

# ENHANCING PRIVACY AND UTILITY IN DIFFERENTIALLY PRIVATE DEEP LEARNING THROUGH INSTANCE-LEVEL SMOOTHING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

In this paper, we address the dual challenge of maintaining high accuracy and ensuring fairness in differentially private (DP) deep learning models. The optimization process is inherently complicated by the necessity of injecting random noise and limiting training iterations, particularly for over-parameterized models. Moreover, DP mechanisms frequently exacerbate accuracy disparities across subpopulations, complicating the balance between privacy and fairness. To tackle these challenges, we introduce a novel framework that systematically addresses the trade-off between privacy and utility in DP deep learning. At the core of our approach is the concept of instance-level smoothing, which enhances privacy protections without compromising performance. Our theoretical contributions include deep insights into sample complexity, instance-level smoothing factors, and error bounds required to achieve a given privacy budget. These insights provide a robust foundation for optimizing the delicate balance between privacy and utility. Our method demonstrates remarkable robustness, independent of iteration counts, model parameters, batch normalization processes, and subpopulation disparities. This flexibility enables an optimal balance between privacy preservation and utility, adaptable to a wide range of scenarios. Through extensive empirical studies on the large-scale medical imaging dataset CheXpert, we validate the effectiveness of our approach. Our findings align with theoretical predictions, showing that our method can effectively meet stringent privacy requirements while maintaining high performance. By bridging the gap between formal privacy guarantees and practical deep learning applications, our work lays the groundwork for future advancements in the field. This research empowers practitioners to protect sensitive data during model training and ensures both data privacy and model generality, paving the way for more secure and equitable AI systems.

## 1 INTRODUCTION

Large neural networks with billions of parameters have achieved state-of-the-art performance across a wide range of machine learning tasks Brown & Mann (2020); Zhai et al. (2022). Despite their success, these models are known to memorize training data Zhang et al. (2021), leading to potential leakage of sensitive information. Privacy attacks capable of extracting memorized training data have been demonstrated on various types of models, including language models Carlini et al. (2021), diffusion models Carlini et al. (2023), and image classification models Balle et al. (2022). Additionally, membership inference attacks, which can determine whether a particular data point was used to train a model, have proven successful across multiple architectures and data modalities Carlini et al. (2022). The implications of these vulnerabilities are profound, particularly in applications involving private, confidential, or proprietary data, such as healthcare, finance, recommendation systems, and mobility. To ensure the safe deployment of models trained on such data, it is crucial to address these privacy concerns effectively.

Differential Privacy (DP) Dwork et al. (2006) has emerged as the gold standard for protecting individual privacy in data processing algorithms. However, achieving strong privacy protections with DP is particularly challenging in deep learning, especially as model sizes and data dimensions increase. The most popular DP training technique in deep learning is Differentially Private Stochastic Gradi-

ent Descent (DP-SGD) Abadi et al. (2016), which privatizes gradients by clipping and adding noise. The strength of the privacy guarantee  $\epsilon$  depends on the noise scale, batch size, number of training samples, and number of iterations. While DP-SGD is a near drop-in replacement for standard SGD, it presents significant challenges that have hindered its widespread adoption.

Achieving high accuracy with differentially private (DP) models faces two major challenges. First, injecting random noise and limiting training iterations complicates optimization, as noisy gradients and iteration constraints hinder effective hyperparameter tuning Anil et al. (2021). The noise’s Euclidean norm scales with model size, degrading performance for larger models Yu et al. (2021). This issue is particularly problematic given the prevalence of over-parameterized neural networks in AI. Second, fairness issues arise in deep neural networks, exacerbated by DP mechanisms. Models can show accuracy disparities across subpopulations, such as race, age, sex, or insurance type in medical tasks Seyyed-Kalantari et al. (2021). DP can worsen these disparities, increasing accuracy gaps due to unbalanced subgroup data Bagdasaryan et al. (2019); Santos-Lozada et al. (2020). DP limits information extraction from individual data points, conflicting with the accurate learning needed for small or underrepresented subgroups, particularly in medical datasets Cummings et al. (2019); Suriyakumar et al. (2021).

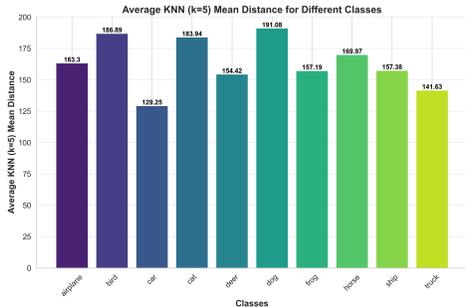


Figure 1: Histogram of the mean distances for k-nearest neighbors (KNN) in CIFAR10. This histogram visualizes the distribution of average distances among data points as computed by the KNN algorithm.

Understanding the trade-off between privacy and utility in differentially private (DP) deep learning models is crucial for both theoretical and practical advancements. This paper introduces a novel framework that aims to bridge the performance gap between DP models and their non-private counterparts by training an end-to-end DP network. Our central hypothesis is that samples with sparse distributions in their feature space, based on pre-trained large models, are more prone to privacy leakage. In contrast, classes with sufficient samples that are compactly clustered in the feature space will be well-represented, incurring minimal privacy sacrifice while maintaining high performance.

To illustrate the difference in feature density among classes in the CIFAR-10 dataset, we present a graph depicting the mean distances within the 10 classes. This graph measures the proximity of each sample to its 5 nearest neighbors. As shown in Figure 1, the mean sample distance for *car* is 129.25, whereas for *dog*, it is 191.08. The observed trade-off between privacy and utility Berrada et al. (August 2023) aligns well with these mean distances: the *car* class exhibits the best classification accuracy and the least privacy leakage, whereas the *dog* class suffers from significant privacy leakage and class disparity issues.

Given that each class has an equal number of 5000 training samples, the within-class compactness and between-class distance largely determine the privacy-utility trade-off. In scenarios where the sample numbers across classes are imbalanced, privacy preservation and class disparity issues worsen. This observation motivates us to consider instance-level probability density function (PDF) estimation and instance-specific kernel smoothing to enhance privacy without degrading performance. We prefer the term *instance smoothing* over *noise*, as it better captures the essence of our approach in providing privacy protection alongside model generality.

Our method remains flexible and can incorporate future advancements in deep learning, including improved network architectures and pre-trained models, without additional constraints. This work

108 represents a significant step towards enabling deep learning practitioners to leverage formal privacy  
 109 guarantees offered by DP while protecting sensitive data during model training.

110 The contributions of this paper are multifaceted and significant, paving the way for a new paradigm  
 111 in differentially private deep learning:  
 112

- 113 • **Innovative Framework:** We introduce a novel framework that systematically addresses  
 114 the trade-off between privacy and utility in differentially private (DP) deep learning models.  
 115 Central to our approach is the concept of instance-level smoothing and PDF representation  
 116 learning, which significantly enhances privacy without compromising model performance.
- 117 • **Theoretical Insights:** Our work provides profound theoretical insights into key aspects  
 118 such as sample complexity, instance-level smoothing factors, and error bounds for achiev-  
 119 ing a given privacy budget. These insights offer a robust foundation for understanding and  
 120 optimizing the balance between privacy and utility.
- 121 • **Robust Methodology:** The proposed method demonstrates remarkable robustness, being  
 122 independent of iteration counts, backbone parameters, batch normalization processes, and  
 123 sub-class disparities. This flexibility allows for an optimal balance between privacy preser-  
 124 vation and utility, making it adaptable to a wide range of scenarios.
- 125 • **Empirical Validation:** Through extensive empirical studies on the large-scale medical  
 126 image dataset CheXpert Irvin et al. (January 2019), we validate the effectiveness of our  
 127 approach. Our findings confirm the theoretical predictions, demonstrating that our method  
 128 can effectively meet privacy requirements while maintaining high performance. This new  
 129 paradigm has the potential to significantly expand the use of large models in real-world  
 130 scenarios where privacy is paramount.

## 131 2 RELATED WORKS

132 **Zero-Concentrated Differential Privacy** Differential Privacy (DP) was initially formalized by  
 133 Dwork et al. (2006) and later adapted to deep learning by Abadi et al. (2016), who operationalized  
 134 DP-SGD Bassily et al. (2014) to train neural networks with differential privacy guarantees. The  
 135 central goal in differentially private learning is to train a classifier while satisfying a rigorous math-  
 136 ematical definition of privacy known as differential privacy Dwork (2006), which ensures that any  
 137 individual training data point cannot be identified using the trained model and any additional side  
 138 information. More formally, we adopt the popular modern variant called zero-centered Concentrated  
 139 Differential Privacy (zCDP), as defined below:

140 **Definition 1.** ZERO-CONCENTRATED DIFFERENTIAL PRIVACY BUN & STEINKE (2016)  
 141 *Two datasets  $D_0$  and  $D_1$  are neighbors if they can be constructed from each other by adding or*  
 142 *removing one data point. A randomized mechanism  $\mathcal{A}$  satisfies  $\rho$ -zero-concentrated differential*  
 143 *privacy ( $\rho$ -zCDP) if, for all neighboring datasets  $D_0$  and  $D_1$ , we have*

$$144 R_\alpha(\mathcal{A}(D_0) \parallel \mathcal{A}(D_1)) \leq \rho\alpha,$$

145 where

$$146 R_\alpha(P \parallel Q) = \frac{1}{\alpha - 1} \log \int \left( \frac{p(x)}{q(x)} \right)^\alpha q(x) dx$$

147 is the Rényi divergence between two distributions  $P$  and  $Q$ .

148 In the above definition,  $\rho \geq 0$  is the privacy loss parameter that measures the strength of the pro-  
 149 tection.  $\rho = 0$  indicates perfect privacy, while  $\rho = \infty$  means no protection at all. The privacy  
 150 protection is considered sufficiently strong in practice if  $\rho$  is a small constant, e.g., 1, 2, 4, 8. For  
 151 readers familiar with standard approximate DP but not zCDP,  $\rho$ -zCDP implies  $(\epsilon, \delta)$ -DP for all  $\delta > 0$   
 152 with  $\epsilon = \rho + 2\sqrt{\rho \log(1/\delta)}$ .

153 While DP-SGD Berrada et al. (August 2023) is foundational for differential privacy in deep learning,  
 154 it has significant drawbacks. It requires a trade-off between privacy budget ( $\epsilon$ ), noise addition,  
 155 and the number of iterations, impacting model utility. The method is sensitive to hyperparameters  
 156 and regularization, and per-sample gradient clipping introduces bias and high variance, degrading  
 157 performance. Additionally, the noise scales with model dimension, complicating its application to  
 158 large models and disrupting batch normalization.

We propose a novel approach that bypasses these issues. Our method adjusts the smoothing bandwidths of the PDF per instance, independent of iterations, enhancing the balance between privacy and accuracy. It is robust to hyperparameter settings and avoids gradient clipping and noise addition, reducing bias and variance. This allows for fine-tuning large pre-trained models in high-dimensional spaces without affecting batch normalization. We also provide theoretical bounds on sample size and classification errors, demonstrating its effectiveness.

Overall, our method offers a practical and efficient alternative to DP-SGD, maintaining privacy while improving model performance.

**Adaptive Kernel Density Fine-tuning** Our work diverges from the DP-SGD method by drawing on adaptive kernel metric representation learning, particularly useful for tasks like anomaly detection. Kernel density-based learning captures the underlying probability density of data, using Gaussian kernels to assess local density and identify anomalies, which are prone to privacy risks in sparse regions Zhang et al. (2018). Kernel-based methods handle uncertainty and class-specific variances well, as demonstrated by Non-isotropic von Mises-Fisher (nivMF) distributions that model complex variances to enhance generalization Kirchof et al. (2022).

We employ a projection network parameterized by  $W \in \mathbb{R}^{p \times d}$ , where  $d$  is the feature dimension from the backbone network, and  $p$  is the projected dimension in the PDF space. This adaptive kernel estimation-based approach emphasizes local data density, crucial for identifying and mitigating privacy risks in sparse regions. It allows fine-tuning of pre-trained models in high-dimensional spaces, compatible with any backbone architecture.

Unlike DP-SGD, which requires per-sample gradient clipping and noise addition, our method adjusts the smoothing bandwidth at the instance level, leading to more efficient training. By focusing on instance-level smoothing, we effectively balance privacy and accuracy, offering a practical and efficient alternative to DP-SGD for differentially private learning.

### 3 METHOD

Our training framework for medical images starts by feeding samples into feature extraction networks, such as pre-trained transformers on ImageNet, to obtain initial feature representations. We then project these raw features into a PDF representation space, modeling each class with an independent PDF using an adaptive kernel density estimation method. This approach assigns a specific smoothing bandwidth to each instance based on its relative distribution density.

As shown in Figure 2, we construct a PDF for each class by applying adaptive kernel density estimation in the projected space, ensuring customized privacy preservation for each instance. To maintain compactness within classes and distinctiveness between different classes, we employ a PDF contrastive loss with a margin. This loss function encourages the model to achieve the desired feature separation. During testing, we compute the likelihood of a test sample belonging to each class by comparing it against the trained PDFs. The predicted label is assigned based on the class with the highest probability. The detailed steps are depicted in Figure 2 and will be elaborated in the following sub-sections.

#### 3.1 INSTANCE LEVEL PRIVACY BUDGET ALLOCATION

This projection network can map the samples in feature space in  $\mathbb{R}^d$  to PDF in  $\mathbb{R}^p$ . For example, in our experiment settings,  $d = 768$  is for transformer and  $p = 64$  for the embedding PDF space. The parameter  $p$  is selected by evaluation process.

Each training sample exhibits different local sensitivity in the PDF space. We should allocate the privacy budget according to their sensitivity. Adaptive smoothing mechanism adjusts the amount of bandwidth based on the local sensitivity of the PDF estimate. This approach enhances privacy while maintaining the utility of the PDF representation.

**Proposition 1.** INSTANCE LEVEL PRIVACY SMOOTHING

Let  $x_i \in \mathbb{R}^p$  be a sample. To protect each sample  $x_i$ , smoothing factor  $N_i$  is applied for each dimension, where  $N_i \sim \mathcal{N}(0, s_i^2 \cdot \sigma^2)$ . Here,  $s_i$  represents the local sensitivity of  $x_i$ , and  $\sigma$  is a global privacy setting applicable to all training samples. The protected sample  $\tilde{x}_i$  is then given

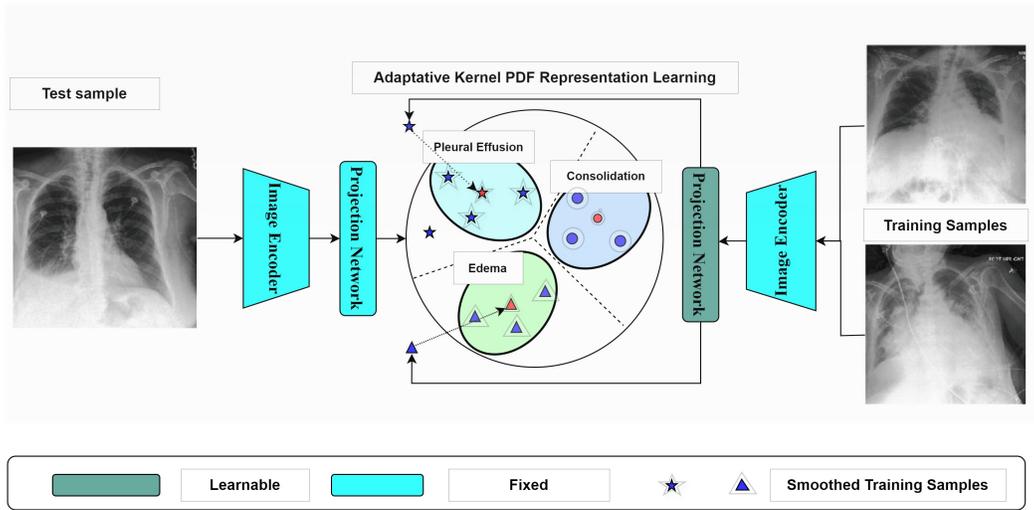


Figure 2: Adaptive Kernel PDF Representation Learning Framework. In the first stage, each training sample is smoothed with a customized bandwidth: densely distributed samples receive less smoothing, while sparsely distributed samples are protected with larger bandwidths. In the second stage, the projection network is trained to optimally separate each class. The mean vectors for each class are represented by red markers: stars, circles, and triangles correspond to the classes *Pleural Effusion*, *Edema*, and *Consolidation*, respectively, positioned at the centers of their corresponding elliptical PDFs. Each class is formed as a compacted PDF with a margin among other classes. The varying sizes of the shaded areas around the samples represent different local bandwidths, corresponding to varied probability densities and different allocations of privacy budgets.

by  $\tilde{x}_i = x_i + N_i$ . This mechanism ensures that the level of smoothing corresponds to the local sensitivity of each data point, balancing privacy protection and data utility by scaling the smoothing variance with  $s_i^2 \cdot \sigma^2$ . The introduction of smoothing  $N_i$  adheres to differential privacy principles, providing a customizable level of privacy suited to the sensitivity of the data.

We will define the  $s_i$  firstly. The PDF for the training set is defined as:

$$\hat{p}(x) = \frac{1}{n} \sum_{j=1}^n K_h(x - x_j)$$

The PDF without sample  $x_i$  is:

$$\hat{p}_{-i}(x) = \frac{1}{n-1} \sum_{j \neq i} K_h(x - x_j)$$

Then the local sensitivity  $s_i$  at  $x_i$  is:

$$s_i = |\hat{p}(x) - \hat{p}_{-i}(x)|$$

For large  $n$ , this can be approximated as:

$$s_i \approx \frac{1}{n} K_h(x - x_i)$$

The  $K_h$  is the kernel function with smoothing bandwidth  $h$ , such as the Gaussian Kernel. However, in this paper, we will define the Adaptive Kernel Function in the next sub-section to best describe each sample's distinctiveness.

The  $\tilde{x}_i = x_i + N_i$  provides instance level privacy protection mechanism. The modified PDF estimate  $\hat{p}(x)$  is:

$$\hat{p}(x) = \frac{1}{n} \sum_{i=1}^n K_h(x - \tilde{x}_i)$$

In the following section, we will not use  $\tilde{x}_i$  but  $x_i$  for brevity to represent the enhanced samples. The customized smoothing factor  $N_i$  is related to instance sensitivity  $s_i$  and an global smoothing variance  $\sigma^2$  independent of sample index.

**Theorem 1. PRIVACY-PRESERVING FOR  $\rho$ -ZCDP**

To achieve  $\rho$ -zero Concentrated Differential Privacy ( $\rho$ -zCDP), the smoothing variance is chosen as  $\sigma^2 = \frac{2 \ln(1/\delta)}{\rho}$ . The privacy-preserving parameter  $\rho$  controls the overall level of privacy, ensuring that the samples in PDF space remains stable and robust while preserving privacy. This formulation provides a balance between privacy protection and data utility by adjusting the smoothing variance according to the specified privacy parameter  $\rho$ .

The proof can be found in Appendix A.1.

### 3.2 ADAPTIVE KERNEL IN HIGHER DIMENSIONS

**Proposition 2. ADAPTIVE KERNEL SMOOTHING**

In the context of the higher-dimensional adaptive kernel, we can define  $h'$  explicitly based on the local density, which provides additional privacy-preserving ability for the sparse training samples given the base bandwidth  $h$ . The adaptive kernel  $K_h(x)$  in  $p$  dimensions is given by:

$$K_h(x) = \begin{cases} \frac{1}{(2\pi h^2)^{p/2}} \exp\left(-\frac{\|x\|^2}{2h^2}\right) & \text{if } N_r(x) \geq \lambda, \\ \frac{1}{(2\pi(h^2+h'^2))^{p/2}} \exp\left(-\frac{\|x\|^2}{2(h^2+h'^2)}\right) & \text{if } N_r(x) < \lambda, \end{cases}$$

where  $h' = h \cdot \frac{c}{N_r(x)}$ . To define sparse and dense regions in high-dimensional space, we use a measure of local density, for which one effective approach is the  $k$ -nearest neighbors ( $k$ -NN) method.

To distinguish between dense and sparse regions, we set a threshold  $\lambda$ . This threshold can be determined based on the empirical distribution of  $k$ -NN distances in the dataset. For example, setting  $\lambda$  to the median or the 75th percentile of  $k$ -NN distances ensures a robust boundary. Formally, we define regions as follows:

$$\begin{cases} \text{Dense region} & \text{if } N_r(x) \geq \lambda \\ \text{Sparse region} & \text{if } N_r(x) < \lambda \end{cases}$$

where  $\lambda$  is a chosen percentile of the distribution of  $k$ -NN distances.

For example, in the CIFAR10 dataset, the mean 5-nearest-neighbor distance for *car* is 129.25. We can define  $\lambda = 129.25$ . For samples with mean distance to their 5-nearest-neighbor beyond this threshold, we should use extra bandwidth to protect the samples.

**Theorem 2. SAMPLE SIZE COMPLEXITY**

To preserve  $\rho$ -zCDP in PDF space, the number of samples  $n$  should satisfy:

$$n \geq \frac{C}{\rho(\alpha - 1)}$$

This theoretical bound ensures that the change in the KDE estimate due to removing a single sample is sufficiently small to maintain the desired privacy level  $\rho$ . The constant  $C$  is given by:

$$C = \frac{1}{\sqrt{2\pi h^2}}$$

This bound provides a theoretical guideline for the minimum number of samples needed to ensure  $\rho$ -zCDP in the context of PDF probability density estimation.

The proof can be found in Appendix A.2.

### 3.3 KERNEL DENSITY-BASED REPRESENTATION LEARNING

In the training framework, the pre-trained baseline network are frozen, which provides feature representation ability thanks to their extensive pre-training on large public datasets such as ImageNet. We define  $\mathbf{W}^\top \in \mathbb{R}^{d \times p}$  serves as the projection matrix that transforms the raw feature space of training samples to the PDF space, facilitating comparisons among different classes.

We propose our learning objective to act as the training loss, replacing the conventional cross-entropy loss and guiding the network training process.

$$\begin{aligned} \mathcal{L}(\mathbf{W}) = \max\left(-\sum_{x \in D_t} \mathbb{1}_{\{y_b=t\}} \mathbf{K}(x - x_b) \right. \\ \left. + \sum_{x \in D_t} \mathbb{1}_{\{y_b \neq t\}} (\mathbf{K}(x - x_b) + m), 0\right) \end{aligned} \quad (1)$$

In Eq. 1, the set of trainable parameters is denoted by  $\mathbf{W}$ , which is implemented as a projection network in our design. In this context,  $x$  represents the PDF representations of all training sample, and  $\mathbf{x}_b$  signifies a batch of instances, while the  $y_b$  is the class label for  $\mathbf{x}_b$ . The symbol  $m$  is introduced as a PDF margin, ensuring that the PDF for positive samples with the same  $y_b$  exceeds that of negative instances with different class label by a safe margin. Eq. 1 is employed as the contrastive loss function within our framework. Notably, the probability values involved in the computation are expressed in logarithmic format. This approach not only stabilizes the training process but also prevents the values from experiencing underflow during back-propagation.

For the PDF-based classification model, if a test sample is nearer to a class in the PDF representation, the sample should be classified to the given class. Assuming there are  $k$  classes, we need to incorporate the privacy-preserving mechanism into the accuracy calculation.

For a sample  $x$ , compute the KDE estimates for each class  $j \in \{1, \dots, k\}$ :

$$\hat{p}_j(x) = \frac{1}{n_j} \sum_{i=1}^{n_j} K_h(x - x_{ij})$$

where  $n_j$  is the number of samples in class  $j$  and  $x_{ij}$  are the samples belonging to class  $j$ . Assign the sample  $x$  to the class with the highest KDE estimate:

$$\hat{y}(x) = \arg \max_j \hat{p}_j(x)$$

#### **Theorem 3.** BOUNDS ON MISCLASSIFICATION ERRORS

*The accuracy of KDE-based classification is affected by the amount of smoothing added to achieve  $\rho$ -zCDP. The trade-off between privacy and accuracy can be analyzed using the following bound for a projection network  $f : \mathbb{R}^d \rightarrow \mathbb{R}^p$  such that the misclassification error:*

$$\text{Err}(f) \leq 2 \exp\left(-\frac{m^2 \rho}{4 \ln(1/\delta)}\right)$$

*Higher  $\rho$  leads to more smoothing and thus lower classification accuracy. Balancing  $\rho$  appropriately ensures a reasonable trade-off between privacy and accuracy. The  $m$  is the PDF margin in the training process.*

The proof can be found in Appendix A.3.

## 4 EMPIRICAL STUDIES

### 4.1 DATASETS

The CIFAR10 dataset Krizhevsky (2009) comprises 32×32 color images across 10 distinct classes, encompassing 50,000 training examples and 10,000 testing examples. We train the model using the proposed privacy-preserving mechanism. This setup is crucial for demonstrating the effectiveness of privacy-preserving techniques in image classification tasks, ensuring that sensitive data is

safeguarded during the training process. We further conduct experiments on the real medical CheXpert dataset Irvin et al. (January 2019), a substantial collection comprising 224,316 chest X-rays from 65,240 patients. This dataset includes five classes representing various thoracic pathologies: (a) Atelectasis, (b) Cardiomegaly, (c) Consolidation, (d) Edema, and (e) Pleural Effusion. For our purposes, we reinitialize only the projection layer while keeping the other layers fixed, ensuring no impact on privacy leakage as described in Abadi et al. (October 2016).

## 4.2 EXPERIMENT SETTINGS

To evaluate both non-private and private deep learning models, we employ the area under the curve (AUC) metric. For our model, we set the projected PDF dimension to 64, as this configuration offers the best balance between prediction accuracy and training efficiency. The PDF margin in our learning process is set to 10 (in logarithmic scale) based on cross-validation on the CIFAR10 dataset Krizhevsky (2009). For the differentially private stochastic gradient descent (DP-SGD) methods, the training settings are configured according to the guidelines provided in Berrada et al. (August 2023).

An algorithm that satisfies  $\rho$ -zCDP also satisfies  $(\epsilon, \delta)$ -DP, where:

$$\epsilon = \rho + 2\sqrt{\rho \log(1/\delta)}$$

Based on this relationship, we will compare our method with other approaches under the same pre-trained models, such as ResNet-50, ViT-B/16, and DenseNet. In our experimental settings, the broken probability  $\delta$  is fixed at  $10^{-5}$ . It is important to note that in this context,  $\delta$  denotes the broken probability, which is distinct from the  $\delta$  used in the previous section to represent the confidence bound.

We will assess how varying  $\epsilon$  values in  $(\epsilon, \delta)$ -DP affect model performance, corresponding to different  $\rho$  values in  $\rho$ -zCDP. This evaluation will provide insights into the trade-offs between privacy guarantees and model efficacy.

## 4.3 RESULTS AND DISCUSSION ON CIFAR10

To validate the effectiveness of the proposed method and the theoretical analysis presented in Section 3, we conducted experiments using the CIFAR-10 dataset. The DP-SGD methods adhere to the settings outlined in Berrada et al. (August 2023), with the primary difference being the backbone models used: ResNet18, ResNet-50, and ViT-B/16. All models were pre-trained on ImageNet-21K. For the Non-Private method, we utilized the optimal settings without enforcing any smoothing. Several observations can be drawn from the results. The batch size was set to 4096 to favor DP-SGD, although we did not explore extremely large batch sizes due to computational limitations.

**Model Capacity and Smoothing Resilience** our method exhibits a narrow gap with the Non-Private model, lagging by just 2-3 percent when  $\epsilon = 8$ . However, for the DP-SGD method, even under the setting of  $\epsilon = 8$ , its performance is still inferior to our  $\epsilon = 3$  counterpart. This indicates that our method provides significantly better utility for a given privacy budget. Our experiments reveal that larger models, such as ResNet50 and ViT-B/16, exhibit higher resilience to the smoothing introduced by differential privacy mechanisms. This resilience can be attributed to the models' enhanced capacity to learn robust feature representations, which can better withstand the perturbations caused by smoothing. This finding aligns with existing literature that suggests a direct correlation between model complexity and its ability to generalize from noisy data Kaplan et al. (2020). The significant performance gap between our method and the traditional DP-SGD approach, even at higher  $\epsilon$  values, underscores the effectiveness of our approach in balancing the privacy-utility trade-off. Traditional DP-SGD methods often suffer from a substantial reduction in utility due to the added smoothing, which is necessary to ensure privacy. Our method, however, achieves competitive accuracy with a much lower privacy budget, demonstrating a more efficient utilization of the privacy budget.

**Impact of  $\epsilon$  on Performance** The parameter  $\epsilon$  in differential privacy quantifies the trade-off between privacy and utility. Our results show that while increasing  $\epsilon$  improves the performance of the DP-SGD method, it still lags behind our method with a lower  $\epsilon$ . This suggests that our method can

432 achieve higher accuracy even under stricter privacy constraints, making it more suitable for applica-  
 433 tions where both high utility and stringent privacy guarantees are required. Our method can allocate  
 434 the budget at instance level, which will not waste privacy on the samples in dense PDF regions and  
 435 degrade the performance. The smoothing effect for the sparse samples is economically allocated as  
 436 the backbone can provide more precise representation, resulting in less performance degradation as  
 437 the privacy budget increases.

#### 438 4.4 RESULT AND DISCUSSION ON CHEXPART

439 In the realm of medical image analysis, ensuring patient privacy while maintaining high model  
 440 performance is a paramount concern. The CheXpert dataset, a comprehensive collection of thoracic  
 441 condition images, serves as an excellent benchmark for evaluating privacy-preserving techniques.  
 442 Our proposed instance-level PDF smoothing model exhibits significant advancements over existing  
 443 methods such as DP-SGD and non-private models.  
 444

445 The experimental results depicted in Figure 2 reveal several critical insights. Our proposed PDF  
 446 smoothing model consistently outperforms DP-SGD under the same backbone and experiment set-  
 447 tings and has a narrow gap with non-private model. This indicates that our method effectively  
 448 balances the trade-off between privacy and utility. As expected, higher values of  $\epsilon$  correlate with im-  
 449 proved model accuracy due to reduced smoothing factor. This is a well-documented phenomenon in  
 450 differential privacy literature. However, our adaptive DP model demonstrates superior performance  
 451 even at lower  $\epsilon$  values, underscoring its robustness. For thoracic conditions with privacy budgets set  
 452 to  $\epsilon = 2$  and  $\epsilon = 8$ , our model achieves an average testing AUC of 85% across all five labels. In  
 453 contrast, DP-SGD lags behind, achieving approximately 70%. This significant margin underscores  
 454 the efficacy of our adaptive approach in maintaining high accuracy while preserving privacy.  
 455

456 **Instance-Level Privacy Allocation** A key strength of our model lies in its instance-level privacy  
 457 allocation. By dynamically adjusting the privacy budget for each data instance based on its prob-  
 458 ability density, we ensure that samples in sparse regions receive higher privacy protection. This  
 459 targeted approach is more economical and effective compared to uniform noise addition strategies.  
 460 Medical images often contain critical and unique information that might appear infrequently (i.e.,  
 461 in sparse regions of the data distribution). Our model enhances privacy for these samples without  
 462 significantly compromising their utility. This is particularly vital in medical contexts where rare con-  
 463 ditions need accurate representation and analysis. Conversely, in dense regions where data points  
 464 are abundant, excessive smoothing injection can severely disrupt the model’s learning process. Our  
 465 method strategically injects less smoothing in these areas, preserving the inherent structure and re-  
 466 lationships within the data. This selective smoothing ensures that the overall performance of the  
 467 model remains high.

468 **Class Disparity Improved by Instance level Probability Smoothing** In Figure 3, we compare  
 469 class disparities between our private models (at  $\epsilon=8$ ) and non-private baselines. Both model types  
 470 show similar disparities across subgroups, with private models not systematically worse in terms of  
 471 AUC disparities. Notably, smaller subgroups, such as *Atelectasis*, suffer more performance degrada-  
 472 tion due to their sparse feature distribution. Contrary to prior works, our private models demonstrate  
 473 significant improvement over the DP-SGD method, achieving better group fairness outcomes than  
 474 non-private baselines. This suggests that increased disparities in private models observed previously  
 475 are not inherent. Our instance-level PDF smoothing model effectively balances privacy and per-  
 476 formance, protecting sparse samples economically while minimizing disruption in dense regions.  
 477 Future work could extend this method to other medical datasets and real-time clinical systems, and  
 478 explore adaptive mechanisms for dynamic privacy budget adjustments.

## 479 5 CONCLUSION

480 In conclusion, we introduce a novel instance-level PDF smoothing model that enhances privacy  
 481 preservation while maintaining high performance in medical image analysis. Our approach of-  
 482 fers profound theoretical insights into sample complexity, smoothing factors, and error bounds for  
 483 achieving privacy budgets. Our experimental results demonstrate the efficacy of our model in balanc-  
 484 ing privacy and utility, dynamically adjusting privacy parameters to protect samples in sparse regions  
 485

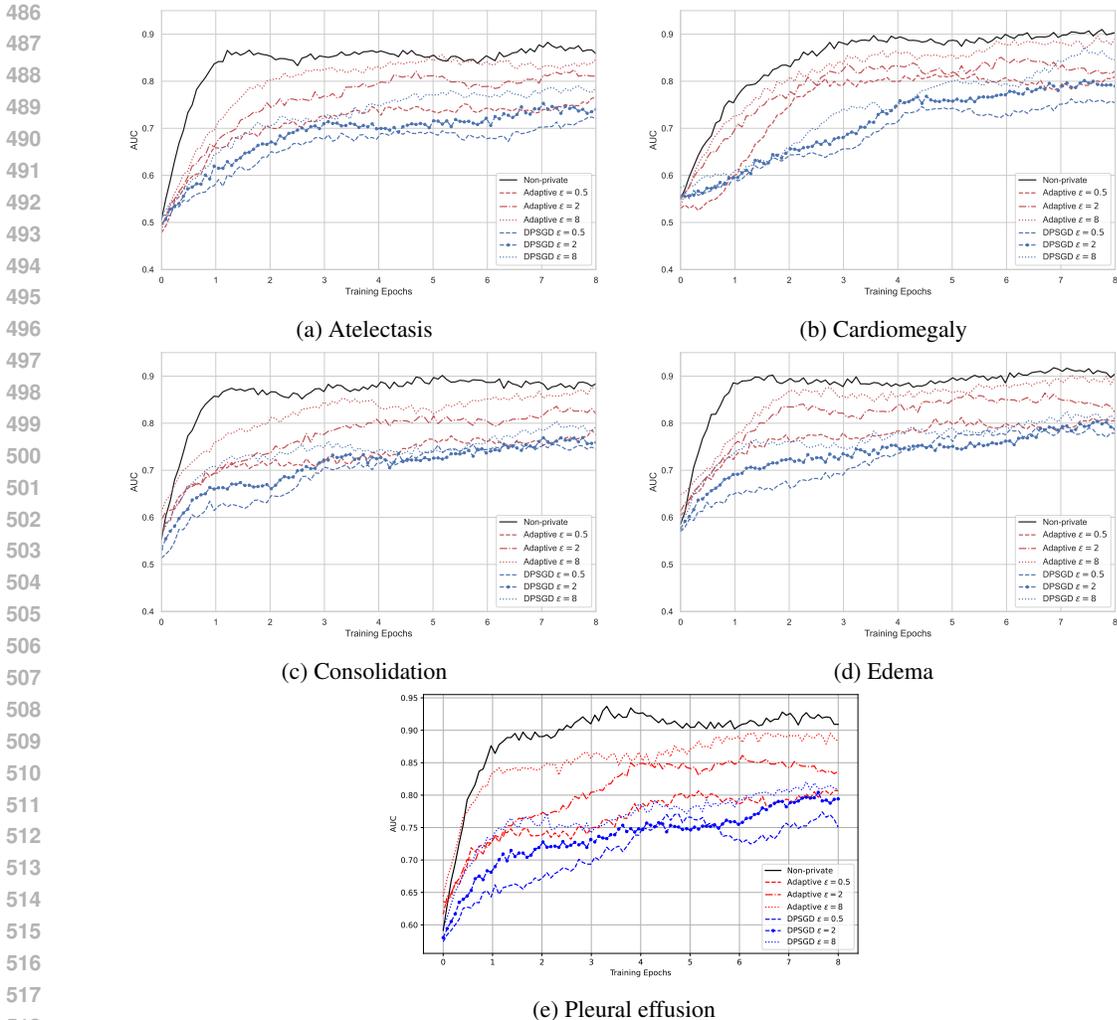


Figure 3: Comparison between non-private model, DP-SGD and our adaptive DP model under different privacy budget settings. The x-axis the training iterations in epochs and the y-axis is the model performance in AUC.

Table 1: Performance of Different Algorithms on CIFAR-10

DP Algorithm	ResNet18	ResNet50	ViT-B/16
DPSGD ( $\epsilon = 3$ )	77.8	81.5	92.4
DPSGD ( $\epsilon = 8$ )	82.9	84.1	94.0
Ours ( $\epsilon = 3$ )	85.2	89.9	94.2
Ours ( $\epsilon = 8$ )	90.2	92.9	95.2
Non-Private	93.2	94.9	97.1

and addressing class disparities in medical datasets. The adaptive DP model not only preserves the utility of critical medical information but also improves overall model fairness. Our findings open several avenues for future research, including extending to other medical datasets and integrating with real-time clinical systems. Our instance-level PDF smoothing model represents a significant step forward in privacy-preserving medical image analysis, offering a robust, comprehensive solution to the challenges posed by differential privacy. We hope our work inspires further innovation, contributing to safer and more effective healthcare technologies. Limitations list in Appendix A.4.

## REFERENCES

- 540  
541  
542 Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and  
543 Li Zhang. Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC  
544 Conference on Computer and Communications Security. ACM, October 2016.
- 545  
546 Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and  
547 Li Zhang. Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC  
548 Conference on Computer and Communications Security, Vienna, Austria, October 2016.
- 549  
550 Rohan Anil, Badih Ghazi, Vineet Gupta, Ravi Kumar, and Pasin Manurangsi. Large-scale differen-  
551 tially private bert. arXiv preprint arXiv:2108.01624 [cs.LG], 2021.
- 552  
553 Eugene Bagdasaryan, Omid Poursaeed, and Vitaly Shmatikov. Differential privacy has disparate  
554 impact on model accuracy. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence  
555 d’Alch’e-Buc, Emily B. Fox, and Roman Garnett (eds.), Advances in Neural Information Processing  
556 Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019,  
557 December 8-14, 2019, Vancouver, BC, Canada, pp. 15453–15462, 2019.
- 558  
559 Borja Balle, Giovanni Cherubin, and Jamie Hayes. Reconstructing training data with informed  
560 adversaries. In Symposium on Security and Privacy (SP). IEEE, May 2022.
- 561  
562 Raef Bassily, Adam D. Smith, and Abhradeep Thakurta. Private empirical risk minimization: Effi-  
563 cient algorithms and tight error bounds. In 55th Annual Symposium on Foundations of Computer  
564 Science. IEEE, October 2014.
- 565  
566 Loubna Berrada, Samin De, Junfeng Shen, Jiajun Hayes, Rory Stanforth, David Stutz, Pushmeet  
567 Kohli, Samuel Smith, and Borja Balle. Unlocking accuracy and fairness in differentially private  
568 image classification. arXiv preprint arXiv:2308.10888, August 2023.
- 569  
570 Tom B. Brown and Benjamin Mann. Language models are few-shot learners. In Hugo Larochelle,  
571 Marc’ Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), Advances  
572 in Neural Information Processing Systems 33: Annual Conference on Neural Information  
573 Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual, 2020.
- 574  
575 Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions,  
576 and lower bounds. In Martin Hirt and Adam D. Smith (eds.), Theory of Cryptography -  
577 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016,  
578 Proceedings, Part I, volume 9985 of Lecture Notes in Computer Science, pp. 635–658. Springer,  
579 2016.
- 580  
581 Nicholas Carlini, Florian Tram’er, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B.  
582 Brown, Dawn Song,  
583 ’Ulfar Erlingsson, Alina Oprea, and Colin Raffel. Extracting training data from large language  
584 models. In 30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021, pp.  
585 2633–2650. USENIX Association, 2021.
- 586  
587 Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tram  
588 ’er. Membership inference attacks from first principles. In 43rd IEEE Symposium on Security  
589 and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022, pp. 1897–1914. IEEE, 2022.
- 590  
591 Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tram  
592 ’er, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion  
593 models. Preprint arXiv:2301.13188 [cs.CR], 2023.
- 594  
595 Rachel Cummings, Varun Gupta, Dhamma Kimpara, and Jamie Morgenstern. On the compatibil-  
596 ity of privacy and fairness. In Adjunct Publication of the 27th Conference on User Modeling,  
597 Adaptation and Personalization, pp. 309–315. ACM, jun 2019.

- 594 Cynthia Dwork. Differential privacy. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and  
595 Ingo Wegener (eds.), Automata, Languages and Programming, 33rd International Colloquium,  
596 ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II, volume 4052 of Lecture Notes  
597 in Computer Science, pp. 1–12. Springer, 2006.
- 598 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity  
599 in private data analysis. In Theory of cryptography conference, volume 7, pp. 17–51, may 2006.
- 600 Tom Farrand, Fatemehsadat Mireshghallah, Sahib Singh, and Andrew Trask. Neither private nor  
601 fair: Impact of data imbalance on utility and fairness in differential privacy. In Proceedings of  
602 the 2020 Workshop on Privacy-Preserving Machine Learning in Practice, pp. 15–19. ACM, nov  
603 2020.
- 604 Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik  
605 Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest  
606 radiograph dataset with uncertainty labels and expert comparison. In Proceedings of the AAAI  
607 Conference on Artificial Intelligence, volume 33, pp. 590–597, Honolulu, HI, January 2019.
- 608 Jared Kaplan et al. Scaling laws for neural language models. arXiv preprint arXiv:2001.08361,  
609 2020.
- 610 Michael Kirchhof, Karsten Roth, Zeynep Akata, and Enkelejda Kasneci. A non-isotropic proba-  
611 bilistic take on proxy-based deep metric learning. In European Conference on Computer Vision,  
612 2022.
- 613 Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University  
614 of Toronto, 2009.
- 615 Alexis R Santos-Lozada, Jeffrey T Howard, and Ashton M Verdery. How differential privacy will  
616 affect our understanding of health disparities in the united states. Proc. Natl. Acad. Sci. U. S. A.,  
617 117(24):13405–13412, jun 2020.
- 618 Amartya Sanyal, Yaxi Hu, and Fanny Yang. How unfair is private learning? In James Cussens  
619 and Kun Zhang (eds.), Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial  
620 Intelligence, volume 180 of Proceedings of Machine Learning Research, pp. 1738–1748. PMLR,  
621 01–05 Aug 2022.
- 622 Laleh Seyyed-Kalantari, Guanxiong Liu, Matthew B. A. McDermott, Irene Y. Chen, and Marzyeh  
623 Ghassemi. Chexclusion: Fairness gaps in deep chest x-ray classifiers. In Biocomputing 2021:  
624 Proceedings of the Pacific Symposium, Kohala Coast, Hawaii, USA, January 3-7, 2021. World-  
625 Scientific, 2021.
- 626 Vinith M Suriyakumar, Nicolas Papernot, Anna Goldenberg, and Marzyeh Ghassemi. Chasing your  
627 long tails: Differentially private prediction in health care settings. In Proceedings of the ACM  
628 Conference on Fairness, Accountability, and Transparency. ACM, mar 2021.
- 629 Florian Tram  
630 ‘er and Dan Boneh. Differentially private learning needs better features (or much more data). In  
631 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria,  
632 May 3-7, 2021. OpenReview.net, 2021.
- 633 Da Yu, Huishuai Zhang, Wei Chen, Jian Yin, and Tie-Yan Liu. Large scale private learning via  
634 low-rank reparametrization. In Marina Meila and Tong Zhang (eds.), Proceedings of the 38th  
635 International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event,  
636 volume 139 of Proceedings of Machine Learning Research, pp. 12208–12218. PMLR, 2021.
- 637 Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers.  
638 In Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, jun 2022.
- 639 Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding  
640 deep learning (still) requires rethinking generalization. Commun. ACM, 64(3):107–115, 2021.
- 641 Liangwei Zhang, Jing Lin, and Ramin Karim. Adaptive kernel density-based anomaly detection for  
642 nonlinear systems. Knowledge-Based Systems, 139:50–63, 2018.
- 643
- 644
- 645
- 646
- 647

## 648 A APPENDIX

649 You may include other additional sections here.

### 652 A.1 PROOF OF THEOREM 1

654 *Proof.* Zero-Concentrated Differential Privacy (zCDP) is defined such that a mechanism  $M$  satisfies  
 655  $\rho$ -zCDP if the privacy loss random variable  $L$  satisfies  $\mathbb{E}[e^L] \leq e^\rho$ . The privacy loss  $L$  is given by  
 656  $L = \log \left( \frac{M(D_1)}{M(D_2)} \right)$ , where  $D_1$  and  $D_2$  are neighboring datasets differing by one element.  
 657

658 Consider the Gaussian mechanism  $M$  that adds smoothing  $N_i \sim \mathcal{N}(0, \sigma^2)$  to each sample  $x_i$ .

660 The privacy loss  $L$  for the Gaussian mechanism is given by  $L = \log \left( \frac{p(M(D_1))}{p(M(D_2))} \right)$ . For two neigh-  
 661 boring datasets  $D_1$  and  $D_2$ , the difference in the KDE estimates is small, and the privacy loss  $L$  can  
 662 be approximated using the Rényi divergence.

663 The Rényi divergence  $D_\alpha$  between two distributions  $P$  and  $Q$  is defined as:

$$665 D_\alpha(P\|Q) = \frac{1}{\alpha - 1} \log \left( \int p(x)^\alpha q(x)^{1-\alpha} dx \right)$$

666 For  $\alpha > 1$ , the Rényi divergence is a measure of the distance between two distributions.

667 The  $\rho$ -zCDP condition can be expressed in terms of Rényi divergence:

$$670 D_\alpha(P\|Q) \leq \rho$$

671 For  $\alpha = 2$ , the Rényi divergence simplifies to:

$$672 D_2(P\|Q) = \log \left( \int p(x)^2 q(x)^{-1} dx \right)$$

673 For the Gaussian mechanism, the smoothing  $N_i \sim \mathcal{N}(0, \sigma^2)$  ensures that the privacy loss  $L$  is  
 674 controlled. The Rényi divergence for the Gaussian mechanism is given by:

$$675 D_2(P\|Q) \leq \frac{\Delta^2}{2\sigma^2}$$

676 where  $\Delta$  is the sensitivity of the function being perturbed.

677 The sensitivity  $\Delta$  is the maximum change in the PDF estimate due to adding or removing one  
 678 sample.

679 To achieve  $\rho$ -zCDP, we need  $D_2(P\|Q) \leq \rho$ . Substituting the Rényi divergence for the Gaussian  
 680 mechanism, we get:

$$681 \frac{\Delta^2}{2\sigma^2} \leq \rho$$

682 Solving for  $\sigma^2$ , we have:

$$683 \sigma^2 \geq \frac{\Delta^2}{2\rho}$$

684 To account for the confidence bound  $\delta$ , we use the fact that the tail probability of the Gaussian  
 685 distribution is controlled by  $\delta$ :

$$686 \Pr[L > \rho] \leq \delta$$

687 Using the Chernoff bound, we obtain:

$$688 \sigma^2 \geq \frac{2 \ln(1/\delta)}{\rho}$$

689  $\square$

## A.2 PROOF OF THEOREM 2

*Proof.* The change in Rényi divergence  $\Delta D_\alpha$  can be bounded using the properties of the PDF estimate:

$$\Delta D_\alpha = \frac{1}{\alpha - 1} \left( \log \left( \int \hat{p}_{-k}(x)^\alpha Q(x)^{1-\alpha} dx \right) - \log \left( \int \hat{p}(x)^\alpha Q(x)^{1-\alpha} dx \right) \right)$$

For large  $n$ , the change in the KDE estimate due to removing a single sample can be approximated as:

$$\hat{p}_{-k}(x) \approx \hat{p}(x) - \frac{1}{n} K_h(x - x_k)$$

The change in Rényi divergence can be bounded using the stability of the KDE estimate:

$$\Delta D_\alpha \leq \frac{1}{\alpha - 1} \log \left( 1 + \frac{C}{n} \right)$$

where  $C$  is a constant that depends on the kernel and the bandwidth  $h$ .

To ensure  $\rho$ -zCDP, we need:

$$\frac{1}{\alpha - 1} \log \left( 1 + \frac{C}{n} \right) \leq \rho\alpha$$

Simplifying this inequality:

$$\log \left( 1 + \frac{C}{n} \right) \leq \rho(\alpha - 1)$$

For small  $C/n$ , we can use the approximation:

$$\frac{C}{n} \leq \rho(\alpha - 1)$$

Solving for  $n$ :

$$n \geq \frac{C}{\rho(\alpha - 1)}$$

To derive the constant  $C$ , we need to consider the maximum possible change in the KDE estimate due to removing a single sample.

Maximum Change in KDE Estimate:

$$\Delta \hat{p}(x) \approx \frac{1}{n} K_h(x - x_k)$$

Bounding the Maximum Change, the maximum value of the kernel function  $K_h(x)$  is:

$$\frac{1}{\sqrt{2\pi h^2}}$$

Therefore, the maximum change in the KDE estimate is:

$$\Delta \hat{p}(x) \leq \frac{1}{n} \cdot \frac{1}{\sqrt{2\pi h^2}}$$

Formulating  $C$ : The constant  $C$  can be formulated as:

$$C = \frac{1}{\sqrt{2\pi h^2}}$$

□

756 A.3 PROOF OF THEOREM 3  
757

758 *Proof.* To achieve  $\rho$ -zCDP, smoothing  $N_i \sim \mathcal{N}(0, s_i^2 \cdot \frac{2\ln(1/\delta)}{\rho})$  is added to each sample  $x_i$ :  
759

$$760 \quad x'_i = x_i + N_i$$

761 The privacy preserved PDF  $\hat{p}'(x)$  is:

$$762 \quad \hat{p}'(x) = \frac{1}{n} \sum_{i=1}^n K_h(x - x'_i)$$

763 The classification error is affected by the smoothing added to the PDF. It can be bounded by the PDF  
764 margin between different classes:  
765

$$766 \quad \text{Err}(f) \leq \Pr(|\hat{p}(x) - \hat{p}'(x)| > m)$$

767 The smoothing variance  $\sigma^2$  is:

$$768 \quad \sigma^2 = \frac{2\ln(1/\delta)}{\rho}$$

769 Thus, the classification error bound is given by:

$$770 \quad \text{Err}(f) \leq \Pr(|\hat{p}(x) - \hat{p}'(x)| > m) \leq \Pr(|N_i| > m)$$

771 Using the properties of the Gaussian distribution:

$$772 \quad \Pr(|N_i| > m) \leq 2 \exp\left(-\frac{m^2}{2\sigma^2}\right)$$

773 Substituting  $\sigma^2 = \frac{2\ln(1/\delta)}{\rho}$ , we get:

$$774 \quad \Pr(|N_i| > m) \leq 2 \exp\left(-\frac{m^2 \rho}{4\ln(1/\delta)}\right)$$

775 Therefore, the classification error bound is:

$$776 \quad \text{Err}(f) \leq 2 \exp\left(-\frac{m^2 \rho}{4\ln(1/\delta)}\right)$$

777 □

778 A.4 LIMITATIONS  
779

780 While our proposed framework demonstrates significant advancements in balancing privacy and  
781 utility in differentially private (DP) deep learning models, several limitations warrant discussion:  
782

- 783 1. **Generalization Across Domains:** Although our empirical studies on the CheXpert dataset  
784 show promising results, the generalization of our method to other domains and data types  
785 remains to be thoroughly validated. Medical imaging datasets, like CheXpert, may possess  
786 unique characteristics that are not representative of other datasets, potentially limiting the  
787 applicability of our findings to different domains such as natural language processing or  
788 other forms of visual data.
- 789 2. **Scalability and Computational Overhead:** Our instance-level smoothing mechanism,  
790 while effective, introduces additional computational overhead that may not be scalable for  
791 extremely large datasets or models with billions of parameters. The increased computa-  
792 tional requirements could hinder the practical deployment of our framework in resource-  
793 constrained environments.
- 794 3. **Hyperparameter Sensitivity:** The success of our method heavily relies on the careful tun-  
795 ing of hyperparameters related to instance-level smoothing and noise injection. Identifying  
796 the optimal settings for these parameters can be challenging and may require extensive  
797 experimentation, which could limit the ease of adoption for practitioners with limited re-  
798 sources or expertise in hyperparameter optimization.

- 810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826
4. **Impact on Training Time:** The necessity to balance privacy and utility often comes with a trade-off in training time. Our method, which includes additional smoothing processes and noise injection, may result in longer training times compared to non-DP models. This could be a critical limitation for applications requiring rapid model development and deployment.
  5. **Potential for Subpopulation Fairness:** Despite our efforts to mitigate accuracy disparities across subpopulations, the inherent complexity of ensuring fairness in DP models means that some level of disparity may still persist. Future work is needed to develop more robust methods for guaranteeing fairness across diverse and heterogeneous subpopulations.
  6. **Theoretical Assumptions:** Our theoretical insights and error bounds are based on certain assumptions regarding sample complexity and instance-level smoothing factors. The practical validity of these assumptions in real-world scenarios needs further empirical validation across a broader range of datasets and model architectures.
  7. **Privacy Budget Allocation:** The allocation of the privacy budget in DP models is a critical factor that influences the overall performance. Our framework provides guidelines for optimizing this balance, but the practical implementation of these guidelines can be complex and may not always achieve the desired outcomes in every context.

827  
828  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863

Addressing these limitations in future research will be crucial for advancing the practical applicability and robustness of differentially private deep learning models, ensuring that they can meet the dual challenges of high accuracy and fairness across a wide range of applications.