

Privacy-Preserving Personalized Recommendation with Federated Learning for Cross-Domain User Modeling

Anton Richards

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
antonr@buffalo.edu

Shane Morales

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
moraless Shane@colostate.edu

Sanjay A. Bether

Department of Computer Science, University of Houston, Houston, TX, USA.
smathur@uh.edu

Clifford A. Jimenez

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL,
USA.
hellocifford@uab.edu

Abstract

Personalized recommender systems increasingly depend on cross-domain user modeling to overcome data sparsity, improve cold-start performance, and produce contextually relevant suggestions across heterogeneous digital services. However, centralized collection and processing of user interaction data expose individuals to privacy risks and create legal tensions with modern data protection regimes. Federated learning offers a promising alternative by enabling collaborative model training without transferring raw user data to a central server. This article presents a system-level examination of privacy-preserving personalized recommendation with federated learning for cross-domain user modeling. It analyzes architectural choices, privacy mechanisms, structural trade-offs, robustness, fairness, governance, and deployment sustainability. The discussion emphasizes that privacy-preserving cross-domain recommendation is not solely an algorithmic challenge; it requires careful coordination among client-side representation learning, edge aggregation, domain-specific adaptation, incentive design, and regulatory compliance. We argue that the design space is shaped by fundamental tensions between personalization depth and communication efficiency, local differential privacy and model utility, cross-domain transfer and negative interference, and participant autonomy and system-wide consistency. The article further considers how popularity and demographic biases can be amplified by federated optimization and how fairness constraints interact with resource heterogeneity. It concludes by outlining forward-looking research directions in hierarchical trust architectures, adaptive privacy budgeting, auditable cross-domain governance, and sustainable deployment models for privacy-preserving personalization infrastructures.

Keywords

federated learning, cross-domain recommendation, privacy preservation, user modeling, distributed systems, data governance, fairness.

1. Introduction

Personalized recommendation systems have become foundational components of digital commerce, media streaming, social platforms, and location-based services. Traditional recommendation architectures rely on centralized collection of user profiles, item catalogs, interaction histories, and contextual signals. Collaborative filtering and latent factor models have historically produced strong personalization results because they can infer preference structures from large matrices of user-item interactions [4]. Nevertheless, the centralized data assumption creates serious privacy vulnerabilities, including unauthorized profiling, data breaches, and secondary use of behavioral traces. Regulatory frameworks such as the European Union General Data Protection Regulation have transformed privacy from a desirable feature into a structural requirement for deployment [6]. These developments motivate a shift toward decentralized learning architectures that preserve user privacy while retaining personalization quality.

Federated learning has emerged as a practical framework for training machine learning models across distributed clients without directly sharing raw data [1]. It was later formalized as a broader paradigm for collaborative model learning under heterogeneous data conditions [2]. The federated approach enables mobile devices, browsers, edge servers, and organizational silos to contribute model updates to a shared coordinator while keeping sensitive interaction data local. Comprehensive analyses of federated learning have identified open challenges related to communication cost, statistical heterogeneity, systems reliability, and adversarial robustness [3]. These challenges become more pronounced when federated learning is applied to personalized recommendation, because recommendation models must capture fine-grained user preferences while maintaining sufficient generalization across user populations.

Cross-domain user modeling adds further complexity. A user may interact with a retail platform, a music streaming service, a navigation application, and a news aggregator. Each domain has distinct item schemas, interaction semantics, temporal dynamics, and privacy sensitivities. Recent work on AI-driven personalization across domains for local categorical query understanding and context-aware retrieval [7] illustrates the need for models that can interpret domain-specific categorical queries without transferring raw identifiers or unstructured behavioral logs across organizational boundaries. Differential privacy provides a formal foundation for limiting information leakage from model updates [5]. However, translating differential privacy guarantees into practical recommendation systems requires careful consideration of trust boundaries, noise injection points, and domain-specific utility constraints.

This article examines privacy-preserving personalized recommendation with federated learning from a system-level perspective. Rather than focusing solely on optimization algorithms, it emphasizes architectural design, structural trade-offs, infrastructure requirements, robustness, fairness, governance, and policy implications. The central argument is that sustainable cross-domain personalization infrastructures must balance competing objectives: local personalization quality, global model utility, communication overhead, privacy protection, auditability, and participant incentives. The remainder of this article develops this argument across six sections, moving from architectural foundations to deployment sustainability and future research directions.

2. System Architecture and Cross-Domain Data Integration

A privacy-preserving cross-domain recommendation architecture typically consists of multiple client nodes, optional edge aggregation layers, a central coordination server, and domain-specific model adapters. Client nodes hold private interaction logs, preference signals, and contextual features. Instead of uploading raw records, clients compute local model updates and transmit only the necessary parameters or gradients. The central server aggregates these updates into a global model, which is then redistributed to clients. Production deployments such as federated mobile keyboard prediction have demonstrated that this pattern can operate at scale across millions of heterogeneous devices [8]. Large-scale federated systems require robust scheduling, dynamic device selection, secure communication channels, and mechanisms for handling intermittent participation [9]. These architectural elements are equally important for recommendation use cases, where user sessions are often short, context-dependent, and spread across multiple applications.

The system design becomes more complex when recommendation spans multiple domains. A single federated global model may be insufficient because user behavior in one domain does not transfer uniformly to another. Therefore, a layered architecture is often needed: a shared federated representation layer captures general user preferences, while domain-specific adapters translate the shared representation into actionable recommendations within each domain. This design preserves cross-domain knowledge transfer while allowing each domain operator to maintain specialized item catalogs, ranking constraints, and business rules. It also creates a natural boundary for privacy governance, because domain-specific adapters can be trained under different privacy policies without exposing raw cross-domain identifiers.

Federated learning systems face well-documented challenges related to statistical heterogeneity, communication bottlenecks, and unreliable clients [10]. Recommendation models are especially sensitive to these challenges because user preferences are highly skewed, item popularity distributions are long-tailed, and interaction signals are sparse. In mobile edge environments, hierarchical aggregation can reduce wide-area communication traffic and improve responsiveness by placing intermediate aggregation nodes closer to users [11]. An edge layer can also perform domain-specific filtering and local context interpretation, which is particularly valuable for location-aware or time-sensitive recommendations. However, introducing edge aggregation nodes expands the attack surface and requires additional trust assumptions.

Cross-domain user modeling intensifies the non-IID data problem. Each client may contribute data from multiple domains with incompatible feature spaces, inconsistent label semantics, and varying update magnitudes [12]. The federated coordinator must decide whether to train a single joint model, separate per-domain models with shared embeddings, or a mixture of experts with domain routing. These architectural choices affect communication cost, convergence speed, personalization quality, and privacy leakage. A well-designed architecture treats domain identity as a first-class structural dimension rather than a nuisance variable. Domain-specific adapters can be combined with secure aggregation and differential privacy to prevent the central server from reconstructing sensitive domain-level behavior while still benefiting from cross-domain knowledge.

3. Privacy Mechanisms and Trust Boundaries

Privacy protection in federated recommendation systems relies on a combination of cryptographic protocols, noise injection, access controls, and architectural trust boundaries.

Although federated learning avoids sharing raw data, model updates can leak information about individual training records through gradient inversion attacks, membership inference, and unintended memorization. Differential privacy has become a standard formal framework for bounding such leakage. Federated learning with differentially private update mechanisms can provide client-level or record-level guarantees while maintaining acceptable model performance in many settings [13]. Deep learning models trained with differential privacy require careful noise calibration and gradient clipping to balance utility and privacy [14]. Recommendation models, which often rely on high-dimensional embedding tables and sparse interaction signals, are particularly sensitive to noise perturbation.

Trust boundaries determine where privacy mechanisms should be enforced. In a basic client-server architecture, the central coordinator is trusted only with aggregate updates. Secure aggregation can further reduce the coordinator's visibility by ensuring that individual updates remain encrypted during the aggregation process. Alternatively, differential privacy can be applied locally on each client before transmission, ensuring that even the coordinator cannot infer individual contributions. The choice between local and central privacy models depends on the threat model, regulatory requirements, and performance constraints. Local mechanisms often introduce more noise because each client must obscure its own update independently. Central mechanisms can achieve better utility under a trusted aggregator, but they require stronger assumptions about the aggregator's integrity.

Cross-domain recommendation introduces additional privacy challenges. A single user may contribute data from multiple applications, and linking those contributions across domains can reveal sensitive behavioral patterns. Therefore, cross-domain federated systems should avoid sharing stable user identifiers across domains. Pseudonymous identifiers, domain-specific user keys, and federated identity mapping can reduce the risk of unintended linkage. At the same time, personalization quality often depends on maintaining a consistent user representation across domains. This creates a structural tension between privacy and utility that cannot be resolved by algorithmic techniques alone. It requires careful system design, including clear data minimization policies, limited retention of linkage keys, and explicit user consent for cross-domain modeling.

Privacy mechanisms must also account for heterogeneous device capabilities and network conditions. Noise injection, secure enclaves, and cryptographic operations consume additional computational resources. On constrained mobile devices, these costs may conflict with battery life, latency budgets, and user experience. As a result, privacy-preserving recommendation systems must support adaptive privacy configurations that match device capabilities and application requirements. The governance layer should document the privacy guarantees provided for each domain and the residual risks that remain after technical controls are applied.

4. Structural Trade-offs in Personalization and Resource Management

The design of privacy-preserving cross-domain recommendation systems is shaped by fundamental structural trade-offs. One central trade-off occurs between global generalization and local personalization. A global federated model can capture broad preference patterns but may fail to represent rare interests, niche communities, or rapidly changing local contexts. Local models can adapt more finely to individual behavior but require more client-side computation and may overfit to sparse interaction histories. Cross-domain modeling adds another dimension: shared representations can improve cold-start performance in one domain

by transferring knowledge from another, but excessive sharing may introduce negative transfer when domain semantics diverge.

Resource heterogeneity across devices and organizations imposes additional constraints. Federated recommendation systems must operate across smartphones, tablets, wearable devices, and edge servers with different storage, memory, bandwidth, and battery capacities. Communication-efficient aggregation methods can reduce bandwidth but may limit the expressiveness of model updates. Compression, quantization, and sparse updates improve efficiency but can distort gradients and slow convergence. These trade-offs are particularly severe for large embedding tables, which are common in collaborative filtering models. Recommendation models trained with federated learning must therefore balance embedding compression, client sampling, and communication frequency against personalization accuracy.

Bias and fairness concerns are also structural rather than purely algorithmic. Recommender systems have been shown to exhibit popularity bias and demographic bias, where mainstream content is overrepresented and minority users receive lower-quality recommendations [15]. Federated learning can amplify these biases when clients with high interaction volumes dominate the global model. Similarly, cross-domain transfer may propagate biases from one domain to another, especially when domain-specific feedback loops differ in density and label quality. Fairness-aware machine learning highlights the limitations of optimizing only aggregate accuracy, because aggregate improvements can mask disparate outcomes for protected groups [16]. In federated settings, fairness must be evaluated across clients, domains, and demographic segments simultaneously.

Therefore, system designers should treat fairness as a first-order constraint in architectural decisions. This includes auditing client sampling strategies, weighting updates to prevent majority dominance, monitoring domain-specific performance disparities, and incorporating fairness metrics into model evaluation. Resource management policies must also be examined for fairness, because requiring all clients to participate equally may disadvantage users with limited connectivity or older devices. A sustainable system should allow heterogeneous participation without systematically excluding already marginalized users.

5. Robustness, Fairness, and Incentive Alignment

Robustness in federated recommendation systems encompasses resilience to adversarial clients, noisy labels, distribution shift, and system failures. Because federated learning relies on contributions from many participants, a small number of malicious or malfunctioning clients can degrade the global model if their updates are not carefully filtered. Deep learning based recommender systems have demonstrated strong capacity for capturing complex user-item relationships, but they also introduce vulnerabilities related to overfitting, adversarial perturbations, and unstable training dynamics [17]. In cross-domain settings, a compromised domain adapter could poison shared representations and thereby affect recommendations in unrelated domains. Robust aggregation procedures, anomaly detection, and isolation mechanisms are necessary to contain such failures.

Cross-domain recommender systems have been studied along dimensions such as user overlap, item overlap, and knowledge transfer mechanisms [18]. These taxonomies are useful for understanding when cross-domain modeling is likely to succeed and when it may harm personalization. For example, domains with high user overlap but low item overlap may benefit from shared user representations, while domains with little semantic alignment may perform better with domain-specific models. Federated learning introduces additional

uncertainty because the central server cannot directly inspect the distribution of local data. Therefore, cross-domain federated systems should include domain-specific validation procedures and dynamic adaptation policies that can disable transfer when negative interference is detected.

Incentive alignment is essential for long-term participation. Users are more likely to contribute data and computational resources if they perceive a clear benefit in personalization quality, privacy protection, and control over their data. Organizations operating different domains may be reluctant to share model updates if they fear losing competitive advantage or exposing proprietary item representations. Federated cross-domain recommendation therefore requires governance mechanisms that allocate benefits fairly, protect domain-specific intellectual property, and provide transparent audit trails. Without such mechanisms, participants may withhold high-quality updates or exit the system, reducing the global model's value for everyone.

Privacy-preserving mechanisms can also affect fairness. Differential privacy noise is typically added uniformly, but uniform noise may disproportionately harm users with sparse data, because their signals are more easily obscured. This can widen the gap between active and inactive users. Adaptive privacy budgets, where high-sensitivity domains receive stronger protection but low-sensitivity domains preserve more utility, may help address this tension. However, adaptive budgets must be governed carefully to avoid creating new forms of disparate treatment. Fairness auditing should therefore accompany privacy auditing throughout the system lifecycle.

6. Governance, Policy, and Deployment Sustainability

The deployment of privacy-preserving cross-domain recommendation systems raises governance challenges that extend beyond technical implementation. A systematic literature review of cross-domain recommender systems highlights considerable fragmentation in evaluation protocols, data assumptions, and deployment contexts [19]. This fragmentation complicates the translation of research results into production systems. Governance frameworks should define clear accountability for model updates, data access, privacy guarantees, and fairness outcomes. They should also specify how cross-domain data flows are approved, monitored, and terminated when necessary.

Regulatory compliance is a central governance concern. Data protection laws impose obligations regarding purpose limitation, data minimization, user consent, and the right to erasure. Federated systems complicate these obligations because user influence is distributed across many devices and domains. The right to erasure, for example, may require removing a user's contribution from aggregated model history, which is difficult when updates have already been aggregated into a global model. Techniques such as approximate unlearning, periodic model retraining, and update lineage tracking can support compliance, but they require system-level support. Cross-domain recommendation systems must also respect sector-specific rules, such as health data protections or financial privacy regulations.

Deployment sustainability depends on operational costs, model maintenance, and organizational alignment. Federated recommendation pipelines often combine deep learning components with traditional models for candidate retrieval, feature engineering, and ranking. Scalable tree boosting models [20] can be used locally for efficient candidate scoring in resource-constrained environments, complementing federated neural representations. This hybrid approach reduces latency and computation while preserving privacy-sensitive features

on the client. However, maintaining multiple model families across domains increases operational complexity. Sustainable deployment requires versioned model artifacts, automated evaluation pipelines, and fallback strategies when federated updates fail or produce degraded outcomes.

Auditability is another critical governance requirement. Regulators, users, and participating organizations need evidence that privacy commitments are being honored. This requires tamper-evident logs of model updates, privacy budget consumption, aggregation decisions, and data access events. Cryptographic techniques can support verifiable aggregation and audit trails, but they introduce computational overhead and key management complexity. A pragmatic governance layer should combine technical audit mechanisms with independent oversight and periodic risk assessments. Cross-domain systems should also provide users with meaningful transparency about how their data influence recommendations across domains, even when raw data remain local.

7. Future Directions

Several future directions merit attention. First, hierarchical trust architectures should be developed to support multiple privacy models within a single cross-domain system. Different domains may require different trust assumptions, and a rigid architecture cannot accommodate such diversity. Second, adaptive privacy budgeting for embedding-based recommender models remains underexplored. Because embedding tables concentrate information in dense vectors, they may leak more than conventional model parameters. Future research should examine privacy accounting for federated embedding updates and cross-domain embedding sharing. Third, cross-domain negative transfer detection requires more robust runtime diagnostics. Systems should be able to identify when a shared representation degrades domain-specific performance and automatically reduce transfer strength.

Fourth, user-centered evaluation should move beyond aggregate accuracy metrics. Privacy-preserving recommendation systems should be evaluated on perceived personalization quality, control, transparency, and fairness. This requires interdisciplinary research involving human-computer interaction, law, and systems engineering. Fifth, sustainability models should address the energy and carbon costs of frequent federated training cycles. Although federated learning can reduce centralized data center workloads, it may increase distributed computation and communication. Cross-domain recommendation systems should optimize for energy efficiency and long-term maintainability rather than short-term model accuracy alone.

Finally, standardized benchmarks and governance protocols for cross-domain federated recommendation would improve reproducibility and regulatory clarity. Real-world deployments currently rely on ad hoc combinations of privacy mechanisms and evaluation practices. Establishing shared protocols for threat modeling, privacy auditing, fairness assessment, and cross-domain data flow documentation would strengthen both research and industrial practice. These future directions underscore the need for continued interdisciplinary collaboration among machine learning researchers, systems engineers, policy experts, and domain practitioners.

8. Conclusion

Privacy-preserving personalized recommendation with federated learning offers a promising path toward cross-domain user modeling that respects user privacy and regulatory constraints. However, realizing this vision requires more than algorithmic innovation. System-level design decisions about architecture, trust boundaries, privacy mechanisms, resource

management, fairness, governance, and deployment sustainability shape the effectiveness and legitimacy of these systems. This article has examined these dimensions and highlighted the structural tensions that arise when personalization depth, cross-domain transfer, communication efficiency, differential privacy, and participant incentives interact. It has argued that sustainable systems must treat privacy, robustness, fairness, and auditability as first-order architectural requirements rather than post hoc constraints. By aligning technical mechanisms with governance structures and policy frameworks, practitioners can build cross-domain recommendation systems that deliver useful personalization while preserving user trust and autonomy. Future research should focus on adaptive privacy architectures, runtime transfer diagnostics, user-centered evaluation, and sustainable deployment models to advance this field responsibly.

References

1. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)* (pp. 1273–1282). PMLR.
2. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19.
3. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., Rouayheb, S. E., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.
4. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8), 30–37.
5. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407.
6. Voigt, P., & von dem Bussche, A. (2017). *The EU General Data Protection Regulation (GDPR): A practical guide*. Springer.
7. Yu, X. (2026, January). AI-Driven Personalization across Domains for Local Categorical Query Understanding and Context-Aware Retrieval. In *Proceedings of the 2nd International Conference on Artificial Intelligence, Digital Media Technology and Social Computing* (pp. 77-82).
8. Hard, A., Rao, K., Mathews, R., Ramaswamy, S., Beaufays, F., Augenstein, S., Eichner, H., Kiddon, C., & Ramage, D. (2018). Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*.
9. Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., Kiddon, C., Konečný, J., Mazzocchi, S., McMahan, H. B., Van Overveldt, T., Petrou, D., Ramage, D., & Roselander, J. (2019). Towards federated learning at scale: System design. *Proceedings of Machine Learning and Systems*, 1, 374–388.
10. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60.

11. Lim, W. Y. B., Luong, N. C., Hoang, D. T., Jiao, Y., Liang, Y.-C., Yang, Q., Niyato, D., & Miao, C. (2020). Federated learning in mobile edge networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(3), 2031–2063.
12. Zhu, H., Xu, J., Liu, S., & Jin, Y. (2021). Federated learning on non-IID data: A survey. *Neurocomputing*, 465, 371–390.
13. Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., Jin, S., Quek, T. Q. S., & Poor, H. V. (2020). Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15, 3454–3469.
14. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 308–318).
15. Ekstrand, M. D., Tian, M., Azpiaz, I. M., Ekstrand, J. D., Anuyah, O., McNeill, D., & Pera, M. S. (2018). All the cool kids, how do they fit in? Popularity and demographic biases in recommender evaluation and effectiveness. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* (pp. 172–186). PMLR.
16. Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and machine learning: Limitations and opportunities*. MIT Press.
17. Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys*, 52(1), 1–38.
18. Cantador, I., Fernández-Tobías, I., Berkovsky, S., & Cremonesi, P. (2015). Cross-domain recommender systems. In F. Ricci, L. Rokach, & B. Shapira (Eds.), *Recommender Systems Handbook* (pp. 919–959). Springer.
19. Khan, M. M., Ibrahim, R., & Ghani, I. (2017). Cross-domain recommender systems: A systematic literature review. *ACM Computing Surveys*, 50(3), 1–34.
20. Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794).