Proceedings of the 6th International Conference on Computing and Data Science
DOI: 10.54254/2755-2721/68/20241418

The evolution and impact of Multi-Armed Bandit algorithms in social media

Yiming Yan

Department of Statistics, The University of Wisconsin-Madison, Madison, United States

yyan227@wisc.edu

Abstract. This paper examines the transformative impact of Multi-Armed Bandit (MAB) algorithms on user experiences across social media platforms. Initially conceptualized in the 1930s and formalized in the 1950s, MAB algorithms have become foundational to the evolution of digital interactions and content personalization. These algorithms adeptly navigate the trade-off between exploration and exploitation to maximize user engagement and satisfaction. By scrutinizing their implementation from early adopters like Yahoo to contemporary giants such as Facebook, Instagram, and TikTok, this analysis elucidates the algorithms' prowess in tailoring content recommendations, refining advertising strategies, and bolstering overall platform engagement. Moreover, this study addresses the ethical dimensions of MAB algorithms, with a particular emphasis on concerns surrounding user privacy and the perpetuation of echo chambers. Through an extensive synthesis of theoretical insights and empirical applications, this paper highlights the pivotal role of MAB algorithms in shaping the digital and social media landscape, advocating for future research focused on improving algorithmic transparency and ethical governance.

Keywords: Multi-Armed Bandit algorithms, social media, content personalization.

1. Introduction

The Multi-Armed Bandit (MAB) algorithm, essential for decision-making under uncertainty, derives its name from a gambler's strategic choice among multiple slot machines, colloquially known as "one-armed bandits." Each machine has a distinct, yet unknown, payout ratio, compelling the gambler to either explore a new "arm" for potentially better returns or exploit a known "arm" based on historical performance [1]. Today, the MAB algorithm is pivotal in various social media applications, ranging from online advertising to personalized content recommendations. These algorithms dynamically learn from user interactions to predict, for instance, which movie might engage a viewer, thereby optimizing recommendations to maximize user satisfaction. This adaptability is crucial for sustaining long-term user engagement and is an integral component of the digital experience on social media platforms [2].

This article delves into how the MAB algorithm has become foundational within social media recommendation systems. It traces the algorithm's evolution, demonstrating its increasing capability to offer more personalized content recommendations that significantly enhance user experiences. By examining the deployment of MAB algorithms across different facets of social media—from boosting

DOI: 10.54254/2755-2721/68/20241418

user engagement to fine-tuning advertising strategies—this study illuminates how these algorithms balance user satisfaction with business objectives. Furthermore, the paper addresses the ethical challenges and implications of using MAB algorithms, particularly their impact on user privacy and their role in fostering echo chambers [3]. Through this analysis, the paper seeks to provide a thorough understanding of how MAB algorithms have become indispensable in the design of recommendation systems and their profound influence on the architecture of social media platforms.

2. Historical Development of MAB Algorithms

The gradual development of the multi-armed bandit algorithm has had a great impact in various fields, especially in the development of early social media platforms [4]. Based on the dilemma of exploration and development, MAB algorithms have been developed from the simple theoretical structure to the complex tools for promoting user participation and content personalization.

2.1. Origin and early models of MAB algorithms and how they play role in early social media platform

While the concept of MAB was initially proposed in the 1930s, its formal mathematical model was not established until the 1950s, initially applied to clinical trials by the Center for Drug Evaluation and Research with the aim of identifying the most effective treatment. Early attempts to resolve the MAB problem, such as the epsilon-greedy algorithm, sought to balance known rewards against potentially more lucrative but untested options [5]. Early social media platforms like Digg and StumbleUpon were among the pioneers in utilizing MAB algorithms for content filtering and recommendation, treating each piece of content as an "arm" to be evaluated by the algorithm, thereby identifying user preferences and potentially increasing time spent on the site [6].

A notable early instance of MAB application in social media was at Yahoo. By the late 2000s, Yahoo had begun incorporating MAB algorithms to personalize the news articles displayed on each user's homepage [7]. This early adoption of MAB algorithms in content customization marked a significant step, demonstrating how these algorithms could substantially enhance user engagement through tailored content recommendations. The fact that MAB algorithms were employed early in the digital era to shape social media interactions underscores their fundamental role in digital platforms. By significantly influencing user experiences, these algorithms have not only maximized user engagement but also established a framework for how future systems might enhance digital interactions [8].

2.2. The improvement over time and branches it developed

Since MAB algorithms emerge, it later have improved into different branches and its applications have expanded. Some of the key algorithm that play important role in nowadays social media platform [9]. Three important algorithms that advances in this area include upper Bound Confidence (UCB), Exploration-then Commit (ETC), and Thompson Sampling (TS) algorithms.

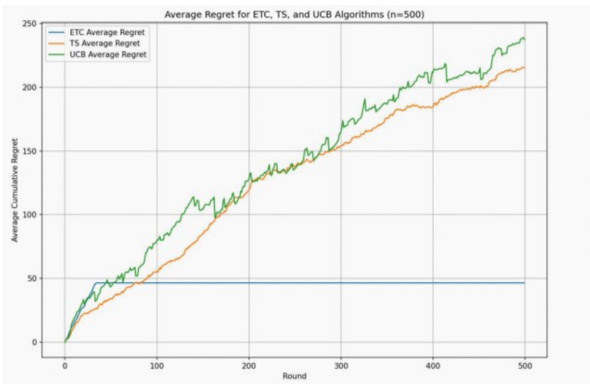

Figure 1. Performance of ETC, TS, UCB algorithm at small number of rounds (Photo/Picture credit: Original).

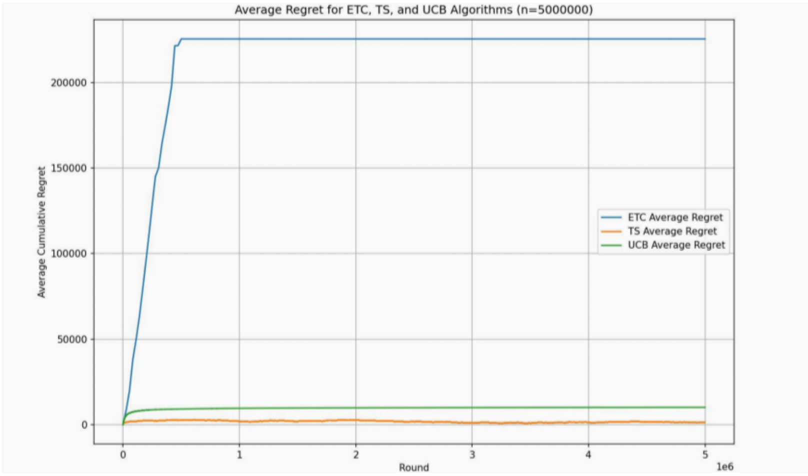

Figure 2. Performance of ETC, TS, UCB algorithm at large number of rounds (Photo/Picture credit: Original).

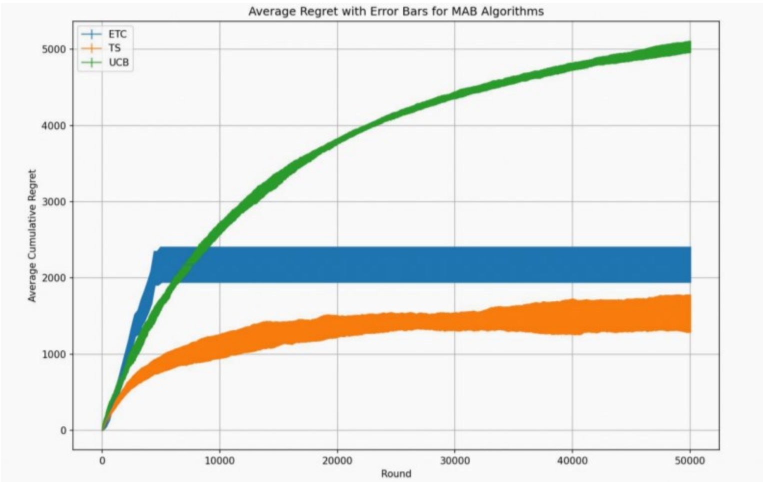

Figure 3. Variance overtime of ETC, TS, UCB algorithm (Photo/Picture credit: Original).

As illustrated in Figure 1, the Explore-Then-Commit (ETC) algorithm demonstrates superior effectiveness during short experimental rounds. Therefore, ETC used most at the starting stage at the social media platform. In contrast, as depicted in Figure 2, the Thompson Sampling (TS) algorithm excels as the number of rounds increases, resulting in lower long-term regret. TS algorithm thereby are used mostly in dynamic environment. Notably, the Upper Confidence Bound (UCB) algorithm, as shown in Figure 3, exhibits the lowest variance, indicating a more stable performance. UCB are used mostly for generating stable inference on social media recommendation platform.

2.3. Upper confidence Boundary (UCB)

UCB_ i (t) = X_ i (t) + sqrt ((2 * ln(t)) / N_ I (t))

-`UCB_ i(t)` is the upper confidence bound of arm i at the t time

- `X_ I (t)` is the average reward from arm i to time t

- `N_ I (t)` is the number of times arm i been selected up to time t

- `ln(t)` is the natural log of time t

The UCB algorithm is in simple words: the algorithm calculates a confidence interval of the expected return for each of the options, and then selects the highest [10]. These serve to resolve the traditionally exploited-explored dilemma, since they balance the need to exploit the most promising

DOI: 10.54254/2755-2721/68/20241418

options available, meanwhile still exploring to find out the underlying better option. This make the UCB algorithm capable with low variance and are hence less variant overtime.

UCB plays good in the ad placement platform of social media. It can give reliable and predicable performance overtime. UCB algorithm can predicts which advertisements are more likely to have a higher user engagement rate. So, the platforms make a good balance for placing ads on their platform without user dissatisfaction and hence ensure the good economics of their profit and high user satisfaction.

2.4. Explore - then submit (ETC)

The difference with ETC is that these latter algorithms decompose decision-making into two phases: the exploration phase, where one information gathers for each option to test; and the commitment phase, where one decides upon the option from which most rewards returned based upon the exploration data. ETC algorithm perform the best at the starting stage. The algorithm is preferably used in cases where preferences by users do not change with time. Social media platforms that use ETC algorithms in their initial understanding greatly influence the long-term perception of the platform for use and improve the long-run user engagement. This could have been done in the login process first, showing the users' varied content to gauge the user's interest.

2.5. Thompson sampling (TS)

Thompson sampling make decision based on each option’s likelihood to be the best choice. It is adaptive to the dynamic real-time environment. It plays a great role in social media platforms where user preferences change overnight. Most of the usage goes to personalization, such as video recommendations. In early social media platforms, it sets down a basic foundation for the MAB of choice to later counterparts, a more complicated and dynamic platform like Facebook, Tiktok, and YouTube.

These new platforms also draw lessons from earlier models and bring in advanced versions of MAB algorithms to manage personalized feeds, recommend connections, and optimize advertising. These new platforms draw lessons from earlier models, and adopt the advanced versions of MAB algorithms to manage personalized feeds, recommend connections and optimize advertising, so as to maintain user participation and increase revenue.

From such an adaptation point, the MAB algorithm has proved very essential in continuous learning and adaptation according to the preferences of the users, who give the world over personalized and fascinating content.

3. MAB Algorithms in Social Media

3.1. The role of MAB algorithms in social media platforms

The Multi-Armed Bandit (MAB) algorithms are significant in improving user experience and improving the function of social media platforms. In the context of social media applications, MAB algorithms gradually adapt to user needs, demonstrating a significant evolution from theoretical underpinnings to practical implementations focused on user centrality. These algorithms excel within dynamic and ever-evolving social media systems by providing a constant flow of data for optimizing user preferences and behaviors.

3.2. Personalized content recommendation

The personalization of content recommendation is one of the key features of MAB algorithms in social media. Parameters of interactivity, including likes, views, comments, and the time users spend on posts, are all influenced by MAB algorithms on popular platforms like Facebook, Instagram, and TikTok. This analysis enables platforms to offer tailored content recommendations for each user. Platforms can explore new types of content that users might like while exploiting known user preferences to optimally present posts, videos, and ads that effectively engage and interest users. This

DOI: 10.54254/2755-2721/68/20241418

approach ensures users are satisfied with more personalized recommendations, even though they spend less time on the platforms.

3.3. Advertising optimization

A significant source of revenue for any platform is advertising. However, striking a balance between ad presence and avoiding user aversion is a substantial challenge. Here, the MAB algorithm plays a crucial role in optimizing advertising on social media platforms. It tries to maximize both click-through rates (CTRs) and conversion rates by dynamically selecting ads based on individual users' interests or CTRs, ultimately benefiting advertisers.

3.4. Engagement and retention

MAB algorithms ensure users always receive content that appeals to them, thus increasing user engagement. The algorithm adapts to changes in user preferences, introduces new content, and maintains the freshness and attractiveness of the platform to minimize user churn while fostering long-term user relationships.

4. Modern Applications of MAB Algorithms in Personalizing Recommendations

4.1. Overview different MAB algorithms are used in personalizing recommendations

Upper Confidence Bound (UCB), Exploration-then Submit (ETC), and Thompson Sampling (TS) are the major multi-armed bandit (MAB) algorithms contributing to the reward-based systems adapted for the continuously changing and diverse needs of social media platforms. They find a balance between exploration and exploitation to optimize different aspects of user experiences, such as content recommendation, advertising, and functional testing. Their roles could differ significantly due to diverse issues and approaches to user engagement on platforms like Facebook, TikTok, or Instagram.

4.2. Application of Upper Confidence Bound (UCB) in social media

The UCB algorithm excels in environments where there is a balance between exploring new content and exploiting known user preferences. For example, social media sites like Facebook favor UCB algorithms due to their emphasis on content diversity and discovery. The UCB algorithm sorts content according to the upper limits of the confidence bounds for the potential engagement level of each piece of content. Higher bounds indicate either great uncertainty or a high level of user engagement. Facebook uses UCB algorithms in its news feed to display new pages or groups within the user's network, ensuring that users receive varied yet relevant content, thus enhancing satisfaction and engagement on the platform.

4.3. Application of Thompson Sampling (TS) in social media

Thompson Sampling is renowned for its effectiveness in dynamic environments where trends change rapidly and user interests evolve, such as on TikTok. TS selects relevant content, ads, or features in a probabilistic manner and infers beliefs based on continuously received user feedback. This is particularly useful for platforms aiming to maximize real-time user engagement. TikTok uses TS to adapt quickly to changes in user behavior and tastes, optimizing video recommendations to capitalize on current trends and viral content.

4.4. Explore - then submit (ETC) applications in social media

The ETC algorithm is crucial during the user login process or when introducing new features, as it allows for an initial exploration phase before making optimal choices. This approach enables platforms to collect data on user preferences in a stepwise manner, initially offering content or features and later customizing them based on these preferences. For instance, if Instagram is testing new features like "Stories" or video volume controls, it could employ the ETC model to first showcase various types of content and later focus on promoting the most successful formats based on engagement metrics.

The choice of MAB algorithms is often dictated by platform-specific goals such as maximizing engagement, exploring novel content types, and personalizing experiences, which stem from the nature of supported content and user interactions. For example, Facebook might prioritize algorithms that support content diversity and long-term user engagement, whereas TikTok focuses on adaptive, real-time decision-making capabilities. Overall, MAB algorithms are selected and applied in a manner that suits the unique ecosystems within each social media platform, playing a significant role in enhancing user experience, engagement, and satisfaction.

4.5. Detailed examples of real-world applications

The deployment of Multi-Armed Bandit (MAB) algorithms across various social media platforms showcases their flexibility and effectiveness in enhancing user experiences with personalized recommendations. Beyond their theoretical applications, algorithms like UCB, TS, and ETC are actively utilized in real-world scenarios by major companies to drive participation and tailor content on platforms.

4.5.1. LinkedIn's Use of TS for Content and Job Recommendations. LinkedIn utilizes Thompson Sampling for personalizing feeds and job recommendations, treating each article and job listing as a unique "arm". This approach dynamically updates the relevance of content or jobs based on user interactions, ensuring the delivery of highly pertinent professional content and opportunities, thereby boosting user engagement and satisfaction on the platform (Kenthapadi et al., 2017).

4.5.2. Spotify's Dynamic Playlist Curation. Spotify employs a variant of the UCB algorithm within its music recommendation engine, particularly for managing personalized playlists like "Weekly Discovery." This method balances users' existing music preferences with the introduction of new tracks, enabling discovery while maintaining a satisfying listening experience. This equilibrium allows users to explore new artists and songs, thereby enhancing their engagement with the service.

4.5.3. Amazon's Personalized Shopping Experience. Amazon implements an Explore-Then-Commit strategy to offer personalized product recommendations. Initially, it explores a range of product recommendations; upon detecting patterns in user preferences for certain brands or categories, it commits to these preferences to optimize its suggestions. This method is clearly demonstrated in how Amazon customizes the homepage and product recommendations for each user over time, enhancing user engagement and sales.

By leveraging these algorithms, companies are able to customize services and content to match individual user preferences, thereby boosting participation, satisfaction, and ultimately, business success.

5. Ethical questions MAB algorithm

Deploying Multi-Armed Bandit (MAB) algorithms on social media and digital platforms significantly enhances personalization and user engagement but also raises ethical considerations and challenges. The two main issues are the impact on user privacy and the potential creation of echo chambers.

5.1. Critics about MAB algorithms on social media recommendation

5.1.1. Privacy Concerns. MAB algorithms optimize user experiences by learning from online behaviors, preferences, and interactions. However, this involves collecting, analyzing, and storing vast amounts of personal data, raising serious privacy concerns. The extensive data collection necessitates transparent consent mechanisms, yet users often lack clear information about what data is collected and how it is used, leading to potential privacy breaches. Additionally, storing user data for algorithm processing makes it a target for cyberattacks, which may lead to unauthorized access to sensitive personal information. Ensuring data security to prevent leaks is crucial.

5.1.2. Echo Chambers. MAB algorithms are designed to maximize user engagement by providing content and recommendations that align with known preferences. While this approach effectively retains users, it inadvertently strengthens echo chambers. As the algorithms become increasingly adept at predicting and delivering content that matches user preferences, the variety of content encountered by users narrows, leading to a homogenization of information and opinions online. This continuous reinforcement of existing preferences contributes to the formation of echo chambers, limiting users' exposure to diverse opinions and potentially exacerbating societal polarization (Peckham, 2024).

5.2. Approach to solve ethical questions to make good use of MAB algorithm

Approaches to Mitigate Ethical Concerns:

Transparency: Increasing transparency about how recommendations are generated can help mitigate the echo chamber effect. However, platforms face the challenge of balancing transparency with the protection of proprietary algorithms.

Data Protection: Implementing robust data protection measures, transparent data collection practices, and empowering users with greater control over their data can alleviate privacy concerns.

Diversity in Recommendations: Intentionally introducing diversity and unexpected elements into recommendation algorithms can counteract the trend towards creating echo chambers. This may involve adjusting MAB algorithms and periodically reviewing recommendation mechanisms to ensure a broader range of content is presented to users.

By addressing these ethical issues, platforms can better harness the benefits of MAB algorithms while minimizing negative impacts on user privacy and societal cohesion.

6. VI. Impact of MAB Algorithms in Social Media

6.1. Impact of MAB algorithms on various aspects of social media

The deployment of Multi-Armed Bandit (MAB) algorithms across social media platforms has fundamentally reshaped how content is curated and presented to users, significantly impacting digital engagement and advertising landscapes. These algorithms, known for their dynamic learning and adaptability, have enhanced user experiences, influenced platform economies, and transformed the broader social media ecosystem. Looking to the future, MAB algorithms hold even greater potential to revolutionize decision-making within social media through heightened personalization, increased efficiency, and integrated ethical considerations.

6.1.1. Personalized User Experiences. MAB algorithms have taken content personalization to new heights by continuously learning from user interactions. This capability allows platforms to tailor content feeds, advertisements, and recommendations to individual preferences, thus boosting user satisfaction and engagement. As a result, platforms see increased user retention rates and more time spent on the platform.

6.1.2. Optimizing Advertisements. MAB algorithms refine ad placements and targeting strategies, enhancing engagement rates and maximizing return on investment for advertisers. By balancing the exploration of new targeting strategies with the exploitation of known high-performers, these algorithms ensure that ads reach the most receptive audiences, thereby optimizing ad revenue for social media companies.

6.1.3. Enhancing Content Diversity. MAB algorithms foster content diversity and discovery, introducing users to new topics, creators, and communities. This not only enriches the user experience but also supports a vibrant ecosystem of content creators by amplifying their reach and influence.

DOI: 10.54254/2755-2721/68/20241418

Future developments could focus on integrating fairness and reducing bias in content recommendations, ensuring a balanced and equitable content distribution to mitigate the risk of reinforcing stereotypes or spreading misinformation.

6.1.4. Integration with Emerging Technologies. As social media platforms increasingly incorporate technologies such as augmented reality (AR) and virtual reality (VR), MAB algorithms are poised to play a crucial role in managing immersive content experiences. These algorithms could tailor content not just based on user preferences but also considering background factors, creating highly personalized and engaging virtual environments.

6.1.5. Addressing Ethical Challenges. The future of MAB algorithms in social media must also tackle ethical challenges, particularly those related to privacy and the formation of echo chambers. Developing more transparent algorithms that empower users to manage their data and content exposure is essential. Additionally, mechanisms that promote content diversity and present broader perspectives can help mitigate the development of echo chambers.

6.1.6. Responsiveness to Global Events. MAB algorithms could be further refined to respond in real-time to global events and shifts in public sentiment, adjusting content recommendations to provide timely and relevant information. This enhancement could solidify the role of social media as a critical source of news and information, ensuring users have access to diverse viewpoints on current events.

These advancements suggest a future where MAB algorithms not only enhance user engagement but also address key ethical issues, supporting a more informed and diverse public discourse.

7. Conclusion

7.1. Evolution and impact of MAB algorithms in social media

The Multi-Armed Bandit (MAB) algorithm has profoundly influenced the social media industry, enhancing personalization, advertising, and content discovery. As we look to the future, these algorithms are expected to further improve user experience, economic efficiency, and moral governance within social media. However, realizing this potential requires continuous research, innovation and a careful consideration of the moral impacts of algorithm-driven decisions.

The evolution of MAB algorithms marks a significant breakthrough in the digital age. Initially inspired by the gambler's dilemma of choosing which slot machine to play, these algorithms have transcended their gambling origins to become essential tools for optimizing online experiences. From simple theoretical models to complex tools driving user engagement and personalization on social media, MAB algorithms highlight the dynamic nature of technological advancement and its profound impact on user interactions and content consumption.

Initially developed to address challenges in clinical trials and economics, MAB algorithms have found their most impactful application within social media. By striking a balance between exploring new content and leveraging known user preferences, these algorithms significantly enhance the personalization of user experiences. Platforms like Facebook, Instagram, and TikTok utilize MAB algorithms to tailor content and ads to individual users, improving user retention, participation, and platform growth. The integration of MAB algorithms has also revolutionized marketing strategies by optimizing ad placement and ensuring ads are relevant to users' interests and behaviors.

7.2. Future developments of MAB algorithms in social media

Looking ahead, the potential for future advancements in Multi-Armed Bandit (MAB) algorithms within social media is vast. As these algorithms incorporate increasingly sophisticated artificial intelligence and machine learning technologies, we anticipate a more refined level of personalization efficiency. This continuous enhancement is poised to render social media platforms more engaging, equitable, and responsive. In conclusion, the evolution of MAB algorithms has dramatically reshaped the social media landscape, driving innovation in personalization, advertising, and user engagement. The ongoing development of these algorithms holds the promise of crafting even more captivating and fair social media ecosystems. However, realizing these advancements necessitates a concerted effort to confront ethical challenges, ensuring that the technology augments rather than undermines the social media experience.

References

[1] Center for Drug Evaluation and Research (no date) Bispecific antibodies: An area of research and clinical applications, U.S. Food and Drug Administration. Available at: https://www.fda.gov/drugs/spotlight-cder-science/bispecific-antibodies-area-research-and-clinical-applications (Accessed: 16 April 2024).

[2] Chu, W. et al. (2009) ‘A case study of behavior-driven conjoint analysis on Yahoo!’, Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining [Preprint]. doi:10.1145/1557019.1557138.

[3] Costacurta, M. (2023) Dynamic pricing with multi-armed bandit: Learning by doing, Medium. Available at: https://towardsdatascience.com/dynamic-pricing-with-multi-armed-bandit-learning- by- doing-3e4550ed02ac (Accessed: 16 April 2024).

[4] Zhu, X., Zhao, Z., Wei, X., & others. (2021). Action recognition method based on wavelet transform and neural network in wireless network. In 2021 5th International Conference on Digital Signal Processing (pp. 60-65).

[5] Kenthapadi, K., Le, B. and Venkataraman, G. (2017) ‘Personalized job recommendation system at linkedin’, Proceedings of the Eleventh ACM Conference on Recommender Systems [Preprint]. doi:10.1145/3109859.3109921.

[6] Peckham, S. (2024) What are algorithms? how to prevent echo chambers, Internet Matters. Available at: https://www.internetmatters.org/hub/news-blogs/what-are-algorithms- how-to-prevent-echo-chambers/ (Accessed: 16 April 2024).

[7] What is upper confidence bound (UCB) (no date) ActiveLoop. Available at: https://www.activeloop.ai/resources/glossary/upper-confidence-bound-ucb/ (Accessed: 16 April 2024).

[8] Elena, G., Milos, K., & Eugene, I. (2021). Survey of multiarmed bandit algorithms applied to recommendation systems. International Journal of Open Information Technologies, 9(4), 12-27.

[9] Bouneffouf, D., Rish, I., & Aggarwal, C. (2020, July). Survey on applications of multi-armed and contextual bandits. In 2020 IEEE Congress on Evolutionary Computation (CEC) (pp. 1-8). IEEE.

[10] Jeunen, O., & Goethals, B. (2023). Pessimistic decision-making for recommender systems. ACM Transactions on Recommender Systems, 1(1), 1-27.