# Securing Social Commerce Infrastructure through Privacy Preserving Federated Learning Synergizing Multi-Modal Large Language Models and Differential Privacy

Miranda Vance
School of Information Systems, University of Maryland Baltimore County
mvance@umbc.edu

Julian Montgomery
Department of Systems and Industrial Engineering, University of Arizona
jmontgomery@arizona.edu

Hugh Brooks
Department of Electrical Engineering, Oregon State University
hbrooks@oregonstate.edu

**Abstract**

The convergence of social media and electronic commerce has birthed a complex socio-technical ecosystem known as social commerce, which relies heavily on the ingestion of heterogeneous user data to drive personalized recommendation engines. However, the centralization of multi-modal data—including text, images, and transactional behaviors—poses significant privacy risks and systemic vulnerabilities. This paper proposes a comprehensive architectural framework for securing social commerce infrastructure by integrating privacy-preserving federated learning with multi-modal large language models and differential privacy mechanisms. By shifting the paradigm from centralized data aggregation to decentralized model training, the proposed system ensures that sensitive user attributes remain on-device while allowing the global model to benefit from collective intelligence. We provide a deep analytical exploration of the system-level trade-offs between privacy guarantees, computational overhead, and recommendation accuracy. The discussion emphasizes the structural requirements for hardware-aware orchestration and the necessity of robust governance frameworks to manage autonomous agentic behaviors in digital marketplaces. Furthermore, we examine the socio-technical implications of this infrastructure, focusing on algorithmic fairness, environmental sustainability, and the evolving global policy landscape regarding data sovereignty. By synergizing the semantic depth of multi-modal transformers with the mathematical rigor of differential privacy, this research offers a resilient blueprint for the next generation of social commerce, ensuring that commercial efficiency does not come at the expense of individual privacy or systemic security.

**Keywords**
Social Commerce Infrastructure, Federated Learning, Multi-Modal Large Language Models, Differential Privacy, Privacy-Preserving Computation, Socio-Technical Governance, Distributed Systems.

## 1. Introduction

The contemporary digital economy is increasingly defined by the erosion of boundaries between social interaction and commercial transaction. Social commerce has emerged as a dominant paradigm, leveraging the massive scale of social networks to facilitate peer-to-peer commerce, influencer-driven marketing, and algorithmically curated product discovery. At the heart of this infrastructure lies the ability to process and interpret multi-modal data streams, where linguistic context, visual aesthetics, and historical purchasing patterns are fused to predict consumer intent. While this integration has unlocked unprecedented levels of commercial efficiency, it has simultaneously created a "privacy-security paradox." To provide high-fidelity personalization, platforms have traditionally required the centralization of vast quantities of sensitive user data, creating lucrative targets for adversarial exploitation and raising profound ethical concerns regarding surveillance and data autonomy.

Securing the infrastructure of social commerce requires a fundamental departure from centralized data architectures toward decentralized, privacy-preserving computation. Federated learning provides a viable pathway for this transition, allowing models to be trained across a distributed network of devices without ever requiring the raw data to leave its point of origin. However, applying federated learning to the complex, high-dimensional space of social commerce introduces significant technical hurdles. Multi-modal large language models, which are essential for understanding the nuances of social commercial interactions, are computationally expensive and sensitive to the noise introduced by privacy-preserving mechanisms. Furthermore, the decentralized nature of the training process makes the system vulnerable to new classes of attacks, such as model poisoning and gradient inversion.

This research addresses these challenges by proposing a synergistic framework that combines the semantic reasoning of multi-modal transformers with the formal privacy guarantees of differential privacy within a federated learning ecosystem. We move beyond the optimization of individual algorithms to focus on the system-level orchestration required to deploy these technologies at scale. This includes an investigation into the structural trade-offs between privacy budget and model utility, the infrastructure requirements for heterogeneous device participation, and the governance frameworks necessary to ensure algorithmic fairness. By treating social commerce as a critical socio-technical infrastructure, this paper aims to provide a robust analytical foundation for a digital marketplace that is both technologically advanced and ethically grounded.

## 2. Theoretical Framework of Multi-Modal Federated Learning

The application of federated learning in social commerce must account for the unique characteristics of the data involved. Unlike traditional horizontal federated learning, which often deals with homogeneous tabular data, social commerce involves multi-modal inputs

where the relationship between a user's text comment and an uploaded product image provides a richer signal than either modality in isolation. The integration of multi-modal large language models into the federated pipeline allows the system to capture these cross-modal dependencies. However, the high dimensionality of transformer embeddings poses a challenge for traditional federated averaging protocols. From a systems perspective, this requires the development of "modality-aware" aggregation strategies that can weight updates based on the quality and diversity of the local data available on each participating device.

Synergizing these models with differential privacy adds a layer of mathematical rigor to the privacy guarantees. Differential privacy ensures that the inclusion or exclusion of a single user's data in the training set does not significantly alter the output of the global model, thereby providing a formal defense against membership inference attacks. In the context of federated learning, this is typically achieved through "local differential privacy," where noise is added to the gradients before they are transmitted to the central aggregator. The core analytical challenge here is the "utility-privacy frontier." For multi-modal LLMs, which rely on subtle semantic patterns, excessive noise can lead to "catastrophic forgetting" or the loss of long-tail personalization capabilities. System designers must therefore implement dynamic noise-scaling mechanisms that adapt to the sensitivity of the specific modality being processed.

Furthermore, the theoretical framework must consider the socio-technical nature of the "participation economy" in social commerce. Devices participating in the federated network range from high-performance smartphones to low-power internet-of-things devices. This heterogeneity requires an infrastructure that can manage "asynchronous federation," where the global model is updated in a non-blocking manner to accommodate the varying latencies of the distributed nodes. This theoretical foundation sets the stage for a resilient architecture that balances the aggressive pursuit of commercial intelligence with the non-negotiable requirement of user privacy.

**3. System Architecture and Distributed Orchestration**

The proposed architecture for a secure social commerce infrastructure is built upon a tiered distributed fabric. At the "User Tier," individual devices host local instances of a multi-modal encoder, typically a distilled version of a large-scale transformer. These local models perform initial feature extraction and gradient calculation using the user's private data. To manage the computational constraints of mobile hardware, we employ "parameter-efficient fine-tuning" (PEFT) techniques, where only a small subset of the model's weights—such as adapter layers—are updated and transmitted. This significantly reduces the bandwidth and energy requirements for participation, addressing the sustainability concerns inherent in large-scale AI deployment.

The "Orchestration Tier" acts as the central hub for model aggregation and global state management. Unlike traditional servers, this tier does not see or store raw data; instead, it utilizes "Secure Multi-Party Computation" (SMPC) to aggregate noisy gradients from thousands of devices simultaneously. The orchestration layer is also responsible for managing

the global privacy budget. By tracking the cumulative epsilon-delta values across multiple training rounds, the system can ensure that the overall privacy guarantee does not degrade over time. This layer must also implement "robustness-by-design," utilizing anomaly detection algorithms to identify and quarantine malicious devices attempting to perform model poisoning attacks by sending divergent gradients intended to bias the global model.

The third tier is the "Infrastructure Backbone," which provides the high-performance computing resources necessary for periodic global model updates and cross-modal latent space alignment. This tier manages the complex task of ensuring that the linguistic and visual encoders remain synchronized across the distributed network. Given the high throughput required for social commerce platforms, the backbone must utilize hardware-aware scheduling to optimize the placement of aggregation tasks based on network topology and energy availability. This three-tiered approach creates a "zero-trust" environment where data remains local, computation is distributed, and the global intelligence of the platform is an emergent property of the federated collective.

**4. Structural Trade-offs: Privacy, Performance, and Fairness**

Every design choice in a federated social commerce system involves a complex web of structural trade-offs. The most prominent is the conflict between "privacy granularity" and "recommendation utility." Implementing strict differential privacy guarantees inherently introduces noise into the model updates, which can obscure the nuanced signals required for high-accuracy product discovery. In the social commerce context, this can lead to a "homogenization of taste," where the model becomes adept at predicting popular items but fails to surface niche products that align with specific subcultures. This has profound implications for market diversity and small-business visibility on the platform, necessitating a socio-technical approach to "privacy-aware diversity."

Another critical trade-off exists between "system throughput" and "security overhead." Protocols like Secure Multi-Party Computation and homomorphic encryption provide robust defenses against curious aggregators, but they incur significant latency and computational costs. In the high-velocity environment of digital commerce, where the window of opportunity for a recommendation may last only a few seconds, excessive security overhead can lead to "stale intelligence." To manage this, our architecture proposes a "risk-adaptive security" model. In this paradigm, the system dynamically modulates the strength of the privacy-preserving protocols based on the sensitivity of the data and the perceived threat level of the current network environment.

Finally, we must address the "fairness-privacy trade-off." Federated learning can inadvertently exacerbate algorithmic bias. Because the model is trained on decentralized data, it may struggle to represent marginalized or under-represented groups whose local updates are drowned out by the noise of the majority or the differential privacy mechanisms. Protecting privacy should not mean sacrificing equity. We advocate for the inclusion of "fairness-constrained aggregation," where the orchestrator monitors the model's performance across various demographic slices (identified through non-private proxy variables) and

adjusts the federated weighting to ensure that the global model remains inclusive. These structural trade-offs are not merely technical problems but reflect the core values of the socio-technical infrastructure.

**5. Deployment Challenges and Infrastructure Sustainability**

Deploying a federated LLM-based infrastructure for social commerce at a global scale introduces unprecedented operational challenges. One of the primary hurdles is "network non-stationarity." Devices frequently drop in and out of the network, and the quality of connections varies across different geographic regions. A robust federated system must be "fault-tolerant," capable of completing training rounds even when a significant percentage of nodes fail to report their gradients. This requires the implementation of "staleness-aware" aggregation, where the central orchestrator can incorporate delayed updates without destabilizing the global model's convergence.

Sustainability is another central pillar of our deployment framework. The continuous training of multi-modal transformers across millions of devices consumes a staggering amount of energy. To mitigate this, we propose "opportunistic federation," where model updates are only triggered when devices are connected to renewable energy sources and are in a charging state. Furthermore, we emphasize the use of "model compression" and "knowledge distillation" to ensure that the local models are as energy-efficient as possible. By reducing the number of active parameters and utilizing low-precision arithmetic, we can minimize the carbon footprint of the social commerce infrastructure, aligning technical innovation with global environmental imperatives.

The physical deployment must also account for "regional data sovereignty." Different jurisdictions, such as the European Union under GDPR or China under the PIPL, have divergent requirements for data processing and privacy. A global federated system must be "governance-aware," allowing for regional aggregation hubs that ensure local model updates comply with specific legal standards before they are integrated into the global model. This "federated-governance" approach allows social commerce platforms to operate seamlessly across borders while respecting the legal and ethical boundaries of each sovereign territory. Building a resilient and sustainable infrastructure is thus a matter of aligning hardware optimization with global policy compliance.

**6. Socio-Technical Governance and Algorithmic Ethics**

The governance of federated social commerce systems represents a significant shift from traditional corporate oversight. In a centralized system, the platform has total visibility and control over the data and the algorithms. In a federated system, the platform delegates the primary data processing to the users, creating a decentralized power structure. This requires a "collaborative governance" model, where the rules of the federated network are transparent and subject to collective verification. We argue for the implementation of "governance-as-code," where the privacy budgets, fairness constraints, and data-usage policies are embedded directly into the software protocols that govern the federated rounds.

Ethics in social commerce also involves the management of "autonomous influence." Multi-modal LLMs are not just recommendation engines; they are persuasive agents capable of generating synthetic content that can influence consumer behavior. In a privacy-preserving environment, it becomes harder for external auditors to detect if these agents are engaging in predatory practices or spreading misinformation. To address this, we propose the development of "auditable federated logs." These are tamper-proof records of the global model's evolution that, while not revealing individual user data, provide sufficient information for independent third parties to verify that the model's reasoning remains within ethical bounds.

Furthermore, we must address the "accountability gap." If a federated model produces an unfair or harmful outcome, determining the locus of responsibility is complex. Is it the fault of the platform that designed the aggregator, or the collective of users whose local updates shaped the global state? We advocate for a "systemic accountability" framework, where the platform takes ultimate responsibility for the behavior of the emergent global model. This involves maintaining a "human-in-the-loop" oversight committee that can override the federated training process if the global model exhibits signs of ethical drift. By integrating these governance mechanisms, we ensure that the social commerce infrastructure remains a tool for human empowerment rather than automated exploitation.

**7. Policy Implications and Global Regulatory Landscapes**

The rise of privacy-preserving federated learning necessitates a fundamental re-evaluation of global privacy policies. Current regulations are largely focused on "data-at-rest" and "data-in-transit," but they often lack clear guidelines for "data-in-use" within a federated ecosystem. Policy-makers must now define what constitutes a "privacy-preserving update." If a model gradient can be used to reconstruct a user's original data, even if the raw data never left the device, does that constitute a data breach? This ambiguity creates a "regulatory vacuum" that can hinder the adoption of secure infrastructures. We argue for a shift toward "output-based regulation," where the legal focus is on the formal privacy guarantees provided by mechanisms like differential privacy rather than the physical location of the data.

Another critical policy dimension is the "interoperability of privacy standards." As social commerce platforms become increasingly interconnected, there is a need for a common language of privacy that allows federated models to share insights across different ecosystems without compromising security. International standards bodies must collaborate to define a "global privacy exchange" framework, ensuring that the benefits of collective intelligence are accessible to all while maintaining local data sovereignty. This is particularly important for preventing "digital protectionism," where nations use privacy regulations as a tool for excluding foreign commercial platforms.

Finally, we must consider the "digital equity" implications of privacy-preserving technologies. The hardware requirements for participating in a federated LLM network may exclude users in developing regions who lack access to high-performance mobile devices. Policy-makers should encourage "inclusive innovation," providing incentives for platforms to develop lightweight, low-power federated protocols that can run on older hardware. This ensures that

the transition to secure social commerce does not create a new "privacy-poverty" divide. By aligning technical architecture with inclusive public policy, we can build a global commerce infrastructure that is truly resilient and equitable.

**8. Forward-Looking Perspectives and Emerging Frontiers**

As we look toward the next decade, the evolution of secure social commerce will be driven by the integration of even more advanced computational paradigms. One such frontier is "quantum-safe federated learning." As quantum computing becomes a reality, traditional cryptographic protocols used in SMPC and homomorphic encryption may become vulnerable. Our infrastructure must be "future-proofed" by incorporating lattice-based cryptography and other post-quantum algorithms to ensure long-term data security. This is especially critical for social commerce platforms that store transactional histories spanning decades.

Another emerging frontier is the move toward "fully decentralized autonomous commerce" (FDAC). In this vision, the central orchestrator is replaced by a decentralized ledger, such as a blockchain, that manages the federated aggregation process through smart contracts. This would eliminate the "central point of trust" entirely, creating a truly peer-to-peer commercial ecosystem. While this offers the ultimate in data autonomy, it introduces massive challenges regarding system throughput and governance. The socio-technical challenge of the future will be to find the optimal balance between the efficiency of centralized orchestration and the security of decentralized ledgers.

Finally, we anticipate a shift from "intent discovery" to "value alignment." Future social commerce agents, powered by multi-modal LLMs, will not just try to sell a product but will attempt to align their recommendations with the user's long-term well-being and ethical values. This "pro-social AI" requires a sophisticated understanding of human psychology and social dynamics, necessitating an even deeper interdisciplinary dialogue between engineers, sociologists, and philosophers. By building the foundations of a secure and privacy-preserving infrastructure today, we are preparing the groundwork for a future where commerce is not just an exchange of goods, but a reflection of our collective values.

**9. Conclusion**

The development of a secure social commerce infrastructure is a landmark challenge at the intersection of large-scale systems engineering and socio-technical governance. By synergizing privacy-preserving federated learning, multi-modal large language models, and differential privacy, we have proposed a framework that addresses the critical vulnerabilities of centralized data architectures. Our analysis has demonstrated that while the path to decentralization involves complex structural trade-offs between utility, performance, and fairness, it offers a resilient blueprint for a digital marketplace that respects individual autonomy.

The success of this infrastructure depends on our ability to look beyond technical optimization and embrace the broader dimensions of sustainability, ethics, and policy. As social commerce continues to expand into every corner of human interaction, the focus must

remain on building systems that are transparent, inclusive, and accountable. Through interdisciplinary collaboration and a commitment to "privacy-by-design," we can ensure that the next generation of social commerce serves as a force for positive socio-economic transformation. The architecture proposed in this paper is a first step toward that future, providing the technical and conceptual tools necessary to secure the digital fabric of our society.

**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 2016 ACM SIGSAC Conference on Computer and Communications Security, 308-318.

2. Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and creates labor. Journal of Economic Perspectives, 33(2), 3-30.

3. Agrawal, R., & Srikant, R. (2000). Privacy-preserving data mining. ACM Sigmod Record, 29(2), 439-450.

4. Bonawitz, K., et al. (2017). Practical secure aggregation for privacy-preserving machine learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 1175-1191.

5. Bommasani, R., et al. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258.

6. Brown, T., et al. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877-1901.

7. Cho, J. H., et al. (2020). Toward a sustainable and resilient social commerce infrastructure: A systems perspective. IEEE Transactions on Systems, Man, and Cybernetics: Systems, 50(12), 4987-5002.

8. Dwork, C. (2008). Differential privacy: A survey of results. International Conference on Theory and Applications of Models of Computation, 1-19.

9. Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing, 169-178.

10. Hard, A., et al. (2018). Federated learning for mobile keyboard prediction. arXiv preprint arXiv:1811.03604.

11. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770-778.

12. Kairouz, P., et al. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1-2), 1-210.

13. Kaplan, J., et al. (2020). Scaling laws for neural language models. arXiv preprint arXiv:2001.08361.

14. Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492.

15. Liu, T. (2026). A Comparative Study of Transformer-Based and Classical Models for Financial Time-Series Forecasting. Journal of Risk and Financial Management, 19(3), 203.

16. 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.

17. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Artificial Intelligence and Statistics, 1273-1282.

18. Mohri, M., Sivek, G., & Suresh, A. T. (2019). Agnostic federated learning. International Conference on Machine Learning, 4615-4625.

19. Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. 2008 IEEE Symposium on Security and Privacy, 111-125.

20. O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown.

21. Pasquale, F. (2015). The Black Box Society: The Secret Algorithms That Control Money and Information. Harvard University Press.

22. Radford, A., et al. (2021). Learning transferable visual models from natural language supervision. International Conference on Machine Learning, 8748-8763.

23. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. 2017 IEEE Symposium on Security and Privacy (SP), 3-18.

24. Stoica, I., et al. (2017). Ray: A distributed framework for emerging AI applications. 13th USENIX Symposium on Operating Systems Design and Implementation, 561-577.

25. Vaswani, A., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.

26. Wu, S., et al. (2023). BloombergGPT: A large language model for finance. arXiv preprint arXiv:2303.17564.

27. Yi, X. (2026). A Federated and Differentially Private Incentive–Marketing Framework for Privacy-Preserving Cross-Channel Measurement in AI-Powered Digital Commerce.

28. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology (TIST), 10(2), 1-19.

29. Zaharia, M., et al. (2012). Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. 9th USENIX Symposium on Networked Systems Design and Implementation, 15-28.

30. Zhang, C., et al. (2021). Survey on differential privacy with machine learning. IEEE Access, 9, 13329-13350.

31. Zhao, Y., et al. (2018). Federated learning with non-iid data. arXiv preprint arXiv:1806.00582.

32. Zuboff, S. (2019). The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. PublicAffairs.

33. Mo, F., Haddadi, H., Katiyar, K., Ansari, R., & Chuah, C. N. (2021). PPFL: Privacy-preserving federated learning with trusted execution environments. Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services, 94-108.

34. Wang, J., et al. (2021). A field guide to federated optimization. arXiv preprint arXiv:2107.06917.

35. Rothchild, D., et al. (2020). FetchSGD: Communication-efficient federated learning with sketching. Proceedings of the 37th International Conference on Machine Learning.

36. Kairouz, P., et al. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1-2), 1-210.