JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

Predictive Ad Targeting Powered by Machine Learning Models in the Cloud

[1]Praveen Gujar, [2]Sriram Panyam

[1] Director of Product @ LinkedIn, [2] CTO @ DagKnows
[1] San Francisco, USA

Abstract : The revolution in digital advertising, spearheaded by predictive ad targeting, marks a significant leap in how businesses engage with current and future buyers. Central to this transformation is the deployment of machine learning (ML) models on cloud platforms, which has significantly enhanced the precision and efficiency of ad targeting. This comprehensive exploration delves into the sequential processes involved in employing ML models for predictive ad targeting, covering data collection, preprocessing, model training, evaluation, deployment, and the overarching necessity for ethical considerations and bias mitigation at each stage. The paper aims to foreground the critical importance of fostering fairness, transparency, and respect for user privacy in leveraging advanced predictive technologies within the advertising domain.

IndexTerms - Predictive Ad Targeting, Machine Learning, Cloud Computing, Bias Mitigation, Ethical AI, Personalization, Data Privacy

I. Introduction

The digital advertising landscape has been profoundly reshaped by the advent of predictive targeting, which utilizes data analysis and ML algorithms to forecast consumer behaviors, intent, and preferences[1]. Predictive targeting enables brands to discover new consumers that are more likely to convert, even if they hadn't considered targeting them before. This enables a more effective advertising approach, significantly increasing engagement, growth, and conversion rates. Predictive audiences are typically computed from brands' assets (eg: product pages), existing consumers and their behaviors, conversion signals, potential consumers web footprint, and more. The scalability, flexibility, and efficiency provided by cloud computing have been instrumental in this advancement, allowing for the processing and analysis of vast datasets and the execution of complex ML models at a competitive cost structure[2]. However, this technological evolution brings to the fore significant ethical concerns, especially regarding data privacy, consent, and the risk of perpetuating biases, which necessitates a careful and considered approach.

II. Organization of the Survey Paper

- Provide an overview of key aspects such as Predictive Ad Targeting is made effective with Machine Learning Models, Cloud Computing, Bias Mitigation, Ethical AI considerations, and Data Privacy.
- Discuss the theoretical framework
- Highlight the key value propositions of predictive ad targeting in digital advertising
- Model selection, training, and pitfalls to avoid for effective implementation of predictive ads targeting
- Ability of Cloud to enable scale the AI/ML models in a cost effective manner
- Future directions of research in this "new world"

3 Theoretical Framework

3.1 Comprehensive Data Collection and Ethical Preprocessing

The foundational element that underpins the effectiveness of any predictive ad targeting system is the robustness and comprehensiveness of the data it leverages[3]. This data, a rich tapestry of user interactions, preferences, and behaviors, is the lifeblood of predictive modeling in advertising. It encompasses a wide array of data starting from first-party signals (product engagement, social media engagement and more), browsing behaviors that reveal interest and intent, conversion signals (eg: server side signals), and lead or sales opportunity signals (eg: CRM). The collection of this multifaceted data is the first critical step in building a predictive ad targeting system that can accurately forecast user intent and enhance ad relevance[4].

3.1.1 Data Collection Strategies

Collecting comprehensive user data involves deploying sophisticated tracking technologies and methodologies. Web tracking tools, such as cookies, pixels and SDKs, are commonly used to monitor browsing behaviors across different websites, capturing data on the pages visited, the duration of visits, and the interactions on those pages[5]. Transaction records are typically gathered by customer CRMs or 3P e-commerce platforms and online retail databases, detailing the items purchased, the transaction amounts, and the frequency of purchases. This is then combined with first-party data of the advertising platform to enrich user behavior and intent. This multi-dimensional data collection strategy is designed to create a holistic view of the user, enabling highly targeted and personalized advertising campaigns.

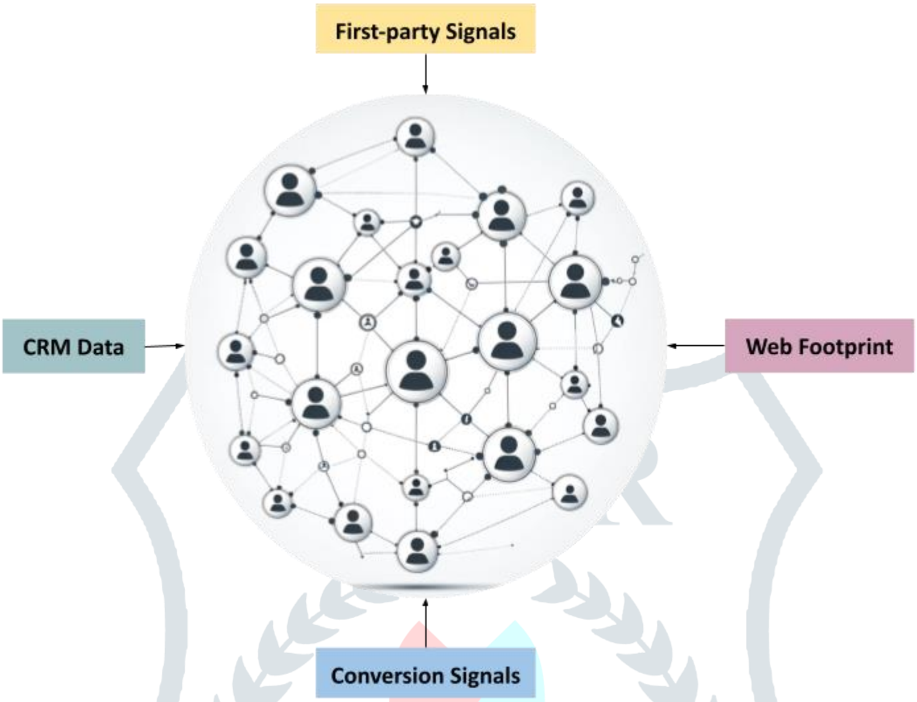

Fig 1: Signals that help build buyer journey

3.1.2 Preprocessing for Data Quality and Privacy

Once collected, the data undergoes a series of preprocessing steps to ensure it is of high quality and ready for analysis. Data cleaning is the first crucial step, involving the removal of inaccuracies, inconsistencies, and duplicate entries to prevent skewed results. Normalization follows, standardizing the range of data values so that no single feature disproportionately influences the model outcomes. Anonymization is another key preprocessing step, particularly vital for maintaining user privacy and data security. It involves stripping away personally identifiable information (PII) from the dataset, ensuring that the data cannot be traced back to an individual user. Techniques such as data masking, pseudonymization, and encryption are employed to safeguard user privacy while still allowing for the meaningful analysis of patterns and trends[6].

3.1.3 Ethical Considerations and Compliance

The ethical dimensions of data collection and preprocessing are paramount, given the sensitive nature of personal data. Obtaining

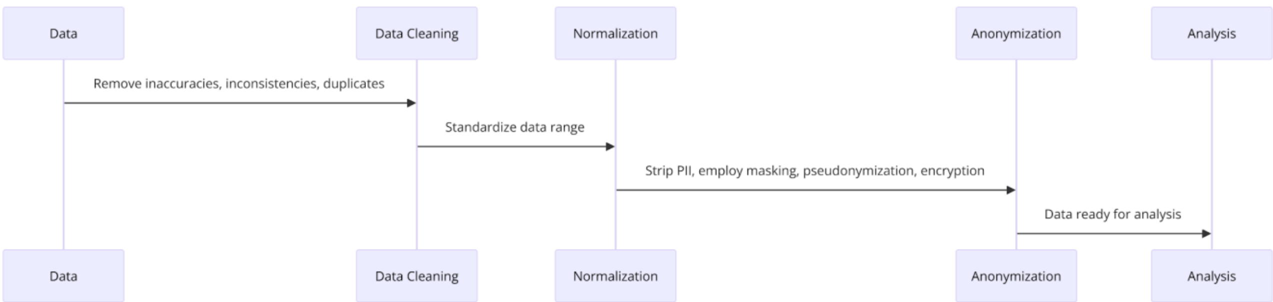

explicit user consent before data collection is not just an ethical imperative but also a legal requirement under data protection laws like the General Data Protection Regulation (GDPR) in the European Union and the California Consumer Privacy Act (CCPA) in the United States[7]. These regulations mandate clear communication with users about the data being collected, the purpose of its collection, and obtaining their explicit consent to do so. Furthermore, they provide users with rights over their data, including the right to access, rectify, delete, or port their data.

A principled approach to data collection and preprocessing is advocated, one that prioritizes building trust and transparency with users. This involves not only complying with legal standards but also going beyond them to establish ethical guidelines that respect user privacy and ensure data security. Transparent data practices, such as providing users with clear privacy policies, options to opt-out of data collection, and controls over their data, are essential for fostering trust. Additionally, regular audits and assessments of data practices help in identifying potential ethical risks and ensuring ongoing compliance with both legal and ethical standards[8].

3.2 Advanced Model Training and Rigorous Evaluation

After meticulous preparation and preprocessing of the dataset, the focus transitions to a critical phase in the development of predictive ad targeting systems: the selection, training, and comprehensive evaluation of suitable machine learning (ML) models. This phase is pivotal, as the chosen models directly influence the system's capacity to accurately predict user behavior and, consequently, the effectiveness of ad targeting. The selection process encompasses a broad spectrum of ML algorithms, ranging from conventional models like logistic regression, known for its simplicity and efficacy in binary classification tasks, to more complex and nuanced methods such as neural networks and deep learning, which offer unparalleled depth and sophistication for capturing intricate patterns in advertising data[9].

3.2.1 Model Selection and Training

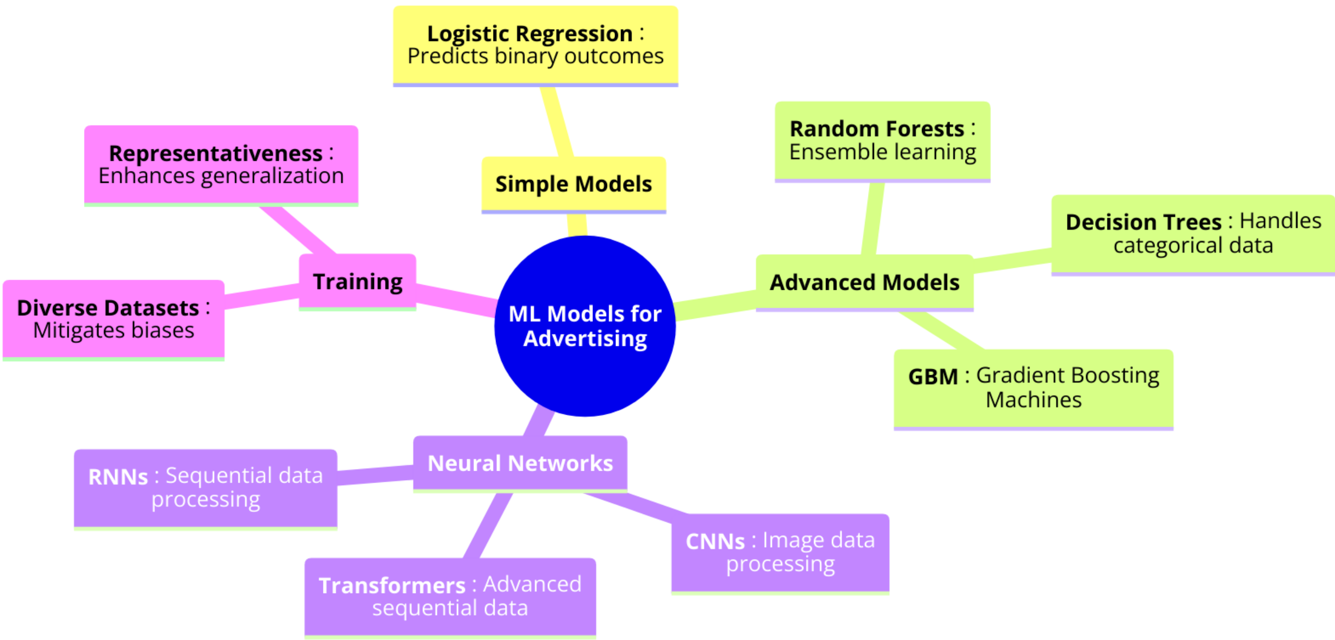

The selection of ML models is guided by the nature and intricacies of the advertising data at hand. Logistic regression, for instance, serves as a powerful tool for predicting binary outcomes, such as whether a user will click on an ad. However, its simplicity might not capture the complex interactions and non-linear relationships present in the data. This is where advanced models like decision trees, random forests, and gradient boosting machines (GBM) come into play, offering more depth in analysis through ensemble learning and the ability to handle a mix of numerical and categorical data effectively[9].

In the realm of neural networks and deep learning, specialized architectures are tailored to specific types of data encountered in advertising. Convolutional Neural Networks (CNNs) are adept at processing image data, making them ideal for analyzing visual content in ads, while Recurrent Neural Networks (RNNs) and transformers excel in sequential data processing, such as textual content in user comments or browsing history. These sophisticated models can unearth deep insights from the data, enabling highly personalized and dynamically targeted advertising strategies.

The training of these models is a meticulous process that hinges on the diversity and representativeness of the datasets. Employing a broad spectrum of data helps in mitigating biases and enhancing the models' ability to generalize across different user demographics and behaviors. This diversity is crucial for ensuring that the models do not inadvertently perpetuate or amplify existing biases, leading to unfair or discriminatory ad targeting practices.

3.2.2 Model Evaluation and Bias Mitigation

Evaluating the performance of ML models goes beyond merely assessing their predictive accuracy. Advanced metrics and methodologies are employed to scrutinize the models' outputs for fairness and ethical implications. Metrics such as precision, recall, and the area under the Receiver Operating Characteristic (ROC) curve provide insights into the models' accuracy and sensitivity.

However, evaluating fairness requires additional layers of analysis, looking into whether the model's predictions are equitable across different groups of users, particularly those defined by sensitive attributes like age, gender, or ethnicity[10].

Techniques for identifying and correcting biases are integral to this phase. Algorithmic fairness approaches, such as fairness constraints and adversarial debiasing, are employed to adjust models in a manner that minimizes discriminatory patterns in their predictions. Furthermore, ethical model auditing—a comprehensive review of the models' development process, data handling, and outcome analyses—ensures adherence to ethical AI principles. This auditing process involves stakeholders from diverse backgrounds, including ethicists, domain experts, and representatives from affected communities, to evaluate the models from multiple perspectives and ensure their fairness and transparency.

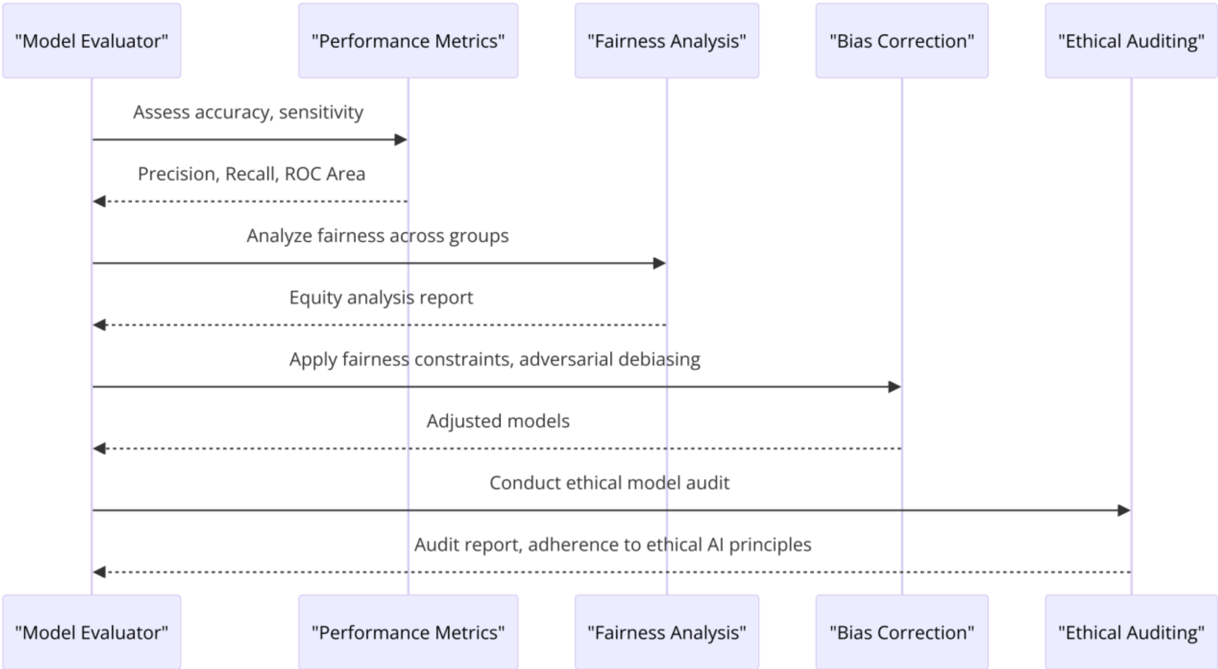

Stressing the role of ethical AI in advertising underscores the industry's responsibility to adopt non-discriminatory practices and uphold high ethical standards. By implementing these rigorous selection, training, and evaluation processes, predictive ad targeting systems can achieve not only high accuracy and efficiency but also fairness and ethical integrity. This balanced approach ensures that the benefits of advanced ML models in advertising are realized without compromising on equity and respect for all users.

3.3 Scalable Deployment and Ethical Real-time Prediction

The deployment of machine learning (ML) models for real-time predictive ad targeting not only heralds a new era of advertising precision but also introduces a complex landscape of technical and ethical challenges. This evolution necessitates a deep dive into the cloud-based infrastructures and cutting-edge technologies that underpin the scalable and efficient deployment of these models, alongside a rigorous examination of the ethical imperatives that govern their operation[9].

3.3.1 Cloud-Based Infrastructures for ML Deployment

Cloud computing platforms offer the computational power and scalability required to process vast datasets and run complex ML models in real time. These platforms provide a suite of services and technologies tailored for ML workloads, including managed data storage solutions, powerful computing resources, and specialized tools for ML model development and deployment. Within this ecosystem, container technologies such as Docker have emerged as pivotal tools[11]. They encapsulate ML models and their dependencies into portable containers, ensuring consistency across different computing environments and facilitating seamless deployment and scaling.

Orchestration tools like Kubernetes further enhance this landscape by managing these containers across a cluster of servers.

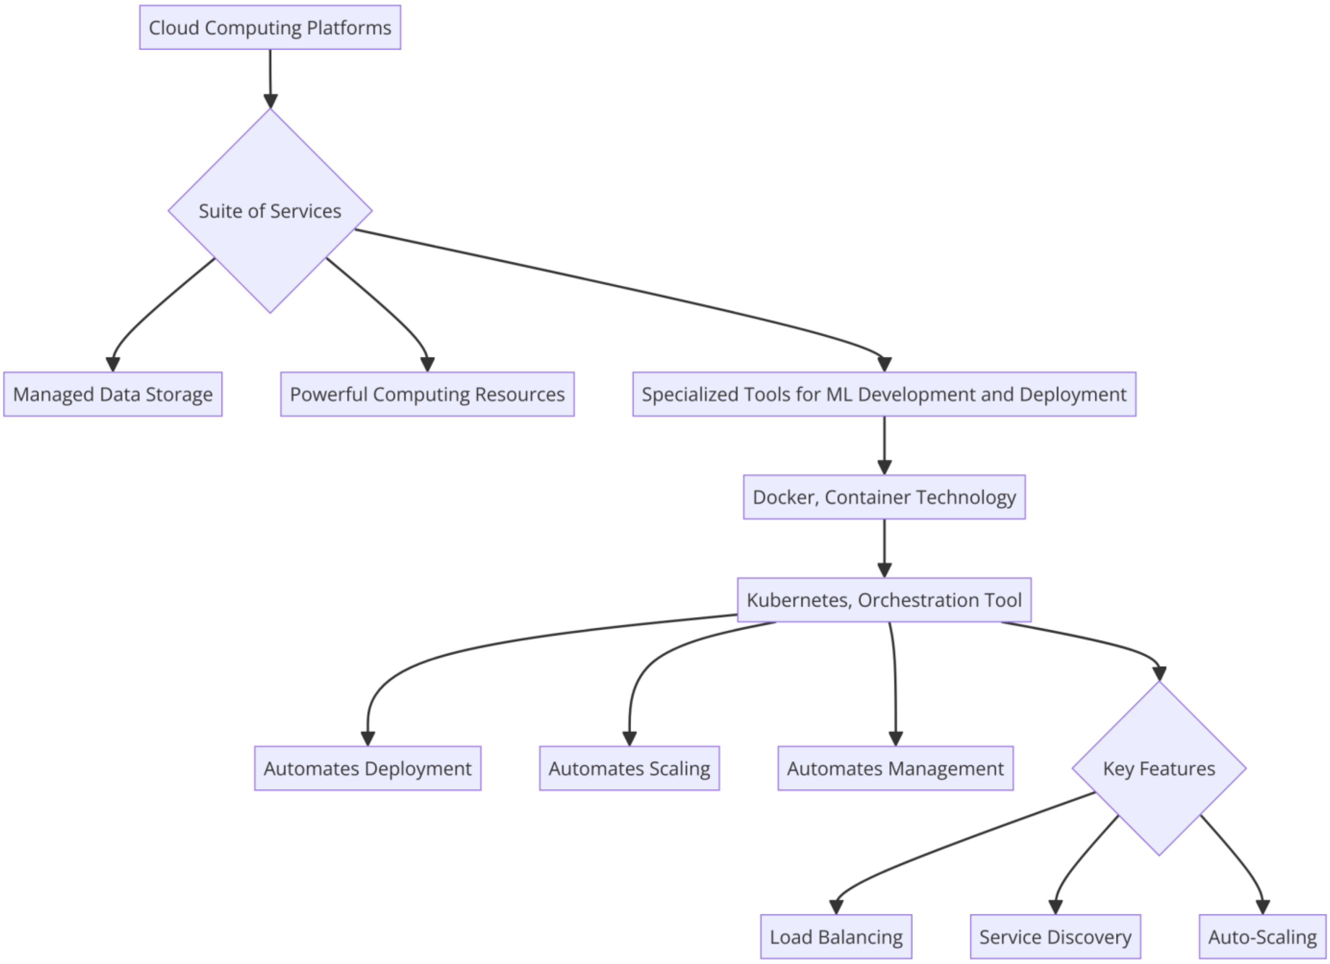

Kubernetes automates the deployment, scaling, and management of containerized applications, addressing key challenges such as load balancing, service discovery, and auto-scaling. This orchestration capability is crucial for real-time predictive ad targeting, where the demand can fluctuate dramatically, and the system must scale efficiently to maintain performance without incurring unnecessary costs[12].

3.3.2 Ethical Considerations in Real-time ML Deployment

Deploying ML models for real-time predictions in ad targeting introduces unique ethical considerations. The dynamic nature of real-time targeting means that models continuously make decisions that directly impact users, necessitating a framework that ensures these decisions are made responsibly. This involves continuous monitoring of model performance to quickly identify and rectify any issues that could lead to inaccurate predictions, which might result in irrelevant or even harmful ads being displayed to users[13].

Moreover, the potential for models to perpetuate or amplify biases presents a significant ethical concern. Continuous model updating becomes essential, not just for maintaining the accuracy and relevance of predictions but also for ensuring that models evolve in response to changing data patterns and societal norms, thereby mitigating biases.

Implementing fail-safes and oversight mechanisms is critical to address potential ethical breaches. Fail-safes can include thresholds and alerts that trigger human review for certain types of predictions or decisions, ensuring that automated systems do not operate unchecked. An ethical oversight mechanism, such as an ethics review board or committee, plays a crucial role in maintaining continuous ethical compliance. This body can oversee the deployment and operation of ML models, conduct regular audits, and review practices and decisions to ensure they align with ethical standards and societal values.

3.3.4 The Responsibility of Advertisers and Technologists

The deployment of ML models for real-time predictive ad targeting bestows a significant responsibility on advertisers and technologists to uphold high ethical standards. This responsibility encompasses not only the technical execution of model deployment but also the ethical implications of how these models interact with users and impact society. Advertisers and technologists must work collaboratively to establish ethical guidelines, implement robust oversight mechanisms, and foster a culture of transparency and accountability[14]. This collaborative effort ensures that the advancements in real-time predictive ad targeting are leveraged ethically, promoting trust and fairness in the dynamic landscape of digital advertising.

In sum, the deployment of ML models for real-time predictive ad targeting in cloud environments represents a convergence of technical innovation and ethical stewardship. By embracing cloud-based infrastructures and technologies like Docker and Kubernetes, advertisers can achieve scalable and efficient model deployment. However, the success of these technological endeavors is inherently tied to the rigorous ethical considerations and frameworks that guide their operation, ensuring that as we advance in our technical capabilities, we also advance in our ethical obligations to users and society at large.

3.4 Ethical Framework and Bias Mitigation Strategies

This comprehensive segment delves into the ethical framework and strategies necessary to navigate the complex ethical landscape of ML-driven predictive ad targeting. It articulates the pressing ethical issues, including the risk of privacy intrusion, the amplification of societal biases through automated processes, and the ethical dilemmas posed by potentially manipulative ad targeting[13]. Proposing a robust ethical framework for AI in advertising, the discussion incorporates principles of accountability, fairness, privacy, and transparency, aiming to establish a balanced approach to predictive advertising. Practical strategies for embedding these ethical principles into the ML lifecycle are explored, from the initial design phase through to deployment and ongoing operation, highlighting the imperative for the advertising industry to adopt a responsible and principled approach to leveraging AI technologies.

4 Methodology of Literature Review

The literature review was undertaken using a step-by-step method to ensure as many important and existing studies around predictive ad-targeting, machine learning models for ad targeting (and others)were represented without bias.

4.1 Search Strategy

Websites like Google Scholar, IEEE, Xplore and JSTOR were searched for terms like “digital advertising”, “brand-consumer interaction”, (and others)?. Only studies published in the last five years were considered as up-to-date and relevant due to rapid changes in the technical and ad-tech landscape.

4.2 Inclusion and Exclusion Criteria

Studies around predictive ads targeting, consumer signals, model selection, model training, minimizing bias, ethical and compliance considerations, privacy of other topics we cared about around ad-tech and marketing data were taken up to understand their problem spaces and the learnings they offered. Studies not in English, or those not peer-reviewed or focussing on topics outside those listed above were excluded from this review.

4.3 Data Extraction and Synthesis

From the studies included in our review, key aspects were extracted, e.g. goals, methods, findings, suggestions and possible areas of future research. This was extremely helpful in synthesizing the key topic areas as well as in identifying gaps in current studies. This guided the journey in this review.

5 Recommendations for Future Research

Looking ahead, the future directions for ML models in predictive ad targeting, including the integration of emerging technologies such as blockchain for enhanced data security and privacy, and the exploration of federated learning approaches to mitigate data privacy concerns. It concludes by reiterating the immense potential of ML models for revolutionizing predictive ad targeting in the cloud, while also underscoring the paramount importance of adhering to ethical standards and actively mitigating biases. By committing to these principles, the field of digital advertising can navigate the challenges posed by advanced technologies, ensuring a future where advertising is not only more effective but also equitable and respectful of user privacy and rights.

6 Conclusion

This journal delves into the evolving landscape of predictive ad targeting, highlighting the significant role machine learning (ML) and cloud computing play in revolutionizing digital advertising. It outlines the process of leveraging extensive user data, from collection and ethical preprocessing to deploying advanced ML models, to enable highly personalized advertising strategies. The discussion emphasizes the critical importance of ethical considerations and bias mitigation throughout this process, advocating for transparency, fairness, and user privacy. The use of cloud-based infrastructures, like Docker and Kubernetes, is explored for their efficiency in deploying scalable and real-time predictive models, while also acknowledging the technical and ethical challenges involved. The narrative stresses the collective responsibility of advertisers, technologists, and regulatory bodies to uphold ethical standards, ensuring that technological advancements in advertising are balanced with the need for privacy and equity. Looking ahead, the journal suggests that the future of digital advertising will continue to be shaped by emerging technologies, underscoring the ongoing need for ethical vigilance and innovation to benefit all stakeholders in a data-driven ecosystem.

III. Conflict of Interest

IV.

The author(s) declare(s) that there is no conflict of interest regarding the publication of this paper.

References

[1] Koneti Chaitanya, Gonesh Chandra Saha, Hasi Saha, Samik Acharya, Manjul Singla - The Impact of Artificial Intelligence and Machine Learning in Digital Marketing Strategies
[2] Adnan Qayyum, Aneeqa Ijaz, Muhammad Usama, Waleed Iqbal, Junaid Qadir, Yehia Elkhatib, Ala Al-Fuqaha - Securing Machine Learning in the Cloud: A Systematic Review of Cloud Machine Learning Security
[3] Michael W. Obal, Wen Lv - Improving banner ad strategies through predictive modeling
[4] Fruergaard, Bjarne Ørum - Statistical learning for predictive targeting in online advertising
[5] Jonathan R. Mayer; John C. Mitchel - Third-Party Web Tracking: Policy and Technology
[6] J. Andrew Onesimu, J. Karthikeyan & Yuichi Sei - An efficient clustering-based anonymization scheme for privacy-preserving data collection in IoT based healthcare services
[7] Lisa V. Ziukovic - The Alignment Between the Electronic Communications Privacy Act and the European Union's General Data Protection Regulation: Reform Needs to Protect the Data Subject
[8] Michael L. Rustad & Thomas H. Koenig - TOWARDS A GLOBAL DATA PRIVACY STANDARD
[9] Jin-A Choi a, Kiho Lim - Identifying machine learning techniques for classification of target advertising
[10] Sule Birim, Ipek Kazancoglu, Sachin Kumar Mangla, Aysun Kahraman & Yigit Kazancoglu - The derived demand for advertising expenses and implications on sustainability: a comparative study using deep learning and traditional machine learning methods
[11] Mikhail M. Rovnyagin; Alexander S. Hrapov; Anna V. Guminskaia; Aleksandr P. Orlov - ML-based Heterogeneous Container Orchestration Architecture
[12] Emiliano Casalicchio, Stefano Iannucci - The state-of-the-art in container technologies: Application, orchestration and security
[13] I. Glenn Cohen, Ruben Amarasingham, Anand Shah, Bin Xie, and Bernard Lo - The Legal And Ethical Concerns That Arise From Using Complex Predictive Analytics In Health Care
[14] Louis Frank - Edge AI: Deploying models directly on edge devices for real-time processing