International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 05 May 2024
Page No: 1525- 1529
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2024.0209

Ad Optimization Via Machine Learning: A Focus on Upper Confidence Bound and Thompson Sampling Algorithms

Ananya Kaul[1], Priyam Aneja[2], Sarthak Tomar[3], Dr Abdul Rahman[4]*
[1,2,3]VIT Bhopal University, Bhopal, Madhya Pradesh, India.
[4]Associate Professor, CSE, VIT Bhopal University, Bhopal, Madhya Pradesh, India.*
Emails: *ananya.kaul2020@vitbhopal.ac.in[1], priyam.aneja2020@vitbhopal.ac.in[2], sarthak.tomar2020@vitbhopal.ac.in[3], abdulrahman@vitbhopal.ac.in[4]*

Abstract

The objective of this project is to improve the effectiveness and efficiency of advertising on various platforms by utilizing advanced algorithms, namely the Upper Confidence Bound and Thompson Sampling Algorithm. The project aims to find a balance between exploring new advertising strategies and exploiting proven high-performing approaches. By implementing these bandit algorithms, the project aims to dynamically optimize ad placements, formats, and targeting to maximize user engagement and ad revenue. The methodology involves an iterative process of data collection, analysis, and adaptation. The initial phases include defining project objectives, understanding the target audience, and reviewing the current ad strategy. The Upper Confidence Bound algorithm enables intelligent decision-making by assigning confidence bounds to different ad strategies, allowing for efficient exploration and exploitation. On the other hand, the Thompson Sampling algorithm, rooted in Bayesian principles, dynamically adapts based on observed outcomes, striking a balance between exploration and exploitation through probabilistic reasoning. In summary, this Ads Optimization Project utilizes the power of the Upper Confidence Bound and Thompson Sampling algorithms to create a data-driven, adaptive, and user-centric approach to advertising. The ultimate goal is to achieve optimal user engagement and ad revenue.

Keywords: *Ads Optimization; Upper Confidence Bound; Thompson Sampling Algorithm.*

1. Introduction

In the contemporary digital landscape, where businesses navigate an intricate web of consumer interactions, online advertising stands out as a linchpin for success. The behemoth that is Facebook, with its vast user base and dynamic content delivery, presents both opportunities and challenges for marketers. This capstone project undertakes a deep dive into the intricate realm of Ads Optimization using the Multi-Armed Bandit (MAB) framework. Specifically, it explores the implementation of the Upper Confidence Bound (UCB) and Thompson Sampling Algorithm (TSA) on Facebook's expansive and diverse ad dataset. The objective of this project is to apply the UCB (Upper Confidence Bound) and TSA (Thompson Sampling Algorithm) algorithms in a dynamic optimization system for Facebook ads. The goal is to seamlessly integrate these algorithms into the system, ensuring adaptability to the ever-changing user preferences and platform dynamics. In addition to implementing the algorithms, the project aims to evaluate their effectiveness in the dynamic context of Facebook's ad platform. This evaluation goes beyond conventional metrics and includes an in-depth exploration of user engagement, conversion rates, and the overall holistic performance of ads. Furthermore, the project seeks to showcase the Multi-Armed Bandit (MAB) Problem in action within the complex landscape of Facebook advertising. It serves as a demonstration of how MAB models can address real-world marketing challenges and adapt to the dynamic nature of the platform. Beyond the technical aspects, the project also aims to generate actionable insights into the advantages and challenges associated with deploying the MAB framework for online advertising. These insights contribute to the broader discourse surrounding algorithmic decision-making in the realm of marketing.

2. Literature Review

Ad optimization plays a crucial role in digital advertising, as it aims to maximize the effectiveness and efficiency of ad campaigns. By utilizing various techniques and strategies, advertisers can improve their targeting, bidding, and creative elements to achieve better results. This literature review aims to explore the key concepts, methodologies, and advancements in ad optimization.

2.1. Ad Optimization Techniques

Targeting Optimization: Targeting optimization involves identifying the most relevant audience segments for an ad campaign. This can be achieved through demographic, geographic, behavioral, or contextual targeting. Several studies have focused on developing advanced targeting algorithms to improve ad relevance and increase conversion rates (Li, H., Yang, Y., (2022) [1].

Bidding Optimization: Bidding optimization aims to determine the optimal bid amount for each ad placement. Various approaches, such as rule-based bidding, machine learning algorithms, and game theory models, have been proposed to optimize bidding strategies. These techniques consider factors like ad quality, competition, and budget constraints to maximize return on investment (ROI) (Li, H., Xu, D., Shmakov, K., Lee, K., Shen, W.) [2].

Creative Optimization: Creative optimization focuses on improving the visual and textual elements of ads to enhance their impact on the target audience. A/B testing, multivariate testing, and dynamic creative optimization are commonly used techniques to optimize ad creatives. These methods help identify the most effective combinations of images, headlines, and call-to-action buttons (Chen, J., Xu, J., Jiang, G., Ge, T., Zhang, Z., Lian, D., Zheng,) [3].

2.2. Advanced Ad Optimization Techniques

Machine Learning and Artificial Intelligence (AI): Machine learning and AI techniques have revolutionized ad optimization by enabling automated decision-making processes. These technologies can analyze vast amounts of data, identify patterns, and make real-time optimizations to improve ad performance (Truong, V., Hoang, V. (2022)) [4].

Predictive Analytics: Predictive analytics leverages historical data to forecast future ad performance. By analyzing past campaign data, advertisers can predict the likelihood of achieving specific goals and optimize their strategies accordingly (Kumar, Vaibhav & L., M., (2018)) [5].

Real-Time Bidding (RTB): RTB is an auction-based ad buying method that allows advertisers to bid on ad impressions in real-time. RTB platforms use algorithms to optimize bidding decisions based on user data, ad relevance, and campaign objectives (Zhang, Chong-rui & Zhang, E. (2014)) [6].

Ad optimization is a critical aspect of digital advertising, enabling advertisers to maximize the effectiveness of their campaigns. By utilizing various techniques such as targeting optimization, bidding optimization, and creative optimization, advertisers can improve their ad performance and achieve better results. Advanced techniques like machine learning, predictive analytics, and real-time bidding further enhance the capabilities of ad optimization. As technology continues to evolve, it is expected that ad optimization will become even more sophisticated, enabling advertisers to reach their target audience more effectively and efficiently.

3. Methodology

3.1. System Design / Architecture

The system's architectural design is not just a technical detail; it combines the MAB framework with the Facebook ads platform seamlessly. The data pre-processing stage goes beyond basic tasks, including not only extracting features but also

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 05 May 2024
Page No: 1525- 1529
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2024.0209

creating a dataset that enhances the algorithms. The UCB and TSA algorithms are not standalone components; they are decision-making engines that adapt ad allocation based on user interactions in real-time.

3.2. Algorithm Implementation and Visualizations

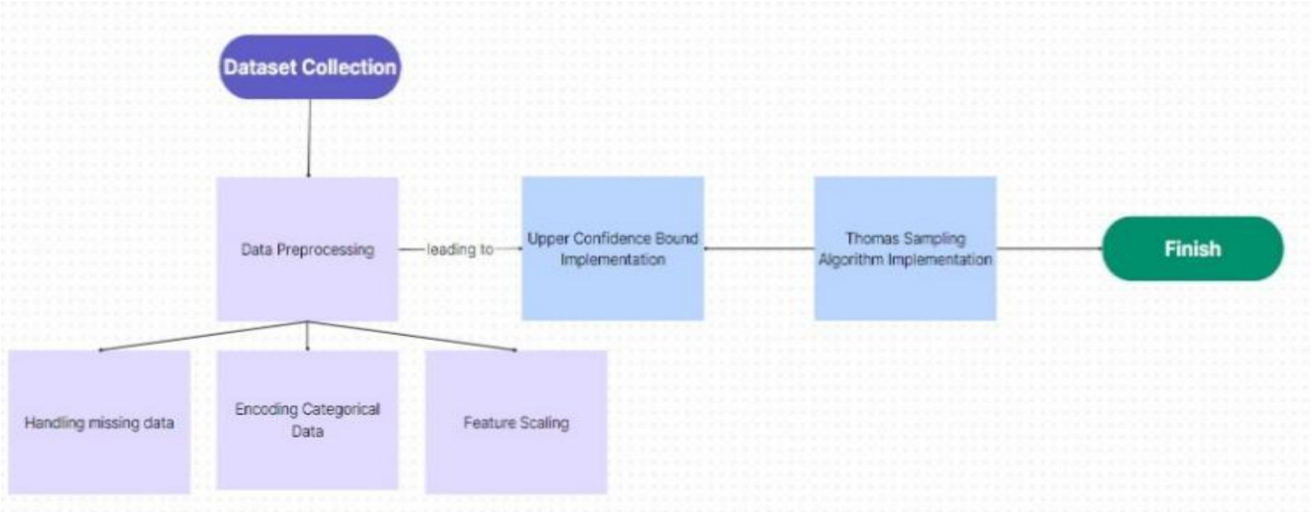

Figure 1 **Block Diagram for Algorithm Implementation**

Algorithm steps are shown in Figure 1.

3.2.1. Data Pre-Processing

Regarding the libraries used for data pre-processing, we utilized numpy, matplotlib, and pandas. We used pandas to read the data, and sklearn to handle missing values. Since our dataset did not have any categorical variables, we did not need to convert them into numeric variables using Encoding Data. Additionally, feature scaling was not necessary as our dataset did not have independent variables that needed to be confined. Furthermore, we did not split the data as our goal was not to predict data, but rather to classify it.

3.2.2. Implementation & Visualization of UCB

The Upper Confidence Bound can be implemented in the following steps:

Step 1. At each round n, we consider two numbers for each ad i:

- $N_i(n)$ - the number of times the ad i was selected up to round n,
- $R_i(n)$ - the sum of rewards of the ad i up to round n.

Step 2. From these two numbers we compute:

- the average reward of ad i up to round n

$$\bar{r}_i(n) = \frac{R_i(n)}{N_i(n)}$$

- the confidence interval $[\bar{r}_i(n) - \Delta_i(n), \bar{r}_i(n) + \Delta_i(n)]$ at round n witl

$$\Delta_i(n) = \sqrt{\frac{3}{2}\frac{\log(n)}{N_i(n)}}$$

Step 3. We select the ad i that has the maximum UCB $\bar{r}_i(n) + \Delta_i(n)$.

Figure 2 **Upper Confidence Bound**

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 05 May 2024
Page No: 1525- 1529
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2024.0209

In Figure 2. the trend can be seen in the histogram given below where 4th ad has the highest number of clicks.

3.2.3. Implementation & Visualizations of TSA

Steps:

Initialize Prior Beliefs: At the start, Thompson Sampling sets up a prior distribution for each machine's reward. Typically, a non-informative prior like the Beta distribution is used, assuming equal probabilities for all possible rewards in Figure 3.

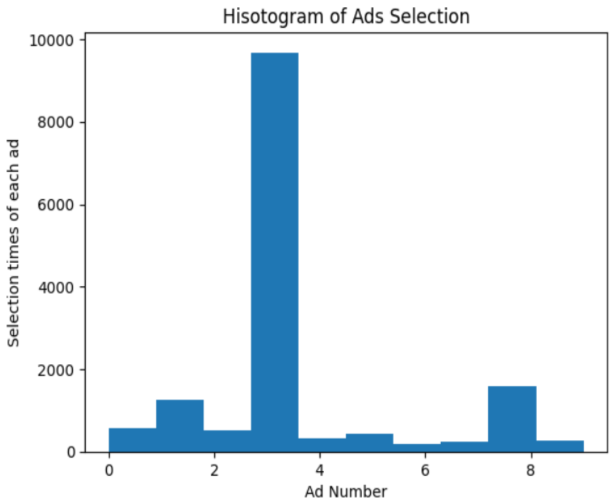

Figure 3 Histogram of ADS Selection After UCB

Action Selection: In each round, Thompson Sampling samples a reward value from each machine's distribution. The machine with the highest sampled value is selected, and its lever is pulled.

Observe Reward: After the selected action is executed, the algorithm observes the actual reward obtained from that machine.

Update Probability Distribution: Using the observed reward, Thompson Sampling updates the probability distribution (posterior) for the selected machine using Bayesian inference. The updated distribution becomes the prior distribution for the next round.

Repeat: The process continues iteratively, with the algorithm balancing exploration (trying out different machines to learn more about their rewards) and exploitation (focusing on the machines that have shown higher rewards so far).

3.3. A/B Testing Vs Multi-Armed Bandit Problem

A/B testing is a simple concept where you split your sample into two groups and provide each group with a different version of an ad. You then measure the impressions on each ad and determine the winner based on the highest number of impressions. However, this approach has a drawback as it isolates exploration from exploitation, leading to wasted resources and limited variation.

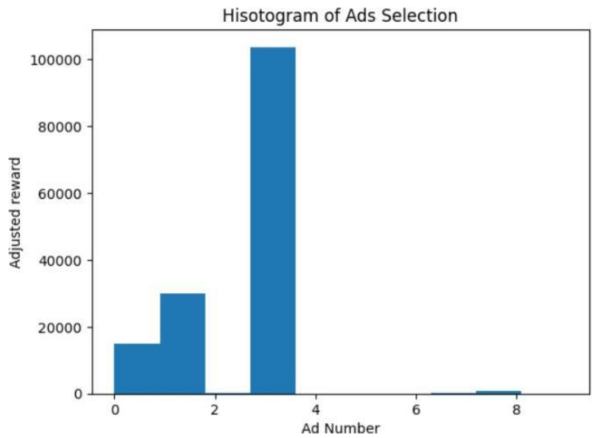

Figure 4 Histogram of ADS Selection After TSA

In Figure 4&5 the trend can be seen in the histogram given below where 4th ad has the highest number of clicks.

Figure 5 A/B Testing

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 02 Issue: 05 May 2024
Page No: 1525- 1529
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2024.0209

4. Results

To address this issue, bandit algorithms can be used. These algorithms, also known as learning, while earning, allow you to make progress in finding the winning ad while also generating revenue. They offer excellent performance in shorter time periods. For example, during a Black Friday sales campaign that lasts only a week, A/B testing may not be feasible due to limited time and insufficient data. In such cases, a multi-armed bandit approach is recommended for its effectiveness in shorter time-frames.

Conclusion

To summarize, this project is not just a technical endeavor; it is an exploration into the uncharted realms of ad optimization. The findings are not simply outcomes; they are revelations that provide insights into the potential benefits and challenges of implementing UCB and TSA in the dynamic and complex advertising landscape of social media. This project does not signify an endpoint, but rather a starting point that propels the journey towards more advanced ad optimization strategies. It opens opportunities for the development of advertising systems that are not only sophisticated but also responsive, systems that can adapt and even anticipate. Future work does not imply a conclusion, but rather a continuation that involves further refining algorithms, exploring their application across various marketing platforms, and integrating real-time user feedback to create marketing strategies that are not just personalized but also forward-thinking.

References

[1]. Li, H., Yang, Y., (2022). Keyword targeting optimization in sponsored search advertising: Combining selection and matching, Electronic Commerce Research and Applications, Volume 56, 101209, ISSN1567-4223, https://doi.org/10.1016/j.elerap.2022.101209.

[2]. Li, H., Xu, D., Shmakov, K., Lee, K., Shen, W., Bid Optimization for Offsite Display Ad Campaigns on eCommerce, https://arxiv.org/abs/2306.10476

[3]. Chen, J., Xu, J., Jiang, G., Ge, T., Zhang,. Z., Lian, D., Zheng, K., Automated Creative Optimization for E-Commerce Advertising, https://arxiv.org/abs/2103.00436

[4]. Truong, V., Hoang, V. (2022). Machine Learning Optimization in Computational Advertising—A Systematic Literature Review. In: Abdul Karim, S.A. (eds) Intelligent Systems Modeling and Simulation II. Studies in Systems, Decision and Control, vol 444. Springer, Cham. https://doi.org/10.1007/978-3-031-04028-3_8

[5]. Kumar, Vaibhav & L., M. (2018). Predictive Analytics: A Review of Trends and Techniques. International Journal of Computer Applications. 182. 31-37. 10.5120/ijca2018917434.

[6]. Zhang, Chong-rui & Zhang, E. (2014). Optimized bidding algorithm of real time bidding in online ads auction. International Conference on Management Science and Engineering - Annual Conference Proceedings.33-42. 10.1109/ICMSE.2014.6930205.