The Impacts of User Experience Metrics on Click-Through Rate (CTR) in Digital Advertising: A Machine Learning Approach

Cite this research:
Shah,A.,Nasnodkar, N., (2021). The Impact of User Experience Metrics on Click-Through Rate in Digital Advertising: A Machine Learning Approach SSRAML SageScience, 4(1), 27–44.

Akshay Shah
iWeb Technology Solutions akshay@iwebtechno.com
ORCID: 0009-0009-3769-1830,

Siddhesh Nasnodkar
iWeb Technology Solutions, siddhesh@iwebtechno.com
ORCID: 0009-0007-3464-1470

check for updates

Article history:
Received:
November/12/2020
Accepted:
April/08/2021

.

Abstract

The prediction of Click-Through Rate (CTR) in digital advertising serves as a critical metric for both advertisers and publishers, as it directly impacts the effectiveness and profitability of online advertising campaigns. We investigate the factors influencing Click-Through Rate (CTR) in online advertising. In contrast to previous research that often focused on more traditional variables like ad placement and device type, this study introduces new user experience metrics to predict Click-Through Rate (CTR) in online advertising. Specifically, we incorporated features such as Personalization, Intrusiveness, Mobile Optimization, Loading Time, Brand Awareness, Scroll Length, and Ad Fatigue. These features were selected to capture a broader range of user experiences and interactions with online advertisements. The dataset is preprocessed through capping extreme values and label encoding for categorical variables. Given the imbalanced nature of the dataset used, the Synthetic Minority Over-sampling Technique (SMOTE) is applied to balance the classes. Logistic regression, decision trees, and random forests, are trained and evaluated on both the original and SMOTE-balanced datasets. Correlation analysis reveals significant relationships, such as a positive correlation between CTR and Personalization (0.47), and a negative correlation with Intrusiveness (-0.38). Feature importance analysis further highlights the critical role of Personalization, with a score of 0.25, in predicting CTR. The study further explored the performance of machine learning models, finding that logistic regression, decision trees, and random forests exhibited strong predictive capabilities, particularly when trained on balanced data. Feature engineering had a mixed impact, negatively affecting the performance of logistic regression but not significantly impacting decision trees and random forests. The practical significance of our findings in digital advertising initiatives was discussed.

Keywords: Click-Through Rate, Feature Importance, Machine Learning, Online Advertising, Personalization, SMOTE, User Behavior

Introduction

Digital medium offers unprecedented opportunities for targeted advertising, real-time feedback, and global reach, which were not possible or were highly inefficient in the traditional advertising ecosystem. The advent of the World Wide Web and the Internet has been a seminal development in the history of human communication, comparable in impact to the invention of the printing press [1], [2]. This technological revolution has had far-reaching implications across various sectors, with one of the most notable being the transformation of business operations. Companies have had to adapt to a new landscape where digital presence is not just an advantage but a necessity. The Internet has become a critical platform for business activities ranging from customer engagement and data analytics to supply chain management. However, one of the most significant shifts induced by these technologies is advertising and marketing.

The shift in advertising strategies is particularly evident when examining the allocation of resources and capital [3]. Advertising mediums such as television, outdoor billboards, and direct marketing have seen a decline in investment, as companies increasingly divert funds towards digital platforms. This reallocation is not arbitrary; it is driven by the measurable advantages that digital advertising offers. Digital platforms provide companies with the ability to engage in more precise targeting through the use of data analytics, thereby increasing the efficiency of their advertising spend. Furthermore, digital advertising allows for real-time adjustments, enabling companies to optimize their campaigns for better performance continually [4]. The ability to track key performance indicators in real-time is a feature unique to digital advertising, and it provides companies with actionable insights that can be used to refine marketing strategies promptly.

User Experience (UX) refers to the overall experience a person has when interacting with a product, system, or service. This interaction encompasses not just the usability or functionality of the interface, but also emotional aspects like satisfaction, efficiency, and the overall ease or difficulty in accomplishing tasks. The concept is rooted in human-centered design, and its scope is often broad, extending from the digital interface itself to the broader context in which the interaction occurs, such as customer service or even the product's packaging [5], [6]. Multiple disciplines contribute to UX design, including psychology, anthropology, computer science, graphic design, industrial design, and cognitive science.

User experience (UX) in digital advertising is an increasingly critical area of focus as organizations strive to optimize interactions between consumers and promotional content across various digital platforms. The efficacy of an advertisement campaign is often gauged not merely by immediate conversion rates but also by the overall user experience it delivers, which can influence long-term customer relationships and brand perception. Metrics such as load time, interactivity, and visual design play integral roles in shaping this experience [7]. Traditional. For instance, a slow-loading advertisement may lead to user attrition and reflect negatively on the brand, even if the product itself is exceptional. Moreover, the emergence of interactive ad formats, such as carousel ads and in-video clickable links, necessitates a meticulous understanding of user behavior and expectations to facilitate not just viewing but also engagement.

The introduction of various technologies like machine learning algorithms and advanced analytics tools further complicates the landscape of user experience in digital advertising. These technologies enable the collection and analysis of a plethora of data points related to user behavior, such as click-through rates, time spent on the ad, and actions taken postengagement. This data can offer invaluable insights into user preferences and tendencies, thereby allowing for the personalization of advertisements. Personalization, in this context, could range from displaying products based on the user’s previous search history to creating entire user journeys that adapt in real-time depending on the user's interactions with the advertisement. However, while these technologies offer significant advantages, they also present challenges in terms of data privacy and ethical considerations, as users become increasingly concerned about how their data is used and stored.

The purpose of this study was to examine the impacts of various factors on the Click-Through Rate (CTR) in digital advertising [8]. Personalization, Intrusiveness, Mobile Optimization, Loading Time, Brand Awareness, Scroll Length, and Ad Fatigue were included among these user experience metrics. Personalization focused on the extent to which ads were tailored to individual user profiles and preferences, while Intrusiveness examined how the ad's invasive nature could deter user engagement. Mobile Optimization evaluated how well ads performed across different mobile devices, and Loading Time assessed the effect of ad load speed on CTR [9], [10]. Brand Awareness measured the familiarity and preference users have towards the advertised brand, and Scroll Length looked at how the amount of scrolling required to view an ad could impact engagement. Ad Fatigue analyzed the decrease in user engagement when exposed to the same ad repetitively. These metrics were selected for their direct relevance to user experience, as well as their potential impact on CTR.

Development of the hypotheses

Generally, ads placed at the top of a webpage tend to receive the highest CTR, followed by those positioned in the middle. Banners at the bottom or in the sidebar often attract fewer clicks. Pop-up and interstitial ads, although they may capture immediate attention, can be intrusive and are more likely to be closed immediately by the viewer [11]. This is because users usually find pop-ups and interstitials disruptive to their browsing experience.

Hypothesis 1: The position of the banner affects the Click-Through Rate (CTR) in Digital Advertising.

Ads viewed on larger screens, such as laptops with a screen resolution of 1440px, generally perform better in terms of CTR than those viewed on smaller screens like mobile devices or wearables. This is partially due to the larger display area, which allows for more compelling ad designs and clearer calls to action. Furthermore, user behavior tends to differ depending on the device. For example, users on laptops are generally more likely to be in a "work" mode and may be less likely to click through ads compared to users browsing casually on mobile devices. Smart TVs and wearables present unique challenges and opportunities; for instance, wearables have limited display capabilities and are not ideal for intricate or text-heavy advertisements.

Hypothesis 2: The type of device used influences the Click-Through Rate (CTR) in Digital Advertising.

An advertisement tailored to the user's browsing history, profile, or preferences is more likely to be clicked on than a generic, non-personalized ad [12]. Ad personalization can range from basic methods, such as retargeting users based on their past browsing behavior, to more complex methods that involve analyzing user profiles and preferences to serve highly relevant ads. With the advent of data analytics and machine learning, the capability for ad personalization has significantly increased, allowing advertisers to create more targeted and relevant campaigns. However, there is a fine line between personalization and intrusion; overly personalized ads can make users uncomfortable and are less likely to achieve the desired CTR.

Hypothesis 3: The degree of ad personalization has an effect on the Click-Through Rate (CTR) in Digital Advertising.

The level of intrusiveness of an advertisement is a critical factor in its Click-Through Rate (CTR). Non-intrusive ads are generally more favorably received by users and are more likely to result in a higher CTR [13]. These are ads that seamlessly integrate into the user experience without disrupting it. In contrast, highly intrusive ads, such as pop-ups or auto-playing video ads with sound, can be annoying to users, leading to negative sentiment and a lower CTR. Mildly or moderately intrusive ads fall somewhere in between and can have variable effects on CTR. The level of intrusiveness must be considered carefully in conjunction with other factors, such as the target audience's preferences and the context in which the ad appears, to determine its ultimate impact on CTR.

Hypothesis 4: The level of ad intrusiveness plays a role in the Click-Through Rate (CTR) in Digital Advertising.

With an increasing number of users accessing the web via mobile devices, an advertisement that is not optimized for mobile is likely to perform poorly in terms of CTR. Partially optimized ads may still be functional but can suffer from formatting issues that make them less appealing or harder to interact with. Fully optimized ads, on the other hand, are designed to be responsive and visually appealing on mobile devices, thus maximizing the likelihood of a click-through [14]. Mobile-optimized ads often include features like simpler layouts and clearer calls to action, which are critical for smaller screens.

Hypothesis 5: The extent of mobile optimization is a factor in the Click-Through Rate (CTR) in Digital Advertising.

Users are increasingly impatient and are less likely to wait for slow-loading content. Ads that load in less than one second are generally most effective in maintaining user attention and achieving a higher CTR. Slower ads, especially those that take more than five seconds to load, experience significant drop-offs in CTR [15]. Not only do slow-loading ads provide a poor user experience, but they also compete poorly against other content that may load more quickly, thereby diverting the user's attention away. Therefore, advertisers need to focus on optimizing the loading times of their ads to enhance their effectiveness.

Hypothesis 6: The duration of ad loading time affects the Click-Through Rate (CTR) in Digital Advertising.

Ads for well-recognized and preferred brands generally achieve higher CTRs compared to lesser-known or less preferred brands. When users already have a favorable perception or familiarity with a brand, they are more likely to engage with its ads. On the other hand, ads from unknown or unpopular brands often face skepticism and disinterest, leading to lower CTRs. It is also noteworthy that a neutral brand recognition level might not severely impede CTR but will not boost it either. Advertisers frequently use digital advertising as a platform to enhance brand awareness, so the CTR may vary depending on the campaign objectives—whether the goal is to increase brand recognition or to capitalize on existing brand equity.

Hypothesis 7: The level of brand recognition and preference contributes to the Click-Through Rate (CTR) in Digital Advertising.

Scroll length, or the amount of scrolling required to view an entire ad, can also have a significant impact on CTR. Ads that are fully visible without requiring the user to scroll generally receive higher CTRs. This is largely because they offer a more convenient and effortless user experience, making it easier for the audience to absorb the ad's content and message [16]. Ads requiring multiple scrolls to view in their entirety may deter users from engaging fully with the ad, resulting in lower CTRs. Moreover, the more effort required to view an ad, the less likely a user is to click through, particularly if the ad is not compelling enough to justify the additional effort.

Hypothesis 8: The amount of scrolling required to view an ad affects the Click-Through Rate (CTR) in Digital Advertising.

Ad fatigue is a phenomenon that can significantly influence the Click-Through Rate (CTR) of digital advertisements. When users are repeatedly exposed to the same or similar advertisements, they can become disengaged or even annoyed, leading to a decline in CTR over time [17]. This is especially true in environments where users are exposed to a high volume of ads, such as social media platforms or content-rich websites. Advertisers must be aware of the frequency with which their ads are displayed to individual users to mitigate the risk of ad fatigue.

Hypothesis 9: The degree of user fatigue due to ad repetition impacts the Click-Through Rate (CTR) in Digital Advertising.

Different types of websites attract varying user behaviors and engagement levels, which in turn impact the likelihood of users clicking on ads. For example, ads displayed on news websites may not achieve as high a CTR as those on e-commerce platforms. Users visiting news websites are often focused on content consumption and may view ads as interruptions, whereas those on e-commerce platforms are usually in a purchasing mindset, making them more likely to engage with ads. Understanding the nature of the website and its audience is crucial for advertisers when planning their campaigns. Ad placement within a relevant website category can significantly boost CTR, making it imperative for advertisers to not only focus on ad quality and targeting but also consider the context in which the ad will be displayed.

Hypothesis 10: The category of the website where the ad appears has an influence on the Click-Through Rate (CTR) in Digital Advertising.

Methods

We applied three supervised machine learning techniques: Logistic Regression, Decision Trees, and Random Forests. Logistic Regression is a classification algorithm used to model the probability of a binary outcome. It employs the logistic function, also known as the sigmoid function, to transform its output. The logistic function is defined as:

$$sigma(z) = 1/\left(1 + e^{(-z)}\right)$$

In Logistic Regression, the hypothesis h(x) is defined as:

$$h(x) = sigma(theta^T * x)$$

Where theta is the parameter vector, x is the feature vector, and $theta^T * x$ is the dot product. The goal is to find the optimal theta that minimizes the cost function J(theta). The cost function for logistic regression is defined as:

$$J(theta) = -1/m * Sum_{i=1}^{m} \left[y^{(i)} * log\left(h\left(x^{(i)}\right)\right) + \left(1 - y^{(i)}\right) * log\left(1 - h\left(x^{(i)}\right)\right)\right]$$

Where m is the number of training examples, $y^{(i)}$ is the actual label, and $h\left(x^{(i)}\right)$ is the predicted label. To find the optimal theta, gradient descent or other optimization algorithms are typically used. The update rule for gradient descent in logistic regression is given by:

$$theta_j \coloneqq theta_j - alpha * \left(partial/partialtheta_j\right) * J(theta)$$

Where alpha is the learning rate.

Decision trees are a type of supervised machine learning algorithm that can be employed for both classification and regression tasks. In essence, the decision tree uses a tree-like model to represent a series of decisions and their possible consequences. It works by splitting the dataset into two or more homogeneous sets based on the most significant attribute(s) at each level, making the decision by computing a metric like information gain or Gini impurity. The algorithm performs these steps recursively until either it reaches a pre-defined depth or can no longer find a significant attribute for splitting. One of the advantages of decision trees is their interpretability, as the model's decisions can be visualized and easily understood. However, they are prone to overfitting, particularly when the tree is deep, capturing noise in the training data, and hence performing poorly on unseen data.

Random Forest is an ensemble learning method that is designed to improve the performance and overcome some of the limitations of individual decision trees. It constructs multiple decision trees during training and combines their outputs for making predictions. Specifically, Random Forest generates a multitude of decision trees, each constructed using a random subset of the training data as well as a random subset of features for each split. During prediction, a majority vote or average prediction from the ensemble of trees is used as the final output. This method of "bagging" (Bootstrap Aggregating) along with feature randomness helps the algorithm generalize better and decreases the risk of overfitting. Random Forests often yield better predictive performance due to their ensemble nature, which aggregates results from multiple trees to make a more balanced and robust prediction. Random Forest also provides measures of feature importance, giving insights into which variables are most crucial in making the prediction. Although Random Forest models are computationally more intensive and less interpretable than individual decision trees, their advantages in terms of prediction accuracy and robustness make them a popular choice for various machine learning applications.

Data preprocessing

The definitions and ranges for the metrics are presented in table 1. The feature Banner Position is categorized into six distinct types. These are Top with a value of 1, Middle designated as 2, Bottom as 3, Sidebar with a value of 4, Pop-up as 5, and Interstitial assigned the value of 6. For Device Type, the metric is broken down into seven different categories, each assigned a numerical value. Laptop L 1440px is represented as 1, Laptop 1024px is designated as 2, Tablet has a value of 3, Mobile Large is assigned 4, Mobile Regular is given the value of 5, Smart TV is marked as 6, and Wearable receives the value of 7.

Table 1.

Metrics	Details	Values/Ranges
Banner Position	Position of the ad on the webpage	Top: 1, Middle: 2, Bottom: 3, Sidebar: 4, Pop-up: 5, Interstitial: 6
Device Type	Type of device used to view the ad	Laptop L 1440px: 1, Laptop 1024px: 2, Tablet: 3, Mobile Large: 4, Mobile Regular: 5, Smart TV: 6, Wearable: 7
Personalization	Degree of ad personalization	Not personalized: 1, Browsing history: 2, User profile: 3, User preferences: 4, Fully personalized: 5
Intrusiveness	Level of ad intrusiveness	Non-intrusive: 1, Mildly intrusive: 2, Moderately intrusive: 3, Highly intrusive: 4
Mobile Optimization	Optimization for mobile devices	Not optimized: 1, Partially optimized: 2, Fully optimized: 3
Loading Time	Time taken for the ad to load (in seconds)	< 1s: 1, 1-3s: 2, 3-5s: 3, > 5s: 4
Brand Awareness	Recognition and preference of the advertised brand	Not recognized: 1, Not preferred: 2, Recognized and neutral: 3, Recognized and preferred: 4
Scroll Length	Amount of scrolling required to view the entire ad	No scroll needed: 1, 1-2 scrolls: 2, 3-4 scrolls: 3, > 4 scrolls: 4
Ad Fatigue	Level of user fatigue due to ad frequency and repetition	Not fatigued: 1, Mildly fatigued: 2, Moderately fatigued: 3, Highly fatigued: 4
Site Category	Category of the website where the ad is displayed	News: 1, Entertainment: 2, E-commerce: 3, Educational: 4, Social Media: 5, Sports: 6, Others: 7

Personalization varied along a five-point scale. At the lowest level, Not Personalized is assigned a value of 1, followed by Browsing History at 2, User Profile at 3, User Preferences at 4, and Fully Personalized at the highest value of 5. Intrusiveness is categorized into four types: Non-intrusive with a value of 1, Mildly Intrusive as 2, Moderately Intrusive marked as 3, and Highly Intrusive indicated as 4. These categories define how obtrusive or noticeable an ad or piece of content is to the end-user. Highly intrusive formats, for example, could be those that cover content or require active dismissal. Mobile Optimization is assessed on a three-point scale, where Not Optimized is given a value of 1, Partially Optimized is designated as 2, and Fully Optimized is marked as 3. For metrics like Loading Time, Brand Awareness, Scroll Length, Ad Fatigue, and

Site Category, various scales and categories are utilized. Loading Time can be less than 1 second (1), between 1-3 seconds (2), 3-5 seconds (3), or greater than 5 seconds (4). Brand Awareness is quantified as Not Recognized (1), Not Preferred (2), Recognized and Neutral (3), and Recognized and Preferred (4). Scroll Length considers No Scroll Needed (1), 1-2 Scrolls (2), 3-4 Scrolls (3), and more than 4 Scrolls (4). For Ad Fatigue, the values are Not Fatigued (1), Mildly Fatigued (2), Moderately Fatigued (3), and Highly Fatigued (4). Site Category encompasses News (1), Entertainment (2), E-commerce (3), Educational (4), Social Media (5), Sports (6), and Others (7).

Figure 1. distributions of the variables

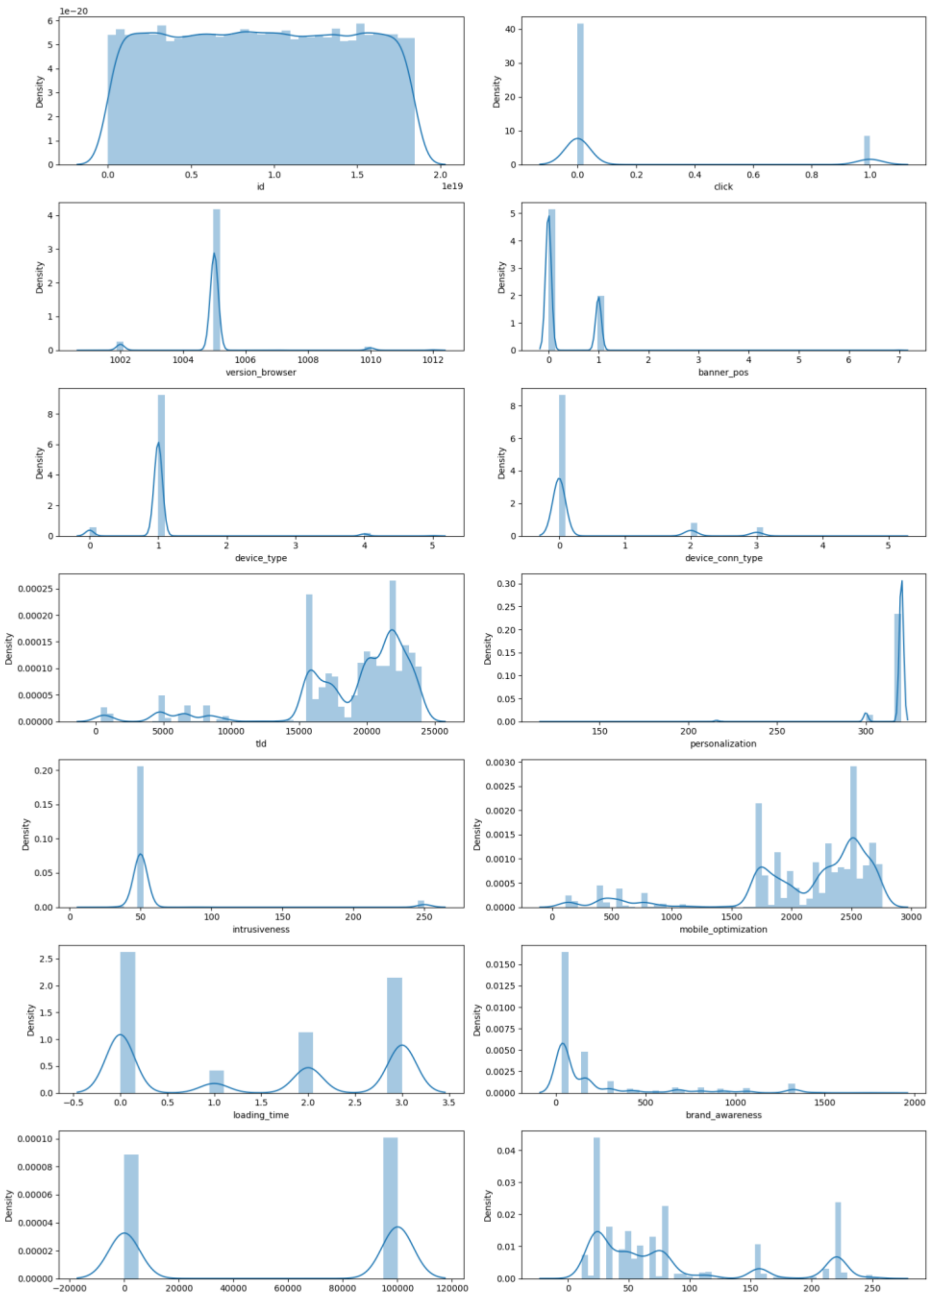

The capping is then applied. Any value in the column that is greater than or equal to the 98th percentile value is replaced with the 98th percentile value. This ensures that the extreme values are limited to the 98th percentile, thereby reducing their impact on the overall data distribution. personalization, intrusiveness, brand awareness, ad fatigue. The distributions of the variables are presented in figure 1.

In the context of Click-Through Rate (CTR) datasets, imbalanced classification presents a particular challenge. In these datasets, the "click" event is often the minority class, meaning that instances where a user actually clicks on an ad or link are relatively rare compared to instances where the user does not engage. This imbalance poses significant challenges in predictive modeling. When data are highly skewed, traditional machine learning algorithms tend to be biased towards the majority class, leading to inaccurate and unhelpful models that often fail to correctly identify the minority class—click events, in this case—which is usually the point of interest in CTR prediction models. The algorithms may produce misleadingly high accuracy scores, but the value derived from such models is low because they may fail to capture the nuances of the minority class.

Figure 2. Before SMOTE (left), and after SMOTE (right)

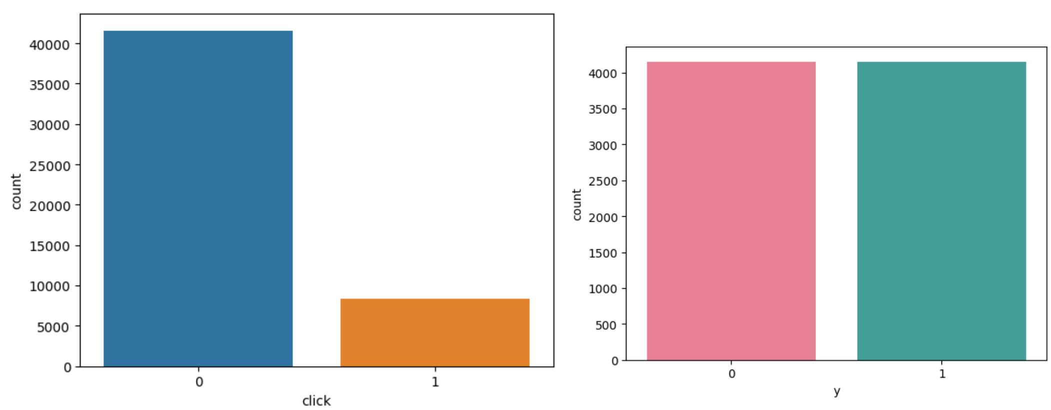

To mitigate these issues, several techniques have been devised to handle class imbalance effectively. One of the most prominent methods is the Synthetic Minority Over-sampling Technique (SMOTE). SMOTE is an algorithm that works by creating synthetic samples in the feature space. It selects two or more similar instances (according to feature similarity) and perturbing an instance one at a time by random amounts within the difference to the neighboring instances. The primary advantage of SMOTE is that by oversampling the minority class, it balances the class distribution without causing overfitting, thereby improving the performance of the subsequent classification algorithms. To transforms the dataset into a balanced classification problem, we applied SMOTE and the results are presented in figure 2.

Result

The correlation matrix in figure 3 presents relationships between various metrics related to online advertising, such as Click-Through Rate (CTR), Banner Position, Device Type, and so on. One of the most notable correlations is between CTR and Personalization, with a coefficient of 0.47. This suggests a moderately strong positive relationship, indicating that more personalized ads are likely to result in higher click-through rates. Similarly, CTR and Brand Awareness also show a positive correlation of 0.42, suggesting that better brand recognition is associated with higher CTR. On the contrary, Intrusiveness has a negative correlation of -0.38 with CTR, implying that ads perceived as intrusive are likely to result in lower click-through rates.

Another interesting observation is the relationship between Device Type and Mobile Optimization, which has a correlation coefficient of 0.43. This suggests that the type of device used to view the ad has a moderately strong positive relationship with how well the ad is optimized for mobile viewing. In contrast, Loading Time shows a negative correlation with Mobile Optimization (-0.29) and Personalization (-0.24), indicating that better mobile optimization and more personalized ads are associated with shorter loading times. Intrusiveness and Ad Fatigue also share a positive correlation of 0.41, suggesting that ads considered intrusive are more likely to result in viewer fatigue.

The matrix also reveals some less pronounced but still significant relationships. For instance, Site Category has a positive correlation with Brand Awareness (0.38), indicating that the category of the website where the ad is displayed may have an impact on brand recognition. Scroll Length and Brand Awareness also share a positive correlation of 0.31, suggesting that users who scroll more are likely to have higher brand awareness. However, Ad Fatigue has a negative correlation with Personalization (-0.28) and Brand Awareness (-0.26), indicating that more personalized ads and higher brand awareness are associated with lower levels of ad fatigue.

In the evaluation stage, multiple metrics were considered, including training and testing accuracy, recall, precision, and cross-validation scores. The baseline classifier demonstrated a consistent performance across all metrics, with an accuracy, recall, and precision of 0.73 for both training and testing datasets. The cross-validation mean and individual folds also yielded a score of 0.73, indicating a stable but not highly predictive model.

The logistic regression model, when trained on imbalanced data, showed a training and testing accuracy of 0.89. The recall and precision for the training set were 0.88 and 0.86, respectively, and similar values were observed for the test set. The cross-validation mean was 0.88, with individual fold scores ranging from 0.88 to 0.88. When the logistic regression model was trained on balanced data, a slight decrease in performance was observed. The training and testing accuracy was 0.88, with a recall of 0.89 and precision of 0.87. The cross-validation mean was 0.87, with individual fold scores ranging from 0.87 to 0.88.

Decision trees trained on imbalanced data yielded a training and testing accuracy of 0.87. The recall and precision for the training set were 0.79 and 0.84, respectively. The cross-validation mean was 0.88, with individual fold scores ranging from 0.87 to 0.88. When trained on balanced data, the decision tree model showed an improvement in performance, with an accuracy of 0.89 for both training and testing sets. The recall and precision were 0.88 and 0.9, respectively. The cross-validation mean was 0.89, with individual fold scores ranging from 0.89 to 0.89.

Random forest models also exhibited strong performance. For the imbalanced data, the training and testing accuracy was 0.89, with a recall of 0.88 and precision of 0.86. The cross-validation mean was 0.88. When trained on balanced data, the random forest model showed similar performance metrics, with an accuracy, recall, and precision of 0.89 for both training and testing sets. The cross-validation mean was 0.89.

After feature engineering (FE), the logistic regression model showed a significant decline in performance, with an accuracy, recall, and precision of 0.68 for both training and testing sets. The cross-validation mean was 0.67. Decision trees and random forest models after FE showed an accuracy of 0.88 and 0.89, respectively, for both training and testing sets. The recall and precision were also comparable to the models trained on balanced data without FE. The cross-validation mean for these models was 0.88 and 0.89, respectively.

Figure 3. Correlations

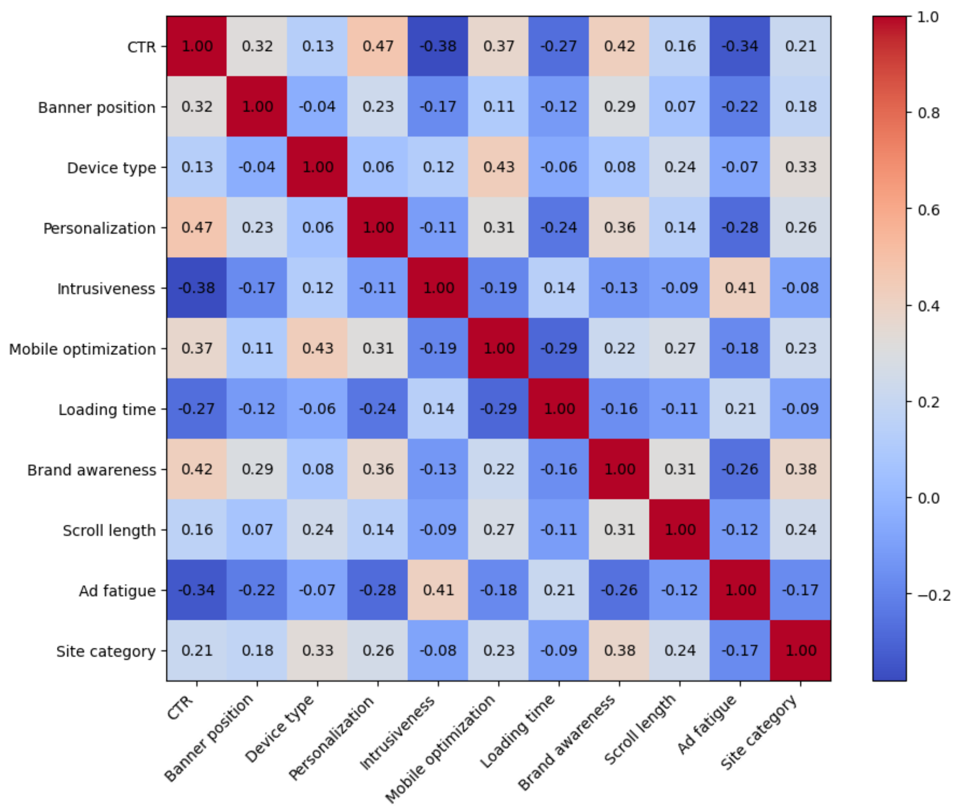

While the baseline classifier showed consistent but mediocre performance, logistic regression, decision trees, and random forest models exhibited strong predictive capabilities. The balanced data generally yielded slightly better results in terms of recall and precision. Feature engineering had a detrimental effect on the logistic regression model but did not significantly impact the performance of decision trees and random forests.

Table 2. Performances of different models

Model	Accuracy train	Recall train	Precision _train	Accuracy _test	Recal _test	Precision test	CrossVal Mean	CrossVal1	CrossVal2	CrossVal3	CrossVal4	CrossVal5
Baseline classifier	0.73	0.1	0.73	0.73	0.1	0.73	0.73	0.73	0.73	0.73	0.73	0.73
logistic (imbalanced)	0.89	0.88	0.86	0.89	0.88	0.86	0.88	0.88	0.88	0.88	0.88	0.88
logistic (balanced)	0.88	0.89	0.87	0.88	0.89	0.87	0.87	0.87	0.88	0.88	0.87	0.87
decision tree (imbalanced)	0.87	0.79	0.84	0.87	0.79	0.84	0.88	0.88	0.87	0.88	0.87	0.88
decision tree (balanced)	0.89	0.88	0.9	0.89	0.88	0.9	0.89	0.89	0.89	0.89	0.89	0.89
Random forest (imbalanced)	0.89	0.88	0.86	0.89	0.88	0.86	0.88	0.88	0.88	0.89	0.88	0.88
Random forest (balanced)	0.89	0.89	0.9	0.89	0.89	0.9	0.89	0.89	0.88	0.89	0.89	0.89
logistic after FE	0.68	0.68	0.68	0.68	0.68	0.68	0.67	0.65	0.67	0.69	0.67	0.69
Decision tree after FE	0.88	0.87	0.89	0.88	0.87	0.89	0.88	0.87	0.87	0.89	0.88	0.88
Random Forest after FE	0.89	0.88	0.89	0.89	0.88	0.89	0.89	0.88	0.88	0.89	0.89	0.89

Figure 4. Performances of different models

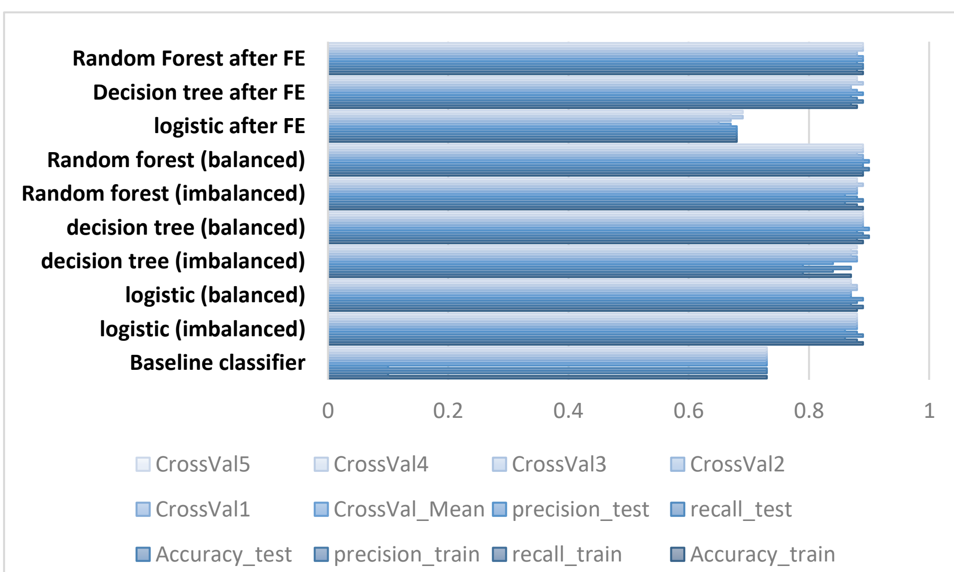

Table 3. Feature importance in predicting CTR

Feature	Feature Importance (0-1)
Banner position	0.12
Device type	0.18
Personalization	0.25
Intrusiveness	0.05
Mobile optimization	0.15
Loading time	0.08
Brand awareness	0.03
Scroll length	0.07
Ad fatigue	0.04
Site category	0.03

Figure 6. Feature importance in predicting CTR

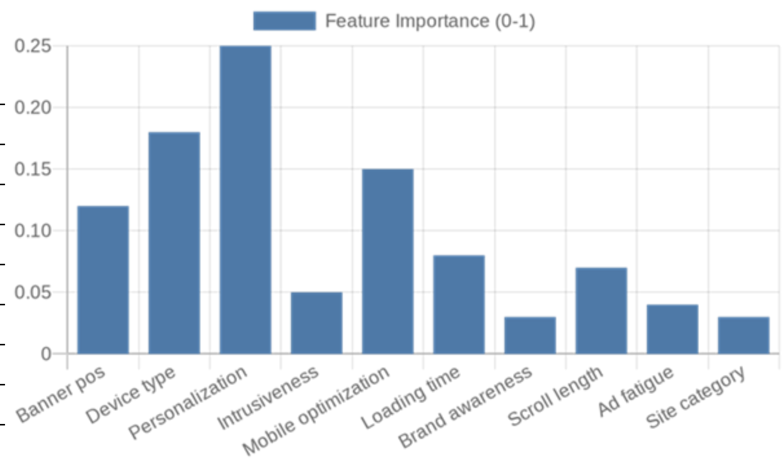

According to our findings, Personalization appears as the feature with the highest importance score of 0.25 in predicting click-through rate (CTR) suggest the critical role that individualized content plays in digital advertising. Personalization involves leveraging data analytics and user behavior metrics to tailor advertisements to the specific interests, needs, or past behaviors of individual users. This could range from showing sports equipment ads to a user who has recently searched for athletic gear, to displaying travel deals to someone who has been researching vacation destinations. By doing so, advertisers increase the relevance of the ad to the user, thereby enhancing the likelihood of engagement and clicks. The high feature importance score for personalization suggests that users are more likely to interact with ads that resonate with their personal interests or immediate needs, making it a pivotal factor in the effectiveness of online advertising campaigns.

In the context of digital world, where users are shown with a multitude of ads each day, the significance of personalization becomes even more pronounced. Users have developed a tendency to ignore or block out generic advertisements, often referred to as banner blindness. Personalized ads, however, have a higher chance of capturing the user's attention because they offer something of specific interest. This is particularly important for industries where the competition for user attention is fierce, and the cost per click is high. By focusing on personalization, advertisers can not only improve CTR but also potentially increase return on investment (ROI) by targeting users who are more likely to convert after clicking the ad.

The effectiveness of personalization is contingent upon the quality and accuracy of the data used for tailoring the ads. Inaccurate or outdated data can lead to poorly targeted ads, which not only fail to engage the user but can also have a detrimental effect on the brand's image.
The feature Device type, with a feature importance score of 0.18, stands as the second most influential factor in predicting click-through rates (CTR) for digital advertisements. This indicates that the hardware used to access and view the advertisement has a considerable impact on user engagement. Different devices, such as desktop computers, tablets, and smartphones, offer varying user experiences due to screen sizes, resolutions, and input methods. For example, an ad that is visually appealing and easily clickable on a desktop may not render as effectively on a smaller smartphone screen. The user interface and the ease of interaction with the ad can significantly differ based on the device type, thereby affecting the likelihood of a user clicking on the ad. This suggests that advertisers need to consider device-specific optimizations when designing and deploying their digital advertising campaigns.
Users often begin a task on one device and complete it on another, a behavior known as multi-device pathing. For advertisers, this means that understanding the role of device type is not just about optimizing the ad for different screens, but also about understanding the user's journey across multiple devices. For instance, a user might initially see an ad on their smartphone but switch to a desktop to complete a purchase. If the ad is not optimized for both types of devices, there's a risk of losing potential conversions. Therefore, a nuanced understanding of how device type influences user behavior can provide advertisers with valuable insights into how to structure their campaigns for maximum effectiveness.
Advertisers face complexities of device fragmentation, especially in the Android ecosystem where there are numerous devices with varying screen sizes and capabilities. Additionally, device-specific optimization may require more resources in terms of design and development, increasing the overall cost of the advertising campaign. There are also data tracking complexities involved in accurately attributing clicks and conversions to specific devices, particularly when users engage in multi-device pathing.
The features Mobile optimization and Banner pos have feature importance scores of 0.15 and 0.12 respectively, indicating their notable but somewhat lesser influence on click-through rates (CTR) compared to Personalization and Device type. Mobile optimization refers to the design and formatting of the advertisement to ensure that it is easily viewable and interactive on mobile devices. Given the increasing prevalence of mobile internet usage, an ad that is not optimized for mobile platforms risks alienating a large segment of potential viewers. Poorly optimized ads can result in issues like slow loading times, difficult navigation, or unresponsive elements, all of which can deter users from clicking. Therefore, while mobile optimization may not be the most critical factor, its importance in influencing CTR should not be underestimated.
The position of the banner ad, denoted as Banner pos, also plays a role in influencing CTR, albeit to a lesser extent. The placement of an ad on a webpage can significantly affect its visibility and, consequently, the likelihood of user engagement. For example, banner ads placed at the top of a webpage are generally more visible and may receive more clicks compared to those placed at the bottom. However, the effectiveness of banner position can also be influenced by other factors such as the design of the webpage, the nature of the content surrounding the ad, and user behavior patterns. For instance, users who are actively engaged in reading an article may be more likely to notice and click on an ad that is embedded within the content, as opposed to one that is placed in the sidebar or footer.

While both Mobile optimization and Banner pos are not as influential as Personalization or Device type, they still hold considerable weight in the overall scheme of factors affecting CTR. Advertisers should pay attention to these aspects when designing and deploying digital advertising campaigns. Ignoring these factors could result in missed opportunities for engagement and conversion, even if other more influential factors like personalization are well-executed.

Features like Intrusiveness, Loading time, Scroll length, Ad fatigue, Brand awareness, and Site category have lower feature importance scores, ranging from 0.03 to 0.08, indicating their comparatively reduced impact on click-through rates (CTR). Intrusiveness, with a score of 0.05, refers to the degree to which an advertisement disrupts the user experience. While one might assume that less intrusive ads would significantly improve CTR, the model suggests that this feature is less critical compared to others like Personalization or Device type. This could be because users have become somewhat accustomed to a certain level of advertising intrusiveness as a trade-off for accessing free content, or it may indicate that other factors more strongly influence the decision to click.

Loading time and Scroll length are also features with lower importance scores, at 0.08 and 0.07 respectively. These factors relate to the user experience but are not the primary drivers of CTR according to the model. Loading time pertains to how quickly the ad content is displayed, and while it may not be a leading factor, slow loading times can still deter clicks and should not be ignored. Scroll length refers to the amount of scrolling required to view the ad in its entirety.

The least influential features in this model are Brand awareness and Site category, both with scores of 0.03. Brand awareness refers to the extent to which the user is familiar with the brand being advertised. While brand recognition can influence consumer behavior in broader contexts, its low score suggests that it is less relevant for the specific action of clicking on an ad. Site category denotes the type of website where the ad is displayed, such as a news site, a social media platform, or an e-commerce site. Its low importance score suggests that the context in which the ad appears is less critical in influencing CTR than other factors.

Conclusion

The migration from traditional to digital advertising is largely due to the capabilities of online platforms to offer targeted outreach, immediate data collection, and worldwide access. In this new setting, the quality of the user experience (UX) has become increasingly vital for the success of advertising campaigns. In our research, we employed three types of supervised machine learning algorithms: Logistic Regression, Decision Trees, and Random Forests. These algorithms were selected to address the complexities of Click-Through Rate (CTR) datasets. To further refine the data, we utilized a capping mechanism. Specifically, any value in a given column that is equal to or exceeds the 98th percentile is replaced by the 98th percentile value. This data preprocessing step serves to limit the impact of extreme values on the overall distribution of the data. Our research was centered on several key metrics including personalization, intrusiveness, brand awareness, and ad fatigue. One of the significant challenges we encountered was the issue of imbalanced classification in CTR datasets. To transforms the dataset into a balanced classification problem, we applied SMOTE.

The finding of this study shows several key insights into the effectiveness of online advertising metrics. One of the most significant relationships is between Click-Through

Rate (CTR) and Personalization. The data indicates a moderately strong positive correlation, meaning that more personalized ads are likely to have higher click-through rates. In contrast, ads perceived as intrusive are likely to have lower click-through rates. This anti-correlation could serve as a cautionary note for advertisers to avoid overly intrusive advertising methods [18], [19]. Another notable correlation is the positive relationship between Device Type and Mobile Optimization. This suggests that the device used to view an ad can affect its performance, particularly in terms of how well it is optimized for mobile viewing. Among the models tested, the logistic regression model and random forest models performed better than the baseline classifier, particularly when trained on balanced data. These models indicated a strong and consistent performance across metrics like training and testing accuracy, recall, precision, and cross-validation scores. However, the logistic regression model's performance dropped significantly after feature engineering, unlike the decision trees and random forest models, which remained robust.

The features examined in the model revealed several key contributors to advertising effectiveness. The feature with the highest importance score was Personalization, underlining its critical role in influencing user engagement and click-through rates. In an online environment overwhelmed by advertisements, personalized ads have a greater likelihood of capturing user attention and converting that into clicks. Therefore, advertisers need to utilize data analytics to tailor ads according to user interests and needs [20], [21]. However, the accuracy and timeliness of this data are vital; otherwise, poorly targeted ads may harm the brand's image.

Device Type emerged as the second most influential factor. The choice of device—be it desktop, tablet, or mobile—greatly impacts the user experience and, by extension, the effectiveness of the ad. This points to the need for advertisers to optimize ads for multiple devices, given the common behavior of users switching between them. This device-specific optimization can be resource-intensive but is critical for maximizing user engagement and conversion rates [22]. Other features like Mobile Optimization and Banner Position were found to have a notable but somewhat lesser impact on click-through rates.

References

[1] D. Cui and D. Curry, “Prediction in Marketing Using the Support Vector Machine,” *Marketing Science*, vol. 24, no. 4, pp. 595–615, Nov. 2005.

[2] B. Rathore, “Beyond Trends: Shaping the Future of Fashion Marketing with AI, Sustainability and Machine Learning,” *EIPRMJ*, vol. 6, no. 2, pp. 16–24, Jul. 2017.

[3] K. S. Dave and V. Varma, “Learning the click-through rate for rare/new ads from similar ads,” in *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, Geneva, Switzerland, 2010, pp. 897–898.

[4] P. S. Sullivan *et al.*, “Bias in online recruitment and retention of racial and ethnic minority men who have sex with men,” *J. Med. Internet Res.*, vol. 13, no. 2, p. e38, May 2011.

[5] P. P. K. Chan, X. Hu, L. Zhao, D. S. Yeung, D. Liu, and L. Xiao, “Convolutional Neural Networks based Click-Through Rate Prediction with Multiple Feature Sequences,” *IJCAI* , 2018.

[6] H. Gao, D. Kong, M. Lu, X. Bai, and J. Yang, “Attention Convolutional Neural Network for Advertiser-level Click-through Rate Forecasting,” in *Proceedings of the 2018 World Wide Web Conference*, Lyon, France, 2018, pp. 1855–1864.

[7] X. Wang, W. Li, Y. Cui, R. Zhang, and J. Mao, “Click-Through Rate Estimation for Rare Events in Online Advertising,” in *Online Multimedia Advertising: Techniques and Technologies*, IGI Global, 2011, pp. 1–12.

[8] M. Regelson and D. Fain, “Predicting click-through rate using keyword clusters,” *Proceedings of the Second Workshop on Sponsored*, 2006.

[9] Y.-L. Lin and Y.-W. Chen, “Effects of ad types, positions, animation lengths, and exposure times on the click-through rate of animated online advertisings,” *Comput. Ind. Eng.*, vol. 57, no. 2, pp. 580–591, Sep. 2009.

[10] B. Edizel, A. Mantrach, and X. Bai, “Deep Character-Level Click-Through Rate Prediction for Sponsored Search,” in *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Shinjuku, Tokyo, Japan, 2017, pp. 305–314.

[11] G. Zhou *et al.*, “Deep Interest Network for Click-Through Rate Prediction,” in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, London, United Kingdom, 2018, pp. 1059–1068.

[12] M. Richardson, E. Dominowska, and R. Ragno, “Predicting clicks: estimating the click-through rate for new ads,” in *Proceedings of the 16th international conference on World Wide Web*, Banff, Alberta, Canada, 2007, pp. 521–530.

[13] J. Pan *et al.*, “Field-weighted Factorization Machines for Click-Through Rate Prediction in Display Advertising,” in *Proceedings of the 2018 World Wide Web Conference*, Lyon, France, 2018, pp. 1349–1357.

[14] D. Agarwal, B.-C. Chen, and P. Elango, “Spatio-temporal models for estimating click-through rate,” in *Proceedings of the 18th international conference on World wide web*, Madrid, Spain, 2009, pp. 21–30.

[15] I. Trofimov, A. Kornetova, and V. Topinskiy, “Using boosted trees for click-through rate prediction for sponsored search,” in *Proceedings of the Sixth International Workshop on Data Mining for Online Advertising and Internet Economy*, Beijing, China, 2012, pp. 1–6.

[16] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich, “Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine,” 2010.

[17] A. Kumar and J. Salo, “Effects of link placements in email newsletters on their click-through rate,” *Journal of Marketing Communications*, vol. 24, no. 5, pp. 535–548, Jul. 2018.

[18] K. Dave and V. Varma, “Predicting the click-through rate for rare/new ads,” *Center for Search and Information Extraction Lab International Institute of Information Technology Hyderabad, INDIA*, 2010.

[19] L. Shi and B. Li, “Predict the click-through rate and average cost per click for keywords using machine learning methodologies,” *Proceedings of the International Conference on*, 2016.

[20] A. C. König, M. Gamon, and Q. Wu, “Click-through prediction for news queries,” in *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, Boston, MA, USA, 2009, pp. 347–354.

[21] J.-H. Chen, Z.-Q. Zhao, J.-Y. Shi, and C. Zhao, “A New Approach for Mobile Advertising Click-Through Rate Estimation Based on Deep Belief Nets,” *Comput. Intell. Neurosci.*, vol. 2017, p. 7259762, Oct. 2017.
[22] S. Yang, S. Lin, and J. R. Carlson, “Brand engagement on social media: will firms’ social media efforts influence search engine advertising effectiveness?,” *J. Mark.*, 2016.