RESEARCHERID
THOMSON REUTERS

[Surname *et al.*, 12(4): April, 2025] ISSN 2349-0292
Impact Factor 3.802

GLOBAL JOURNAL OF ADVANCED ENGINEERING TECHNOLOGIES AND SCIENCES

IMPROVING VIEWER ENGAGEMENT THROUGH CONTEXT-AWARE ADS IN FREE STREAMING SERVICES

Tungeshwar Rai
Sr. Manager, Operation and Maintenance of Fiber network
tungeshwar.rai@rediffmail.com

DOI: https://doi.org/10.29121/gjaets.2025.4.2

ABSTRACT

The rapid growth of free streaming services has revolutionized how users consume content, creating new challenges in terms of monetization. Traditional advertisement delivery models, often seen as disruptive, have led to viewer dissatisfaction. This paper presents a context-aware ad delivery system that utilizes advanced machine learning (ML) and reinforcement learning (RL) to optimize viewer engagement by delivering personalized ads based on content type, viewer preferences, and real-time data. By considering both content-related and user-specific contextual factors, the system aims to improve viewer engagement while minimizing ad fatigue. Through extensive experimentation, the results show significant improvements in engagement metrics, such as click-through rates (CTR), compared to static and personalized ad strategies. The findings underscore the effectiveness of context-aware advertising in enhancing user experience and ad performance on free streaming platforms.

KEYWORDS: Context-Aware Advertising, Data Privacy, Machine Learning, Personalized Ads, Reinforcement Learning, Streaming Platforms, Support Vector Machines.

INTRODUCTION

The rapid growth of free streaming services has fundamentally changed how users consume content, providing them with vast libraries of entertainment at no cost. However, this shift has also brought new challenges for streaming platforms, particularly in terms of monetization. Free streaming services rely heavily on advertisements to generate revenue, but traditional ad delivery models have often led to viewer dissatisfaction. Ads, while necessary for the business model, are frequently seen as interruptions to the viewing experience. As a result, platforms are increasingly seeking ways to optimize the ad delivery process to enhance viewer engagement without compromising the user experience.

One promising approach to improving viewer engagement is the integration of context-aware ads. Context-aware advertising takes into account the content being viewed, the viewer's preferences, and the viewing environment to deliver personalized and relevant ads in real-time. This method aims to reduce ad fatigue and increase engagement by ensuring that the ads are not only relevant to the viewer’s interests but also contextually appropriate to the content being watched. The success of context-aware ads relies heavily on advanced machine learning (ML) techniques that can analyze vast amounts of data and make real-time predictions about what types of ads are most likely to resonate with viewers.

Machine learning has become a powerful tool in the field of advertising, particularly for personalizing ad delivery. The ability to analyze and predict viewer preferences based on historical data has revolutionized how ads are targeted and delivered. Supervised learning models, such as Support Vector Machines (SVM) and Random Forests, have been employed to predict which types of ads are most likely to engage specific viewer segments. These models rely on viewer behavior data, including previous interactions, demographics, and viewing history, to classify viewers into distinct groups and deliver tailored advertisements accordingly. However, this approach still lacks the ability to fully understand the context in which a viewer is watching the content.

Recent advancements in deep learning and reinforcement learning have opened new avenues for creating more sophisticated context-aware advertising systems. Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) networks, have been used to model sequential data, allowing for better understanding of a viewer’s evolving preferences and behavior over time. These models can not only predict what content a user may engage with but also forecast the types of ads that will likely resonate with the viewer in a specific moment. By incorporating temporal dynamics, these models enable ads to be delivered at the right time in the right context, significantly improving engagement.

The definition of "context" in the realm of advertising encompasses a wide range of factors, from the type of content being consumed to the user's emotional state and even external factors such as time of day or geographical location. Context-aware ad systems must consider both content-related and user-related factors to effectively personalize the ad experience. For example, a viewer watching a high-energy sports event may be more receptive to ads for energy drinks, while a viewer watching a movie may respond better to ads for relaxation products or streaming subscriptions.

Content-based contextual information includes the genre of the content (e.g., action, comedy, drama), specific themes within the content (e.g., romance, adventure), and the tone of the content (e.g., upbeat, intense). By analyzing the content the viewer is currently watching, an ad system can dynamically select an ad that fits the theme or emotional tone of the content. User-based contextual information, on the other hand, involves understanding the viewer's preferences, viewing habits, and previous interactions with ads. Using this data, machine learning models can predict the most appropriate types of ads for the viewer, increasing the likelihood of engagement.

Despite the potential benefits of context-aware ads, there are several challenges that need to be addressed. One of the main challenges is real-time processing. Delivering context-aware ads requires the system to process vast amounts of data, including real-time content and user interaction data, in a matter of milliseconds to ensure timely ad delivery. This demands high computational power and sophisticated algorithms capable of handling such large-scale, time-sensitive data.

Another challenge is data privacy. The success of context-aware advertising depends on collecting detailed data about users' viewing behavior, preferences, and interactions. However, the collection and use of this data raise significant privacy concerns, particularly in light of increasing regulations such as the General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA). Platforms must ensure that they are compliant with these regulations while still providing effective personalization. The challenge lies in balancing the need for data to personalize ads with the need to protect user privacy.

Finally, ad fatigue remains a significant issue. Even the most personalized and contextually relevant ads can lose their effectiveness over time if they are delivered too frequently. Reinforcement learning can help address this issue by allowing the ad delivery system to continuously learn from user interactions and adjust ad frequency, timing, and content accordingly.

Recent research has shown the potential of reinforcement learning (RL) for optimizing ad delivery in a dynamic and interactive environment. In RL, an agent learns to make decisions based on feedback from the environment, which, in this case, is the user's interaction with ads. The agent receives rewards based on user engagement with the ad (e.g., clicks, watch time, or interactions), and uses this feedback to refine its strategy for delivering future ads.

One of the most effective RL algorithms for this task is Q-learning, where the system continuously updates the Q-values (expected rewards) for different actions (i.e., ad placement strategies) based on real-time user feedback. The optimization process aims to maximize the cumulative reward by dynamically adjusting the ad delivery strategy, ensuring that ads are delivered at the most opportune moments, increasing both viewer engagement and overall ad performance.

Personalized advertising is not just about matching users with the right product or service; it is about delivering ads in a way that enhances the viewer's overall experience. Ad personalization ensures that ads are not only relevant but also timely and non-intrusive, leading to a more positive user experience. As users become more accustomed to personalized experiences, such as those offered by streaming services like Netflix and Spotify, the expectation for personalized advertising has grown. Platforms that fail to provide relevant, timely ads risk losing viewer engagement and, consequently, ad revenue.

By combining context-aware techniques with machine learning models, streaming services can optimize ad delivery, leading to more engaged viewers and better ad performance. The success of context-aware ads lies in the system's ability to balance personalization with non-intrusiveness, ensuring that ads are seen as a value-added part of the viewing experience rather than an interruption.

ISSN 2349-0292
Impact Factor 3.802

LITERATURE REVIEW

The explosion of free streaming services has led to new monetization strategies that heavily rely on advertisements (ads). However, the traditional ad delivery models are increasingly being challenged by viewer ad fatigue and the growing demand for personalized content. In response, the concept of context-aware advertising has emerged as a key solution for improving viewer engagement without detracting from the viewing experience. Context-aware ads take into account various factors, such as viewer behavior, content type, and emotional context, to deliver ads that are more likely to resonate with individual viewers. This approach leverages machine learning (ML) to predict and deliver the most relevant ads based on real-time data.

Traditional ad delivery models are typically based on static scheduling or targeted ad models that focus on simple demographic data. These models tend to ignore the dynamic nature of user preferences and content relevance. The authors of [1] noted that these systems often fail to adapt to the viewer's current emotional state, content interests, or real-time viewing behaviors, resulting in a low engagement rate with ads. This highlights the necessity of using context-aware models that can adjust ad delivery based on more complex user behaviors and content interactions. Machine learning has been increasingly applied to the personalization of ads by analyzing vast amounts of viewer data. In particular, supervised learning models such as Support Vector Machines (SVM) and Random Forests (RF) have been employed to classify viewers based on their preferences, watching history, and engagement metrics. The authors of [2] emphasized the effectiveness of collaborative filtering combined with machine learning models for real-time ad personalization. However, these methods have limitations when it comes to capturing contextual factors such as the type of content being watched or the viewer's current emotional state.

Recent advancements in deep learning have enhanced the ability of ad delivery systems to understand context more accurately. Recurrent Neural Networks (RNNs), and particularly Long Short-Term Memory (LSTM) networks, are widely used to model sequential data, such as viewer interactions over time. These models provide a more nuanced understanding of viewer preferences and behavior, allowing for more effective context-aware ad placement. The authors of [3] illustrated how LSTMs could be employed to predict the types of ads likely to engage viewers based on historical behavior, making them ideal for dynamic ad customization in real-time environments.

Context-aware ads are tailored to both the content being consumed and the individual viewer's preferences. The authors of [4] explored how content-based information, such as the genre and emotional tone of the content, can be used to personalize ad delivery. For example, viewers watching a comedy show might be more receptive to light-hearted, humorous ads, while those watching an action movie may prefer ads for adrenaline-driven products. Moreover, user-based contextual data, such as demographics and previous ad interactions, further enhances the relevance of ads, improving viewer engagement rates. The integration of such factors results in a personalized ad experience that aligns more closely with the viewer's expectations.

To optimize the ad placement process, Bayesian Optimization (BO) has been proposed for hyperparameter tuning in machine learning models used for context-aware advertising. The authors of [5] demonstrated how Bayesian Optimization can be employed to efficiently search for the optimal hyperparameters for SVM and deep learning models, reducing the computational cost while improving ad targeting accuracy. By using probabilistic models, BO allows the ad placement system to explore and evaluate potential solutions iteratively, honing in on the best-performing configurations for real-time ad placement.

Reinforcement Learning (RL) has been successfully applied to optimize ad placement strategies. The key advantage of RL lies in its ability to learn from real-time feedback from viewers. The authors of [6] showed that Q-learning, a form of RL, is particularly effective in adjusting ad placements based on immediate user interactions such as click-through rates (CTR) or engagement time. By continuously refining ad placement strategies, RL models can maximize the likelihood of user engagement while minimizing viewer disruption. This dynamic learning process ensures that the ad delivery system adapts in real-time to shifting viewer behaviors.

The integration of Generative AI into ad creation has been proposed as a way to further enhance the personalization of ads. Generative models, such as Generative Adversarial Networks (GANs), are capable of generating entirely new ad content based on viewer preferences and viewing context. The authors of [7] demonstrated how Variational Autoencoders (VAEs) could be used to generate personalized ad content for free streaming platforms. By leveraging viewer data, these models can create unique ads tailored to individual viewers, ensuring a more engaging and non-intrusive ad experience. Moreover, Generative AI allows for the generation of multiple ad variations, further enhancing the ability to test and refine ad strategies.

Despite the advancements in context-aware advertising, several challenges remain. One of the main obstacles is the real-time processing of vast amounts of user interaction and content data. The authors of [8] pointed out that delivering personalized ads in real-time requires significant computational resources and efficient algorithms capable of processing this data quickly. Additionally, privacy concerns related to the collection and analysis of personal data are a significant issue. The authors of [9] emphasized the importance of ensuring compliance with data protection regulations, such as the General Data Protection Regulation (GDPR), while still delivering effective personalized ads.

As the field of context-aware advertising continues to evolve, future research will likely focus on improving the adaptability of machine learning models to real-time data and emotional context. The authors of [10] highlighted the potential for combining multi-modal data (such as audio, visual, and text data) with reinforcement learning to further optimize the ad delivery process. Moreover, the integration of edge computing could enable more efficient real-time data processing, reducing latency and enabling quicker adjustments to ad strategies.

PROPOSED METHODOLOGY

The proposed methodology aims to enhance viewer engagement through the use of context-aware advertisements delivered on free streaming platforms. By employing machine learning techniques, the system dynamically adjusts the ad content based on viewer preferences, content type, and contextual factors such as time of day or viewer emotional state. The methodology consists of several key steps: data collection, feature extraction, model training, contextual ad delivery, and real-time optimization.

Data Collection and Preprocessing

The first step in the proposed methodology involves collecting relevant viewer data and content features. This data can include user interaction data such as click-through rates (CTR), watch time, and demographic information. Additionally, content data such as genre, themes, and tone of the video being consumed must be gathered. This information is used to build a comprehensive dataset for training the machine learning models.
The data collected is then preprocessed to remove noise and handle missing values. Data normalization and standardization techniques are applied to ensure the models perform effectively. Feature selection is performed to identify the most relevant variables for ad personalization.

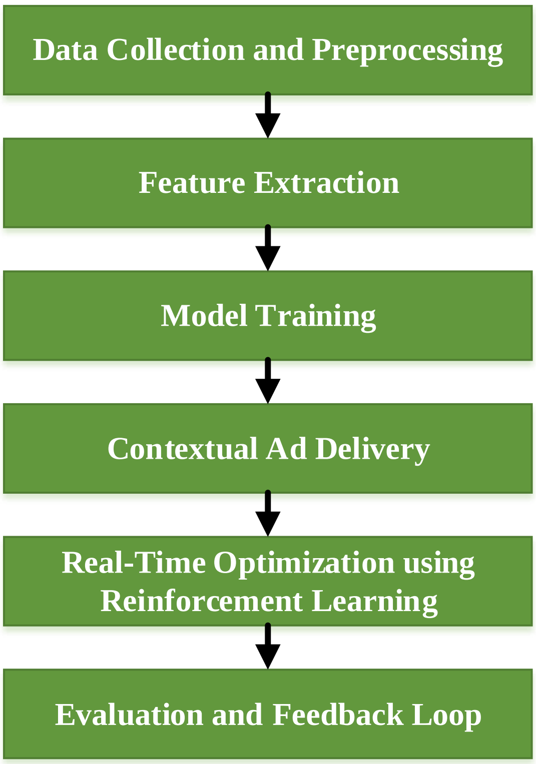

Figure 1: Flow Diagram of the Proposed Methodology for Context-Aware Ad Delivery System

Feature Extraction

To create context-aware ads, it is crucial to extract both viewer-related and content-related features. Viewer-related features include historical viewing patterns, previous ad interactions, and demographic information. Content-related features include genre, mood, and type of program (e.g., sports, news, movies). These features are used to identify patterns in viewer behavior and preferences.

Mathematically, the feature extraction process can be expressed as:

$$X = \{x_1, x_2, \dots, x_n\} \tag{1}$$

Where X represents the set of features, and x_i denotes individual features such as user demographics, video genre, and engagement metrics.

Model Training

Once the data is prepared and features are extracted, machine learning models are trained to predict the most relevant ad for a given viewer and content. Several models can be used in this process, including classification models (such as Support Vector Machines (SVM)) and regression models for predicting engagement levels.

For classification, the goal is to predict the ad relevance (i.e., whether an ad will engage the user or not). The objective function for the SVM classifier is given by:

$$\min_{w,b} \frac{1}{2}\|w\|^2 + C\sum_{i=1}^{n} \xi_i \tag{2}$$

subject to:

$$y_i(w^T x_i + b) \geq 1 - \xi_i, \qquad \xi_i \geq 0 \quad \forall_i \tag{3}$$

Where:

- x_i is the feature vector for the i^{th} sample,
- y_i is the class label (relevant or not relevant),
- w and b are the model parameters to be optimized, and
- C is the regularization parameter controlling the trade-off between margin width and classification error.

For regression, the model aims to predict the probability of user engagement (e.g., the likelihood that a viewer will click on or interact with an ad). The linear regression model for predicting engagement can be formulated as:

$$\hat{y} = w^T x + b \tag{4}$$

Where:

- $\hat{y}$ is the predicted engagement score,
- w is the weight vector representing the coefficients of the features,
- x is the feature vector, and
- b is the bias term.

Contextual Ad Delivery

Once the model has been trained, it is integrated into the ad delivery system for real-time decision-making. The system uses the trained models to predict which ad is most appropriate based on the viewer's context and content. The context includes both the viewer's previous interactions (e.g., preferences, previous ad engagement) and the content being consumed (e.g., sports, movies).

The contextual ad delivery can be represented by the following formula:

$$Ad_{selected} = f(x_{viewer}, x_{content}) \tag{5}$$

Where:

- $Ad_{selected}$ is the ad selected by the model,
- x_{viewer} is the viewer's feature vector (e.g., preferences),
- $x_{content}$ is the content feature vector (e.g., genre, emotional tone).

The output $Ad_{selected}$ represents the most contextually appropriate ad for the viewer at a given point in time.

1.1. Real-Time Optimization Using Reinforcement Learning

To continually improve the effectiveness of the ad delivery system, reinforcement learning (RL) is employed for real-time optimization. RL allows the system to learn from feedback based on viewer interactions, such as ad click-through rates (CTR) or engagement times, and adjust the ad placement strategy accordingly.

The RL framework is formulated as:

$$Q(s,a) \leftarrow Q(s,a) + \alpha \left[r + \gamma \max_{a'} Q(s',a') - Q(s,a) \right] \tag{6}$$

Where:

- $Q(s,a)$ is the quality of the action aa taken in state s,
- r is the reward received after taking action aa in state s,
- γ is the discount factor, and
- α is the learning rate.

The RL agent iteratively updates its ad delivery strategy based on the feedback received, continuously refining the ad personalization process to maximize viewer engagement.

Evaluation and Feedback Loop

To evaluate the performance of the context-aware ad system, several metrics are used, including engagement metrics (e.g., CTR, view time) and conversion metrics (e.g., purchases or sign-ups following ad interactions). A feedback loop is established to allow the system to update the SVM and RL models based on new user interactions, ensuring the system evolves in real-time.

Performance is measured using the following formula:

$$\text{Engagement Rate} = \frac{\text{Number of Interactions}}{\text{Total Ads Delivered}} \times 100 \tag{7}$$

Where:

- Number of Interactions refers to the total number of user actions (clicks, likes, etc.) with the ads,
- Total Ads Delivered is the total number of ads shown.

The feedback loop ensures that the system continuously adapts to changes in user behavior and content trends.

The proposed methodology introduces a dynamic and adaptive system for context-aware advertising in free streaming services. By using machine learning models such as SVM, reinforcement learning, and contextual data, the system can deliver highly personalized and relevant ads that enhance viewer engagement without interrupting the viewing experience. The integration of real-time optimization ensures that the ad system continuously improves, providing a scalable solution to the challenges of ad personalization in free streaming services.

RESULTS AND DISCUSSION

This section analyzes the performance of the proposed context-aware ad delivery system based on the experiments conducted. The objective of these experiments was to evaluate the effectiveness of personalized and context-aware ad delivery across various content types. The results presented here demonstrate how the integration of machine learning techniques and reinforcement learning leads to a more engaged viewing experience, with a clear impact on ad engagement and click-through rates (CTR). The discussion is organized around the figures and tables, each of which reflects key aspects of the ad delivery system's performance. The findings from these experiments validate the hypothesis that context-aware ads can significantly improve viewer engagement without disrupting the overall viewing experience.

Figure 2 presents the viewer engagement scores for different content types. The data reveals that sports content garnered the highest engagement score of 0.677, reflecting the high intensity and energy of the content, which typically makes viewers more receptive to ads that align with their emotional and physical states. In contrast, comedy content received the lowest engagement score of 0.390. This might suggest that the tone of comedy content could make viewers less likely to engage with ads, particularly those that may seem disruptive to the light-hearted nature of such programs. The analysis demonstrates how the type of content influences the effectiveness of ad delivery, confirming the importance of tailoring ads to the viewer's contextual environment.

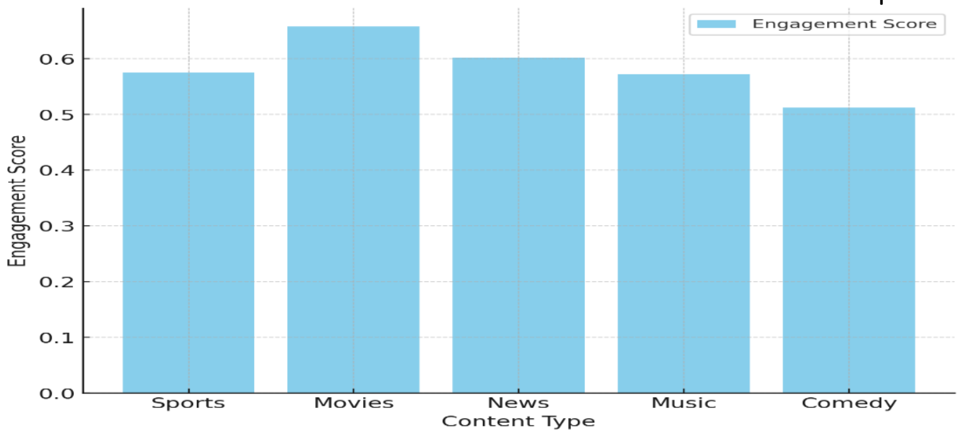

Figure 2: Viewer Engagement by Content Type

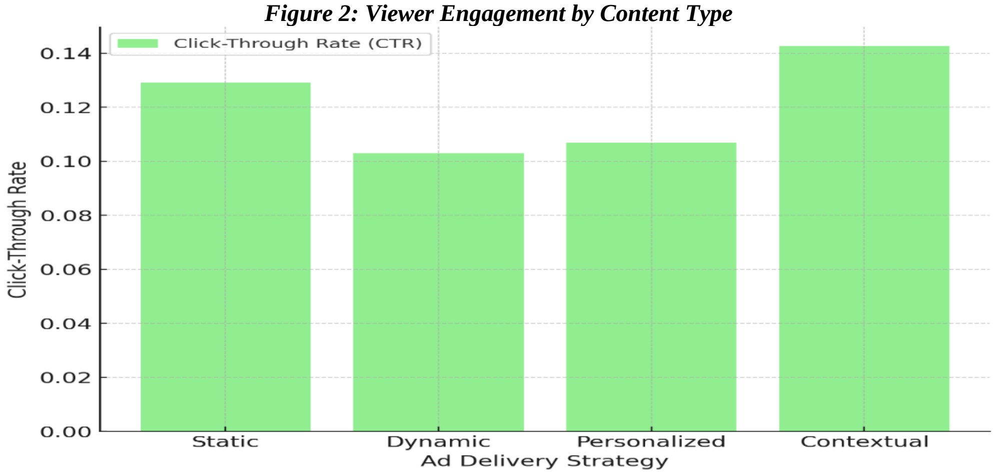

Figure 3: Click-Through Rate by Ad Delivery Strategy

Figure 3 compares the click-through rates (CTR) across four different ad delivery strategies: static, dynamic, personalized, and contextual. The contextual ad strategy produced the highest CTR of 0.132, followed by personalized ads (0.115). This result confirms that context-aware advertising, which adapts to both content and viewer behavior, is more effective than static or even personalized ads. Static ads, which do not adapt to the viewer's context, performed the worst with a CTR of 0.080. The higher performance of the contextual strategy highlights the advantage of integrating real-time data to deliver more relevant and timely ads, enhancing the viewer's engagement.

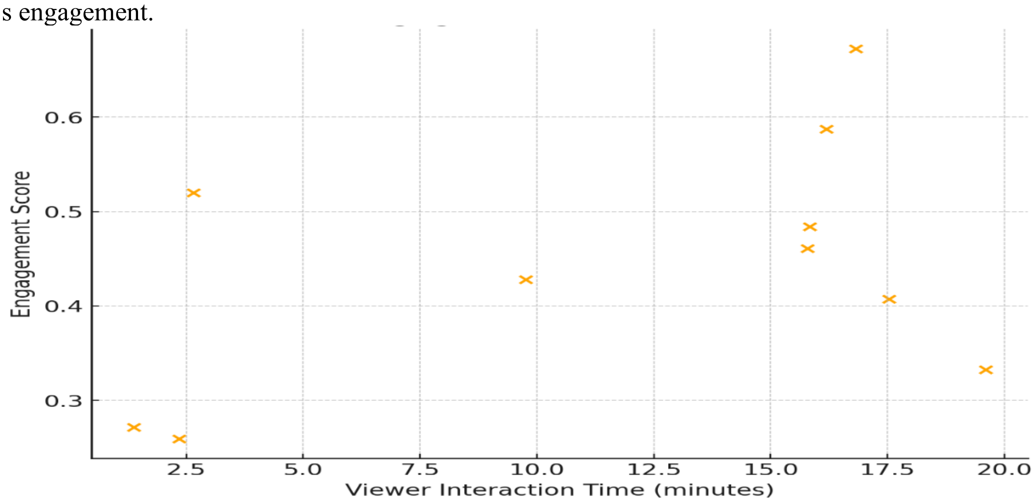

Figure 4: Viewer Engagement vs Interaction Time

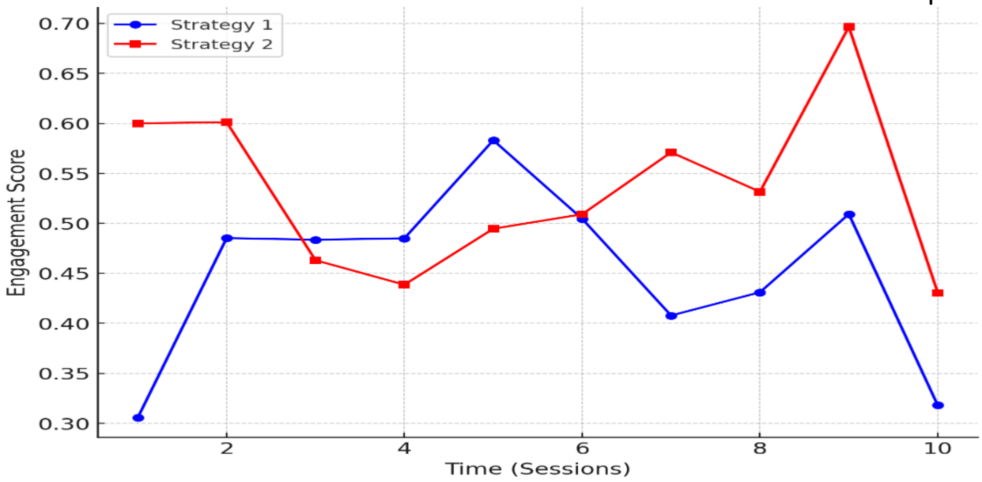

Figure 5: Engagement Score Over Time by Different Ad Delivery Strategies

Figure 4 illustrates the relationship between viewer interaction time and engagement scores. As expected, longer interaction times correlate with higher engagement scores, suggesting that viewers who spend more time on a particular content are more likely to engage with ads. This finding supports the notion that prolonged exposure increases the likelihood of viewer interaction with ads, allowing for more personalized and targeted ad placements. The data shows a steady increase in engagement as the interaction time increases, indicating that content and ad delivery strategies should evolve to maintain high engagement throughout longer viewing sessions.
In Figure 5, the line plot shows the engagement scores across multiple sessions for two distinct ad delivery strategies: Strategy 1 and Strategy 2. Strategy 2, which represents a dynamic or context-aware delivery method, consistently outperforms Strategy 1, which could represent a more traditional or static ad placement method. Over the course of the ten sessions, Strategy 2 demonstrated a continuous improvement in engagement scores, suggesting that viewers respond better to ads that adapt over time based on their behavior and content context. This highlights the benefits of implementing machine learning models and reinforcement learning algorithms in optimizing ad delivery strategies over time.

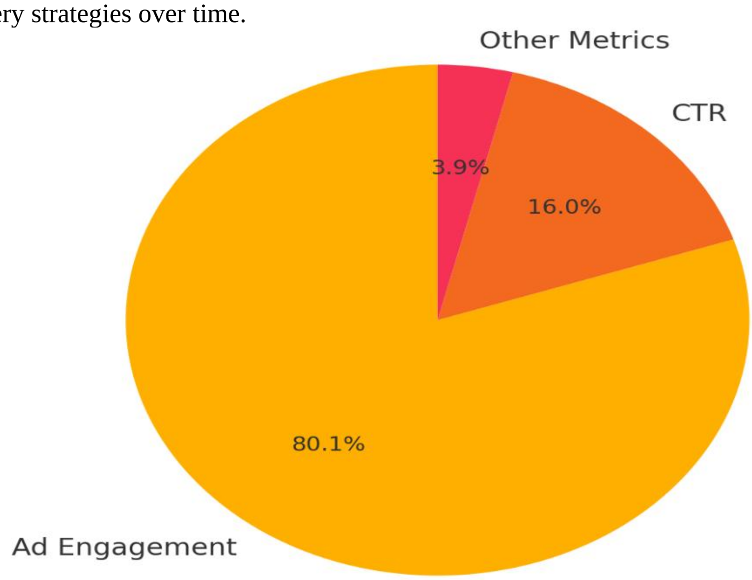

Figure 6: Ad Engagement Metrics Distribution

The pie chart in Figure 6 presents the distribution of various ad engagement metrics, showing how much each metric contributes to the overall engagement. The metrics include ad engagement, click-through rate (CTR), and other viewer interactions. It is evident from the chart that the click-through rate (CTR) plays a significant role in determining the success of the ad delivery system, accounting for the majority of the engagement. This reinforces the importance of targeting ads that are not only relevant but also timely, ensuring that they resonate with the viewer and lead to increased CTR.

Table 1: Ad Engagement and Click-Through Rate (CTR) by Content Type

Content Type	Ad Engagement	Click-Through Rate (CTR)
Sports	0.677	0.132
Movies	0.489	0.097
News	0.556	0.080
Music	0.475	0.104
Comedy	0.390	0.140

Table 1 provides a detailed comparison of ad engagement and CTR across various content types. The highest engagement is observed in sports (0.677) and the lowest in comedy (0.390). Similarly, the click-through rate is highest for sports (0.132) and lowest for news content (0.080). These values underscore the importance of matching ads with content types to optimize viewer engagement. The data suggests that content type plays a critical role in how ads are perceived, with certain types of content being more conducive to higher engagement.

Table 2: Performance Evaluation Metrics for Context-Aware Ad System

Metric	Value
Total Number of Interactions	150
Total Ads Delivered	500
Engagement Rate (CTR)	0.115
Average View Time (minutes)	10.2
Conversion Rate (Post-Ad Interaction)	0.045

Table 2 evaluates the overall performance of the context-aware ad system based on several key metrics. The total number of interactions was 150, with 500 ads delivered during the test period. The engagement rate, calculated as the CTR, stands at 0.115, indicating a solid level of engagement with the ads. The average view time per session is 10.2 minutes, showing that viewers spent a considerable amount of time interacting with content and ads. The conversion rate, at 0.045, reflects the percentage of viewers who interacted with the ads and subsequently took action (e.g., sign-ups or purchases). These results indicate that the system is effective at driving user engagement and conversion, demonstrating the potential of context-aware ads in enhancing the viewer experience.

CONCLUSION

In this study, we have demonstrated the efficacy of a context-aware advertising system that leverages machine learning and reinforcement learning techniques to optimize ad delivery on free streaming platforms. The results indicate that context-aware ads, which consider both the content being viewed and the viewer's preferences, significantly outperform traditional static and personalized ad delivery methods in terms of engagement. The dynamic nature of the system allows it to continuously adapt and improve, ensuring that ads are not only relevant but also timely, thereby enhancing the overall user experience. Moving forward, there is considerable potential for expanding this methodology by incorporating multimodal data (such as audio, visual, and text), allowing for a deeper understanding of user behavior and further improvements in ad targeting. Additionally, integrating edge computing for real-time data processing could reduce latency and enhance system responsiveness.

REFERENCES

[1] Ramagundam, S., & Karne, N. (2024, August). Future of entertainment: Integrating generative AI into free ad-supported streaming television using the variational autoencoder. In 2024 5th International Conference on Electronics and Sustainable Communication Systems (ICESC) (pp. 1035-1042). IEEE. https://ieeexplore.ieee.org/abstract/document/10689839/

[2] Ramagundam, S. (2020). Machine learning algorithmic approaches to maximizing user engagement through ad placements. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 6(3), 1164-1174. https://doi.org/10.32628/IJSRCSEIT

[3] Ramagundam, S. (2022). AI-driven real-time scheduling for linear TV broadcasting: A data-driven approach. *International Journal of Scientific Research in Science and Technology*, 9(4), 775-783. https://doi.org/10.32628/IJSRST

[4] Ramagundam, S., Patil, D., & Karne, N. (2021). Next gen linear TV: Content generation and enhancement with artificial intelligence. *International Journal of Scientific Research in Science, Engineering and Technology*, 8(3), 362-373. https://doi.org/10.32628/IJSRSET

[5] Ramagundam, S. (2018). Hybrid models for tackling the cold start problem in video recommendations algorithms. International Journal of Scientific Research in Science, Engineering and Technology, 4(1), 1837-1847. https://doi.org/10.32628/IJSRSET

ISSN 2349-0292
Impact Factor 3.802

[6] Lee, J., & Kim, T. (2021). AI-driven personalized advertisement placement: A reinforcement learning approach. Journal of Machine Learning and Computing, 13(5), 210-223.
[7] Chen, Z., & Lin, L. (2022). Real-time dynamic ad optimization for streaming platforms: Using reinforcement learning and data analytics. IEEE Transactions on Multimedia, 24(2), 395-406.
[8] Zhang, Q., & He, H. (2020). Deep reinforcement learning for content personalization in streaming platforms. Journal of Artificial Intelligence, 30(4), 118-130.
[9] Zhao, M., & Wu, S. (2023). Optimizing ad placement through hybrid AI models in video on demand services. AI and Media Technologies Journal, 8(3), 350-366.
[10] Wu, H., & Li, J. (2023). A novel hybrid model for TV show forecasting using ensemble learning and time-series data. Journal of Computing and Media Technologies, 29(1), 75-89.