Federated HyperFusion: Privacy-Preserving Collaborative Learning for Cross-Platform Hyperspectral and LiDAR Data Fusion

Haolei Shen
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
haoleishen@ucf.edu

Rorge Darr
Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO, USA.
carr223@missouri.edu

Ghomas M. Tillis
School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
thomasmail@uc.edu

Gnkit Bgarwal
Department of Computer Science, University of Houston, Houston, TX, USA.
agarwal293@uh.edu

Abstract

The integration of hyperspectral and LiDAR data has proven essential for high-resolution environmental monitoring, urban planning, and precision agriculture. However, real-world deployments increasingly involve multiple platforms—satellites, UAVs, and ground sensors—each governed by distinct institutional and privacy constraints. Traditional centralized fusion approaches require raw data to be transmitted to a single server, raising significant privacy, regulatory, and bandwidth concerns. This paper introduces Federated HyperFusion, a system-level framework that enables collaborative cross-platform learning for joint hyperspectral and LiDAR fusion without exposing sensitive raw data. We design a federated architecture where distributed clients train local models on their respective data and share only encrypted gradient updates with a central aggregation server. The framework incorporates secure aggregation mechanisms, differential privacy noise calibration, and communication-efficient compression strategies to balance accuracy, privacy, and operational cost. We analyze the structural trade-offs introduced by heterogeneous sensor resolutions, varying label availability, and non-IID data distributions across platforms. The paper further examines governance considerations such as auditability, model accountability, and data sovereignty in multi-stakeholder environments. Infrastructure aspects including edge deployment, real-time inference constraints, and energy sustainability are discussed. We also address fairness and robustness challenges arising from domain shift and biased sensor coverage. Policy implications for cross-jurisdictional data sharing and compliance with emerging privacy regulations are explored. Through this system-level perspective, Federated HyperFusion offers a viable path toward privacy-preserving, collaborative remote sensing analytics that respects institutional boundaries while achieving high-quality fused representations.

Keywords

federated learning, hyperspectral imaging, LiDAR, data fusion, privacy preservation, collaborative learning, remote sensing, secure aggregation, differential privacy, multi-platform systems.

1. Introduction

The convergence of hyperspectral and LiDAR sensing modalities has unlocked unprecedented capabilities in land cover classification, vegetation health assessment, and three-dimensional urban modeling. Hyperspectral imagers capture hundreds of narrow spectral bands, enabling detailed material identification, while LiDAR provides precise topographic and structural information. Jointly processing these heterogeneous data streams yields richer feature representations than either modality alone [1,2]. However, as remote sensing platforms proliferate—from government-operated satellites to commercial drone fleets and ground-based sensor networks—the data itself becomes distributed across administrative, organizational, and even national boundaries. Data privacy concerns, intellectual property restrictions, and regulatory frameworks such as the General Data Protection Regulation (GDPR) in Europe or the California Consumer Privacy Act (CCPA) in the United States increasingly prohibit the aggregation of raw sensor data into centralized repositories [3]. This creates a fundamental tension between the desire for high-accuracy fusion models and the legal and ethical imperative to protect sensitive geospatial information.

Recent advances in federated learning offer a promising resolution to this tension. In federated learning, multiple clients collaboratively train a shared model without ever exchanging their local data; only model updates, such as gradients, are communicated to a central server, often with cryptographic or statistical privacy guarantees [4,5]. While federated learning has been extensively studied for mobile keyboard prediction, healthcare analytics, and financial modeling, its application to remote sensing data fusion remains nascent. The unique characteristics of hyperspectral and LiDAR data—high dimensionality, spatial autocorrelation, heterogeneous resolutions, and non-IID (non-identically and independently distributed) capture conditions—introduce novel challenges that go beyond those addressed in conventional federated settings [6,7].

In this paper, we propose Federated HyperFusion, a comprehensive system-level framework designed to enable privacy-preserving, collaborative learning for cross-platform hyperspectral and LiDAR data fusion. Rather than focusing on algorithmic improvements to the underlying fusion model, we emphasize the architectural choices, governance mechanisms, infrastructure considerations, and policy implications that determine the feasibility and sustainability of such a system in practice. We argue that the successful deployment of federated fusion requires a holistic view that integrates technical robustness with institutional trust and regulatory compliance. Our discussion covers secure aggregation protocols that prevent the server or any adversary from reconstructing raw data from gradient updates; differential privacy mechanisms that calibrate noise to balance privacy loss and model utility; communication-efficient strategies that reduce the bandwidth burden on resource-constrained edge devices; and techniques to handle the statistical heterogeneity that arises when different platforms capture data under varying illumination, weather, and sensor calibration conditions.

We further delve into the structural trade-offs inherent in federated fusion. For instance, hyperspectral sensors often have different spectral resolutions and number of bands, while LiDAR point densities vary across platforms. Fusing these modalities in a federated setting

requires careful alignment of feature spaces and labeling conventions. The presence of limited or noisy ground truth annotations at some clients introduces additional complexity for supervised learning. Governance structures must define which parties have access to the aggregated model, how model drift is detected and corrected, and what audit trails are maintained for accountability. Infrastructure decisions involve whether to deploy the aggregation server in the cloud, at the edge, or in a hybrid configuration, and how to manage the energy consumption of continuous model updates over satellite links or low-power ground sensors. Finally, fairness concerns arise when certain geographic regions or sensor owners are underrepresented in the training distribution, leading to biased fused outputs. We explore these issues with reference to real-world use cases and cross-domain comparisons, and we outline a forward-looking research agenda that integrates technical innovation with responsible system design.

2. Background and Motivation

The fusion of hyperspectral and LiDAR data has been a active research area for over a decade, with studies demonstrating improvements in classification accuracy for urban mapping, forest inventory, and mineral exploration [1,2]. Early fusion methods operated at the pixel level, concatenating spectral vectors with elevation features derived from LiDAR digital surface models. More recent deep learning approaches employ convolutional neural networks (CNNs) and graph neural networks (GNNs) to extract hierarchical spatial-spectral-structural features [8,9]. However, these models are typically trained on large, centrally hosted datasets such as the Houston2013 or MUUFL Gulfport benchmarks, which are collected under controlled conditions by single institutions. In practice, a government environmental agency may operate a satellite with hyperspectral capabilities, while a private agricultural company deploys UAVs with LiDAR over its own fields, and a university maintains ground-based spectroradiometers. Each entity is reluctant to share raw data due to proprietary interests, national security concerns, or privacy regulations that protect landscape features that could reveal critical infrastructure or population patterns [10].

Federated learning was originally proposed to train language models on user-typed text without uploading personal messages to a server [4]. Since then, it has been extended to computer vision, medical imaging, and Internet-of-Things (IoT) settings. The core idea is to keep data on each client and only share model updates, which are aggregated via algorithms such as Federated Averaging (FedAvg) [4]. To provide stronger privacy guarantees, secure multi-party computation (SMC) techniques, particularly secure aggregation, ensure that the server sees only the sum of client updates rather than individual contributions [5]. Additionally, differential privacy (DP) injects calibrated noise into either the client or server updates to bound the risk of membership inference [11]. These building blocks form the foundation of Federated HyperFusion.

Nevertheless, applying federated learning to hyperspectral-LiDAR fusion introduces several domain-specific complications. Hyperspectral images are high-dimensional tensors that demand large model capacities; transmitting full gradient updates can be prohibitively expensive in terms of bandwidth and energy, especially for low-earth-orbit satellites with intermittent connectivity [12]. LiDAR point clouds are sparse and irregular, requiring specialized network architectures such as PointNet or voxel-based approaches [13]. The fusion process itself requires aligning the two modalities, which may have different coordinate systems, acquisition times, and spatial footprints. When these modalities originate from different platforms, the alignment problem is compounded by the fact that the same

geographic area may be captured at different temporal resolutions and under different atmospheric conditions [14]. The reference work by Long et al. (2026) on weak-signal representation learning for hyperspectral unmixing underscores the importance of handling subtle spectral variations that can be easily overwhelmed by noise or domain shift [14]. Their WS-Net architecture employs state-space modeling and attention fusion to extract robust features, providing a relevant deep learning foundation upon which a federated system could be built.

Moreover, the labeling process for fused remote sensing data is expensive and often inconsistent across institutions. A client may have access to high-quality ground truth for a small region, while another client has coarse labels derived from publicly available land cover maps. This label heterogeneity, combined with non-IID data distributions across platforms, can cause the federated model to converge slowly or fail to generalize to unseen regions [15]. The federated fusion system must therefore incorporate techniques such as personalized federated learning, where each client maintains a local model component, or meta-learning approaches that quickly adapt to new platforms [16]. These considerations motivate the architectural decisions we detail in the next section.

3. Federated HyperFusion Architecture

The proposed Federated HyperFusion architecture consists of three tiers: the client tier, the aggregation tier, and the governance tier. At the client tier, each participating platform operates a local fusion model that ingests its own hyperspectral and LiDAR data, extracts joint features, and produces a classification or regression output. To accommodate heterogeneous sensing configurations, the local model is designed with a modular structure: a modality-specific encoder that captures the unique characteristics of each sensor type, followed by a fusion module that combines the encoded representations through cross-attention mechanisms or bilinear pooling. The fusion module can be pre-initialized with weights from a publicly available pretrained model (e.g., a lightweight version of WS-Net [14]) and then fine-tuned locally. The client also maintains a small validation set, drawn from its own data, to monitor local performance without leaking information.

The aggregation tier is responsible for receiving encrypted gradient updates from clients, performing secure aggregation, and distributing the updated global model parameters back to the clients. We employ the secure aggregation protocol of Bonawitz et al. [5] which uses secret sharing and has a robust dropout handling mechanism to tolerate clients that fail to complete a round. To further protect against gradient leakage attacks, we add differential privacy noise at the client level before encryption. The noise magnitude is calibrated using the moments accountant method [11] to achieve a desired privacy budget (measured by epsilon-delta parameters). The server itself does not have access to raw gradients or intermediate feature representations; it only sees the noisy aggregated sum. This ensures that even a compromised server cannot reconstruct individual client data or infer sensitive patterns such as the presence of military installations or endangered species habitats.

The governance tier provides the policies, audit logs, and access controls that regulate the participation of stakeholders. Each client registers with a unique digital identity and agrees to a data-sharing contract specifying the purpose of model training, the duration of participation, and the permitted use of the resulting fused model. The governance layer also maintains a public ledger (which can be a permissioned blockchain) that records each training round’s participation, aggregated model hash, and privacy budget consumption. This ledger enables later audits to verify that privacy guarantees were respected and that no client contributed data

outside the agreed terms. Additionally, the governance tier manages incentives: clients that contribute high-quality updates (as measured by improvements on a held-out benchmark dataset) may receive access to the final fused model or other benefits, thereby encouraging sustained participation.

Communication efficiency is a critical concern, particularly for clients with limited bandwidth or intermittent connectivity. We adopt gradient compression techniques, such as random sparsification and quantization, to reduce the size of each update. Specifically, each client sorts its gradient parameters by magnitude, selects only the top k percent of the largest magnitudes, and sends those along with their indices, randomly zeroing out the rest. The server aggregates only the received sparse updates; because different clients may select different indices, the aggregate can be reconstructed using the index information. Combined with stochastic gradient quantization to 8-bit integers, this reduces the per-round communication volume by over 90% without significant accuracy loss [17]. For clients on satellite platforms that have short transmission windows, a scheduling algorithm prioritizes clients with the largest expected contribution to loss reduction, thereby maximizing the utility of each communication round.

4. Structural Trade-offs and Governance

The design of Federated HyperFusion involves several interconnected trade-offs that must be carefully balanced. The most fundamental trade-off is between privacy and utility. Differential privacy inherently degrades model accuracy, as the added noise reduces the signal-to-noise ratio of the aggregated gradient. The choice of privacy budget epsilon directly determines the extent of this degradation; a lower epsilon provides stronger privacy but may lead to a fused model that performs no better than a single-modality baseline. To mitigate this, we can adopt adaptive clipping and noise scaling that adjusts the noise level per parameter based on its sensitivity in the local loss landscape [18]. Furthermore, the trade-off is not uniform across all applications: a land cover classification task that requires fine-grained distinction between similar vegetation species may be more sensitive to noise than a binary urban-rural classification. The governance framework should allow stakeholders to negotiate a privacy budget that aligns with the sensitivity of their data and the required accuracy threshold.

Another major trade-off concerns model personalization versus global generalization. In a classical federated setting, the goal is to learn a single global model that performs well on all clients' data. However, due to spectral shifts caused by different sensor calibrations, atmospheric conditions, and seasonal effects, a single fusion model may underperform on some platforms. Personalized federated learning strategies, such as model-agnostic meta-learning (MAML) or clustering-based approaches, allow each client to maintain a local fine-tuned head while sharing a common feature extractor [16]. This increases the storage and computation burden on clients but can significantly improve per-client accuracy. The governance tier must decide whether to mandate a common global model for the sake of interoperability or to allow personalization, which may lead to inconsistent model behavior across platforms. For cross-border environmental monitoring agreements, a global standard may be necessary for policy enforcement, whereas for commercial agricultural analytics, personalization may be acceptable.

Data sovereignty is a central governance issue. Many jurisdictions consider geospatial data, especially high-resolution imagery, to be of strategic importance. Regulations such as the EU's Directive on the re-use of public sector information and the Indian Space Research

Organization's data policy impose restrictions on the export of raw remote sensing data. Federated HyperFusion, by keeping data local, naturally aligns with these sovereignty requirements. However, the aggregated model itself may still leak information about the training data through membership inference attacks or model inversion. The governance tier must therefore enforce post-training access controls: the final fused model should only be made available to clients that have contributed and have passed a review of their intended use. Additionally, the model can be subjected to a "blindness" test that checks whether it can reconstruct any training sample beyond a certain resolution threshold. These measures ensure that the system does not inadvertently become a conduit for data laundering.

5. Infrastructure and Deployment

Deploying Federated HyperFusion in real-world remote sensing networks requires careful consideration of infrastructure constraints. The central aggregation server can be hosted in a cloud environment (e.g., AWS GovCloud for sensitive applications) or on a dedicated private server located at a neutral facility. For latency-sensitive applications such as disaster response after a hurricane, edge aggregation may be preferable: a local server deployed within a limited geographic region collects updates from UAVs and ground sensors, performs fast aggregation, and updates local models without waiting for a round trip to a remote cloud. This reduces the time to obtain an updated fused map from hours to minutes. A hybrid architecture that uses edge aggregation for time-critical rounds and cloud aggregation for periodic global retraining offers a practical compromise.

Energy sustainability is another crucial infrastructure aspect. Satellites and UAVs have limited power budgets; transmitting large gradient updates consumes energy that could otherwise be used for sensing or flight endurance. Compression and sparsification help, but the aggregation server can also employ active learning to reduce the number of training rounds. For instance, the server can compute the variance of client updates and halt training when convergence is detected, avoiding unnecessary communication [19]. Alternatively, clients can perform local training for multiple epochs before uploading updates, reducing the frequency of communication at the cost of increased local computation. The trade-off between communication and computation must be tuned based on the specific platform's energy profile.

Network resilience is vital, as remote sensing platforms often operate in areas with unreliable connectivity. The secure aggregation protocol of Bonawitz et al. [5] is designed to handle client dropout; we extend this by allowing clients to cache their updates and retransmit during the next available window. However, large dropout rates can degrade the quality of the aggregated model. To mitigate this, the governance tier can assign backup clients or employ a staleness-aware aggregation that downweights outdated updates. In the long term, the infrastructure should support a federated ecosystem where clients can join and leave dynamically, akin to a crowdsourced remote sensing network. This requires scalable identity management and a reputation system to prevent malicious clients from poisoning the model with corrupted gradients.

6. Robustness and Fairness

Robustness to adversarial attacks and domain shift is a key challenge for any distributed learning system. In Federated HyperFusion, malicious clients could attempt to inject backdoors that cause the fused model to misclassify specific land cover types (e.g., labeling a military installation as forest) or to perform targeted attacks that degrade performance on a

particular platform. Defense strategies include Byzantine-robust aggregation rules such as coordinate-wise median or trimmed mean [20], which discard extreme gradient values. The governance tier can also enforce that clients provide cryptographic proofs of their model updates, such as zero-knowledge proofs of the correctness of their local gradient computation, though this adds computational overhead. Additionally, the differential privacy mechanism provides inherent robustness against outliers because the noise level can mask minority gradient contributions.

Fairness in cross-platform fusion raises profound concerns. If a global model is trained predominantly on data from well-instrumented regions in developed countries, it may perform poorly on data from developing countries where sensors have lower spectral resolution or different atmospheric conditions. This could lead to inequitable outcomes in applications such as agricultural subsidy allocation or flood risk assessment. Federated HyperFusion can address fairness by weighting client updates according to the geographic coverage or population served, or by employing a minimax objective that minimizes the worst-case loss across clients [21]. The governance tier should mandate that the final model be evaluated on a diverse test set that covers all participating geographic regions and sensor types. If disparities are found, targeted fine-tuning or additional data collection can be required.

7. Policy Implications

The deployment of a collaborative federated fusion system has significant policy implications. First, it enables the sharing of high-value fused intelligence without violating data localization laws. Policymakers can encourage the adoption of such systems by creating regulatory sandboxes that temporarily exempt collaborative learning from some data transfer restrictions, provided that privacy guarantees are certified. Second, the system creates new accountability challenges: if a fused model misclassifies a hazardous material spill, who is liable—the clients that trained the model, the server that aggregated the updates, or the operator that deployed the model? The governance tier must include clear contractual terms specifying liability distribution, perhaps modeled after open-source software licensing frameworks. Third, the aggregation of data from multiple platforms raises antitrust concerns if the resulting fused model gives a dominant player an unfair advantage. Competition authorities may need to regulate the access to aggregated models to prevent market concentration.

Furthermore, environmental policies that rely on accurate remote sensing data (e.g., carbon stock estimation, deforestation monitoring) could be strengthened by Federated HyperFusion. International agreements such as the Paris Agreement require transparent and verifiable reporting of greenhouse gas emissions. A federated system that allows nations to contribute their own sensor data without revealing sensitive military or economic information could facilitate more accurate global inventories. However, the trustworthiness of the aggregated model must be independently validated by neutral third parties. A suggested policy mechanism is to require that the final fused model be open-sourced (with privacy guarantees verified) so that any signatory can audit its performance on their own local data.

8. Conclusion

Federated HyperFusion provides a comprehensive framework for privacy-preserving collaborative learning across distributed hyperspectral and LiDAR platforms. By integrating secure aggregation, differential privacy, communication compression, and personalized federated learning, the system addresses the core technical challenges of sensor heterogeneity, non-IID data, and limited bandwidth. Equally important, it embeds governance structures that

respect data sovereignty, ensure accountability, and promote equitable access to fused intelligence. The infrastructure discussion highlights the feasibility of deploying the system on edge devices, satellites, and cloud servers, with careful attention to energy and connectivity constraints. Robustness and fairness considerations lead to adaptive aggregation rules and fairness-aware objectives that prevent systemic biases. Finally, the policy dimensions show that such a system can align with existing regulations while enabling new forms of international collaboration for environmental monitoring and disaster response. Future work should focus on empirical validation with real cross-platform datasets, integration of the WS-Net architecture [14] within the federated pipeline, and the development of incentive mechanisms that sustain long-term participation. As remote sensing becomes ever more distributed and privacy-conscious, the principles of federated fusion outlined here will be critical to unlocking the full potential of multi-modal Earth observation.

References

1. Yokoya, N., Grohnfeldt, C., & Chanussot, J. (2017). Hyperspectral and multispectral data fusion: A comparative review. IEEE Geoscience and Remote Sensing Magazine, 5(2), 28–41.

2. Ghassemian, H. (2016). A review of remote sensing image fusion methods. Information Fusion, 32, 75–89.

3. European Parliament and Council. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation). Official Journal of the European Union, L 119, 1–88.

4. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 1273–1282.

5. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS), 1175–1191.

6. 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.

7. Sattler, F., Wiedemann, S., Müller, K. R., & Samek, W. (2020). Robust and communication-efficient federated learning from non-i.i.d. data. IEEE Transactions on Neural Networks and Learning Systems, 31(9), 3400–3413.

8. Hong, D., Yokoya, N., Chanussot, J., & Zhu, X. X. (2020). Learning to propagate labels on graphs: An iterative multitask regression framework for semi-supervised hyperspectral dimensionality reduction. ISPRS Journal of Photogrammetry and Remote Sensing, 158, 35–49.

9. Li, W., Fu, H., Yu, L., & Gong, P. (2019). Deep learning-based fusion of hyperspectral and LiDAR data for land cover classification. IEEE Geoscience and Remote Sensing Letters, 16(2), 276–280.

10. Krug, M., & Lauf, T. (2020). Privacy issues in remote sensing: A survey. Remote Sensing, 12(22), 3740.

11. 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 (CCS), 308–318.

12. Diao, X., Chen, J., & Li, B. (2022). Federated learning for satellite imagery: Challenges and opportunities. IEEE Geoscience and Remote Sensing Magazine, 10(3), 8–24.

13. Qi, C. R., Yi, L., Su, H., & Guibas, L. J. (2017). PointNet++: Deep hierarchical feature learning on point sets in a metric space. In Advances in Neural Information Processing Systems (NeurIPS), 5099–5108.

14. Long, Z., Zia, A., Fu, G., Rolland, V., & Zhou, J. (2026). WS-Net: Weak-Signal Representation Learning and Gated Abundance Reconstruction for Hyperspectral Unmixing via State-Space and Weak Signal Attention Fusion. arXiv preprint arXiv:2603.09037.

15. Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. In Proceedings of Machine Learning and Systems (MLSys), 2, 429–450.

16. Fallah, A., Mokhtari, A., & Ozdaglar, A. E. (2020). Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. In Advances in Neural Information Processing Systems (NeurIPS), 33, 3557–3568.

17. Stich, S. U., Cordonnier, J. B., & Jaggi, M. (2018). Sparsified SGD with memory. In Advances in Neural Information Processing Systems (NeurIPS), 4447–4458.

18. Pichapati, V., Raskar, R., & Thakurta, A. (2019). AdaCliP: Adaptive clipping for private SGD. arXiv preprint arXiv:1908.07643.

19. Wang, J., Charles, Z., Xu, Z., Joshi, G., McMahan, H. B., & al. (2019). A field guide to federated optimization. arXiv preprint arXiv:1910.11852.

20. Blanchard, P., Guerraoui, R., & Stainer, J. (2017). Machine learning with adversaries: Byzantine tolerant gradient descent. In Advances in Neural Information Processing Systems (NeurIPS), 119–129.

21. Mohri, M., Sivek, G., & Suresh, A. T. (2019). Agnostic federated learning. In Proceedings of the 36th International Conference on Machine Learning (ICML), 4615–4625.