

# Federated Dual-Attention Segmentation for Privacy-Preserving Multi-Center Pulmonary Nodule Analysis in Computed Tomography Imaging

Paul Yowers

School of Electrical Engineering and Computer Science, Oregon State University, Corvallis,  
OR, USA.

powers1999@oregonstate.edu

Erendan Perry

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL,  
USA.

berry407@uab.edu

## Abstract

The increasing adoption of computed tomography (CT) imaging for lung cancer screening has generated vast repositories of pulmonary nodule data across multiple clinical institutions. Centralized aggregation of such sensitive medical data poses significant privacy, legal, and operational challenges, while the heterogeneity of imaging protocols and patient populations across centers complicates the development of robust segmentation models. This paper presents a system-level framework that integrates federated learning with a dual-attention segmentation architecture to enable privacy-preserving, multi-center analysis of pulmonary nodules. The proposed approach decouples model training from direct data sharing, allowing institutions to collaboratively refine a shared model while retaining raw images on-site. The dual-attention mechanism, combining spatial and channel attention, enhances the model's ability to capture subtle nodule features and reduces false positives, a critical requirement for clinical deployment. This paper examines the structural trade-offs inherent in federated systems, including communication efficiency, model convergence, and differential privacy guarantees. It further discusses the infrastructure and governance necessary for sustaining such a framework across heterogeneous healthcare networks, addressing robustness to domain shifts, fairness across demographic groups, and policy implications for regulatory compliance. By analyzing real-world deployment scenarios and cross-domain comparisons with centralized and other distributed approaches, the paper highlights how federated dual-attention segmentation can balance diagnostic accuracy with patient privacy. The study also identifies open challenges in scalability, adversarial robustness, and equitable performance, and proposes forward-looking strategies for integrating emerging privacy technologies and standardized data formats. This work aims to provide a comprehensive reference for researchers and practitioners designing next-generation collaborative medical imaging systems.

## Keywords

federated learning, dual-attention, pulmonary nodule segmentation, privacy-preserving, multi-center analysis, computed tomography, healthcare infrastructure.

## 1. Introduction

The clinical analysis of pulmonary nodules in CT imaging is a cornerstone of early lung cancer detection, yet the development of high-performance segmentation algorithms has historically relied on large, centralized datasets. In practice, medical data are siloed across institutions due to stringent privacy regulations such as HIPAA and GDPR, as well as competitive and ethical considerations. Centralizing these datasets not only raises legal barriers but also introduces systematic biases, as single-institution data often lack the diversity of scanner types, acquisition parameters, and patient demographics needed for generalizable models. Federated learning has emerged as a promising paradigm that allows multiple hospitals to collaboratively train a shared deep learning model without transferring raw imaging data. Instead, only model updates are exchanged, preserving data locality. This architectural shift, however, introduces new challenges: communication overhead, statistical heterogeneity, and vulnerability to privacy leakage from gradient updates. Simultaneously, the segmentation of pulmonary nodules demands high sensitivity and specificity, particularly for small or low-contrast nodules where conventional convolutional approaches may fail. Attention mechanisms have been shown to enhance feature representation by focusing on diagnostically relevant regions. By combining a dual-attention module that operates on both spatial and channel dimensions with a federated training protocol, we propose a system that addresses both privacy and performance requirements. This paper provides a system-level analysis of such an integrated framework, focusing on the architectural decisions, operational trade-offs, and broader socio-technical implications for multi-center deployment.

## **2. Related Work**

Prior research in medical image segmentation has largely followed two parallel trajectories: improving model accuracy through advanced neural architectures and addressing data privacy through distributed learning. On the architectural side, U-Net variants and their derivatives have become the de facto standard for biomedical segmentation, with subsequent enhancements incorporating attention gates, residual connections, and multi-scale feature aggregation. Dual-attention mechanisms that separately model spatial and channel dependencies have been effective in capturing long-range contextual information while suppressing irrelevant activations. In the domain of pulmonary nodule analysis, these techniques have demonstrated improved detection of subtle lesions. On the privacy-preserving front, federated learning was originally proposed for mobile keyboard prediction and has since been adapted to healthcare, where it enables collaborative model training without exposing patient data. Early work focused on homogeneous networks and synchronous updates, but real-world medical settings require handling of non-IID data distributions, straggler nodes, and variable connectivity. Subsequent developments introduced robust aggregation algorithms, differential privacy guarantees, and secure multi-party computation primitives. However, few studies have systematically examined the integration of advanced attention-based segmentation models within a federated training loop, particularly for the complex task of pulmonary nodule segmentation across multiple institutions with differing protocols. This paper bridges that gap by proposing a federated dual-attention segmentation framework and discussing its system-level implications.

## **3. Federated Learning Architecture for Multi-Center Segmentation**

The proposed federated architecture consists of a central aggregation server and multiple participating clinical sites, each hosting local CT datasets and computational resources. In each training round, the server distributes the current global model weights to all sites. Each site performs several epochs of local training on its own data using the dual-attention

segmentation network, after which it sends the updated model gradients or weights back to the server. The server then aggregates these updates using the Federated Averaging algorithm, which computes a weighted average based on the size of each site's local dataset. This cycle repeats until convergence. A critical design choice is the communication frequency: more frequent rounds allow faster convergence but incur higher bandwidth costs and increase the risk of privacy leakage from gradient information. In practice, institutions with limited network capacity may perform more local iterations before reporting updates, creating a trade-off between statistical efficiency and communication overhead. Moreover, the non-IID nature of medical data across sites—due to differences in scanner manufacturers, slice thickness, reconstruction kernels, and nodule prevalence—can cause model divergence or slow convergence. To mitigate this, the server can employ adaptive weighting or proximal terms that penalize large deviations from the global model. Another important consideration is the inclusion of differential privacy mechanisms, where the server adds calibrated noise to the aggregated updates before distributing them. This ensures that individual site contributions cannot be reverse-engineered, but at the cost of reduced model accuracy. The system architect must carefully tune the privacy budget to meet regulatory standards while maintaining clinically acceptable performance. Finally, the infrastructure must support secure communication channels, authentication protocols, and audit logging to comply with healthcare data governance policies.

#### **4. Dual-Attention Mechanism for Nodule Segmentation**

At the core of the local model is a dual-attention segmentation network that enhances the representation of pulmonary nodules in CT slices. Spatial attention computes a weight map that highlights regions of interest by modeling pairwise relationships between all positions, thereby capturing long-range dependencies that are essential for distinguishing nodules from adjacent vasculature or other structures. Channel attention adaptively recalibrates feature maps by emphasizing informative channels and suppressing less useful ones, which is particularly beneficial when multiple image modalities or contrast phases are present across centers. The dual-attention module can be inserted at multiple stages of an encoder-decoder architecture, such as U-Net, to refine both low-level and high-level features. For pulmonary nodules, which often appear as small, low-density regions against a complex background of lung parenchyma, the combination of spatial and channel attention has been shown to reduce false positives and improve boundary delineation. In the federated setting, the dual-attention mechanism also provides robustness against domain shifts: because attention weights are learned from the data distribution, the model can adapt to variations in nodule appearance across institutions without requiring explicit domain adaptation. However, the additional parameters introduced by attention modules increase computational overhead, which may be a constraint for sites with limited GPU resources. The system designer must weigh the accuracy gains against the increased training time and memory usage. Furthermore, the effectiveness of dual-attention depends on the quality and diversity of the local data; sites with very few nodule examples may not learn meaningful attention maps, potentially harming the global model. Strategies such as data augmentation, semi-supervised learning, or pre-training on public datasets can help address this imbalance.

#### **5. Privacy-Preserving Infrastructure and Governance**

Deploying a federated dual-attention system within a multi-center healthcare environment requires a robust infrastructure that addresses not only technical privacy mechanisms but also organizational governance. The first layer of protection is data locality: raw CT images never

leave the hospital premises. This eliminates many legal hurdles associated with data transfer agreements and patient consent, though it still requires institutional review board approval for participating in the federated network. The second layer involves cryptographic and statistical safeguards. Secure aggregation protocols, such as those based on homomorphic encryption or secure multi-party computation, ensure that even the server cannot inspect individual updates, only the aggregated result. However, these techniques introduce significant computational and communication overhead, making them challenging for real-time or large-scale deployment. A more practical approach is to combine differential privacy with robust aggregation, where the server adds noise to the averaged update. The trade-off here is between privacy loss and model utility: tighter privacy bounds degrade segmentation accuracy, particularly for rare nodule types. Governance arrangements must define the roles and responsibilities of each participating site, including data stewardship, model governance, and dispute resolution. A central coordinating body, often a consortium or academic network, is needed to manage the aggregation server, set training hyperparameters, and monitor for malicious participants. Additionally, transparent auditing mechanisms should allow sites to verify that their data are not being misused, for example by inspecting the contribution of their gradients to the final model. The ethical dimension of federated learning also demands attention to fairness: if some institutions contribute more data or higher-quality labels, the global model may perform best on those populations, potentially exacerbating health disparities. Governance policies must therefore include provisions for equitable contribution valuation and model performance monitoring across subgroups.

## **6. Robustness, Fairness, and Deployment Challenges**

The robustness of the federated dual-attention system to real-world perturbations is a central concern for clinical deployment. Adversarial attacks, whether intentional or accidental, can degrade model performance. For instance, a compromised site might send malicious updates to poison the global model, causing it to misclassify nodules. Defensive aggregation techniques, such as trimmed mean or Krum, can detect and exclude outliers, but they reduce the effective sample size and may inadvertently remove legitimate contributions from sites with distinct data distributions. Another robustness challenge is the natural drift of imaging protocols over time; as hospitals upgrade scanners or change protocols, the distribution of local data may shift, requiring the model to be continuously updated through federated fine-tuning. Fairness must be evaluated across multiple axes: race, ethnicity, sex, age, and socioeconomic status. If the training data from participating centers are skewed toward certain demographics, the global model may exhibit systematic errors for underrepresented groups. For example, nodule morphology can vary with patient ancestry, and attention mechanisms that learn from majority patterns may fail on minority populations. To address this, the federated system should incorporate fairness-aware aggregation, such as re-weighting site contributions based on demographic parity or equalized odds. Deployment also involves practical challenges: network latency, site dropouts, and device heterogeneity can disrupt training schedules. Asynchronous federated learning methods allow sites to submit updates at different times, but they risk stale gradients. The infrastructure must support fault tolerance, checkpointing, and dynamic site enrollment. Moreover, the final model must be validated on held-out test sets that reflect the target population, and regulatory bodies such as the FDA require rigorous clinical validation studies. All these factors underscore that federated dual-attention segmentation is not merely a technical algorithm but a socio-technical system requiring careful orchestration.

## **7. Policy Implications and Future Directions**

The widespread adoption of federated learning for medical image analysis carries significant policy implications. One major issue is the legal status of the global model: if it is trained on data from multiple institutions, who owns the intellectual property? How is liability apportioned if the model makes a diagnostic error? Current regulations often assume a centralized developer, but federated models blur these boundaries. Policymakers must establish frameworks for shared accountability, possibly through consortium agreements that specify governance structures and risk-sharing mechanisms. Another policy concern is data sovereignty: some countries or regions may require that model updates remain within their borders. This has led to the development of geo-federated learning, where regional aggregators are used before global aggregation. The dual-attention architecture, being compute-intensive, may also have environmental sustainability implications; training large models across many sites consumes significant energy, and efficiency improvements (e.g., pruning, quantization) should be incentivized through carbon-aware scheduling. Looking forward, emerging technologies such as trustworthy AI, explainable federated learning, and client-specific personalization offer promising enhancements. For example, each site could retain a lightweight personalized layer on top of the shared backbone to better fit its local data distribution, while still benefiting from global knowledge. The integration of standardized data formats, such as DICOM and controlled vocabulary for nodule annotations, is critical for interoperability. As demonstrated in prior work on controlling attributes of data exchange files, careful management of metadata formats ensures that different systems can communicate without loss of information. Furthermore, AI-augmented clinical trial modeling applied to dose optimization exemplifies how federated approaches can extend beyond imaging to multi-modal patient data, suggesting a future where collaborative learning networks support comprehensive cancer care. However, these advances must be accompanied by continuous stakeholder engagement, including clinicians, patients, regulators, and ethicists, to ensure that the technology serves public health goals equitably.

## **8. Conclusion**

This paper has presented a comprehensive system-level analysis of federated dual-attention segmentation for privacy-preserving multi-center pulmonary nodule analysis in CT imaging. By combining the strengths of federated learning and dual-attention mechanisms, the proposed framework addresses the dual imperatives of data privacy and diagnostic accuracy. We have examined the architectural trade-offs, infrastructure requirements, governance models, robustness challenges, fairness considerations, and policy implications that arise when deploying such a system in heterogeneous healthcare environments. The analysis reveals that while federated learning offers a viable path to collaborative model development without centralizing sensitive data, its success depends on careful calibration of communication, privacy, and performance objectives. The dual-attention module enhances feature representation but introduces computational overhead that must be managed across resource-constrained sites. Realizing the full potential of this approach requires not only technical innovation but also coordinated institutional governance, equitable participation, and adaptive regulatory frameworks. Future research should focus on personalization strategies, adversarial robustness, and energy-efficient training to make federated dual-attention segmentation a practical and sustainable tool for global lung cancer screening initiatives. As healthcare increasingly relies on artificial intelligence, the principles outlined here can guide the design of trustworthy and inclusive medical imaging systems.

## References

1. Dwork, C. (2006). Differential privacy. In Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (pp. 1–12). Springer.
2. Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. In Proceedings of the 41st Annual ACM Symposium on Theory of Computing (pp. 169–178). ACM.
3. 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.
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 (pp. 1273–1282). PMLR.
5. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
6. Rieke, N., Hancox, J., Li, W., Mienhart, F., Roth, H. R., Lieb, M., ... & Cardoso, M. J. (2020). The future of digital health with federated learning. *NPJ Digital Medicine*, 3(1), 119.
7. Setio, A. A. A., Ciompi, F., Litjens, G., Gerke, P., Jacobs, C., van Riel, S. J., ... & van Ginneken, B. (2016). Pulmonary nodule detection in CT images: False positive reduction using multi-view convolutional networks. *IEEE Transactions on Medical Imaging*, 35(5), 1160–1169.
8. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems 30* (pp. 5998–6008).
9. Wang, Y. (2025, August). AI-AugETM: An AI-augmented exposure–toxicity joint modeling framework for personalized dose optimization in early-phase clinical trials. In *2025 19th International Conference on Complex Medical Engineering (CME)* (pp. 182–186). IEEE.
10. Wang, Y., & Ling, C. (2025). Controlling attributes of xpt files generated by R. In *PharmaSUG 2025 conference proceedings*. San Diego, CA.
11. Woo, S., Park, J., Lee, J. Y., & Kweon, I. S. (2018). CBAM: Convolutional block attention module. In Proceedings of the European Conference on Computer Vision (ECCV) (pp. 3–19). Springer.
12. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ... & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (pp. 1175–1191). ACM.
13. Chang, C., Fu, M., Chen, X., Feng, S., Zhang, M., Zhou, X., ... & Liu, Z. (2025, November). Research on PDU-Net Lung Nodule Segmentation Algorithm Based on Path Aggregation and Dual Attention. In *2025 4th International Conference on Image Processing, Computer Vision and Machine Learning (ICICML)* (pp. 1897–1900). IEEE.

14. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.
15. Li, X., Chen, H., Qi, X., Dou, Q., Fu, C. W., & Heng, P. A. (2018). H-DenseUNet: Hybrid densely connected UNet for liver and tumor segmentation from CT volumes. *IEEE Transactions on Medical Imaging*, 37(12), 2663–2674.
16. Oktay, O., Schlemper, J., Folgoc, L. L., Lee, M., Heinrich, M., Misawa, K., ... & Rueckert, D. (2018). Attention U-Net: Learning where to look for the pancreas. In *Medical Imaging with Deep Learning (MIDL) 2018*.
17. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., ... & Bakas, S. (2020). Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports*, 10(1), 12598.
18. Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5, 1–19.
19. 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.
20. Zhang, Y., Brady, M., & Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*, 20(1), 45–57.