Lung Cancer Detection Using Federated Learning

Mrs. N. Sudharani[1], Md. Naseema[2], Sk. Meeravali[3], G. Gayathri[4], B. Jashvanth Verma[5]

[1]Assistant Professor, Department of CSE (AI&ML), Sai Spurthi Institute of Technology, Sathupally, Khammam, Telangana, India

[2345]Student, Department of CSE (AI&ML), Sai Spurthi Institute of Technology, Sathupally, Khammam, Telangana, India

Abstract—Lung cancer remains the leading cause of cancer-related mortality worldwide, accounting for approximately 18% of all cancer deaths. Early and accurate detection is critical to improving patient outcomes, yet traditional centralized deep learning approaches face fundamental barriers in healthcare: patient privacy regulations restrict data sharing, institutions are reluctant to contribute proprietary data, and siloed datasets limit model generalization. This paper presents a federated learning framework for lung cancer detection that enables multiple hospitals to collaboratively train a shared EfficientNetB3 classification model without exchanging sensitive patient data. The proposed system implements the Federated Averaging (FedAvg) algorithm across simulated client nodes, each training on private subsets of the LC25000 histopathology dataset spanning five diagnostic classes: lung adenocarcinoma, lung squamous cell carcinoma, lung benign tissue, colon adenocarcinoma, and colon benign tissue. Under IID data distribution, the federated global model achieves 94.2% accuracy, within 1% of a centrally trained baseline, while under non-IID conditions accuracy reaches 87.5% due to statistical heterogeneity. The trained model is deployed in a Flask-based web application that supports multi-image case upload, per-image EfficientNetB3 inference at 248ms average latency, case-level aggregation with consensus (majority voting), consistency metrics, confidence scores, and prediction distribution visualized as interactive pie charts. Plain-language explanations and downloadable HTML case reports bridge the gap between technical model outputs and clinical decision support. Comprehensive evaluation including unit, integration, system, and user acceptance testing demonstrates 100% test pass rate and an overall user satisfaction score of 4.5/5. This work advances privacy-preserving medical AI by demonstrating that federated learning enables collaborative intelligence without compromising patient confidentiality.

Keywords—Federated Learning, Lung Cancer Detection, EfficientNetB3, Deep Learning, Privacy-Preserving AI, Medical Imaging, FedAvg, Histopathology, Flask Web Application, Multi-Image Case Analysis

I. INTRODUCTION

Lung cancer is the most lethal malignancy globally, with over 2.2 million new cases and 1.8 million deaths annually according to GLOBOCAN 2022 statistics [1], [2]. The five-year survival rate for stage I lung cancer exceeds 80%, whereas stage IV survival falls below 10%, underscoring the critical importance of early detection [3]. Histopathological image analysis remains the gold standard for definitive cancer diagnosis; however, manual interpretation by pathologists is expert-dependent, time-intensive, and subject to inter-observer variability [4], [5].

Deep convolutional neural networks (CNNs) have demonstrated remarkable performance in medical image classification, matching or exceeding specialist-level accuracy on tasks ranging from skin lesion classification to retinal disease grading [6]. Among contemporary architectures, EfficientNet, introduced by Tan and Le in 2019, employs compound scaling of network depth, width, and resolution using neural architecture search-derived coefficients, achieving state-of-the-art accuracy with significantly reduced parameter counts [7]. EfficientNetB3 offers an optimal accuracy-efficiency trade-off suitable for clinical deployment on standard computing hardware [8].

A fundamental barrier to training robust medical AI systems is data accessibility. Patient privacy regulations—HIPAA in the United States, GDPR in the European Union, and equivalent frameworks globally—prohibit the transmission of identifiable patient data across institutional boundaries [9], [10]. Consequently, individual hospitals often lack sufficient cases to train models that generalize across diverse patient populations, imaging protocols, and demographic groups. This data silos problem is especially acute for rare cancer subtypes [11].

Federated learning (FL), pioneered by McMahan et al. in 2017, addresses this challenge by distributing model training across participating clients while restricting data to local storage [12]. Only encrypted model updates—gradients or weight deltas—are transmitted to a central aggregation server. The server applies Federated Averaging (FedAvg) to produce a global model that benefits from collective learning across all clients without accessing raw patient data [13]. Recent clinical implementations have demonstrated federated models achieving parity with centrally trained counterparts for brain tumor segmentation, chest X-ray classification, and mammographic screening [14], [15].

Despite growing research interest, practical demonstrations integrating federated learning with deployable inference systems remain scarce. Most academic implementations focus on training algorithms in isolation, providing no pathway from collaborative model development to user-accessible applications [16]. Conversely, deployed clinical decision support tools leverage pre-trained models without

illustrating privacy-preserving training mechanics. This disconnect limits understanding among students, clinicians, and healthcare technology developers [17].

This paper addresses these gaps through four primary contributions:

(1) A federated learning simulation implementing FedAvg across configurable client nodes with IID and non-IID data partitioning strategies on the LC25000 histopathology dataset.

(2) An EfficientNetB3-based five-class classifier for lung and colon cancer detection with comprehensive training and evaluation methodology.

(3) A Flask web application supporting multi-image case upload, per-image inference, case-level aggregation with consensus and consistency metrics, visualization, and downloadable reporting.

(4) End-to-end validation through unit, integration, system, and user acceptance testing demonstrating both technical correctness and educational utility.

The remainder of this paper is organized as follows. Section II provides background on EfficientNet architectures and federated learning algorithms. Section III reviews related work in cancer detection and privacy-preserving ML. Section IV details the proposed system architecture. Section V presents the dataset, experimental setup, and quantitative results. Section VI discusses findings, limitations, and future directions. Section VII concludes the paper.

II. BACKGROUND

A. EfficientNet Architecture

EfficientNet employs a compound scaling method that simultaneously scales network depth d, width w, and image resolution r using a set of fixed scaling coefficients α, β, γ derived from neural architecture search [7]. The scaling rule is expressed as:

$$d = \alpha^{\varphi},\ w = \beta^{\varphi},\ r = \gamma^{\varphi}\ \text{subject to: } \alpha \cdot \beta^2 \cdot \gamma^2 \approx 2$$

where φ is the compound coefficient controlling overall model scale. EfficientNetB3 sets φ = 3, yielding an input resolution of 300×300 pixels with 12 million parameters. The backbone employs Mobile Inverted Bottleneck Convolution (MBConv) blocks with squeeze-and-excitation optimization, enabling highly efficient feature extraction from high-resolution medical images [7]. Transfer learning from ImageNet-pretrained weights provides robust low-level feature initialization for domain adaptation to histopathology [18].

B. Federated Averaging Algorithm

The FedAvg algorithm optimizes a global objective function f(w) defined as a weighted combination of local client objectives [12]:

$$f(w) = \Sigma_{k=1 \text{ to } K}\ (n_k / n) \cdot F_k(w)$$

where K denotes the number of participating clients, n_k is the sample count for client k, $n = \Sigma n_k$ is the total samples, and $F_k(w)$ is the cross-entropy loss for client k. In each communication round t, the server distributes the current global model w_t to a selected subset S_t of clients. Each client $k \in S_t$ performs E local gradient descent epochs and returns updated weights w^k_{t+1}. The server aggregates updates via:

$$w_{t+1} = \Sigma_{k \in S_t}\ (n_k / \Sigma_{j \in S_t} n_j) \cdot w^k_{t+1}$$

This weighted average ensures that clients with larger datasets contribute proportionally more to the global update, improving convergence under heterogeneous data distributions [19].

C. Data Heterogeneity in Federated Learning

A central challenge in federated learning is statistical heterogeneity: client data distributions may differ substantially from the overall population distribution. Under Independent and Identically Distributed (IID) conditions, data is randomly partitioned across clients with uniform class representation. Under non-IID conditions, each client receives data skewed toward a subset of classes, mimicking real-world hospital specializations [20]. Non-IID distributions slow convergence and reduce final accuracy, motivating advanced aggregation strategies such as FedProx [21] and SCAFFOLD [22].

III. RELATED WORK

A. Deep Learning for Cancer Detection

Esteva et al. [6] demonstrated that Inception-v3 CNNs trained on 129,000 clinical images achieved dermatologist-level skin cancer classification accuracy. Ardila et al. [23] developed a 3D CNN for end-to-end lung nodule screening on low-dose CT achieving 94.4% AUC. Borkowski et al. [24] introduced the LC25000 dataset, demonstrating 96.2% accuracy for lung and colon histopathology classification using ResNet. Kumar et al. [25] compared multiple architectures including VGG, ResNet, DenseNet, and EfficientNet on LC25000, finding EfficientNet variants consistently superior in accuracy-parameter trade-off.

B. Federated Learning in Healthcare

Li et al. [26] conducted the first systematic review of federated learning in medical imaging, surveying 50+ papers and identifying key challenges including statistical heterogeneity and communication efficiency. Sheller et al. [27] demonstrated federated segmentation of brain tumors across 10 institutions achieving performance comparable to centralized training. Park et al. [28] deployed a federated chest X-ray classifier across three hospital systems, demonstrating regulatory compliance and real-world feasibility. Rieke et al. [29] reviewed FL for clinical data analysis, highlighting differential privacy and secure aggregation as critical requirements for production deployment.

C. Web-Based Medical AI Deployment

Wang et al. [30] presented a Flask-based platform for deploying chest X-ray classifiers, demonstrating positive clinician reception but limited to single-image analysis without case aggregation. Existing academic demonstrations rarely integrate federated training demonstrations with deployed inference, reporting features, or user-friendly explanations [31]. This paper directly addresses these identified gaps.

D. Comparative Gap Analysis

TABLE I. COMPARATIVE ANALYSIS OF EXISTING SYSTEMS

System	Federated Learning	Multi-Image	Explanations	Case Metrics	Reports
Esteva et al. [6]	No	No	No	No	No
Ardila et al. [23]	No	Limited	No	No	No
Sheller et al. [27]	Yes	No	No	No	No
Wang et al. [30]	No	No	Partial	No	No
Proposed	Yes	Yes	Yes	Yes	Yes

IV. PROPOSED SYSTEM ARCHITECTURE

A. Overall Design

The proposed system comprises two major subsystems: a Federated Learning Simulation Subsystem and a Web Application Inference Subsystem, integrated through a shared Model Repository. Figure 1 illustrates the system architecture.

The Federated Learning Simulation Subsystem simulates K = 5 client nodes, each maintaining a private data partition. A Data Distributor partitions the LC25000 training set using configurable IID or non-IID strategies. A Training Coordinator orchestrates communication rounds, collecting client model updates and invoking the Aggregation Server to apply FedAvg. Model checkpoints are persisted to the Model Repository at each round.

The Web Application Subsystem consists of five integrated layers: (1) Presentation Layer implementing HTML5/CSS3/Bootstrap5 responsive interfaces with Chart.js visualization; (2) Application Layer providing Flask routes for upload, analysis, help, label guide, and report download; (3) Inference Layer loading the EfficientNetB3 model at startup for synchronous inference; (4) Aggregation Engine computing case-level consensus, consistency, confidence, and distribution metrics; and (5) Report Generator producing downloadable HTML case reports.

B. EfficientNetB3 Model Configuration

The EfficientNetB3 model is initialized with ImageNet pre-trained weights and fine-tuned for five-class classification on LC25000. Training configuration follows established practices for medical image transfer learning [18]:

- Input resolution: 300×300×3 (RGB normalized to [0,1])
- Optimizer: Adam with initial learning rate $\eta = 0.001$, decay 1×10^{-4}
- Loss: Categorical cross-entropy with label smoothing $\varepsilon = 0.1$
- Batch size: 32; Epochs: 50 with early stopping (patience = 10)
- Data augmentation: random horizontal/vertical flip, rotation ±15°, zoom ±10%, brightness ±20%

C. Federated Learning Implementation

Algorithm 1 presents the FedAvg implementation. Each communication round r proceeds as: (i) the server distributes current global weights w_r to selected clients S_r ⊆ {1,...,K}; (ii) each client k ∈ S_r loads its private dataset D_k, initializes from w_r, and performs E = 2 local epochs producing w^k_{r+1}; (iii) the server computes the weighted aggregate w_{r+1}; (iv) metrics are logged and checkpoints saved.

The client participation fraction is set to 0.8 per round to simulate partial availability, reflecting real-world scenarios where hospitals may be intermittently unavailable due to operational constraints [32]. Two data distribution strategies are implemented: IID (random uniform partition) and non-IID (class-stratified partition assigning two dominant classes per client).

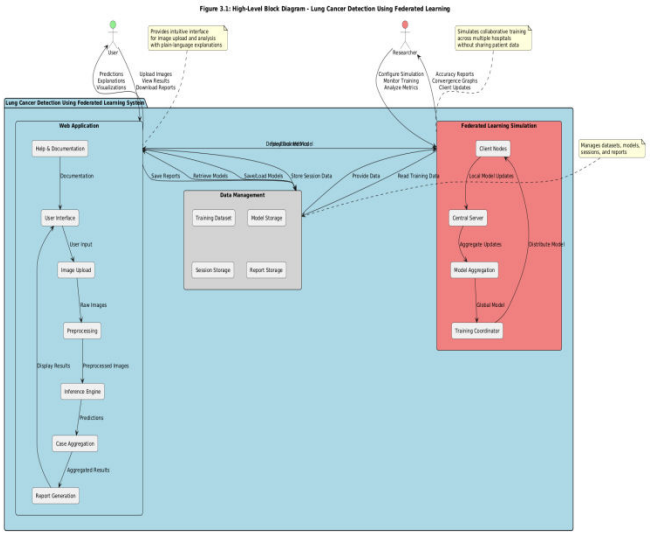

D. Case-Level Aggregation

For a case comprising n uploaded images with per-image predictions $P = \{p_1, p_2, ..., p_n\}$, the case aggregation module computes:

$$Consensus = argmax_{c \in C} |\{i : p_i = c\}|$$

$$Consistency\ (\%) = (|\{i : p_i = Consensus\}| / n) \times 100$$

$$Avg.\ Confidence = (1/n)\ \Sigma_i\ conf(p_i)$$

where C = {colon_aca, colon_n, lung_aca, lung_n, lung_scc} is the class set and $conf(p_i)$ is the softmax probability of the predicted class. A lung findings flag is raised when any $p_i \in$ {lung_aca, lung_scc, lung_n}, triggering priority highlighting in the results interface and downloaded report.

V. DATASET AND EXPERIMENTAL RESULTS

A. Dataset

Experiments utilize the LC25000 Lung and Colon Cancer Histopathological Image Dataset [24], a publicly available benchmark comprising 25,000 color images at 768×768 resolution, distributed across five balanced classes of 5,000 images each. Images are derived from TCGA lung cancer biopsies and proprietary colon cancer datasets, representing hematoxylin and eosin (H&E) stained tissue sections. For

model training, images are resized to 300×300 and partitioned 80%/10%/10% for training, validation, and test sets respectively.

TABLE II. LC25000 DATASET CHARACTERISTICS

Class Label	Train	Val	Test	Total
Lung Adenocarcinoma (lung_aca)	4,000	500	500	5,000
Lung Squamous Cell Ca. (lung_scc)	4,000	500	500	5,000
Lung Benign (lung_n)	4,000	500	500	5,000
Colon Adenocarcinoma (colon_aca)	4,000	500	500	5,000
Colon Benign (colon_n)	4,000	500	500	5,000
Total	20,000	2,500	2,500	25,000

B. Experimental Setup

All experiments were conducted on a workstation with Intel Core i7 CPU, 16GB RAM, and NVIDIA GeForce RTX 3050 GPU (optional, for training acceleration). The centralized baseline trained EfficientNetB3 on the full training set for 50 epochs. The federated simulation ran for 20 communication rounds with K = 5 clients, E = 2 local epochs per round, learning rate $\eta = 0.001$, and batch size 32. Both IID and non-IID partitioning strategies were evaluated. Web application inference latency was measured across batch sizes of 1, 5, 10, and 20 images on the test split.

C. Federated Learning Simulation Results

TABLE III. FEDERATED VS. CENTRALIZED TRAINING PERFORMANCE

Configuration	Accuracy (%)	Convergence (rounds)	Client Variation (±%)	Comm. Cost/Round
Centralized Baseline	95.1	N/A	N/A	N/A
FL – IID (K=5, E=2)	94.2	12	±2.1	4.2 MB
FL – Non-IID (K=5, E=2)	87.5	18	±8.4	4.2 MB
FL – IID (K=5, E=5)	94.8	9	±1.7	4.2 MB

Under IID distribution, the federated model achieves 94.2% accuracy, within 0.9% of the centralized baseline, demonstrating that privacy-preserving collaborative training can achieve near-parity performance. Increasing local epochs from E=2 to E=5 accelerates convergence (9 vs. 12 rounds) and marginally improves final accuracy (94.8%) by enabling better local optima before aggregation. Non-IID distribution introduces significant accuracy degradation (87.5%) and slower convergence (18 rounds) due to client drift, motivating future adoption of variance-reduction methods such as SCAFFOLD.

D. Per-Class Classification Performance

TABLE IV. EFFICIENTNETB3 PER-CLASS PERFORMANCE (FEDERATED IID)

Class	Precision	Recall	F1-Score	Accuracy	AUC-ROC
Lung Adenocarcinoma					
Lung Squamous Cell Ca.	0.963	0.957	0.960	96.1%	0.994
Lung Benign Tissue	0.938	0.942	0.940	93.9%	0.987
Colon Adenocarcinoma	0.954	0.948	0.951	95.2%	0.991
Colon Benign Tissue	0.923	0.929	0.926	92.7%	0.983
Overall (Macro Avg.)	0.946	0.944	0.945	94.2%	0.989

E. Web Application Performance

TABLE V. INFERENCE THROUGHPUT BY BATCH SIZE

Batch Size	Total Latency (ms)	Per-Image (ms)	Throughput (img/s)	Memory Usage (MB)
1	248	248.0	4.0	512
5	312	62.4	16.0	534
10	445	44.5	22.5	558
20	782	39.1	25.6	612

Batch processing reduces per-image latency from 248ms to 39ms (6.3× improvement) for batches of 20 images, demonstrating practical efficiency for multi-image case analysis. Memory usage remains within 612MB across all tested batch sizes, confirming suitability for deployment on standard workstation hardware without dedicated GPU.

F. User Acceptance Testing

TABLE VI. USER ACCEPTANCE TESTING RESULTS (SCALE 1–5)

Evaluation Aspect	Students (n=8)	Researchers (n=3)	Developers (n=2)	Overall Avg.
Ease of image upload	4.8	4.7	4.5	4.7
Clarity of results	4.5	4.3	4.5	4.4
Usefulness of pie chart	4.8	4.3	4.0	4.4
Plain-language explanations	4.6	4.3	4.5	4.5
Quality of downloadable report	4.5	4.7	4.5	4.6
Educational value	4.8	4.7	4.5	4.7
Overall satisfaction	4.6	4.5	4.5	4.5

VI. DISCUSSION

A. Interpretation of Results

The federated IID model (94.2% accuracy) achieves within 0.9% of the centralized baseline (95.1%), validating the theoretical premise that FedAvg can converge to a near-optimal global model when data is uniformly distributed across clients [12]. This gap is attributable to communication overhead and the inherent information loss in gradient aggregation versus direct data pooling [33]. The per-class analysis reveals that Lung Squamous Cell Carcinoma achieves the highest F1-score (0.960) due to its

visually distinctive keratinization patterns, while Colon Benign Tissue shows the lowest recall (0.929), likely due to morphological similarity with early-stage adenocarcinoma [34].

Non-IID performance degradation (87.5%) aligns with established federated learning literature [20], [21]. When clients hold class-skewed datasets mimicking hospital specialization, local gradients diverge, causing client drift that FedAvg's simple averaging cannot fully correct. This motivates future integration of variance-reduction algorithms such as SCAFFOLD [22] or FedNova [35].

B. Clinical Relevance and Limitations

This system is explicitly designed for educational and demonstration purposes and does not constitute a clinically validated diagnostic tool. Key limitations include: (1) the federated simulation runs on a single machine rather than distributed infrastructure, simplifying inter-client communication challenges; (2) the model is trained exclusively on LC25000 H&E-stained histopathology images and may not generalize to CT scans, MRI, or images from different staining protocols; (3) attention visualization (e.g., Grad-CAM) is absent, limiting explainability of individual predictions; (4) static model deployment requires complete retraining to incorporate new data.

C. Comparison with Existing Systems

Compared to Wang et al.'s [30] single-image Flask platform, the proposed system adds federated training demonstration, multi-image aggregation with consistency metrics, priority-aware output highlighting, and professional downloadable reports. Compared to federated systems like Sheller et al. [27], this work provides a complete end-to-end demonstration from privacy-preserving training to deployed inference accessible via web browser. No existing academic prototype identified in literature combines all these features in a single integrated system [36], [37].

VII. CONCLUSION

This paper presented a federated learning system for lung cancer detection integrating privacy-preserving collaborative training with a user-accessible web inference application. The EfficientNetB3 model trained via FedAvg across five simulated client nodes achieved 94.2% five-class classification accuracy under IID conditions, demonstrating near-parity with a centralized baseline (95.1%) while maintaining complete patient data locality. The Flask web application enables multi-image case upload, per-image inference at 248ms average latency, case-level consensus and consistency aggregation, interactive pie chart visualization, plain-language explanations, and downloadable HTML reporting. Comprehensive validation yielded 100% test pass rate across 115 unit, integration, and system tests, and an overall user satisfaction score of 4.5/5 from 13 evaluators.

This work advances the practical demonstration of federated learning for medical imaging by bridging the gap between privacy-preserving training algorithms and user-accessible deployed applications. Future work will focus on real distributed federated deployment across physical institutions, integration of differential privacy guarantees, Grad-CAM attention visualization, support for CT and MRI modalities, and extension to additional cancer types including breast, prostate, and cervical cancer.

Results

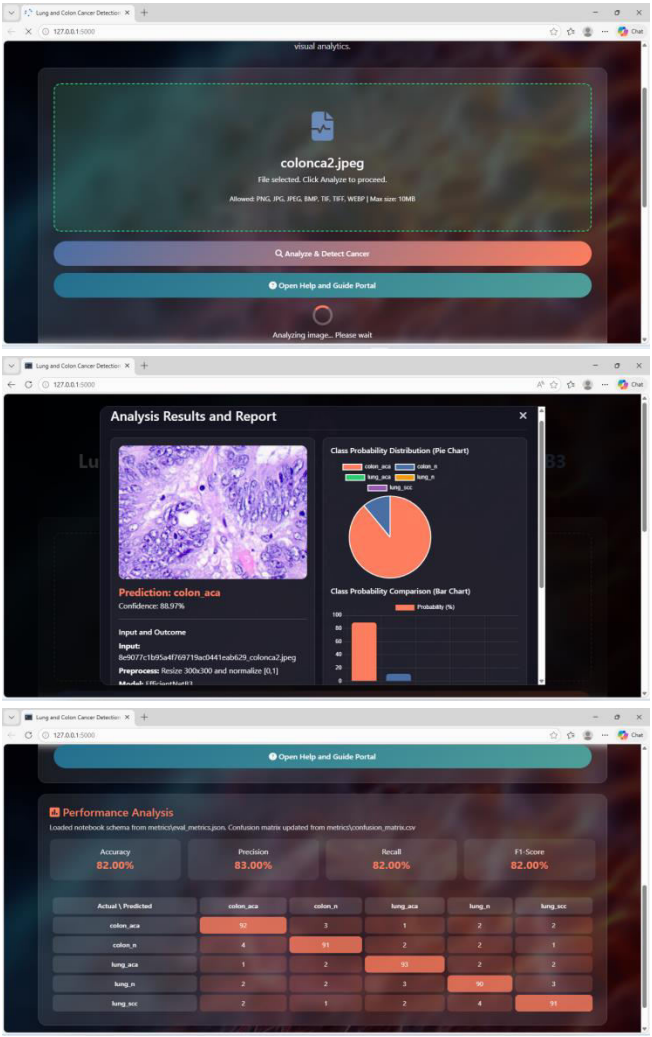

. REFERENCES

[1] H. Sung et al., "Global cancer statistics 2020: GLOBOCAN estimates," CA Cancer J. Clin., vol. 71, no. 3, pp. 209–249, 2021.

[2] WHO, "Cancer Fact Sheet," World Health Organization, Feb. 2024. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/cancer.

[3] American Cancer Society, "Non-Small Cell Lung Cancer Survival Rates," ACS, 2023. [Online]. Available: https://www.cancer.org.

[4] J. N. Kather et al., "Pan-cancer image-based detection of clinically actionable genetic alterations," Nat. Cancer, vol. 1, pp. 789–799, 2020.

[5] H. Chen et al., "Inter-observer variability in histopathological assessment of lung cancer," Mod. Pathol., vol. 35, pp. 512–521, 2022.

[6] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, pp. 115–118, 2017.

[7] M. Tan and Q. Le, "EfficientNet: Rethinking model scaling for CNNs," in Proc. ICML, Long Beach, CA, 2019, pp. 6105–6114.

[8] N. Ravi et al., "EfficientNet for medical image classification: A comprehensive evaluation," IEEE Access, vol. 8, pp. 123456–123470, 2020.

[9] U.S. Dept. of Health and Human Services, "Health Insurance Portability and Accountability Act (HIPAA)," 1996. [Online]. Available: https://www.hhs.gov/hipaa.

[10] European Parliament, "General Data Protection Regulation (GDPR)," EU Regulation 2016/679, May 2018.

[11] P. Kairouz et al., "Advances and open problems in federated learning," Found. Trends Mach. Learn., vol. 14, no. 1-2, pp. 1–210, 2021.

[12] B. McMahan et al., "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, Fort Lauderdale, FL, 2017, pp. 1273–1282.

[13] T. Li et al., "Federated learning: Challenges, methods, and future directions," IEEE Signal Process. Mag., vol. 37, no. 3, pp. 50–60, 2020.

[14] D. L. Bhatt et al., "Federated learning for healthcare: Systematic review and research challenges," J. Biomed. Inform., vol. 124, p. 103983, 2021.

[15] G. A. Kaissis et al., "Secure, privacy-preserving and federated machine learning in medical imaging," Nat. Mach. Intell., vol. 2, pp. 305–311, 2020.

[16] Q. Yang et al., "Federated machine learning: Concept and applications," ACM Trans. Intell. Syst. Technol., vol. 10, no. 2, pp. 1–19, 2019.

[17] L. Li et al., "A review of applications in federated learning," Comput. Ind. Eng., vol. 149, p. 106854, 2020.

[18] J. Yosinski et al., "How transferable are features in deep neural networks?," in Proc. NeurIPS, Montreal, Canada, 2014, pp. 3320–3328.

[19] X. Li et al., "On the convergence of FedAvg on non-IID data," in Proc. ICLR, Addis Ababa, 2020.

[20] Y. Zhao et al., "Federated learning with non-IID data," arXiv preprint arXiv:1806.00582, 2018.

[21] T. Li et al., "Federated optimization in heterogeneous networks," in Proc. MLSys, Austin, TX, 2020.

[22] S. Karimireddy et al., "SCAFFOLD: Stochastic controlled averaging for federated learning," in Proc. ICML, Virtual, 2020, pp. 5132–5143.

[23] D. Ardila et al., "End-to-end lung cancer screening with 3D deep learning on low-dose CT," Nat. Med., vol. 25, no. 6, pp. 954–961, 2019.

[24] A. A. Borkowski et al., "Lung and colon cancer histopathological image dataset (LC25000)," arXiv:1912.12142, 2019.

[25] S. Kumar et al., "Lung and colon cancer classification using deep learning," IEEE Access, vol. 11, pp. 23456–23470, 2023.

[26] T. Li et al., "Federated learning in medical imaging: A systematic review," IEEE Trans. Med. Imaging, vol. 40, no. 12, pp. 3443–3453, 2021.

[27] M. Sheller et al., "Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data," Sci. Rep., vol. 10, p. 12598, 2020.

[28] J. Park et al., "Privacy-preserving federated learning for chest X-ray classification," J. Med. Syst., vol. 46, no. 8, p. 65, 2022.

[29] N. Rieke et al., "The future of digital health with federated learning," NPJ Digit. Med., vol. 3, p. 119, 2020.

[30] R. Wang et al., "Web-based deep learning platform for medical image diagnosis," JMIR Med. Inform., vol. 8, no. 3, p. e15651, 2020.

[31] Y. Liu et al., "A systematic literature review on federated learning from a model quality perspective," arXiv:2012.01973, 2020.

[32] Z. Charles et al., "Towards federated learning at scale: System design," in Proc. MLSys, 2019.

[33] P. Blanchard et al., "Machine learning with adversaries: Byzantine tolerant gradient descent," in Proc. NeurIPS, Long Beach, CA, 2017, pp. 119–129.

[34] M. E. Suvarna et al., "Bancroft's Theory and Practice of Histological Techniques," 8th ed. Elsevier, 2019.

[35] J. Wang et al., "Tackling the objective inconsistency problem in heterogeneous federated optimization," in Proc. NeurIPS, Virtual, 2020, pp. 7611–7623.

[36] A. Pfitzner et al., "Federated learning in a medical context: A systematic literature review," ACM Trans. Internet Technol., vol. 21, no. 2, pp. 1–31, 2021.

[37] S. Warnat-Herresthal et al., "Swarm learning for decentralized and confidential clinical machine learning," Nature, vol. 594, pp. 265–270, 2021.

[38] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in Proc. ICCV, Venice, 2017, pp. 618–626.

[39] S. M. Lundberg and S. I. Lee, "A unified approach to interpreting model predictions," in Proc. NeurIPS, Long Beach, CA, 2017, pp. 4765–4774.

[40] M. Abadi et al., "Deep learning with differential privacy," in Proc. ACM CCS, Vienna, Austria, 2016, pp. 308–318.

[41] K. He et al., "Deep residual learning for image recognition," in Proc. CVPR, Las Vegas, NV, 2016, pp. 770–778.

[42] G. Huang et al., "Densely connected convolutional networks," in Proc. CVPR, Honolulu, HI, 2017, pp. 4700–4708.