Privacy-Aware AI Advertising Systems: A Federated Learning Framework for Cross-Platform Personalization

Ethan Caldwell
School of Information, University of Michigan
ethancal@umich.edu

Sofia Bennett
Department of Computer Science, Rice University
smbennett@rice.edu

Abstract

Digital advertising ecosystems increasingly rely on large-scale artificial intelligence infrastructures that personalize marketing messages, optimize bidding strategies, and allocate attention across millions of users and advertisers. Traditional advertising architectures depend heavily on centralized data aggregation, where behavioral logs from multiple platforms are combined to train large predictive models. While this approach enables highly accurate personalization, it also raises significant concerns related to privacy protection, regulatory compliance, data governance, and systemic concentration of informational power. As privacy regulations expand globally and user expectations regarding data protection intensify, the advertising industry faces increasing pressure to develop new system architectures capable of preserving personalization capabilities while minimizing direct data collection and centralized storage.

This paper proposes a privacy-aware advertising framework based on federated learning for cross-platform personalization. Rather than treating federated learning solely as a distributed optimization technique, the framework conceptualizes it as a socio-technical infrastructure that redistributes data custody, computational responsibilities, and governance accountability across multiple actors in the advertising ecosystem. The study examines how decentralized model training can enable collaborative personalization across advertisers, publishers, and devices without requiring raw behavioral data to leave local environments. Particular attention is given to system-level design challenges including heterogeneous data distributions, delayed feedback signals, adversarial manipulation risks, fairness constraints, and cross-jurisdictional regulatory compliance.

The paper develops a multi-layer architectural model integrating local representation learning, secure aggregation protocols, differential privacy mechanisms, and policy-aware governance structures. It further explores the implications of federated advertising systems for market competition, algorithmic fairness, and institutional accountability. The analysis demonstrates that federated learning can significantly reduce centralized data risks while maintaining

effective personalization performance when combined with robust coordination protocols and transparent governance frameworks. The paper concludes that privacy-aware federated infrastructures represent a promising direction for the future evolution of digital advertising ecosystems.

Keywords：

Federated learning; digital advertising; privacy-aware AI; cross-platform personalization; recommender systems; algorithmic governance; differential privacy; large-scale machine learning

1. Introduction

Digital advertising has become one of the most sophisticated applications of artificial intelligence in modern information economies. Advertising platforms process massive streams of behavioral data to estimate user preferences, predict response probabilities, optimize campaign budgets, and allocate promotional messages in real time. These processes rely on complex machine learning systems capable of analyzing billions of interactions daily. Over the past decade, advances in deep learning, large-scale recommendation systems, and real-time bidding algorithms have dramatically improved the accuracy and efficiency of personalized advertising.

Despite these technological achievements, the dominant architecture of advertising personalization remains heavily dependent on centralized data aggregation. In traditional systems, behavioral signals from search engines, social media platforms, mobile applications, e-commerce websites, and advertising exchanges are consolidated into unified data warehouses. Machine learning models trained on these centralized datasets generate predictions about user interests, click probabilities, and purchase intent. While such architectures offer significant analytical advantages, they also produce serious privacy risks because large volumes of sensitive behavioral data become concentrated within a small number of corporate infrastructures.

Recent developments in privacy regulation and public policy have intensified scrutiny of centralized data practices. Governments around the world have introduced regulations that impose strict requirements on data collection, processing, and sharing. Simultaneously, technology platforms have implemented new restrictions on third-party tracking and cross-site identifiers. These changes have disrupted traditional advertising measurement practices and forced the industry to reconsider how personalization systems should operate in privacy-sensitive environments.

Federated learning has emerged as a promising alternative paradigm for large-scale machine learning under data protection constraints. Instead of transferring raw data to a centralized training environment, federated learning enables model training to occur across decentralized data sources. Each participant trains a local model using its own data and shares only model

updates with a central coordinator. These updates are aggregated to produce an improved global model without requiring the underlying datasets to be centrally stored.

Although federated learning has attracted considerable attention in fields such as mobile computing and healthcare, its implications for advertising systems remain underexplored. Advertising personalization presents unique challenges including non-independent data distributions, delayed reward signals, adversarial incentives, and complex market interactions. Successfully applying federated learning to advertising therefore requires careful consideration of both technical and institutional design.

This paper investigates how federated learning can be integrated into digital advertising infrastructures to enable privacy-aware cross-platform personalization. The analysis focuses on system architecture, governance mechanisms, optimization challenges, and market implications. By examining these dimensions together, the study aims to provide a comprehensive framework for understanding the role of federated learning in the future evolution of advertising technologies.

2. Evolution of Digital Advertising Infrastructure

The architecture of digital advertising systems has evolved significantly over the past two decades. Early online advertising relied primarily on contextual matching between webpage content and advertising messages. These systems required minimal user data because advertisements were selected based on page content rather than individual behavioral history. As internet usage expanded and tracking technologies improved, advertising platforms began incorporating behavioral targeting techniques that relied on user browsing histories and demographic profiles.

The emergence of real-time bidding platforms further transformed advertising infrastructure. Real-time bidding allows advertisers to compete for individual ad impressions through automated auctions that occur within milliseconds. Machine learning models predict the likelihood that a particular user will respond to a given advertisement, enabling advertisers to adjust bids dynamically based on predicted value. This process requires continuous analysis of behavioral data streams to maintain accurate predictions.

Large technology platforms gradually built sophisticated data pipelines capable of collecting signals from multiple sources including search queries, social interactions, location data, and purchase histories. By integrating these signals into centralized data warehouses, platforms could construct detailed user profiles and train highly accurate predictive models. Deep learning architectures such as wide-and-deep models and neural collaborative filtering further improved prediction performance by capturing complex relationships among sparse behavioral features.

However, the benefits of centralized data aggregation have been accompanied by growing concerns regarding surveillance, user autonomy, and market concentration. Critics argue that

large advertising platforms possess disproportionate influence over digital markets because they control both data infrastructure and algorithmic decision systems. These concerns have motivated policymakers and technologists to explore alternative architectures that distribute data control more broadly across the ecosystem.

Federated learning represents one such architectural alternative. By enabling collaborative model training without centralized data storage, federated learning allows multiple organizations to contribute to shared predictive models while retaining control over their own datasets. This approach aligns with emerging privacy norms that emphasize data minimization and local data governance.

3. Fundamentals of Federated Learning

Federated learning is a distributed machine learning paradigm designed to train models using decentralized datasets located across multiple devices or organizations. The core idea is to move computation closer to the data rather than transferring data to a centralized server. In a typical federated training process, a global model is initialized and distributed to participating clients. Each client performs local training using its own dataset and produces a set of model updates. These updates are then transmitted to a central aggregator, which combines them to produce an improved global model.

One of the most important characteristics of federated learning is its ability to operate under heterogeneous data distributions. In many real-world scenarios, datasets stored on different devices or organizational servers exhibit significant statistical differences. For example, user behavior on an e-commerce platform may differ substantially from behavior on a news website or social media application. Federated learning algorithms must therefore accommodate non-independent and non-identically distributed data.

Communication efficiency also plays a critical role in federated learning systems. Because model updates must be transmitted across networks during each training round, efficient update compression and selective participation strategies are often necessary. Techniques such as gradient sparsification, adaptive learning rates, and partial client participation help reduce communication overhead while maintaining convergence performance.

Privacy preservation represents another central motivation for federated learning. Since raw data remain local to the environment where they were generated, the risk of large-scale data breaches or unauthorized data sharing is significantly reduced. Additional privacy mechanisms such as secure aggregation and differential privacy can further protect participants by ensuring that individual updates cannot be reverse engineered to reveal sensitive information.

While federated learning offers substantial advantages, it also introduces new challenges related to system coordination, model robustness, and governance. These challenges become particularly complex in advertising environments where economic incentives and competitive

dynamics shape system behavior.

4. Cross-Platform Personalization Challenges

Cross-platform personalization aims to provide consistent advertising experiences across multiple digital environments such as mobile applications, web browsers, streaming platforms, and e-commerce websites. Achieving this goal traditionally requires the construction of unified identity graphs that link user activities across different contexts. Identity graphs allow advertising systems to attribute conversions, control ad frequency, and optimize campaign performance using integrated behavioral histories.

However, the creation of centralized identity graphs raises significant privacy concerns because it enables extensive tracking of user behavior across digital environments. Many users are unaware of the extent to which their online activities are linked together for advertising purposes. In response to these concerns, technology companies and regulators have begun limiting the use of persistent cross-site identifiers.

Federated learning offers an alternative approach to cross-platform personalization that avoids direct identity consolidation. Instead of building centralized identity graphs, federated systems can coordinate predictive models across platforms using shared representations of user behavior. Each platform maintains its own local data and user identifiers, but contributes model updates that capture statistical patterns relevant to advertising performance.

For example, a mobile application may train a local model that captures patterns of product interest based on in-app browsing behavior. An e-commerce platform may train a separate model based on purchase histories. Through federated aggregation, these models can share generalized representations of consumer preferences without revealing specific user identities or raw behavioral data.

This approach reduces privacy risks while preserving the benefits of collaborative learning. However, implementing cross-platform federation requires careful design to ensure that shared representations do not inadvertently enable re-identification or sensitive inference.

5. System Architecture for Federated Advertising

A privacy-aware federated advertising system can be conceptualized as a multi-layer architecture integrating decentralized data environments, local model training modules, secure coordination mechanisms, and governance infrastructure. Each layer plays a distinct role in maintaining system performance while protecting user privacy.

At the lowest layer, data remain within the local environments where they were originally collected. These environments may include mobile devices, advertiser databases, publisher servers, or platform-specific analytics systems. Local preprocessing pipelines transform raw interaction logs into feature representations suitable for machine learning training.

The next layer consists of local model training components that operate within each environment. These models learn patterns specific to their local datasets, such as user engagement behaviors or purchase tendencies. Training occurs periodically based on locally available data and computational resources.

A coordination layer orchestrates communication among participating nodes. Secure aggregation protocols ensure that model updates are combined in ways that prevent individual contributions from being exposed. Differential privacy mechanisms add controlled noise to updates to further protect participant confidentiality.

Above the coordination layer, policy-aware governance mechanisms manage participation rules, privacy budgets, fairness monitoring, and system auditing. Governance frameworks ensure that federated training processes adhere to regulatory requirements and ethical standards.

Finally, deployment infrastructure distributes updated global models to participating platforms, where they are used to generate real-time advertising predictions. Continuous monitoring systems evaluate model performance and trigger retraining cycles when necessary.

6. Privacy Engineering in Federated Advertising

Protecting user privacy requires more than simply decentralizing data storage. Even in federated systems, model updates can potentially reveal sensitive information through gradient leakage or inference attacks. Privacy engineering therefore requires the integration of multiple protective mechanisms.

Secure aggregation protocols ensure that the central coordinator receives only aggregated model updates rather than individual contributions. This prevents the coordinator from reconstructing data associated with specific participants. Differential privacy mechanisms further limit the influence of individual data points by adding noise to updates during training.

Another important privacy consideration involves purpose limitation. Advertising models should be trained only on data relevant to clearly defined objectives, and data should not be reused for unrelated purposes without user consent. Federated architectures facilitate this principle because data remain within local environments where consent preferences can be directly enforced.

Privacy-preserving measurement techniques also play an important role. Advertising effectiveness metrics must often be computed across multiple platforms without exposing detailed user-level data. Aggregated conversion reporting and secure multiparty computation methods allow performance evaluation while maintaining confidentiality.

7. Optimization Challenges in Federated Advertising Systems

Training advertising models using federated learning presents several unique optimization challenges. Behavioral data are highly unevenly distributed across platforms and user populations. Some nodes may generate large volumes of interactions, while others contribute relatively sparse datasets.

Delayed reward signals further complicate training. Advertising conversions often occur hours or days after initial exposure, making it difficult to align training labels with observed impressions. Federated systems must therefore incorporate mechanisms for asynchronous label integration and delayed gradient updates.

Concept drift represents another significant challenge. Consumer interests, seasonal trends, and marketing campaigns change continuously, requiring models to adapt quickly. Federated systems must balance stability and adaptability by coordinating update schedules across diverse environments.

Communication constraints also influence optimization performance. Frequent communication between nodes may be impractical due to network limitations or computational costs. Efficient aggregation strategies and update compression techniques help reduce communication overhead while preserving model accuracy.

8. Robustness and Security Considerations

Digital advertising ecosystems are vulnerable to manipulation because engagement metrics directly influence financial outcomes. Malicious actors may attempt to manipulate federated learning systems through data poisoning, fake traffic generation, or adversarial model updates.

Robust aggregation algorithms help mitigate these risks by identifying anomalous updates and limiting their influence on global models. Reputation-based participation weighting and anomaly detection techniques can further enhance system resilience.

Security governance also requires procedures for incident response and forensic analysis. Federated systems must provide mechanisms for investigating suspicious activity without violating privacy guarantees. Transparent auditing processes and encrypted logging systems support responsible oversight.

9. Fairness and Market Implications

Advertising algorithms influence the distribution of economic opportunities among advertisers, publishers, and users. Federated learning may reduce concentration of data power by allowing multiple entities to contribute to shared models without surrendering full control of their data assets.

However, fairness challenges remain. Large platforms with extensive datasets may still exert disproportionate influence over model updates. Ensuring equitable participation requires careful design of aggregation rules and fairness-aware optimization metrics.

Fairness considerations also extend to user-level outcomes. Personalization systems must avoid reinforcing harmful stereotypes or excluding vulnerable groups from economic opportunities. Monitoring exposure distributions and implementing fairness constraints in ranking models can help address these concerns.

10. Governance and Policy Implications

The successful deployment of federated advertising systems depends on robust governance frameworks. Technical safeguards must be complemented by institutional policies that define accountability, transparency, and compliance procedures.

Lifecycle accountability ensures that model updates are traceable and auditable. Transparency reports describing personalization practices can help build trust among users and regulators. Consent management systems allow users to control how their data contribute to advertising models.

International regulatory diversity also requires flexible system design. Federated architectures can support jurisdiction-specific policies by allowing local nodes to enforce regional privacy requirements while still participating in global model training.

11. Future Research Directions

Several important research challenges remain in the development of privacy-aware federated advertising systems. Future work should focus on improving optimization algorithms for highly heterogeneous data environments. Advances in representation learning may enable more effective cross-platform knowledge transfer without requiring centralized identity graphs.

Privacy-preserving measurement techniques also require further development to support reliable advertising attribution in decentralized environments. In addition, scalable communication protocols will be essential for coordinating millions of distributed training nodes.

Interdisciplinary collaboration among engineers, economists, policymakers, and ethicists will play a critical role in shaping the future of federated advertising systems.

12. Conclusion

The digital advertising industry stands at a critical transition point as privacy concerns reshape technological and regulatory landscapes. Traditional centralized data architectures are

increasingly difficult to sustain in environments where users demand stronger data protection and governments impose stricter regulatory oversight.

Federated learning offers a promising pathway toward privacy-aware advertising infrastructures capable of maintaining personalization capabilities without relying on centralized behavioral surveillance. By decentralizing data storage and enabling collaborative model training, federated systems redistribute computational responsibilities and governance accountability across the advertising ecosystem.

However, federated learning is not a complete solution on its own. Successful deployment requires integration with privacy engineering techniques, fairness-aware optimization strategies, robust security safeguards, and transparent governance frameworks. Only through the combination of technical innovation and responsible institutional design can privacy-aware advertising systems achieve both effectiveness and legitimacy.

As digital markets continue to evolve, federated architectures may play an increasingly important role in balancing personalization benefits with societal expectations for privacy, fairness, and accountability.

References

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. Proceedings of the ACM SIGSAC Conference on Computer and Communications Security.

2. Bottou, L., Curtis, F. E., & Nocedal, J. (2018). Optimization methods for large-scale machine learning. SIAM Review.

3. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science.

4. Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. Advances in Neural Information Processing Systems.

5. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T. (2017). Neural collaborative filtering. Proceedings of the World Wide Web Conference.

6. Kairouz, P., McMahan, H. B., et al. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning.

7. Li, T., Sahu, A., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine.

8. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017).

Communication-efficient learning of deep networks from decentralized data. Proceedings of AISTATS.

9. Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. IEEE Symposium on Security and Privacy.

10. Richardson, M., Dominowska, E., & Ragno, R. (2007). Predicting clicks: Estimating the click-through rate for new ads. Proceedings of the World Wide Web Conference.

11. Rendle, S. (2012). Factorization machines with libFM. ACM Transactions on Intelligent Systems and Technology.

12. Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. Proceedings of the ACM Conference on Computer and Communications Security.

13. Sun, F., Liu, J., Wu, J., et al. (2019). BERT4Rec: Sequential recommendation with bidirectional encoder representations. Proceedings of CIKM.

14. Varian, H. R. (2007). Position auctions. International Journal of Industrial Organization.

15. Zhou, G., Zhu, X., Song, C., et al. (2018). Deep interest network for click-through rate prediction. Proceedings of KDD.

16. Yi, X. (2025, October). Real-Time Fair-Exposure Ad Allocation for SMBs and Underserved Creators via Contextual Bandits-with-Knapsacks. In Proceedings of the 2025 2nd International Conference on Digital Economy and Computer Science (pp. 1602-1607).

17. Tang, Y., Kojima, K., Gotoda, M., Nishikawa, S., Hayashi, S., Koike-Akino, T., ... & Klamkin, J. (2020, February). InP grating coupler design for vertical coupling of InP and silicon chips. In Integrated Optics: Devices, Materials, and Technologies XXIV (Vol. 11283, pp. 33-38). SPIE.
18. Li, B. (2025). GIS-Integrated Semi-Supervised U-Net for Automated Spatiotemporal Detection and Visualization of Land Encroachment in Protected Areas Using Remote Sensing Imagery.
19. Chang, C., Fu, M., Chen, X., Feng, S., Zhang, M., Zhou, X., ... & Liu, Z. (2025, November). Research on PDU-Net Lung Nodule Segmentation Algorithm Based on Path Aggregation and Dual Attention. In 2025 4th International Conference on Image Processing, Computer Vision and Machine Learning (ICICML) (pp. 1897-1900). IEEE.

20. Tang, Y., Kojima, K., Gotoda, M., Nishikawa, S., Hayashi, S., Koike-Akino, T., ... & Klamkin, J. (2020). Design and Optimization of Shallow-Angle Grating Coupler for Vertical Emission from Indium Phosphide Devices.

21. HOU, R., JEONG, S., WANG, Y., LAW, K. H., & LYNCH, J. P. (2017). Camera-based triggering of bridge structural health monitoring systems using a cyber-physical system framework. Structural Health Monitoring 2017, (shm).

22. Qi, R. (2025). AUBIQ: A Generative AI-Powered Framework for Automating Business Intelligence Requirements in Resource-Constrained Enterprises. Frontiers in Business and Finance, 2(01), 66-86.

23. Qi, R. (2025, June). Enterprise financial distress prediction based on machine learning and SHAP interpretability analysis. In Proceedings of the 2025 International Conference on Artificial Intelligence and Digital Finance (pp. 76-79).

24. Qi, R. (2025, July). DecisionFlow for SMEs: A Lightweight Visual Framework for Multi-Task Joint Prediction and Anomaly Detection. In Proceedings of the 2025 International Conference on Economic Management and Big Data Application (pp. 899-903).

25. Yang, D. (2022). An Investigation on English Translations of Culture-Loaded Words in The Analects of Confucius from the Eco Perspective: A Case Study of the English Translation of Lectures on China's Traditional Political Thoughts. Editorial Board, 7.

26. Dan, Y. A. N. G. AN ANALYSIS OF THE IN-DEPTH TRANSLATION STRATEGY OF THE ENGLISH EDITION OF LECTURES ON CHINA'S TRADITIONAL POLITICAL THOUGHTS.

27. YANG, D., & WANG, Z. A Study on Evaluation of the Integration of Chinese and Foreign Cultures into Oxford Junior High School English Textbooks on the Basis of Multicultural Education. Editorial Board, 33.

28. Tian, Y., Xu, S., Cao, Y., Wang, Z., & Wei, Z. (2025). An Empirical Comparison of Machine Learning and Deep Learning Models for Automated Fake News Detection. Mathematics, 13(13), 2086.

29. Li, B. (2025). GIS-Integrated Semi-Supervised U-Net for Automated Spatiotemporal Detection and Visualization of Land Encroachment in Protected Areas Using Remote Sensing Imagery.

30. Zhang, T. (2025). A Neuro-Symbolic and Blockchain-Enhanced Multi-Agent Framework for Fair and Consistent Cross-Regulatory Audit Intelligence.