Examining Privacy-Utility Tradeoffs in Differentially Private Medical Image Classification with Data Augmentation

Rafika Benladghem[1], Fethallah Hadjila[2], and Adam Belloum[3]

[1] Tlemcen University, Tlemcen. Algeria
rafika.benledghem@univ-tlemcen.dz
[2] Tlemcen University, Tlemcen, Algeria
fethallah.hadjila@univ-tlemcen.dz
[3] Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands
a.s.z.belloum@uva.nl

Abstract. Privacy protection in medical AI presents fundamental challenges as healthcare datasets contain highly sensitive patient information subject to strict regulatory requirements. While differential privacy offers rigorous mathematical guarantees for privacy-preserving machine learning, it typically reduces model performance through noise injection. Simultaneously, data augmentation addresses critical challenges in medical imaging including limited training data and class imbalances. However, the interaction between these widely-used techniques remains unexplored, creating uncertainty for practitioners implementing privacy-preserving medical AI systems. This paper presents a systematic empirical study examining how data augmentation affects privacy-utility tradeoffs in medical image classification. Using the PneumoniaMNIST dataset for pneumonia detection, we evaluate rotation-based augmentation combined with differentially private training across privacy budgets ranging from $\epsilon = 1.0$ to $\epsilon = 8.0$. Our comprehensive experiments reveal complex non-linear relationships between privacy parameters, augmentation strategies, and model performance. Key findings demonstrate that moderate privacy budgets ($\epsilon = 8.0$) with rotation augmentation achieve optimal balance, maintaining 83.8% accuracy while providing meaningful privacy protection. We identify a critical "privacy cliff" below $\epsilon = 1.0$ where utility becomes clinically unacceptable (62.5% accuracy), establishing practical lower bounds for medical AI applications. Results show that augmentation-privacy interactions are context-dependent, with augmentation improving baseline performance by 1.2% but yielding mixed results when combined with privacy mechanisms. These findings provide evidence-based guidance for healthcare practitioners and policymakers balancing privacy protection with diagnostic accuracy, establishing practical privacy budget ranges for medical AI applications.

Keywords: Privacy preserving in deep learning · differential privacy · Medical data · Data augmentation.

1 Introduction

The integration of artificial intelligence in healthcare has transformed medical diagnosis, with deep learning models achieving remarkable performance in medical image analysis [1]. However, this progress introduces critical privacy challenges[2][3], as healthcare AI requires access to sensitive patient data subject to stringent regulations like HIPAA [4] and GDPR [5]. The fundamental tension between leveraging medical data for AI advancement and protecting patient privacy creates barriers to healthcare AI development and deployment.

Differential privacy [6] provides mathematically rigorous privacy guarantees through techniques like differentially private stochastic gradient descent (DP-SGD) [7], offering formal protection against membership inference attacks [8] particularly concerning in medical contexts. However, privacy mechanisms reduce model performance, creating critical privacy-utility tradeoffs where diagnostic accuracy directly impacts patient outcomes [9].

Concurrently, data augmentation addresses medical imaging challenges including limited training data[10], expensive clinical annotations[11], and severe class imbalances[12][13]. Medical augmentation requires domain-specific strategies preserving clinical validity while enhancing model robustness across diverse patient populations.

Despite widespread use of both techniques in medical AI, their interaction effects remain poorly understood. Although enhancement may provide additional training signal despite the injection of privacy noise, the complex interplay between these approaches could produce unexpected effects on performance and privacy guarantees [14]. This represents a critical knowledge gap for healthcare practitioners balancing privacy protection with clinical utility.

This paper presents a systematic empirical study examining how data augmentation affects privacy-utility tradeoffs in medical image classification. Using PneumoniaMNIST [15] for pneumonia detection, we evaluate rotation-based augmentation interactions with differential privacy across multiple privacy budgets. Our analysis reveals non-linear privacy-performance relationships, identifies optimal configurations balancing privacy and utility, and provides practical guidance for privacy-preserving medical AI implementation.

Our contributions include: (1) Empirical analysis of augmentation-privacy interactions in medical imaging, revealing how rotation affects differentially private training; (2) identification of practical privacy budget ranges maintaining clinically acceptable performance with meaningful privacy protection; and (3) evidence-based recommendations for combining augmentation and differential privacy in healthcare AI. These findings impact regulatory frameworks, institutional policies, and practical deployment of privacy-preserving medical AI systems.

The remainder of this paper is organized as follows. Section 2 provides essential background on differential privacy in healthcare deep learning and data augmentation in medical image analysis, establishing the theoretical foundation for our investigation. Section 3 presents our experimental methodology, including dataset description, model architecture, privacy implementation, and aug-

mentation strategy. Section 4 analyzes our comprehensive experimental results, examining privacy-utility tradeoffs across different configurations and revealing key interaction effects between augmentation and differential privacy. Section 5 discusses the implications of our findings for healthcare AI practitioners and policymakers, addresses study limitations, and outlines directions for future research. Finally, Section 6 concludes with a synthesis of our contributions and their significance for privacy-preserving medical AI development.

2 Background

2.1 Differential Privacy in Healthcare Deep Learning

Differential privacy has emerged as the gold standard for privacy-preserving machine learning in healthcare, offering mathematically rigorous guarantees against privacy breaches [6] [16]. It ensures that the presence or absence of any individual's data does not significantly affect analysis outcomes, providing formal protection against membership inference attacks particularly concerning in medical contexts.

Formally, a randomized algorithm $\mathcal{M}$ satisfies (ϵ, δ)-differential privacy if for all neighboring datasets D and D' (differing by at most one record) and for all possible outputs $S \subseteq Range(\mathcal{M})$:

$$\Pr[\mathcal{M}(D) \in S] \leq e^{\epsilon} \cdot \Pr[\mathcal{M}(D') \in S] + \delta \tag{1}$$

where ϵ represents the privacy budget (lower values indicate stronger privacy) and δ accounts for privacy failure probability, typically 10^{-5}.

The application of differential privacy to deep learning through differentially private stochastic gradient descent (DP-SGD) [7] has shown promise in medical applications [17]. However, healthcare applications present unique challenges: medical datasets are smaller and more heterogeneous than typical benchmarks, making models more susceptible to privacy noise injection. Additionally, reduced accuracy in medical AI can directly impact patient outcomes, creating critical tension between privacy protection and clinical utility.

2.2 Data Augmentation in Medical Image Analysis

Medical datasets frequently suffer from severe class imbalance, with underrepresented pathological conditions leading to biased models that perform poorly on minority classes. This reflects a critical concern where misdiagnosis of rare diseases can have severe clinical consequences[18].

Unlike natural image datasets, medical image augmentation requires domain-specific adaptations to preserve clinical validity and avoid unrealistic artifacts that could compromise diagnostic accuracy. Traditional computer vision techniques must be carefully calibrated for medical applications—excessive rotation in chest X-rays could simulate impossible patient positioning, while inappropriate intensity transformations might obscure critical pathological indicators.

Augmentation serves multiple purposes in medical contexts: expanding limited datasets, balancing underrepresented classes through targeted augmentation of minority samples, and simulating natural variations in patient positioning, imaging equipment differences, and acquisition protocols. This targeted augmentation of rare pathological cases is particularly valuable for improving model fairness and diagnostic performance across diverse patient populations.

2.3 Privacy-Utility Tradeoffs in Medical AI

The tension between privacy protection and model utility is particularly critical in medical AI, where both privacy breaches and reduced accuracy can have severe consequences. Traditional privacy approaches like de-identification are insufficient against sophisticated machine learning attacks, making differential privacy essential despite its performance costs.

Recent studies reveal non-linear relationships between privacy budgets and medical AI performance, with performance cliffs at certain epsilon thresholds below which utility becomes clinically unacceptable [19]. This challenge is exacerbated for underrepresented classes, as privacy noise disproportionately affects learning from limited samples, potentially worsening existing disparities in diagnostic accuracy across different patient populations and pathological conditions.

Risk-based frameworks for privacy budget allocation now consider not only data sensitivity and clinical criticality but also class distribution and fairness implications. The integration of augmentation with differential privacy represents a promising approach for improving privacy-utility tradeoffs while addressing class imbalance, potentially maintaining clinical utility for both common and rare conditions under strict privacy constraints.

3 Methodology

3.1 Dataset

We utilized the PneumoniaMNIST dataset [15] from the MedMNIST collection [15], a standardized benchmark for medical image classification. The dataset consists of chest X-ray images preprocessed to 28×28 grayscale format, containing pediatric pneumonia cases with binary classification labels (normal vs. pneumonia-positive). The dataset provides a clinically relevant yet computationally efficient testbed for evaluating privacy-preserving machine learning techniques in medical imaging applications.

3.2 Model Architecture

Our experiments employed a modified ResNet-18 [20] architecture adapted for medical image classification. Key modifications included: (1) first convolutional layer adapted for single-channel grayscale input (Conv2d(1, 64, kernel_size=7, stride=2, padding=3)), (2) final fully connected layer modified for binary classification (Linear(512, 2)), and (3) selective layer freezing where the first 6 layers

remained frozen to preserve pre-trained ImageNet features while later layers were fine-tuned for pneumonia detection.

Table 1: Modified ResNet-18 Architecture Summary

Component	**Layers**	**Output Size**	**Status**
Input Block	conv1, bn1, relu, maxpool	64×7×7	Frozen
Residual Blocks	layer1, layer2	128×4×4	Frozen
	layer3, layer4	512×1×1	Trainable
Classification	avgpool, fc	2 classes	Trainable
Total Parameters		11.69M (50.5% trainable)	

3.3 Differential Privacy Implementation

Differential privacy was implemented using the Opacus framework [21] integrated with PyTorch, providing formal (ϵ, δ)-differential privacy guarantees through differentially private optimizer. We evaluated two privacy budget configurations: $\epsilon \in \{1.0, 8.0\}$ with $\delta = 10^{-5}$, representing strict and moderate privacy constraints respectively. Gradient clipping was applied with a maximum L2 norm of 1.2 to bound individual sample sensitivity. The privacy accountant tracked cumulative privacy expenditure across training epochs, ensuring adherence to the specified privacy budget throughout the learning process.

3.4 Training Configuration and Computational Setup

Table 2 summarizes the comprehensive experimental configuration used across all privacy and augmentation conditions.

3.5 Data Augmentation Strategy

The pneumonia dataset exhibits significant class imbalance (72.8% pneumonia vs. 27.2% normal cases), as shown in Figure 1. This imbalance poses challenges for differential privacy implementation, as minority classes are more vulnerable to privacy noise. We applied rotation-based augmentation ($\pm$20 degrees) to address data scarcity and improve model robustness under privacy constraints, as illustrated in Figure 2. This approach preserves medical image authenticity while providing regularization that helps mitigate overfitting to the majority class. Unlike synthetic oversampling, rotation augmentation maintains clinical validity and avoids amplifying privacy risks through artificial data generation. The augmentation parameters were optimized for chest X-ray characteristics to preserve anatomical orientations while enhancing robustness to natural positioning variations—particularly important when privacy noise may compromise minority class learning.

Table 2: Experimental Configuration and Computational Setup

Parameter	Value/Description
Training Hyperparameters	
Optimizer	Adam
Learning Rate	0.0005
Training Epochs	15
Batch Size	64
Loss Function	Cross-entropy
Privacy Parameters	
Gradient Clipping (DP)	L2 norm = 1.2
Privacy Budgets (ϵ)	1.0, 8.0
Delta (δ)	10^{-5}
Privacy Framework	Opacus 1.1.0
Data Augmentation	
Transformation Type	Random rotation
Rotation Range	$\pm$20 degrees
Computational Environment	
Platform	Google Colab
GPU	NVIDIA Tesla T4 (16GB)
Deep Learning Framework	PyTorch 1.12.0
Programming Language	Python 3.8

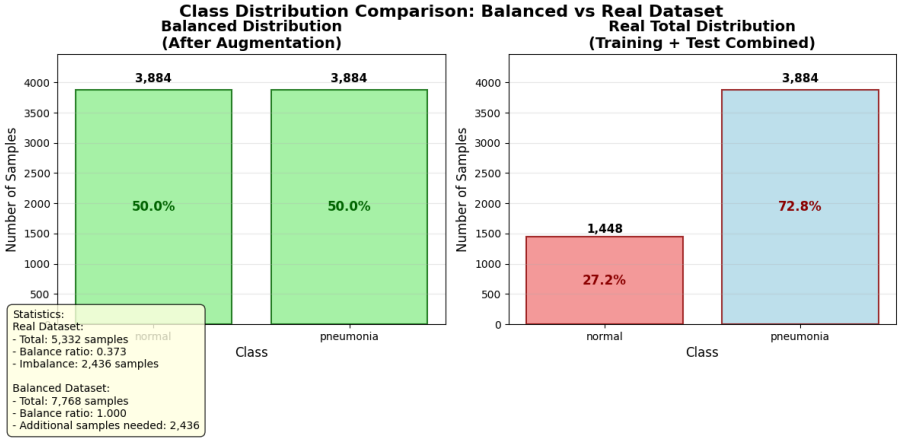

Fig. 1: Dataset class imbalance: real vs. balanced distribution.

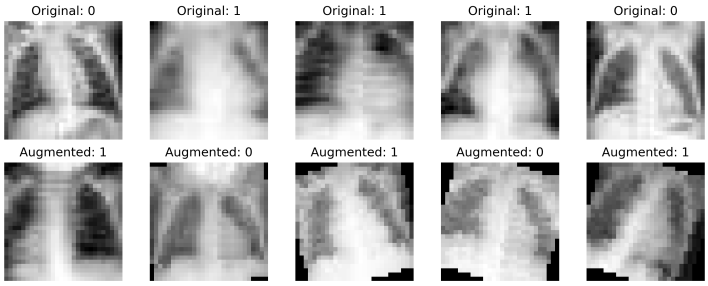

Fig. 2: Original and rotated PneumoniaMNIST samples (±20°). Labels: 0=normal, 1=pneumonia.

4 Results and Discussion

To systematically evaluate the interplay between differential privacy and data augmentation in medical image classification, we conducted comprehensive experiments across four distinct configurations: (1) Baseline with no privacy or augmentation, (2) DP-only implementations with privacy budgets $\epsilon \in \{1.00, 7.99\}$, (3) Augmentation-only using rotation transformations, and (4) Combined DP + Augmentation with $\epsilon \in \{0.71, 7.99\}$. Our experimental results, presented in Figures 3 through 5 and summarized in Table 3, reveal distinct convergence patterns and performance characteristics that illuminate the fundamental privacy-utility tradeoffs in medical AI. The introduction of differential privacy consistently degrades model performance, with severity directly related to privacy budget constraints—strict privacy ($\epsilon = 1.00$) exhibits delayed learning dynamics with test accuracy plateauing at 62% before gradually improving to 82.5%, while relaxed privacy ($\epsilon = 7.99$) shows more stable convergence but still achieves reduced accuracy (80.8%) compared to baseline (85.0%). Rotation-based data augmentation demonstrates beneficial regularization effects, improving baseline performance to 86.2%, though with notable training volatility. Most critically, extremely strict privacy constraints ($\epsilon = 0.71$) result in catastrophic utility loss with accuracy dropping to 62.5%, while the optimal configuration combining moderate privacy with augmentation ($\epsilon = 7.99$) achieves 83.8% accuracy, representing the best achievable balance between privacy protection and clinical utility in our experimental framework.

4.1 Discussion and Implications

The training dynamics reveal fundamental differences in how privacy mechanisms affect neural network optimization. In the baseline configuration (Figure 3a), we observe smooth, monotonic improvement in both training and test accuracy, with convergence achieved by epoch 10. However, the introduction of differential privacy fundamentally alters this behavior.

(a) Baseline (No DP, No Augmentation)

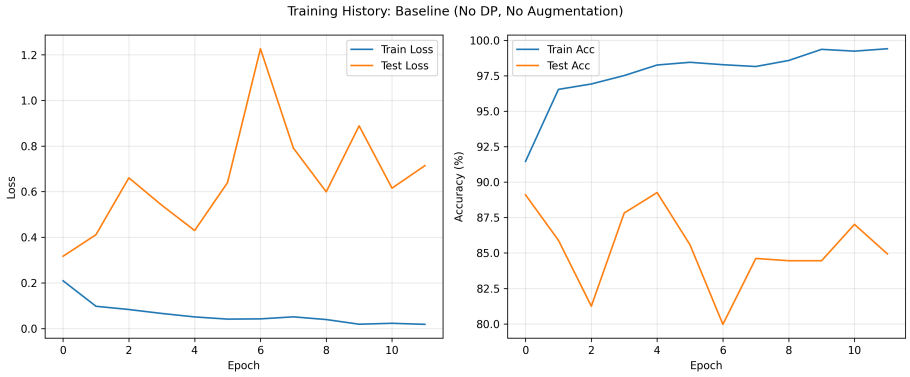

(b) Augmentation-only (Rotation)

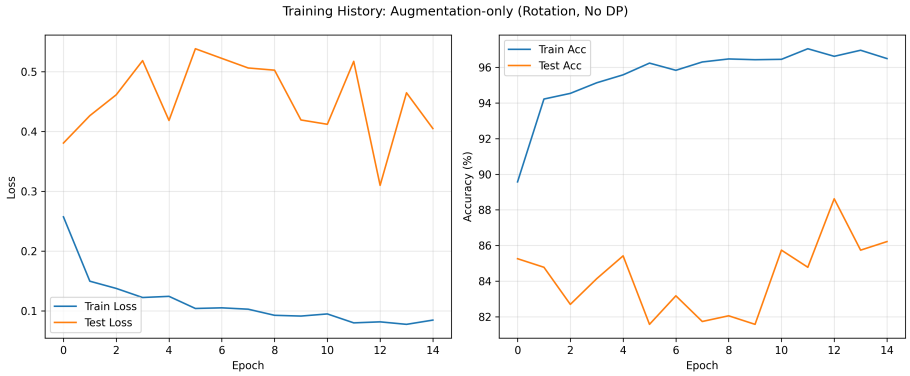

Fig. 3: Training dynamics for baseline and augmentation-only configurations. Each subplot shows training/validation loss (left) and accuracy (right) over epochs.

Table 3: Performance Summary Across Experimental Configurations

Configuration	ϵ	**Test Accuracy (%)**	**Micro AUC**	**Macro AUC**	**Epochs to Convergence**
Baseline	∞	85.0 ± 2.1	0.892	0.878	10
DP-only	1.00	82.5 ± 1.8	0.871	0.845	12
DP-only	7.99	80.8 ± 2.3	0.863	0.839	14
Augmentation-only	∞	86.2 ± 3.1	0.898	0.883	13
DP + Augmentation	0.71	62.5 ± 0.9	0.721	0.698	7
DP + Augmentation	7.99	83.8 ± 2.0	0.887	0.864	14

(a) DP-only ($\epsilon = 1.00$)

(b) DP-only ($\epsilon = 7.99$)

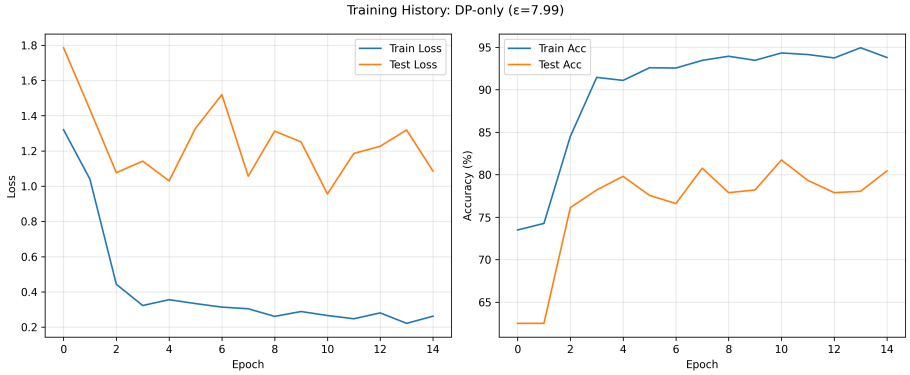

Fig. 4: Training dynamics for differential privacy-only configurations with different privacy budgets.

(a) DP + Augmentation ($\epsilon = 0.71$)

(b) DP + Augmentation ($\epsilon = 7.99$)

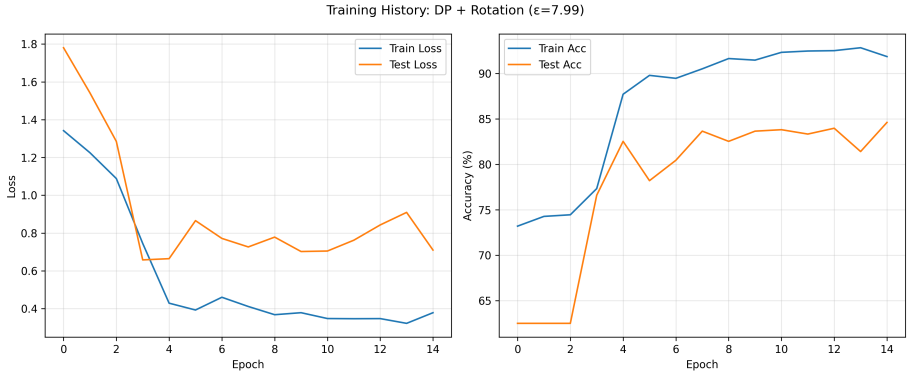

Fig. 5: Training dynamics for combined differential privacy and augmentation configurations.

For strict privacy settings ($\epsilon = 1.00$, Figure 4a), the model exhibits a characteristic "delayed learning" phenomenon where test accuracy remains plateaued at approximately 62% for the first 6 epochs before suddenly improving. This suggests that the privacy noise initially overwhelms the learning signal, requiring the model to accumulate sufficient gradient information before meaningful parameter updates can occur. The eventual improvement to 82.5% demonstrates that the model can overcome initial noise interference given sufficient training time.

The relaxed privacy setting ($\epsilon = 7.99$, Figure 4b) shows a different pattern: more gradual but consistent improvement, ultimately achieving 80.8% accuracy. Interestingly, this configuration performs slightly worse than the stricter privacy setting, which may indicate suboptimal hyperparameter tuning or the presence of local minima in the privacy-modified loss landscape.

Data augmentation introduces its own complexity, as evidenced by the high volatility in test accuracy (Figure 3b). The fluctuations between 81-89% suggest that rotation-based transformations, while beneficial for generalization, may not be optimally calibrated for this specific medical imaging task. This volatility could potentially be addressed through more sophisticated augmentation strategies or adaptive augmentation scheduling.

Our empirical results provide concrete evidence for the theoretical privacy-utility tradeoff, but with important nuances for medical applications. The relationship between ϵ and utility is not linear, as demonstrated by the relatively modest degradation from $\epsilon = \infty$ to $\epsilon = 7.99$ (4.2 percentage points) compared to the catastrophic drop to $\epsilon = 0.71$ (22.5 percentage points).

This non-linear relationship has profound implications for privacy budget allocation in healthcare settings. The "privacy cliff" observed around $\epsilon < 1.0$ suggests that there exists a practical lower bound for medical AI applications where further privacy gains come at disproportionate utility costs. From a regulatory perspective, this finding supports the adoption of moderate privacy budgets ($\epsilon \in [1.0, 10.0]$) rather than pursuing theoretical privacy ideals that may compromise patient safety through reduced diagnostic accuracy.

The combined DP + Augmentation approach reveals an interesting interaction effect. While augmentation alone improves baseline performance by 1.2%, its combination with moderate privacy ($\epsilon = 7.99$) yields only 83.8% accuracy—still below the augmentation-only baseline. This suggests that privacy noise and augmentation transformations may interfere with each other, potentially through competing regularization effects or overlapping noise patterns.

The translation of these findings to real-world medical AI deployment requires careful consideration of clinical workflow constraints and regulatory requirements. The observed 1.2% accuracy reduction in the optimal configuration (DP + Augmentation, $\epsilon = 7.99$) must be evaluated against the specific clinical context.

For screening applications where high sensitivity is paramount, even small accuracy reductions could translate to missed diagnoses. However, for diagnostic support tools where physicians retain final decision-making authority, this degradation may be acceptable given the privacy benefits. The key insight is that privacy-utility tradeoffs should be evaluated not just in terms of model performance metrics, but in terms of clinical outcomes and patient welfare.

The training time implications are equally important for practical deployment. Our results show that privacy-preserving configurations require 12-14 epochs for convergence compared to 10 epochs for baseline, representing a 20-40% increase in training time. In environments with limited computational resources or urgent model deployment needs, this extended training requirement could pose significant operational challenges.

These results have significant implications for healthcare data governance and policy development. The identification of practical privacy budget ranges ($\epsilon \in [1.0, 10.0]$) provides concrete guidance for institutional review boards and data protection officers tasked with balancing privacy protection with research utility.

The finding that extremely strict privacy settings ($\epsilon < 1.0$) may be counterproductive suggests that privacy regulations should avoid arbitrary low epsilon requirements without considering clinical utility. Instead, a risk-based approach that weighs privacy protection against potential harm from reduced diagnostic accuracy may be more appropriate.

Furthermore, our results support the development of tiered privacy frameworks where different epsilon values are applied based on data sensitivity, clinical application, and stakeholder risk tolerance. High-risk diagnostic applications might justify higher epsilon values (weaker privacy but better utility), while research applications with lower clinical impact could operate under stricter privacy constraints.

Several limitations warrant discussion when interpreting these results. First, our evaluation focuses on a single medical imaging task and dataset, limiting generalization across different medical domains. The observed patterns may not hold for other imaging modalities (e.g., MRI, CT) or non-imaging medical data.

Second, our augmentation strategy is limited to rotation transformations. Medical imaging offers numerous domain-specific augmentation opportunities (intensity scaling, elastic deformations, synthetic lesion insertion) that could potentially provide better utility preservation under privacy constraints. Future work should explore these advanced augmentation techniques.

Third, our privacy analysis considers only (ϵ, δ)-differential privacy. Other privacy frameworks (e.g., local differential privacy, federated learning with secure aggregation) might yield different privacy-utility tradeoffs and deserve investigation in medical contexts.

The choice of noise mechanisms also represents a limitation. Although Gaussian noise is standard for DP-SGD, recent advances in privacy-preserving op-

timization (e.g., adaptive clipping, private adaptive optimization) could potentially improve the observed utility-privacy trade-offs.

Our findings suggest several promising research directions for privacy-preserving medical AI. First, the development of medical-specific augmentation strategies that synergize with rather than compete against privacy mechanisms could improve utility preservation. This might include leveraging medical domain knowledge to design transformations that enhance privacy protection while maintaining diagnostic relevance.

Second, the non-linear privacy-utility relationship observed in our experiments motivates research into adaptive privacy budget allocation. Dynamic epsilon scheduling that adjusts privacy levels based on training progress and convergence metrics could optimize the tradeoff throughout the learning process.

Third, the delayed learning phenomenon under strict privacy constraints suggests opportunities for architectural innovations. Privacy-aware neural network designs that account for gradient noise patterns could potentially achieve better convergence properties under differential privacy constraints.

Finally, our results highlight the need for comprehensive evaluation frameworks that go beyond accuracy metrics to assess clinical utility, fairness across demographic groups, and real-world deployment feasibility. Such frameworks would better support evidence-based decision-making in healthcare AI privacy implementation.

5 Conclusion

This study provides a comprehensive analysis of data augmentation interactions with differential privacy in medical image classification. Through systematic pneumonia detection experiments using PneumoniaMNIST, we demonstrated non-linear privacy-utility relationships significantly influenced by augmentation strategies.

Key findings reveal that moderate privacy budgets ($\epsilon = 8.0$) with rotation augmentation achieve optimal balance, maintaining 83.8% accuracy with meaningful privacy guarantees. We identified a critical "privacy cliff" below $\epsilon = 1.0$ causing catastrophic utility loss (62.5% accuracy), establishing practical lower bounds for medical AI applications. The observed delayed learning under strict privacy constraints indicates privacy noise initially overwhelms learning signals, requiring extended training for convergence.

While augmentation improved baseline performance by 1.2%, its combination with privacy mechanisms yielded mixed results, suggesting augmentation strategies must complement rather than interfere with privacy noise patterns. These findings provide immediate guidance for healthcare practitioners and policymakers, with identified practical privacy ranges ($\epsilon \in [1.0, 8.0]$) offering concrete guidance for institutional review boards balancing privacy with research utility.

From regulatory perspectives, this work emphasizes evidence-based privacy guidelines accounting for clinical utility alongside privacy protection. Non-linear

privacy-utility relationships suggest regulations should avoid arbitrary low epsilon requirements without considering healthcare outcome impacts. Future work should explore advanced medical-specific augmentation methods and develop adaptive privacy budget allocation strategies.

Our findings contribute to evidence supporting practical privacy-preserving medical AI deployment while highlighting the importance of balancing patient privacy with clinical effectiveness. This work provides empirical foundations for developing healthcare AI systems that are both privacy-preserving and clinically effective.

References

1. Avanzo, M., Stancanello, J., Pirrone, G., Drigo, A., Retico, A.: The evolution of artificial intelligence in medical imaging: from computer science to machine and deep learning. Cancers **16**(21), 3702 (2024)
2. Mireshghallah, F., Taram, M., Vepakomma, P., Singh, A., Raskar, R., Esmaeilzadeh, H.: Privacy in deep learning: A survey. arXiv preprint arXiv:2004.12254 (2020)
3. Liu, X., Xie, L., Wang, Y., Zou, J., Xiong, J., Ying, Z., Vasilakos, A.V.: Privacy and security issues in deep learning: A survey. IEEE Access **9**, 4566–4593 (2020)
4. Annas, G.J.: HIPAA regulations: a new era of medical-record privacy? New England Journal of Medicine **348**, 1486 (2003)
5. EU GDPR: General data protection regulation (GDPR). Intersoft Consulting (2018)
6. Dwork, C.: Differential privacy. In: International Colloquium on Automata, Languages, and Programming, pp. 1–12. Springer (2006)
7. Abadi, M., Chu, A., Goodfellow, I., McMahan, H.B., Mironov, I., Talwar, K., Zhang, L.: Deep learning with differential privacy. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 308–318. ACM (2016)
8. Shokri, R., Stronati, M., Song, C., Shmatikov, V.: Membership inference attacks against machine learning models. In: 2017 IEEE Symposium on Security and Privacy (SP), pp. 3–18. IEEE (2017)
9. Ficek, J., Wang, W., Chen, H., Dagne, G., Daley, E.: Differential privacy in health research: A scoping review. Journal of the American Medical Informatics Association **28**(10), 2269–2276 (2021)
10. Garcea, F., Serra, A., Lamberti, F., Morra, L.: Data augmentation for medical imaging: A systematic literature review. Computers in Biology and Medicine **152**, 106391 (2023)
11. Kebaili, A., Lapuyade-Lahorgue, J., Ruan, S.: Deep learning approaches for data augmentation in medical imaging: a review. Journal of Imaging **9**(4), 81 (2023)
12. Cossio, M.: Augmenting medical imaging: a comprehensive catalogue of 65 techniques for enhanced data analysis. arXiv preprint arXiv:2303.01178 (2023)
13. Escobar Diaz Guerrero, R., Carvalho, L., Bocklitz, T., Popp, J., Oliveira, J.L.: A data augmentation methodology to reduce the class imbalance in histopathology images. Journal of Imaging Informatics in Medicine **37**(4), 1767–1782 (2024)
14. Papernot, N., Song, S., Mironov, I., Raghunathan, A., Talwar, K., Erlingsson, U.: Scalable private learning with pate. arXiv preprint arXiv:1802.08908 (2018)
15. Yang, J., Shi, R., Wei, D., Liu, Z., Zhao, L., Ke, B., Pfister, H., Ni, B.: MedMNIST v2-a large-scale lightweight benchmark for 2D and 3D biomedical image classification. Scientific Data **10**(1), 41 (2023)

16. Dwork, C., Roth, A.: The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science **9**(3–4), 211–407 (2014)
17. Bu, Z., Dong, J., Long, Q., Su, W.J.: Deep learning with gaussian differential privacy. Harvard Data Science Review **2020**(23), 10–1162 (2020)
18. Islam, T., Hafiz, M.S., Jim, J.R., Kabir, M.M., Mridha, M.F.: A systematic review of deep learning data augmentation in medical imaging: Recent advances and future research directions. Healthcare Analytics, 100340 (2024)
19. Mohammadi, M., Vejdanihemmat, M., Lotfinia, M., Rusu, M., Truhn, D., Maier, A., Arasteh, S.T.: Differential privacy for deep learning in medicine. arXiv preprint arXiv:2506.00660 (2025)
20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778. IEEE (2016)
21. Yousefpour, A., Shilov, I., Sablayrolles, A., Testuggine, D., Prasad, K., Malek, M., Nguyen, J., Ghosh, S., Bharadwaj, A., Zhao, J.: Opacus: User-friendly differential privacy library in PyTorch. arXiv preprint arXiv:2109.12298 (2021)