

Privacy-Preserving Ad Targeting in the Age of Data Protection Legislation: Leveraging Federated Learning and Clean Room Solutions

Subhash Vinnakota

The Walt Disney Company, USA

Abstract: The digital advertising ecosystem faces unprecedented transformation driven by global privacy regulations that restrict traditional data collection and targeting practices. As regulations like GDPR and CCPA redefine permissible data handling, advertisers and publishers must adopt privacy-preserving technologies to maintain targeting effectiveness while ensuring compliance. This article examines two promising solutions: data clean rooms and federated learning. Data clean rooms provide secure environments where multiple parties can collaborate on analytics without exposing raw data, utilizing differential privacy, secure multi-party computation, and homomorphic encryption to enable valuable insights while protecting user information. Federated learning offers a distributed computational framework where models train across decentralized devices without centralizing sensitive data, keeping information at its source while enabling personalization. The comparative assessment reveals distinct advantages for each approach: clean rooms excel in cross-dataset analysis while federated learning provides superior privacy for individual-level data with on-device processing. Future directions indicate potential integration with blockchain technologies, advancement of standardization efforts, and development of formal optimization frameworks that balance privacy protection with targeting effectiveness. As the advertising landscape evolves, these privacy-preserving technologies will reshape industry practices, enabling personalized advertising that respects user privacy and regulatory requirements.

Keywords: Privacy-preserving advertising, data clean rooms federated learning, differential privacy, blockchain integration.

INTRODUCTION AND REGULATORY LANDSCAPE

The digital advertising ecosystem has undergone a profound transformation in recent years, driven largely by the proliferation of data privacy regulations across global jurisdictions. The European Union's General Data Protection Regulation (GDPR), implemented in May 2018, established a comprehensive framework for personal data protection, introducing strict requirements for data collection, processing, and transfer. Similarly, the California Consumer Privacy Act (CCPA), which took effect in January 2020, empowered California residents with unprecedented control over their personal information, including the right to know what data businesses collect and the right to request deletion of personal information. These landmark regulations have created a ripple effect, spurring similar legislation worldwide including Brazil's Lei Geral de Proteção de Dados (LGPD), China's Personal Information Protection Law (PIPL), and Virginia's Consumer Data Protection Act (CDPA), forming a complex global patchwork of privacy requirements that demands attention from multinational advertisers (Rao, S. 2024). The implementation of these regulations reflects a growing consensus among legislators that consumers deserve greater transparency and control over their personal data, particularly as the volume and granularity of data collection have expanded exponentially with the growth of digital platforms.

The regulatory surge has fundamentally disrupted traditional ad targeting models that previously relied on unrestricted access to user data. Third-party cookies, once the backbone of cross-site tracking and behavioral advertising, are being phased out by major browsers in response to privacy concerns and regulatory pressure. Mobile device identifiers have become increasingly restrictive, dramatically reducing their availability for targeting purposes. Furthermore, the legal basis for processing personal data for advertising purposes now requires explicit consent under many regulations, significantly limiting the scope and scale of data collection activities. These changes have created substantial operational challenges for advertisers, who must navigate an increasingly fragmented identity landscape while ensuring compliance with divergent regulatory requirements that can vary significantly by jurisdiction. As regulatory authorities gain experience in enforcement, they have demonstrated increasing willingness to impose substantial penalties for non-compliance, creating significant financial risk for organizations that fail to adapt their advertising practices to the new privacy landscape (Fitzgerald, A. 2024). This regulatory environment has effectively rendered obsolete many of the data practices that previously underpinned digital advertising's targeting precision.

Table 1: Comparison of Global Privacy Regulations and Their Impact on Digital Advertising. (Rao, S. 2024; Fitzgerald, A. 2024)

Regulatory Aspect	Traditional Ad Targeting Approach	Privacy-Preserving Requirements	Key Challenges
User Consent	Implicit or broad consent	Explicit, specific, informed consent	Implementation of consent mechanisms without disrupting user experience
Data Processing	Centralized collection and storage	Data minimization and purpose limitation	Maintaining targeting effectiveness with limited data access
Cross-Border Transfers	Minimal restrictions	Adequacy decisions or safeguards required	Navigating varying requirements across jurisdictions
User Rights	Limited transparency and control	Access, deletion, and portability rights	Technical infrastructure for responding to user requests

The imperative for privacy-preserving technologies in digital advertising has never been more evident. As traditional identifiers disappear and privacy regulations tighten, advertisers and publishers face mounting pressure to develop alternative approaches that balance effective targeting with robust privacy protections. The industry stands at a critical inflection point, where continued reliance on conventional data practices risks substantial legal penalties—with violations potentially resulting in significant fines proportional to global annual revenue according to multiple regulatory frameworks (Fitzgerald, A. 2024). Beyond regulatory compliance, consumer trust has emerged as a crucial competitive factor, with privacy-conscious audiences increasingly favoring brands that demonstrate responsible data stewardship. The growing consumer awareness about data privacy issues, catalyzed by high-profile data breaches and media coverage of regulatory actions, has created market incentives that align with regulatory requirements, further accelerating the adoption of privacy-enhancing technologies in the advertising ecosystem. This convergence of regulatory, technical, and market forces has necessitated innovation in privacy-enhancing technologies that promise to maintain advertising effectiveness while protecting user privacy.

This leads to our central research question: How can advertisers maintain targeting effectiveness while adhering to privacy regulations? The challenge is multifaceted, involving technical, legal, and operational dimensions that must be harmonized to create viable solutions. As third-party identifiers become increasingly restricted, advertisers must pivot toward approaches that leverage first-party data in privacy-compliant ways (Rao, S. 2024). The complexity of this challenge is compounded by the global nature of digital advertising, where campaigns often cross jurisdictional boundaries with differing regulatory requirements. Technologies such as clean room

solutions and federated learning have emerged as promising alternatives that enable sophisticated data analysis and model training without exposing raw user data. These approaches represent a fundamental paradigm shift from centralized data collection toward distributed computation models that keep sensitive information closer to its source, thereby reducing privacy risks while still enabling effective targeting. The remainder of this article examines these technologies in detail, evaluating their potential to reconcile the seemingly competing objectives of effective ad targeting and robust privacy protection in the post-cookie era, while acknowledging the ongoing evolution of the regulatory landscape that continues to shape industry practices and technological innovation.

CLEAN ROOM SOLUTIONS: ARCHITECTURE AND IMPLEMENTATION

Data clean rooms represent a transformative approach to collaborative data analytics in digital advertising, fundamentally reimagining how organizations can derive insights while maintaining stringent privacy protections. At their conceptual core, clean rooms establish secure computational environments where multiple parties can contribute proprietary data without exposing the underlying raw information to other participants. This architectural paradigm addresses the inherent tension between data utility and privacy preservation by creating a controlled ecosystem that facilitates specific analytical operations while implementing robust safeguards against unauthorized data access or extraction. The conceptual framework of data clean rooms is predicated on the principle of data minimization—limiting exposure to only what is necessary for the intended analytical purpose—and establishes clear boundaries between data ownership, access, and computational privileges. Within this framework, participants typically maintain control over their contributed data, retaining the authority to approve

specific analytical operations and to revoke access when necessary. Recent research has highlighted how clean rooms implement governance mechanisms through both technical controls and contractual agreements, establishing a multi-layered approach to privacy protection that addresses both regulatory compliance and ethical data use considerations. The technical implementation of these governance structures typically includes cryptographic protocols that mathematically enforce data access limitations, creating verifiable technical guarantees that complement traditional legal safeguards (Jindal, P. 2024). Studies examining the architectural evolution of clean rooms have identified a progression from early implementations that relied primarily on contractual protections to contemporary solutions that incorporate cryptographic guarantees, differential privacy, and distributed computation models to establish more robust privacy preservation. This evolution reflects the growing recognition that effective privacy protection requires technical mechanisms that cannot be circumvented through administrative means, particularly in contexts involving sensitive advertising data that could reveal consumer behaviors or preferences.

The technical foundations of data clean rooms encompass multiple layers of privacy-enhancing technologies, beginning with sophisticated encryption protocols and anonymization techniques that transform raw data before it enters the shared environment. These include deterministic and probabilistic encryption methods that protect personally identifiable information while preserving the analytical utility of the

underlying data patterns. Advanced anonymization approaches are frequently deployed to ensure that individual records cannot be re-identified through inference attacks or correlation with external datasets. Building upon this protective foundation, secure data matching methodologies enable the crucial function of audience overlap analysis without exposing raw identifiers. Recent research has demonstrated how these methodologies have evolved to address the fundamental challenge of entity resolution across organizational boundaries without compromising privacy, incorporating techniques such as secure multi-party computation (MPC) for threshold-based matching that prevents any single party from accessing the complete matching criteria or results (Ahessin, A. & Ali, A. 2025). The technical implementation often incorporates differential privacy models that introduce calibrated statistical noise into query results, providing mathematical guarantees against re-identification while maintaining acceptable levels of analytical accuracy. Contemporary research has explored the application of homomorphic encryption in advertising contexts, revealing both the profound privacy benefits and the computational challenges associated with performing operations on encrypted data without decryption. These studies have identified strategies for optimizing homomorphic operations for specific advertising use cases, including audience segmentation, conversion measurement, and frequency capping, demonstrating how this advanced cryptographic approach can be tailored to the particular requirements of digital advertising while maintaining stringent privacy protections.

Table 2: Privacy-Enhancing Technologies in Data Clean Rooms. (Jindal, P. 2024; Ahessin, A. & Ali, A. 2025)

Technology	Primary Function	Privacy Guarantees	Implementation Considerations
Differential Privacy	Statistical noise addition	Formal mathematical privacy guarantees	Privacy-utility trade-off calibration
Secure Multi-Party Computation	Joint computation without data sharing	Cryptographic protection of inputs	Computational overhead and protocol complexity
Homomorphic Encryption	Computation on encrypted data	End-to-end encryption	Performance limitations for complex operations
Privacy-Preserving Record Linkage	Secure identity matching	Protection of identifier information	Accuracy versus privacy trade-offs

Case analyses across industry verticals demonstrate the versatility and effectiveness of clean room implementations in addressing diverse advertising challenges while maintaining privacy compliance. In the retail sector, clean room deployments have enabled merchants to collaborate with consumer packaged goods

manufacturers, matching transaction data with product information to optimize merchandising strategies and measure campaign effectiveness without exposing customer identifiers or purchase histories. The implementation architecture in these cases typically includes rigorous access controls, granular permission structures, and output

restrictions that prevent extraction of individual-level data while still delivering actionable business intelligence. Media and entertainment companies have similarly leveraged clean room technology to align content consumption patterns with advertiser data, creating more relevant targeting opportunities while protecting viewer privacy through sophisticated data transformation and query limitation mechanisms that prevent reconstruction of viewing histories. Research examining these implementations has identified critical success factors including clearly defined use cases, purpose limitation enforcement, and comprehensive audit mechanisms that create accountability without compromising privacy protections (Jindal, P. 2024). Studies analyzing financial services implementations have documented how sector-specific regulatory requirements necessitate additional technical safeguards beyond general-purpose clean room architectures, including enhanced encryption standards, more restrictive query interfaces, and augmented consent management capabilities. These implementations demonstrate how the baseline clean room architecture can be extended to address specialized compliance requirements while maintaining core functionality. Scholarly examination of healthcare-adjacent implementations has revealed particularly instructive lessons regarding the balance between utility and privacy, documenting how these deployments incorporate additional technical measures including advanced anonymization techniques, stricter access controls, and more conservative privacy budget management to address the heightened sensitivity of health-related information while still enabling essential advertising and marketing functions.

FEDERATED LEARNING FOR PRIVACY-PRESERVING AD PERSONALIZATION

The theoretical foundations of federated learning represent a paradigm shift in machine learning methodology, particularly pertinent to privacy-sensitive domains such as digital advertising. Unlike traditional centralized approaches that require consolidation of data for model training, federated learning establishes a distributed computational framework where models are trained across multiple decentralized devices or servers holding local data samples, without exchanging the raw data itself. This theoretical architecture originated from the recognition that data locality and privacy are increasingly critical

considerations in contemporary machine learning applications, especially those involving personal information related to consumer behavior and preferences. The foundational principle of federated learning can be formalized as a distributed optimization problem, wherein the objective is to minimize a global loss function across decentralized data without centralizing the underlying information. Recent research has expanded this theoretical foundation to incorporate parallel training methodologies that optimize both privacy protection and computational efficiency, enabling more effective deployment in resource-constrained environments such as mobile devices where much of digital advertising engagement occurs. These parallel training approaches distribute the computational workload across participating devices while maintaining strict data locality, effectively addressing both privacy concerns and performance limitations that might otherwise impede adoption in advertising contexts. The theoretical framework has been further enhanced through the development of specialized loss functions and optimization algorithms specifically designed to handle the unique characteristics of advertising-related data, including its temporal sensitivity, sparse interaction patterns, and heterogeneous distribution across user populations (Cao, T. D. *et al.*, 2025). These advancements enable federated learning systems to effectively model consumer preferences and predict advertising responsiveness without requiring the centralization of sensitive behavioral data, establishing a theoretical foundation that aligns technical capabilities with emerging privacy regulations while preserving the targeting effectiveness essential to advertising business models.

The technical architecture of federated learning for ad personalization encompasses multiple specialized components designed to enable privacy-preserving model training and deployment. At the foundation of this architecture are on-device model training methodologies that execute the learning process directly on user devices rather than in centralized servers. These methodologies typically employ lightweight versions of standard machine learning algorithms, optimized for execution in resource-constrained environments such as mobile devices or browsers. The training process often utilizes techniques such as mini-batch gradient descent with locally available data, allowing the device to compute model updates without transmitting the underlying user interaction data. Building upon this

foundation, secure aggregation protocols enable the privacy-preserving combination of model updates from multiple devices without exposing individual contributions. These protocols have evolved significantly in recent implementations, incorporating advanced cryptographic techniques that provide formal privacy guarantees while minimizing computational overhead, making them increasingly practical for deployment in production advertising systems. The federated averaging algorithm represents a core component of this architecture, providing a mechanism to combine locally trained models into a global model that benefits from the collective learning across all participating devices. Contemporary research has enhanced this approach through adaptive weighting mechanisms that account for varying data quality and relevance across different devices, improving the resulting model's performance for advertising applications where data heterogeneity is particularly pronounced (Cao, T. D. *et al.*, 2025). Recent architectural advances have also focused on update compression and communication efficiency mechanisms that address the bandwidth constraints inherent in distributed learning across consumer devices. These innovations include novel quantization techniques specifically optimized for advertising models, which tend to have distinct characteristics compared to other application domains. The technical architecture increasingly incorporates specialized components for detecting and mitigating potential privacy leakage, including mechanisms that analyze model updates for unintended information disclosure before transmission, ensuring robust privacy protection beyond what conventional federated learning approaches might provide.

Despite its promising theoretical foundations and sophisticated technical architecture, federated learning for ad personalization faces several significant implementation challenges that must be addressed to achieve widespread adoption in production advertising systems. Statistical heterogeneity across user populations represents a fundamental challenge, as the non-identical distribution of data across devices can lead to convergence difficulties and reduced model performance. In advertising contexts, this heterogeneity manifests as variation in user behaviors, preferences, and interaction patterns across different demographic groups, geographic regions, and device types. Recent research has identified specialized adaptation techniques that can mitigate these effects in advertising

applications, including personalized regularization approaches that adjust model training based on local data characteristics without compromising privacy protections. Systems heterogeneity and computational limitations present further implementation hurdles, as federated learning must operate across a diverse ecosystem of devices with varying computational capabilities, memory constraints, network connectivity, and energy limitations. Contemporary studies have documented significant variations in training performance across device types, with implications for both model quality and user experience in advertising applications, leading to the development of adaptive participation mechanisms that optimize system-wide performance while accommodating device limitations (Narula, M. *et al.*, 2024). Security against adversarial attacks represents perhaps the most critical implementation challenge, as federated learning systems are vulnerable to various malicious behaviors including data poisoning, model inversion, and inference attacks. Recent analyses have demonstrated that advertising models may be particularly susceptible to certain attack vectors due to the inherent characteristics of behavioral data and the commercial sensitivity of the resulting insights. Advanced defensive measures have been developed specifically for advertising applications, including specialized anomaly detection systems that identify potentially malicious updates based on patterns particularly relevant to behavioral modeling, though these systems must balance security requirements with privacy preservation to avoid inadvertently compromising the very protections that federated learning aims to provide.

COMPARATIVE ANALYSIS AND PERFORMANCE METRICS

The evaluation of privacy-preserving ad technologies necessitates a comprehensive framework that addresses their multifaceted nature, balancing privacy protection, targeting effectiveness, technical feasibility, and regulatory alignment. Contemporary evaluation approaches have evolved beyond simplistic performance assessments to incorporate sophisticated multi-dimensional analyses that acknowledge the inherent trade-offs between these competing objectives. An effective evaluation framework begins with clear articulation of stakeholder requirements and constraints, recognizing that advertisers, publishers, technology providers, consumers, and regulatory bodies may have divergent priorities and success criteria. The framework must establish standardized testing

methodologies that enable objective comparison across different privacy-preserving approaches, including controlled A/B testing protocols, benchmark datasets that represent realistic advertising scenarios while respecting privacy considerations, and reproducible evaluation procedures that facilitate meaningful cross-study comparisons. Drawing from established data mining evaluation principles, contemporary frameworks typically incorporate five essential dimensions: data utility (how effectively the protected data serves its intended analytical purpose), disclosure risk (the likelihood of sensitive information being revealed), algorithmic complexity (computational and implementation requirements), scalability (performance characteristics with increasing data volumes), and resilience to adversarial attacks (robustness against attempts to compromise privacy protections). These dimensions form an evaluation taxonomy that enables systematic assessment of different privacy-preserving technologies beyond superficial feature comparison, creating a structured approach to measurement that considers both technical performance and privacy outcomes (Bertino, E. *et al.*, 2005). Such frameworks have been adapted specifically for advertising contexts to address unique considerations including time-sensitivity of data, heterogeneity of user behaviors, and the particular commercial value of behavioral insights. Recent advancements have extended these frameworks to incorporate contextualized privacy evaluation that acknowledges the varying sensitivity levels of different data elements within advertising datasets, recognizing that attributes such as health-related interests may warrant stronger protection than general content preferences, thereby enabling more nuanced privacy-utility optimization.

Metrics for assessment of privacy-preserving ad technologies span multiple dimensions, beginning with privacy preservation efficacy as measured through formal privacy guarantees and empirical resistance to various attack vectors. Formal privacy guarantees typically include mathematical properties such as differential privacy with quantifiable privacy budgets, information-theoretic measures of data leakage, cryptographic security guarantees based on computational hardness assumptions, and formal verification of protocol security properties. Complementing these theoretical guarantees, empirical privacy assessments evaluate resistance against practical attack methodologies including membership inference attacks, model inversion, attribute

reconstruction, and side-channel exploits, often through adversarial evaluations that attempt to extract protected information under realistic constraints. Contemporary privacy efficacy metrics increasingly incorporate information-theoretic approaches that quantify the amount of sensitive information potentially leaked through model outputs or query responses, providing more rigorous mathematical foundations for privacy evaluation compared to earlier heuristic approaches. Ad targeting performance metrics retain their traditional importance even in privacy-preserving contexts, with engagement indicators, conversion measurements, and audience reach serving as critical indicators of business value. The evaluation challenge lies in comparing these performance metrics across different privacy-preserving approaches while controlling for confounding variables such as audience composition, creative quality, and temporal factors. Computational efficiency and scalability metrics address the practical feasibility of implementation, encompassing processing requirements, communication costs, storage needs, and scaling characteristics with respect to user population size and data dimensionality. These technical metrics have particular relevance in advertising contexts where real-time decision-making is often required and deployment environments range from high-performance data centers to resource-constrained consumer devices. Regulatory compliance assurance metrics evaluate alignment with relevant privacy legislation, including automated compliance verification, audit capabilities, consent management effectiveness, and adaptability to evolving regulatory requirements across different jurisdictions (Bertino, E. *et al.*, 2005). Contemporary assessment frameworks recognize the inherent tension between these different metric domains and increasingly utilize visualization techniques such as radar charts and trade-off curves to present multidimensional performance characteristics in ways that facilitate informed decision-making among stakeholders with different priorities.

Empirical comparison between federated learning and clean room approaches reveals distinct advantages and limitations that inform implementation decisions in different advertising contexts. Federated learning generally demonstrates superior privacy protection for individual-level data, as raw user information never leaves the local device, while clean rooms typically require some form of data sharing (albeit in protected environments with various technical

safeguards). This fundamental architectural difference results in different privacy threat profiles—federated learning primarily must defend against inference attacks on model updates and potential device compromise, while clean rooms must address risks of re-identification through query patterns, collusion among participants, and potential security vulnerabilities in the clean room infrastructure. Research examining information flow in these systems has developed formal information-theoretic models that quantify the potential privacy leakage under various threat models, demonstrating that the distributed nature of federated learning intrinsically limits certain forms of exposure while creating other unique vulnerabilities that must be specifically addressed. Performance comparisons across multiple advertising use cases indicate that clean room approaches currently deliver superior targeting accuracy for complex audience modeling scenarios that benefit from cross-dataset analysis, while federated learning shows competitive or superior performance for personalization use cases that can leverage rich on-device data without requiring

cross-user comparisons. Latency characteristics differ significantly, with clean rooms generally supporting complex analytical queries with moderate latency suitable for campaign planning and optimization, while federated learning exhibits higher initial training latency but can enable near-instantaneous on-device inference for real-time ad selection. Deployment flexibility also differs substantially, with clean rooms requiring significant infrastructure and typically involving extended implementation timelines, while federated learning can be deployed with varying degrees of sophistication depending on available resources and performance requirements (Shalabi, E. *et al.*, 2025). Recent studies comparing these approaches in real-world advertising scenarios have developed standardized benchmark tasks that evaluate performance across dimensions including privacy protection, targeting accuracy, computational efficiency, and implementation complexity, providing a more structured basis for technology selection than earlier anecdotal comparisons.

Table 3: Federated Learning vs. Clean Room Approaches: Comparative Analysis (Shalabi, E. *et al.*, 2025)

Aspect	Federated Learning	Clean Room Solutions	Key Considerations
Data Movement	Data remains on device/server of origin	Data moves to protected environment	Network bandwidth and data sensitivity
Privacy Model	No raw data sharing	Protected sharing with technical safeguards	Threat models and trust assumptions
Computational Distribution	Distributed across participating devices	Centralized within secure environment	Resource requirements and scalability
Use Case Suitability	On-device personalization, real-time adaptation	Cross-dataset analysis, audience insights	Business objectives and data characteristics

Cost-benefit analysis for implementation of privacy-preserving ad technologies requires structured evaluation of both quantifiable expenses and less tangible strategic advantages across multiple time horizons. Direct implementation costs include technology infrastructure (computational resources, storage, networking, security systems), specialized expertise (data scientists, privacy engineers, legal consultants), integration expenses (API development, system modifications), and ongoing operational costs (maintenance, updates, monitoring, compliance verification). These direct expenses vary significantly across different privacy-preserving approaches, with clean room implementations typically requiring larger upfront investment in secure infrastructure and governance frameworks, while federated learning deployments often entail higher ongoing costs for model maintenance and performance optimization across heterogeneous device ecosystems. Beyond these explicit costs,

organizations must consider opportunity costs associated with potential performance differentials compared to traditional approaches, as well as transition costs related to workflow disruption, retraining requirements, and temporary efficiency reductions during implementation phases. The benefit side of the equation encompasses both defensive advantages (regulatory compliance, breach risk mitigation, consumer trust preservation) and offensive opportunities (competitive differentiation, access to privacy-sensitive markets, partnership expansion with privacy-conscious organizations). Recent research has developed structured methodologies for quantifying these benefits, including information-theoretic approaches to valuing privacy protection and econometric models that relate privacy investments to measurable business outcomes such as user trust indicators and regulatory risk reduction (Shalabi, E. *et al.*, 2025). Comprehensive cost-benefit frameworks

specifically adapted for privacy technologies in advertising contexts now incorporate multi-criteria decision analysis techniques that enable organizations to weight different factors according to their strategic priorities, risk tolerance, and existing infrastructure constraints. These frameworks typically distinguish between immediate compliance benefits, medium-term operational efficiencies, and long-term strategic advantages, recognizing that privacy technology investments yield different types of returns across varying time horizons and should therefore be evaluated using appropriate temporal discounting models rather than simplistic return-on-investment calculations focused solely on short-term performance impacts.

FUTURE DIRECTIONS AND RESEARCH OPPORTUNITIES

The landscape of privacy-preserving advertising continues to evolve rapidly, with emerging technologies offering novel approaches to the fundamental challenge of balancing personalization with privacy protection. Trusted execution environments (TEEs) represent a promising technological frontier, enabling computation on sensitive data within hardware-protected enclaves that shield processing from unauthorized observation or interference, even by the system operators themselves. These secure enclaves potentially offer stronger security guarantees than pure cryptographic approaches, though they introduce different trust assumptions regarding hardware manufacturers and attestation mechanisms. Zero-knowledge proofs (ZKPs) are gaining traction in privacy-preserving advertising contexts, allowing one party to prove possession of certain information or satisfaction of specific criteria without revealing the underlying data. In advertising scenarios, ZKPs could enable verification of audience membership or conversion attribution without exposing individual user identities or behaviors. Synthetic data generation represents another emerging approach, using advanced generative techniques to create artificial datasets that preserve statistical properties and patterns of original user data while eliminating the privacy risks associated with using actual user information. Edge computing architectures are increasingly being explored as infrastructure for privacy-preserving advertising, pushing computation closer to data sources and reducing the need for data centralization. The concept of distributed asynchronous privacy-preserving technologies is evolving beyond basic implementations to incorporate novel approaches

for cross-device synchronization despite intermittent connectivity and heterogeneous computational capabilities, addressing practical deployment challenges in mobile advertising environments. Recent research has proposed reference architectures that formalize the integration of these diverse technologies within comprehensive privacy-preserving advertising frameworks, establishing standardized component interfaces and communication protocols that enable interoperability while maintaining privacy guarantees across system boundaries (Rivadeneira, J. E. *et al.*, 2023). These emerging approaches collectively represent a shift toward privacy-by-design principles rather than retrospective privacy protections, embedding privacy safeguards into the fundamental architecture of advertising systems rather than applying them as supplementary controls.

Integration potential with blockchain technologies and decentralized identifiers presents intriguing opportunities for enhancing both the effectiveness and transparency of privacy-preserving advertising systems. Blockchain infrastructures can provide immutable audit trails for data access, processing activities, and consent management, creating verifiable records of privacy-related operations that enhance accountability while preserving confidentiality of the underlying data. Smart contracts deployed on blockchain networks can encode and automatically enforce privacy policies and data usage restrictions, potentially reducing reliance on trust between advertising ecosystem participants. Decentralized identifiers (DIDs) offer a promising alternative to traditional identity management approaches, enabling users to maintain sovereignty over their digital identities while selectively disclosing attributes relevant to advertising personalization without exposing comprehensive profiles. Recent research has demonstrated how blockchain-based systems can effectively implement privacy-preserving access control mechanisms through cryptographic techniques that verify authorization without revealing sensitive identity attributes, creating potential applications for selective advertising content delivery based on verified but undisclosed user characteristics. The integration of blockchain with federated learning systems has shown particular promise, with novel consensus mechanisms specifically designed to validate model updates while preventing privacy leakage, effectively addressing the dual challenges of ensuring model quality and protecting sensitive information in distributed learning environments.

These approaches typically implement reputation systems and stake-based participation mechanisms that create economic disincentives for adversarial behavior without compromising the privacy guarantees of the underlying federated learning protocol (Li, L. 2024). The decentralized nature of blockchain architectures aligns conceptually with the distributed processing model of privacy-preserving advertising technologies, potentially creating synergistic integration opportunities that enhance both privacy protection and system resilience against manipulation or censorship. While significant research challenges remain regarding scalability, energy efficiency, and regulatory compliance, the convergence of blockchain technologies with privacy-preserving advertising represents a promising direction for addressing trust deficits in digital advertising while maintaining effective personalization capabilities.

Standardization efforts and industry collaboration represent critical enablers for the widespread adoption of privacy-preserving advertising technologies, addressing the fragmentation that currently limits interoperability and increases implementation complexity. Technical standardization initiatives are emerging across multiple domains, including cryptographic protocols for secure data sharing, API specifications for privacy-preserving measurement, data schemas for clean room implementations, and model architecture definitions for federated learning deployments. These standardization efforts aim to create common technical foundations that reduce implementation barriers while ensuring consistent privacy protections across different platforms and providers. Recent research has identified significant challenges in standardization processes, including tensions between rapid innovation and stable reference implementations, competing commercial interests that complicate consensus formation, and the technical complexity of precisely specifying privacy guarantees across diverse implementation contexts. Despite these challenges, collaborative standardization initiatives have made notable progress in establishing common vocabularies, reference architectures, and evaluation methodologies that facilitate meaningful comparison between different privacy-preserving approaches. Cross-organization research collaborations have emerged as valuable mechanisms for addressing particularly complex technical challenges that exceed the resources or expertise of individual organizations, creating shared infrastructures for privacy-preserving

advertising research that accelerate innovation while avoiding duplicative efforts (Rivadeneira, J. E. *et al.*, 2023). These collaborative endeavors frequently leverage open-source development models to enhance transparency, enable community contribution, and facilitate independent security assessment, reducing the barriers to adoption while building trust in implementation quality. Industry self-regulatory frameworks have begun incorporating privacy-preserving technologies as core components rather than supplementary controls, establishing technical requirements and certification mechanisms that provide meaningful assurance regarding privacy protections. The convergence of formal standardization efforts with more informal collaboration mechanisms potentially creates powerful network effects that accelerate adoption across the advertising ecosystem, as interoperability reduces integration costs while the growing participant network increases the value proposition for each organization.

The research agenda for addressing current limitations in privacy-preserving advertising technologies encompasses multiple dimensions, including technical capabilities, performance optimization, usability enhancements, and integration with existing advertising infrastructure. Privacy-utility trade-off optimization remains a central research challenge, requiring more sophisticated mathematical frameworks for quantifying and balancing these competing objectives across different advertising contexts and privacy sensitivity levels. Current approaches often rely on heuristic parameter tuning rather than principled optimization methods, creating opportunities for more rigorous analytical approaches that provide clearer guidance for implementation decisions. Computational efficiency represents another critical research direction, particularly for cryptographic approaches that currently impose significant performance overhead relative to non-privacy-preserving alternatives. The challenge of balancing network latency, computational complexity, and privacy guarantees requires novel approaches that optimize across these multiple dimensions simultaneously rather than treating them as independent considerations. Security research focused on privacy-preserving advertising has identified unique vulnerability patterns that differ from traditional systems, necessitating specialized threat modeling approaches and defensive techniques specifically designed for distributed privacy-preserving architectures. Recent advances

in privacy-preserving machine learning have demonstrated promising techniques for addressing statistical challenges in advertising applications, including methods for handling non-uniform data distributions, mitigating selection bias effects, and quantifying uncertainty in analytical results derived from privacy-protected data sources. These statistical approaches must be further adapted to the specific characteristics of advertising data, including its temporal sensitivity, sparse interaction patterns, and complex attribution challenges (Li, L. 2024). User control mechanisms

represent another important research direction, developing interfaces and interaction models that effectively communicate privacy properties and enable meaningful choice without imposing prohibitive cognitive burdens. This user-centered research stream intersects with blockchain-enhanced transparency mechanisms that could provide verifiable evidence of compliance with stated privacy policies, potentially addressing trust deficits that currently limit consumer engagement with personalized advertising.

Table 4: Research Priorities for Privacy-Preserving Ad Technologies. (Rivadeneira, J. E. *et al.*, 2023)

Research Area	Current Limitations	Future Direction	Potential Impact
Privacy-Utility Optimization	Heuristic approaches, undefined metrics	Formal mathematical frameworks, standardized benchmarks	Improved decision-making for implementation parameters
Blockchain Integration	Limited scalability, high resource requirements	Lightweight consensus protocols, specialized advertising applications	Enhanced transparency and user control mechanisms
Computational Efficiency	High overhead for cryptographic methods	Hardware acceleration, advertising-specific optimizations	Broader adoption across resource-constrained environments
Cross-Technology Standardization	Fragmented implementations, limited interoperability	Unified reference architectures, common APIs	Ecosystem-wide adoption and implementation efficiency

Implications for advertising ecosystem stakeholders extend across technological, operational, economic, and strategic dimensions, with varying impacts for different participants in the advertising value chain. For advertisers, privacy-preserving technologies necessitate fundamental reconsideration of targeting and measurement strategies, shifting from individual-level tracking toward cohort-based or on-device personalization approaches. This transition requires not only technical adaptation but also organizational capability development, including expertise in privacy-preserving analytical techniques, modified workflow processes, and revised performance expectations that acknowledge the inherent limitations of privacy-protected data environments. Publishers face similarly profound implications, balancing potential revenue impacts from reduced targeting precision against opportunities to leverage first-party relationships and contextual relevance. The implementation of privacy-preserving technologies potentially rebalances value exchange dynamics within the advertising ecosystem, with entities holding direct audience relationships potentially gaining advantage relative to intermediaries whose value propositions centered on cross-site tracking capabilities. Advertising technology providers face perhaps the most significant transformation

requirements, necessitating fundamental reconsideration of data processing architectures, analytical methodologies, and value propositions in a privacy-centric ecosystem. Recent research has examined how these ecosystem dynamics might evolve under different privacy technology adoption scenarios, identifying potential equilibrium states that balance commercial viability with privacy protection (Rivadeneira, J. E. *et al.*, 2023). For users, the implications include enhanced privacy protections, potentially modified advertising experiences, and evolving value exchange propositions that may include more explicit consent mechanisms or compensation models for data sharing. The diversity of these stakeholder implications underscores the systemic nature of the transition toward privacy-preserving advertising, requiring coordinated adaptation across the ecosystem rather than isolated technological implementation. Organizations that proactively develop both the technical capabilities and strategic positioning for privacy-preserving advertising may gain significant advantages in an advertising landscape increasingly shaped by privacy considerations, with research suggesting that early adopters can establish competitive differentiation while building internal expertise that becomes increasingly valuable as privacy

regulations continue to evolve and expand globally.

CONCLUSION

Privacy-preserving advertising technologies represent a fundamental paradigm shift for digital advertising, simultaneously addressing regulatory compliance requirements and evolving consumer privacy expectations. Data clean rooms and federated learning each offer distinct advantages that can be strategically deployed based on specific use cases, data sensitivity, and business objectives. The implementation challenges—including computational efficiency limitations, statistical heterogeneity across user populations, and security concerns—require continued technical innovation and interdisciplinary collaboration. The integration potential with blockchain technologies presents particularly promising opportunities for enhancing transparency and trust throughout the advertising ecosystem, though standardization efforts remain essential for widespread adoption and interoperability. As these technologies mature, the economic and strategic implications will reshape competitive dynamics, creating advantages for organizations that proactively develop privacy-preserving capabilities. The fundamental privacy-utility trade-off at the heart of these technologies demands more rigorous mathematical frameworks and standardized evaluation methodologies. Ultimately, the transition toward privacy-preserving advertising necessitates coordinated adaptation across the entire ecosystem, balancing the legitimate business needs of advertisers and publishers with the privacy rights of consumers. Through thoughtful implementation of these technologies, the digital advertising industry can establish a sustainable future that maintains personalization effectiveness while respecting privacy boundaries—ensuring that advertising continues to provide value to all participants in an increasingly privacy-conscious digital landscape.

REFERENCES

1. Rao, S. The Evolution of Privacy Rights in the Digital Age: A Comparative Analysis of GDPR and CCPA. *Indian Journal of Law*, 2.4, (2024) 52-56.
2. Fitzgerald, A. "Non-Compliance Fines and Sanctions: Why It's More Expensive Not to Comply with Regulations," *Secure frame blogs*, (2024). <https://secureframe.com/blog/sanctions-non-compliance-fine>
3. Jindal, P. "Privacy-Preserving Data Analysis: Implications of Clean Rooms." *International Journal of Management, IT & Engineering* 14.6 (2024).
4. Ahessin, A. & Ali, A. "Privacy-Preserving Technologies: Homomorphic Encryption and Secure Multi-Party Computation". *International Journal for Research in Applied Science and Engineering Technology*. 13. 740-745, (2025).
5. Cao, T. D., Truong-Huu, T., Tran, H., & Tran, K. "A federated learning framework for privacy-preserving and parallel training." *arXiv preprint arXiv:2001.09782* (2020).
6. Narula, M., Meena, J., & Vishwakarma, D. K. "A comprehensive review on Federated Learning for Data-Sensitive Application: Open issues & challenges," *Engineering Applications of Artificial Intelligence*, 133 (2024): 108-128.
7. Bertino, E., Fovino, I. N., & Provenza, L. P. "A framework for evaluating privacy preserving data mining algorithms." *Data Mining and Knowledge Discovery* 11.2 (2005): 121-154.
8. Shalabi, E., Khedr, W., Rushdy, E., & Salah, A. "A comparative study of privacy-preserving techniques in federated learning: A performance and security analysis." *Information* 16.3 (2025): 244.
9. Rivadeneira, J. E., Silva, J. S., Colomo-Palacios, R., Rodrigues, A., & Boavida, F. "User-centric privacy preserving models for a new era of the Internet of Things." *Journal of network and computer applications* 217 (2023): 103695.
10. Li, L. "Blockchain-Enhanced Privacy Protection in Social Network Research and Applications". *Highlights in Science, Engineering and Technology*. 85. (2024) 487-493.

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:

Vinnakota, S. "Privacy-Preserving Ad Targeting in the Age of Data Protection Legislation: Leveraging Federated Learning and Clean Room Solutions" *Sarcouncil Journal of Multidisciplinary* 5.7 (2025): pp 718-728.