy Strategy, Epsilon-Greedy, UCB1, Thompson Sampling	Applies MAB algorithms to pricing, evaluates exploration-exploitation balance	Efficiently balances exploration (new pricing strategies) and exploitation (known rewards)	Greedy approach may overlook optimal long-term options, uncertainty in real-world application	Thompson Sampling and UCB1 are most effective for balancing exploration and maximizing revenue
Dynamic Pricing on E-commerce platform with Deep Reinforcement Learning: A Field Experiment	Deep Reinforcement Learning (DRL)	MDP, DRCR reward function, pre-training with historical data	Introduces end-to-end framework for continuous pricing, solves cold-start problem with pre-training	Provides effective continuous pricing strategies, outperforms traditional methods	Computation-heavy, requires historical data for pre-training	Continuous pricing with DRCR outperforms traditional methods, showing increased revenue

Table 1: Comparison of Result of different Study

5. Conclusion:

Dynamic pricing strategies using deep reinforcement learning (DRL) and reinforcement learning (RL) offer significant potential for real-time adaptability in e-commerce and energy sectors.

The proposed approaches demonstrate superior performance in revenue generation, efficiency, and optimization, outperforming traditional methods and manual strategies. These advancements enhance scalability, adaptability, and profitability across industries, particularly in high-demand, dynamic environments. Several works, particularly in the e-commerce domain, have explored RL frameworks like DDPG and DQN for real-time pricing decisions, showcasing their superiority over traditional manual pricing strategies. The application of RL in managing large-scale, high-frequency pricing tasks—such as markdowns and daily pricing for FMCGs—has proven highly effective. These approaches address challenges like insufficient training data for low-volume products, with potential solutions involving clustering, transfer learning, and meta-learning.

Key improvements include leveraging k-means clustering, transfer learning, and novel reward functions to address sparse data and cold-start problems.

Acknowledgements

I would like to express my sincere gratitude to all those who have supported and contributed to the completion of this research. First and foremost, I am deeply grateful to my advisor, Dr. J.B Simha, for his invaluable guidance, insightful feedback, and continuous encouragement throughout this research. His expertise and mentorship were crucial to the development of this work. I am thankful to RACE Dept. for providing the resources and facilities necessary for conducting this study. Special thanks to Dr. Shinu Abhi for her guidance in paper publication. Lastly, I would like to acknowledge the anonymous reviewers whose constructive comments and suggestions significantly improved this paper. Thank you all for your contributions to this work.

References

[1] Afshar, R. R., Rhuggenaath, J., Zhang, Y., & Kaymak, U. (2022). An Automated Deep Reinforcement Learning Pipeline for Dynamic Pricing. IEEE Transactions on Artificial Intelligence.

[2] Jiaxi Liu, Yidong Zhang, Xiaoqing Wang, Yuming Deng, Xingyu Wu. (2021 August). "Dynamic Pricing On E-Commerce Platform With Deep Reinforcement Learning: A Field Experiment" https://arxiv.org/pdf/1912.02572.

[3] Yuchen Liu, Ka Lok Man, Gangmin Li, Terry R. Payne, (August 2023) Enhancing Sparse Data Performance in E-Commerce Dynamic Pricing with Reinforcement Learning and Pre-Trained Learning". 2023 International Conference on Platform Technology and Service (PlatCon).

[4] Maestre, R., Duque, J., Rubio, A., & Arévalo, J. (2019). Reinforcement learning for fair dynamic pricing. In Intelligent Systems and Applications: Proceedings of the 2018 Intelligent Systems Conference (IntelliSys) Volume 1 (pp. 120-135). Springer International Publishing.

[5] Kastius, A., & Schlosser, R. (2021). Dynamic pricing under competition using reinforcement learning. Journal of Revenue and Pricing Management.

[6] Durga B , Comparative Analysis of Robotic Path Planning using Reinforcement learning. Aug 2024.

[7] Raju, C. V. L., Narahari, Y., & Ravikumar, K. (2003, June). Reinforcement learning applications in dynamic pricing of retail markets. In EEE International Conference on E-Commerce, 2003. CEC 2003. (pp. 339-346). IEEE.

[8] *B.V.S.S. Chandana, K. Kiranmai, (*2023, December*).* Dynamic Pricing Using Reinforcement Learning, International Journal of Research Publication and Reviews, Vol 4, no 12, pp 4522-4528 December 2023.

[9] *Ravikumar M L. & Nagashree J,* A Comprehensive Review in Exploring the Role of Reinforcement Learning, Vol 06, Issue 09 Sept 2024

[10] Deqian Kong et al. "Dynamic pricing of demand re- sponse based on elasticity transfer and reinforcement learning". In: *2019 22nd International Conference on Electrical Machines and Systems (ICEMS)*. IEEE. 2019, pp. 1–5.

[11] Palensky P, Dietrich D. Demand side management: demand response, intelligentenergy systems, and smart loads. IEEE Trans Ind Inform 2011;7(3):381–8. https://doi .org /10 .1109 /TII .2011 .2158841.

[12] Huang H, Ning Y, Jiang Y, Tang Z, Qian Y, Zhang X. Dynamic pricing optimization for commercial subcontracting power suppliers engaging demand response considering building virtual energy storage.

[13] Könönen, V. (2006). Dynamic pricing based on asymmetric multiagent reinforcement learning. International journal of intelligent systems, 21(1), 73-98.

[14] Kim, B. G., Zhang, Y., Van Der Schaar, M., & Lee, J. W. (2014, April). Dynamic pricing for smart grid with reinforcement learning. In 2014 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS) (pp. 640-645). IEEE.