



**HAL**  
open science

# Exploring Machine Learning Privacy/Utility Trade-Off from a Hyperparameters Lens

Ayoub Arous, Amira Guesmi, Muhammad Abdullah Hanif, Ihsen Alouani,  
Muhammad Shafique

► **To cite this version:**

Ayoub Arous, Amira Guesmi, Muhammad Abdullah Hanif, Ihsen Alouani, Muhammad Shafique. Exploring Machine Learning Privacy/Utility Trade-Off from a Hyperparameters Lens. International Joint Conference on Neural Networks (IJCNN 2023), Jun 2023, Gold Coast, Australia. 10.1109/IJCNN54540.2023.10191743 . hal-04202435

**HAL Id: hal-04202435**

**<https://hal.science/hal-04202435v1>**

Submitted on 11 Sep 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Exploring Machine Learning Privacy/Utility trade-off from a hyperparameters Lens

Ayoub Arous<sup>1</sup>, Amira Guesmi<sup>1</sup>, Muhammad Abdullah Hanif<sup>1</sup>, Ihsen Alouani<sup>2</sup>, and Muhammad Shafique<sup>1</sup>

**Abstract**—Machine Learning (ML) architectures have been applied to several applications that involve sensitive data, where a guarantee of users’ data privacy is required. Differentially Private Stochastic Gradient Descent (DPSGD) is the state-of-the-art method to train privacy-preserving models. However, DPSGD comes at a considerable accuracy loss leading to sub-optimal privacy/utility trade-offs. Towards investigating new ground for better privacy-utility trade-off, this work questions; (i) if models’ hyperparameters have any inherent impact on ML models’ privacy-preserving properties, and (ii) if models’ hyperparameters have any impact on the privacy/utility trade-off of differentially private models. We propose a comprehensive design space exploration of different hyperparameters such as the choice of activation functions, the learning rate and the use of batch normalization. Interestingly, we found that utility can be improved by using Bounded RELU as activation functions with the same privacy-preserving characteristics. With a drop-in replacement of the activation function, we achieve new state-of-the-art accuracy on MNIST (96.02%), FashionMnist (84.76%), and CIFAR-10 (44.42%) without any modification of the learning procedure fundamentals of DPSGD.

**Index Terms**—DPSGD, bounded activation functions, deep learning, privacy, LayerNormalization, BatchNormalization

## 1. Introduction

Deep Learning (DL) models [1] have gone mainstream in recent years, due to their ability to achieve state-of-the-art performance on a variety of tasks [2] [3]. These models have been used in a variety of applications such as image processing and classification tasks, which require large amounts of data and computational resources to train. However, models also have the potential to compromise the privacy of data used in the training due to this complexity [4]. Furthermore, Machine Learning (ML) models are increasingly applied to solving problems that involve sensitive personal information, such as medical records, which leads to serious privacy concerns. Researchers have proposed several attacks on deep learning models such as membership inference attacks [5], model stealing attacks [6] and inversion attacks [7]. Along with these attacks and to address these privacy issues, a plethora of works have proposed new defenses to enhance the models privacy-preserving such as work in [8] which focus on providing a framework to detect photo privacy with

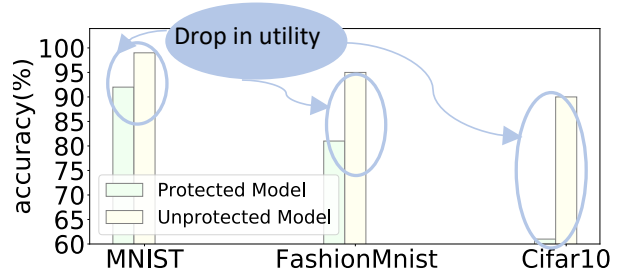


Figure 1: Comparison of accuracy between protected model with DPSGD and unprotected model for benchmark datasets.

a convolutional neural network using hierarchical features, and research in [9] which propose some Adversarial learning techniques for privacy preservation or the use of Generative adversarial networks (GANs) to design models to protect data privacy such as work in [10]. This race between attacks and defenses has been a challenge and many prior works that provide empirical defenses have been broken by these attacks. However, one of the most important defenses used in the state-of-the-art works and that shows great effectiveness in enhancing the model’s privacy is differential privacy [11] which is an emerging defense concept that offers us certain mathematical assurances regarding data privacy, and the most efficient algorithm that incorporates this notion is DPSGD [12], which is an extension of SGD (see section 2.1 for more details). Figure 1 shows the results of DPSGD algorithm on benchmark datasets.

As Figure 1 illustrates, there is a major drop in utility for the protected model with DPSGD compared to the unprotected model. So it is obvious that this privacy gain comes with a price of utility loss. Several related works have tried to find better privacy/utility sweet spots such as work in [13] which tries to enhance the convergence of DPSGD and research in [14] which provides a privacy-enhanced matrix factorization with local differential privacy. However, little attention has been directed to the exploration of the potential inherent impact of the hyperparameter choice on the model’s privacy-preserving properties. To the best of our knowledge, only a few previous studies explored the importance of model hyperparameters in terms of privacy-preserving and resistance to privacy attacks such as MIA.

### 1.1. State-of-the-Art and Their Limitations

Providing a private model comes with the price of a drop in accuracy. So much work has been done to tackle this

<sup>1</sup>Ayoub Arous, Amira Guesmi, Muhammad Abdullah Hanif, and Muhammad Shafique are with eBrain Laboratory, NYU Abu Dhabi, UAE

<sup>2</sup>Ihsen Alouani is with CSIT, Queen’s University Belfast, UK

problem and try to improve the overall trade-off between privacy and accuracy and push the model to converge to higher accuracy with a lower privacy budget. To tackle this problem, prior work tried to understand the cause of a drop in utility with the private model. For example, [15] provides a theoretical explanation of this drop in accuracy and shows the impact of the added steps in the differential private model in terms of convergence and highlight specifically the effect of per simple clipping on the convergence and try to provide a new variety of clipping called global clipping based on batch clipping rather than per simple clipping and show how their method improved this trade-off between privacy and accuracy.

Besides providing new private algorithms some analyses have been conducted to investigate the effects of some parameters on model’s outputs which is our case, such as research in [16] and in [17] which investigates only one activation function, and only one metric is used to evaluate its effectiveness, which is DPSGD. Another work in [18] investigates the impact of batch size and clipping bound on the model’s utility but their analysis has been conducted only on a small model with one dataset (MNIST). Another research in [19] shows how a hyperparameter tuning for model parameters could dramatically help DPSGD but this work focused more on the optimization part and provided three optimization algorithms to tune the model and did not focus on providing a wide framework analysis for the model’s hyperparameter. In our scenario compared to work in [16] and in [17], we will provide a set of activation functions and contrast them in the contexts of DPSGD and an unprotected model, the impact in the case of an unprotected model will be evaluated using a membership inference attack. Additionally, we offer a theoretical justification for why different activation functions are more effective, more specifically the case of the bounded activation function in DPSGD. By doing this, we are providing a more thorough view of those situations, which will aid more people in trying to implement those concepts. Moreover, we will try to conduct our analysis of what is commonly used as datasets and models in a state-of-the-art deep learning analysis.

Finally, the most recent state-of-the-art evaluation of hyperparameters and their effects on model performance lacked a broad framework analysis to clarify those effects and give a more thorough review of the topic. We are motivated to try to provide a better framework analysis and to try to better explore this topic by those limitations and restricted analysis in this context.

## 1.2. Our Novel Contributions

Motivated by the importance of investigating new ground for better privacy/utility trade-off, in this paper, we ask the following 2 research questions:

- **Q1** – Do models’ hyperparameters have an inherent impact on privacy?
- **Q2** – If yes, can this be leveraged for better privacy/utility trade-offs in differentially private models?

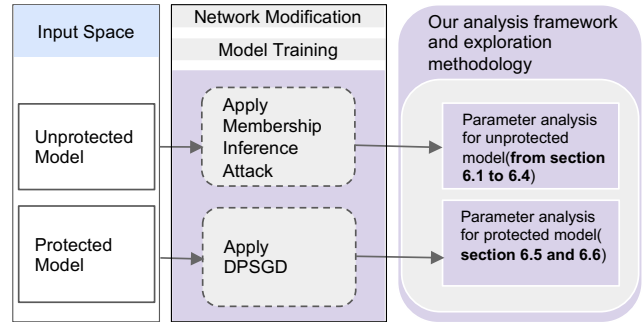


Figure 2: An overview of our major contributions.

To answer these questions, we propose a comprehensive exploratory study on a set of hyperparameters and investigate their impact on baseline models privacy, as well as their impact on differentially-private models’ accuracy under the same privacy budget.

We discover that several hyperparameters, as well as their combination, assure a greater gain in utility than the typically employed CNN design, thus we present a thorough analysis and demonstrate which parameter is appropriate in which case. Finally, our contributions will be:

- 1) We use membership inference attack to evaluate the impact of hyperparameter space on attack success rate with deep learning models.
- 2) we did a wide analysis in the case of unprotected model, we demonstrate the effect of hyperparameter space on the attack success rate of a membership inference attack and compare the baseline model to the model with the modified parameter.
- 3) We did an exploration of the hyperparameter space and its impact in the case of a protected model with DPSGD.
- 4) We demonstrate how models with certain hyperparameters could outperform state-of-the-art work when applied with DPSGD.

Finally Figure 2 summarizes all the major contributions.

## 2. Background

### 2.1. Differential Privacy

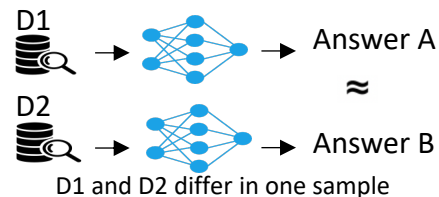


Figure 3: An overview of Differential Privacy.

A method for protecting privacy when training machine learning models on aggregate datasets is differential privacy

highlighted briefly in Figure 3. A survey of the significance of privacy in deep learning can be found in the work [20]. It is defined by the application-specific notion of neighboring databases. For instance, each training dataset used in our study contains a set of image-target pairs. If two of these sets only differ in one item, i.e., if one image-label pair is present in one set but not the other, then we say that the two sets are neighbors.

Definition 1. An algorithm  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$  which run on the input space  $\mathcal{D}$  and range  $\mathcal{R}$  is called  $(\epsilon, \delta)$ -differential privacy if for any two adjacent datasets  $d, d' \in \mathcal{D}$  and for any subset of outputs  $S \subseteq \mathcal{R}$  it satisfies the following expression:

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

$\epsilon$  is a measure of privacy that is inversely correlated to the degree of model’s privacy. Therefore, with higher epsilon values, protection decreases and user data exposure becomes more likely. DPSGD is a well-known algorithm that uses this idea of differential privacy.

Similar to conventional stochastic gradient descent, DPSGD [21] begins by selecting a batch of data samples at random, computing their gradients, and then updating the models’ parameters. The main differences are where Gaussian noise is injected and when the gradient norms are clipped. DPSGD operations include clipping gradients, adding Gaussian noise, and averaging the gradients. As we train, we aggregate the gradients from several mini-batches, and this, too, is differentially private due to the composition properties of differential privacy.

The Moments Accountant system then gets to work to better keep track of privacy loss. The moments accountant [22] yields a substantially smaller value of and, hence, a much greater assurance on privacy for a given noise level and number of training steps. The moments accountant, however, allows DPSGD to be run for a certain privacy budget for much more iterations. No matter how we quantify the privacy loss, the learnt model is the same.

## 2.2. Membership inference attack

To determine whether a data sample  $x$  was used to train a target model  $f_{target}$ , membership inference attacks (MIAs) are applied against deep learning models. Consequently, MIA causes a privacy leakage right away, giving the adversaries access to vital training data. For instance, in the real world,  $x$  could be a patient’s medical file or a person. Because of MIA, the attackers can detect whether another person or a specific patient’s clinical file was used to create a model for a certain condition. This clearly violates both privacy and confidentiality. When the adversary has black-box access to the target model, as is the case in the most common attack scenario, the attackers first train a shadow model  $f_{shadow}$  using a shadow dataset  $\mathcal{D}_{train}^{shadow}$ , which achieves the same objective as  $f_{target}$  (e.g., classification). The attackers then use  $\mathcal{D}_{train}^{shadow}$  (training data) and  $\mathcal{D}_{test}^{shadow}$  (non-training data) to query  $f_{shadow}$ , and they

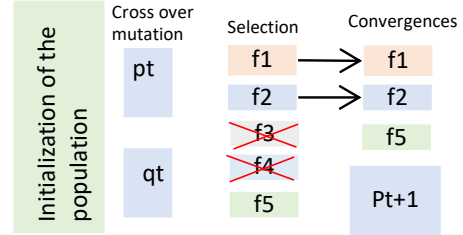


Figure 4: An overview of genetic algorithm. A random starting population of chromosomes is generated. Then results for those chromosomes are determined using a performance metric. The best values (with low cost) are subsequently put through reproduction, crossover, and mutation processes.

successfully retrieve the results. They can create an attack model, in which answers from training data are labeled as 1, and responses from non-training are labeled as 0. Attackers use data sample  $x$  to query  $f_{target}$  at attack time, then use  $f_{attack}$  to determine whether  $x$  is used in training the target model or not based on the results from  $f_{target}$ . In our exploration this technique is going to be used and we will provide later implementation details, to evaluate privacy leakage via attack success rate which is an indicator in  $[0, 1]$  that indicates more privacy leakage with higher values.

## 2.3. Genetic Algorithm

The genetic algorithm [23] used in our approach to tune BoundedRELU added parameter (i.e., the threshold), which is based on natural selection, the mechanism that propels biological evolution, is a technique for resolving both limited and unconstrained optimization issues. At first, chromosomes are generated at random in a population. These parameters values are then submitted to the system model. Results for each set of parameters inside the population are acquired by running a simulation using a performance metric that is based on a cost function. Once each cost value has been determined, it is arranged in ascending order with the corresponding chromosomes. The best values are chosen once more based on which has the lowest cost, and they are then put through processes of reproduction, crossover, and mutation. Figure 4 shows the basic steps of the genetic algorithm.

## 3. Analysis framework

To conduct our analysis, we use the framework described in Figure 5. As the diagram illustrates, our work is divided into two major parts. In the first part, we conduct the analysis on an unprotected model. The structure of the models are mentioned in table 3 and 2, these models will be referred to as baseline models. We define the model parameters and perform the network modification (**step1**), then we train our target model (**step2**) and we use membership inference attack to evaluate attack success rate on those unprotected

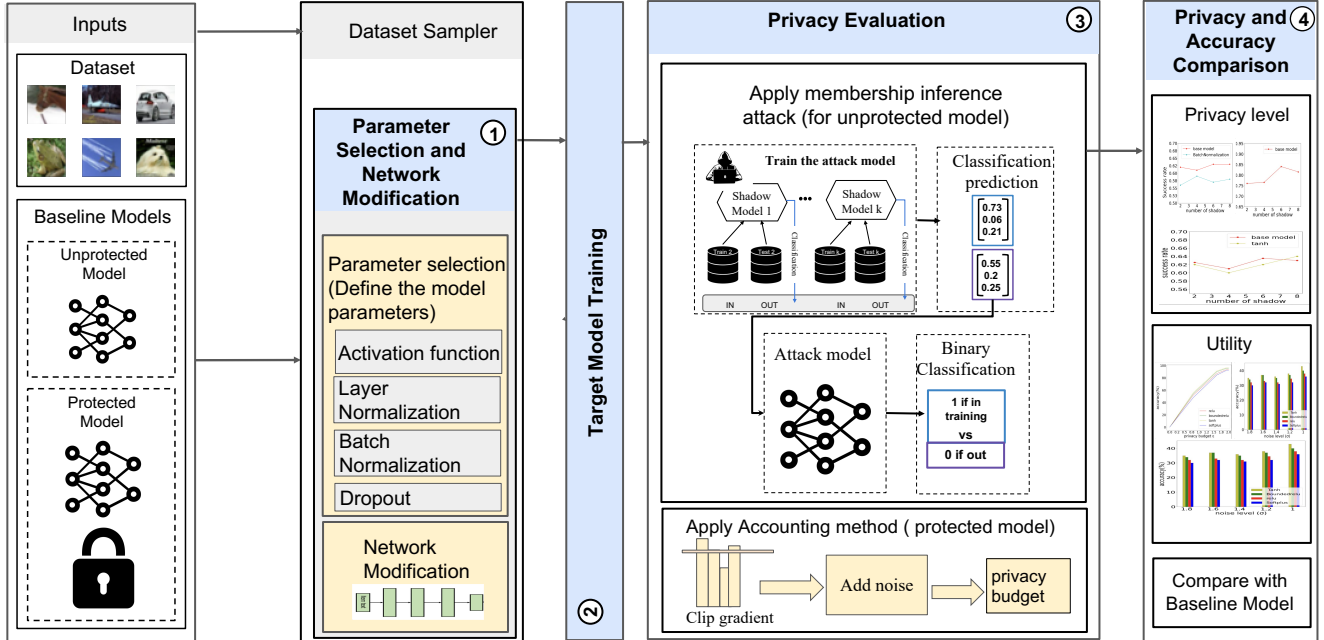


Figure 5: An overview of our analysis framework. For each model (protected/unprotected), we start by selecting the parameter to modify our network. Then, we train these modified networks and we apply the privacy evaluation metric (membership inference attack for the unprotected model and DPSGD for the protected model). Finally, we compare the output (privacy/utility) of those modified networks with the output of the baseline models to see the impact of the added parameter.

List of Parameters	
Parameter	Values
Learning rate	$\{e^{-2}, e^{-3}, e^{-4}\}$
Batch Normalization	$axis = 0$ and $axis = 1$
Layer Normalization	$axis = 0$ and $axis = 1$
Dropout	0.5
Activation functions	{RELU, BoundedRELU, Tanh}

TABLE 1: Hyperparameter space for conducting our analysis.

models (**step3**), finally for each experiment we evaluate the parameter performance by comparing it to the baseline models (**step4**). Table 1 represents the space of hyperparameters that are used in our analysis. In the second part, we conduct our analysis on the protected model with DPSGD and moment accountant as a tool to assess privacy leakage, and we follow the same methodology mentioned for the unprotected model to assess the impact of hyperparameter space on the trade-off between privacy and utility.

## 4. Experimental Methodology

The experiments are conducted on three datasets, MNIST [24] that has 60,000 training data and 10,000 test data, The CIFAR-10 [25] dataset consists of 60000  $32 \times 32$  color images in 10 classes, with 6000 images per class. There are 50000 training images and 10000 test images and the FashionMnist [26] dataset comprising  $28 \times 28$

grayscale images of 70000 fashion products from 10 categories, with 7000 images per category. In the case of the protected model, we implemented DPSGD for evaluating our protected model using Pytorch [27]. The experiments are performed by training the model for  $n$  number of epochs at each noise level and then drawing the privacy budget and accuracy of the last epoch, the noise levels added to the clipped gradient are in the following range:  $list = [1.8, 1.6, 1.4, 1.2, 1, 0.8]$ . The learning rate in the case of DPSGD is fixed to  $0.1$ ,  $delta = 1e - 5$  and  $max - per - sample - grad - norm = 1.0$ . In the case of the unprotected model we implemented a membership inference attack using Tensorflow [28] as follows:

The architecture of the target model is shown in Table 2 for FashionMnist and in Table 3 for CIFAR-10. We then construct the shadow models which inherit the same property of the target model to reveal its pattern to the attack model. We construct the shadow models with the same architecture as the target model for different datasets, and the train and test data of shadow models are sampled from the same target model dataset with the same size. After training the shadow models which we know their ground truth of membership inference attack, we label the output vector of shadow models with 0 if it is not in the training data and with 1 if it is in the training data, then we feed those output vectors with their labels to the attack model which is a binary classifier in our case SVM model. After training the attack model we conduct our attack on the target model and report the

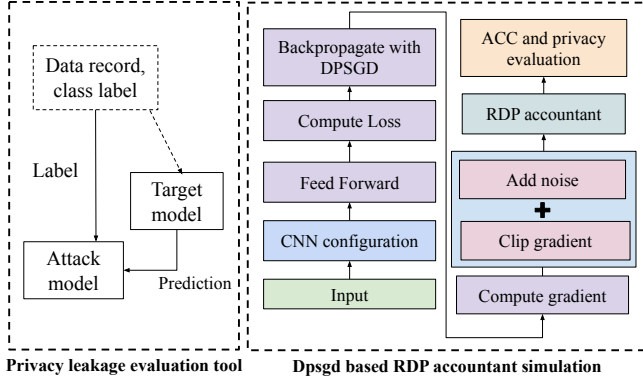


Figure 6: Experimental setup and tool.

attack success rate. The experimental setup for evaluating our methodology is shown in Figure 6.

For MNIST, we set the threshold of BoundedRELU  $a = 2$  for a protected model. For FashionMnist, we set  $a = 3$  and for CIFAR-10 the best performance is for  $a = 4$ . For DPSGD, we use the Opacus package with the RDP accountant method for privacy budget evaluation. For the CNN configuration, the architecture is shown in Table 2.

Layer	Parameters
Convolution	16 filters of 8x8, strides 2
Max-Pooling	2x2
Convolution	32 filters of 4x4, strides 2
Max-Pooling	2x2
Fully connected	32 units
Softmax	10 units units

TABLE 2: Convolutional model architecture.

The architecture mentioned in Table 2 is used for MNIST while the architecture mentioned in Table 3 is used to generate the results for CIFAR-10. For the genetic algorithm, the basic block for CIFAR-10 is highlighted in Table 3 while we maintain the same architecture in Table 2 for MNIST dataset. The results for the baseline models are shown in Figure 7.

Layer	Parameters
Convolution	32 filters of $3 \times 3$ , strides 1
Max-Pooling	2x2
Convolution	64 filters of $3 \times 3$ , strides 1
Max-Pooling	2x2
Convolution	128 filters of $3 \times 3$ , strides 1
Max-Pooling	2x2
Convolution	256 filters of $3 \times 3$ , strides 1
Max-Pooling	2x2
Fully connected	32 units
Softmax	10 units units

TABLE 3: CIFAR-10 model architecture for the Genetic Algorithm.

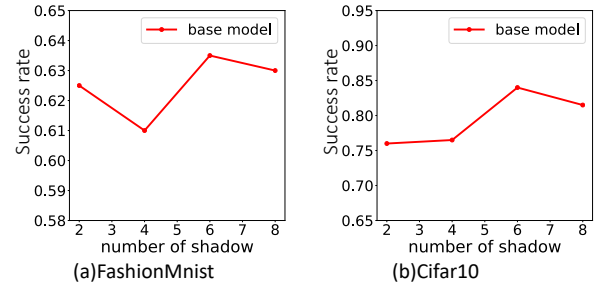


Figure 7: Results for the baseline models.

## 5. Experimental Results

### 5.1. Impact of varying activation functions on attack success rate for an unprotected model

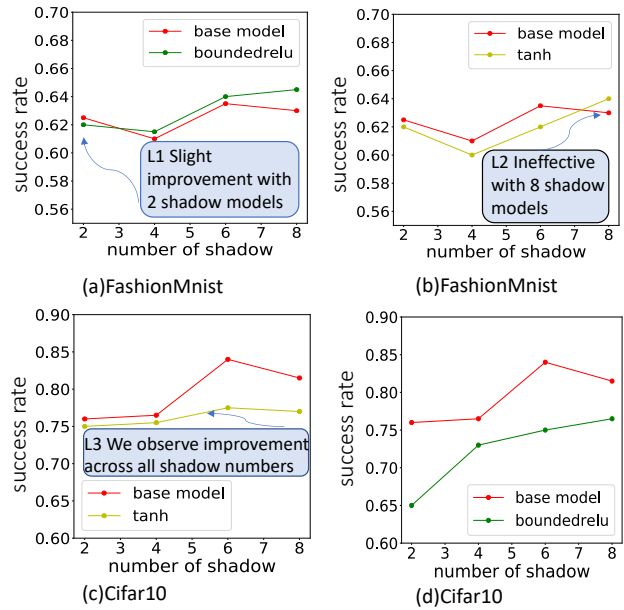


Figure 8: Impact of varying activation functions on attack success rate of membership inferences attack for an unprotected model.

Figure 8 illustrates the impact of varying activation functions in the case of unprotected models on the attack success rate of membership inference attack. To assess the impact of activation functions on the attack success rate, we conducted our experiments on the models mentioned in Table 2 for FashionMnist and Table 3 for CIFAR-10. We replace RELU with BoundedRELU and tanh and then compare attack success rate by varying each time in the number of shadow models. We notice improvement for those activation functions in robustness to the attack for the large model (See L2) compared to the baseline model with RELU

while the simple model shows no improvement in robustness with respect to those activation functions (See L1). The significant improvement is in the case of BoundedRELU with CIFAR-10 which reduced the attack success rate by almost 10% for a  $numberofshadow = 2$ .

### 5.2. Impact of normalization on attack success rate for an unprotected model

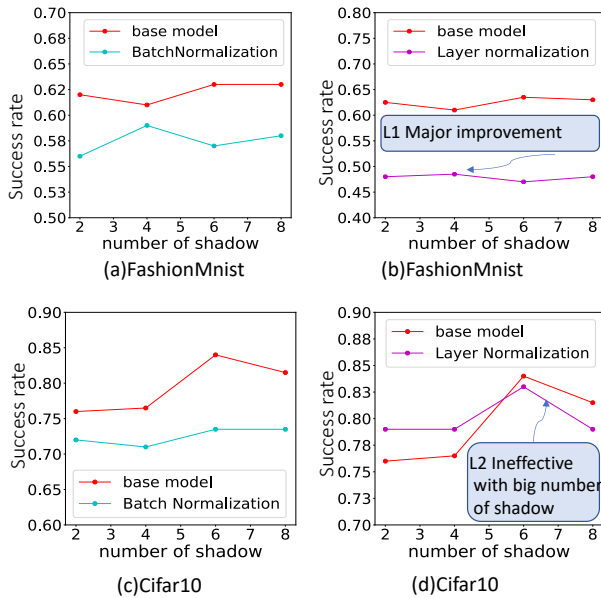


Figure 9: Impact of normalization layer on attack success rate of membership inferences attack for an unprotected model.

Except for the slight drop in robustness for a number of shadows equal to 6 and 8 in Figure 9 compared to the baseline model due to the experimental bias. All the results show a gain in robustness in the case of batch and layer normalization compared to the baseline model, which was expected since work in [29] explains the connection between overfitting and membership inference attack effectiveness, the bigger the gap between training and test accuracy the more the attack model can identify which data is used in the training set, and since the normalization technique help reduce overfitting as mentioned in [30] so this is why these techniques mitigate the effect of membership inference attack.

### 5.3. Impact of varying learning rate on attack success rate for an unprotected model

Figure 10 shows the impact of learning rate on attack success rate. We notice an improvement in target model robustness against membership inference attack with smaller learning rate compared to baseline model which is trained

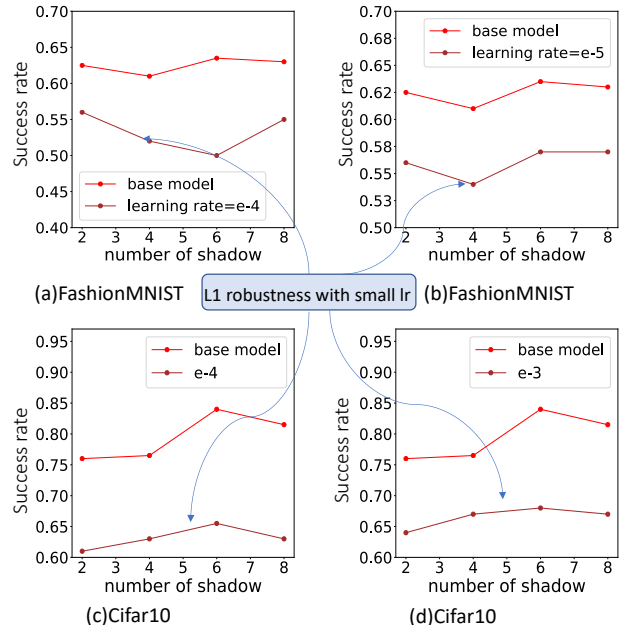


Figure 10: Impact of varying learning rate on attack success rate of membership inferences attack for an unprotected model.

with a learning rate= $e^{-2}$ , decreasing learning rate diminish roughly attack success rate. This is may be due to the fact that when learning rate is too large often moves too far in the correct direction resulting in an instability of the convergence of our model which causes poor generalization accuracy as mentioned in [31] while a small learning rate improves generalization accuracy due to the limitation of those overcorrection with a large learning rate.

### 5.4. Impact of dropout on attack success rate for an unprotected model

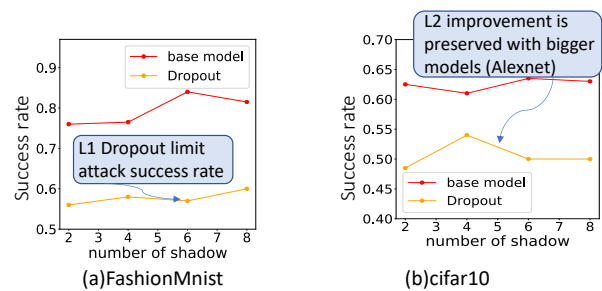


Figure 11: Evaluating Impact of dropout on attack success rate for an unprotected model.

Figure 11 shows the impact of dropout on attack success rate. As mentioned previously with normalization layer and their effect on limiting overfitting, applying dropout in the training phase help also limiting overfitting which reduces

the effectiveness of membership inference attack this is why we see this major improvement for dropout technique especially in the case of CIFAR-10 with a gain of 20% in robustness against the attack.

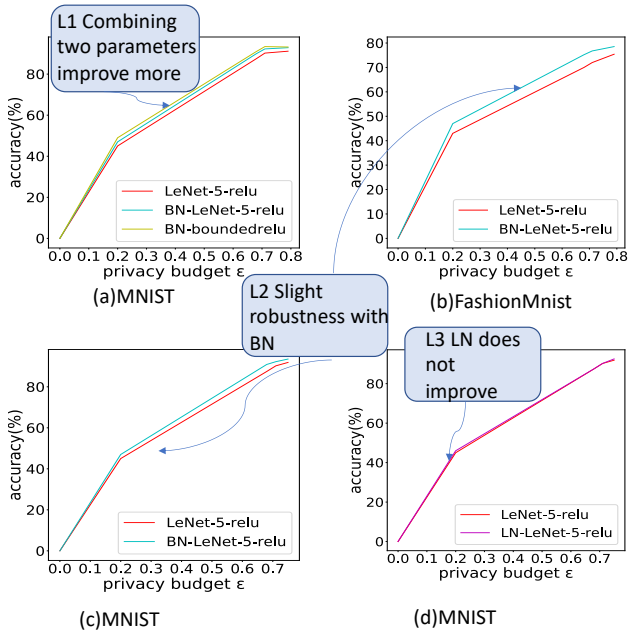


Figure 12: Impact of normalization layer on privacy vs utility trade-off for a protected model with DPSGD.

### 5.5. Impact of normalization layer on privacy vs utility trade-off for a protected model

As Figure 12 shows Batch normalization improves the utility for the same level of privacy (See L2 in Figure 12(b)) due to the fact of stabilizing the learning process and reducing the training time. We also see in Figure 12 (a) that combining BatchNormalization with BoundedRELU improves the overall trade-off (See L1). For Layer Normalization, we notice no improvement for all privacy levels.

### 5.6. Impact of varying activation functions on privacy vs utility trade-off for a protected model

The bar chart 14 includes the average accuracy at each noise level for 10 epochs, the different noise levels values are  $list = [1.8, 1.6, 1.4, 1.2, 1, 0.8]$ . The learning rate is fixed to 0.1,  $delta = 1e - 5$  and  $max - per - sample - grad - norm = 1.0$  and we set the value  $a = 2$  which is the parameter of BoundedRELU.

For Figure 13, we trained our model at each noise level for 4 epochs the last noise level is trained for 10 epochs, the different noise levels values are  $list_2 = [1.4, 1.2, 1, 0.8]$  and each drawn point is the privacy budget and accuracy of the last epoch. The learning rate is fixed to 0.1,  $delta = 1e - 5$ , and  $max - per - sample - grad - norm = 1.0$  and we set

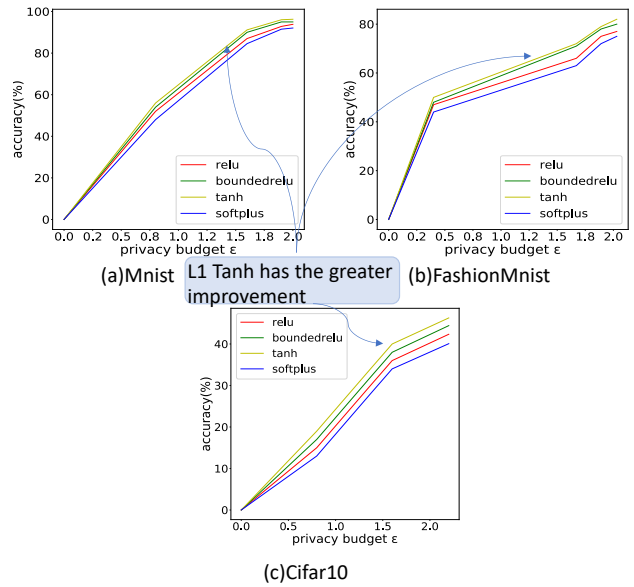


Figure 13: Accuracy vs privacy budget for different level of noise.

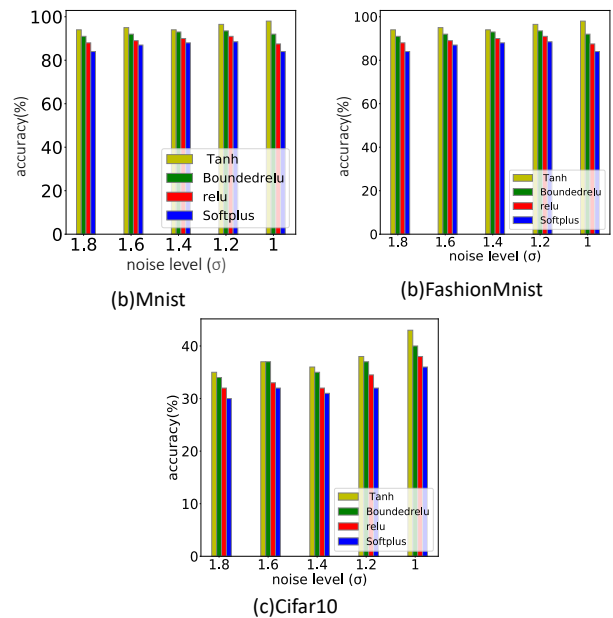


Figure 14: Average accuracy over the ran epochs for different level of noise.



$a = 4$  which is the parameter of BoundedRELU. As Figure 14 and 13 show, bounded activation function outperform unbounded activation functions and this is due to the fact of reducing the impact of clipping gradient and the information discarded at this step which lead to more utility for the same level of privacy. The best performance as the figures show is for Tanh (See L1 in Figure 13) which led to new state-of-the-art results.

## 6. Why do bounded activation functions help privacy?

To respond to this question after the major improvement by bounded activation function seen in section 5.1 more specifically in the case of protected model with DPSGD, we use as an example of bounded activation function BoundedRELU which is a simple variant of RELU that differs from it just by putting a threshold  $a$  so that if the input exceeds this value it is replaced by  $a$ . So its expression is as follows:

$$f(x) = \begin{cases} x & \text{if } a > x > 0 \\ a & \text{if } x > a \\ 0 & \text{otherwise} \end{cases}$$

### 6.1. Impact of activation functions on the gradient magnitude

To highlight the impact of BoundedRELU on the gradient size we perform an analytical comparison between RELU and BoundedRELU.

- 1) gradient size with RELU :

$$\frac{\partial L}{\partial w_i} = y_i \cdot \gamma_{i+1}$$

$$\gamma_i = \begin{cases} (y_i - \text{Label}) \cdot \text{RELU}'(x_i), & \text{if } i \text{ the last} \\ (\sum \gamma_{i+1} \cdot w_i) \cdot \text{RELU}'(x_i), & \text{else} \end{cases}$$

where  $x_i, y_i$  and  $w_i$  are respectively the input, the output and the weights of the  $i$ -th layer.

To highlight the difference between RELU and BoundedRELU in terms of gradient size we consider the case where  $x_i > a$  so  $\text{RELU}'(x_i) = 1$  and consequently due to this unconstrained value of the partial derivative the value of the gradient will explode in this case to  $\frac{\partial L}{\partial w_i} = y_i \cdot \gamma_{i+1}$

- 2) Gradient size with BoundedRELU :

$$\text{BRELU}'(x) = \begin{cases} 1 & \text{if } a > x > 0 \\ 0 & \text{if } x > a \\ 0 & \text{otherwise} \end{cases}$$

$$\frac{\partial L}{\partial w_i} = y_i \cdot \gamma_{i+1}$$

$$\gamma_i = \begin{cases} (y_i - \text{Label}) \cdot \text{BRELU}'(x_i), & \text{if } i \text{ the last} \\ (\sum \gamma_{i+1} \cdot w_i) \cdot \text{BRELU}'(x_i), & \text{else} \end{cases}$$

Taking the case where  $x_i > a$ ,  $\text{BRELU}'(x_i) = 0$  so the value of the gradient will be :

$\frac{\partial L}{\partial w_i} = y_i \cdot \gamma_{i+1} = 0$ . We can notice the importance of BoundedRELU compared to RELU in the case where  $x_i > a$  in constraining the value of the gradient compared to RELU which stays intact since its partial derivative is 1 which will have no impact on constraining the derivative in the chain rule.

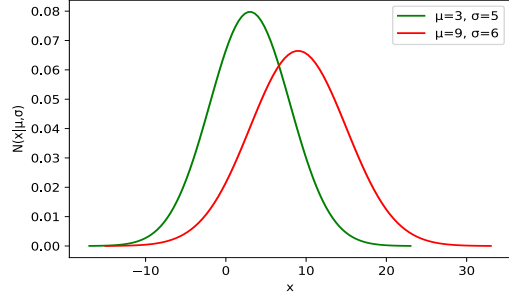


Figure 15: Norm2 of conv output for CIFAR-10.

Figure 15 shows a comparison between the norm2 of the output of the first RELU conv block vs the norm2 of the output of the first BoundedRELU conv block to show the impact of BoundedRELU in our LeNet model for the FashionMnist dataset we can see that the mean value for RELU is higher than the mean value for BoundedRELU which implies the restriction of clipping gradient in DPSGD.

### 6.2. DPSGD with BoundedRELU

Input : Training dataset  $x_1, \dots, x_N$ , Loss function  $L$ , Learning rate  $\mu_t$ , Noise scale  $\sigma$ , Batch size  $b$ , Parameter of RELU  $a$  and clipping threshold  $C$ .

Initialize  $\theta_0$  with uniform initialization;

$a = \text{TuneBound}(L, \theta, a)$ ;

**for**  $t = 0$  **to**  $T - 1$  **do**

**for each**  $i$  **in**  $S_b$  **do**

$g_r(x_i) \leftarrow \nabla_{\theta_t} L(\theta_t, x_i)$

$\bar{g}_t(x_i) \leftarrow g_r(x_i) + G(0, \sigma^2 C^2)$

**end**

$\theta_{t+1} \leftarrow \theta_t - \eta_t \bar{g}_t$

**end**

**Algorithm 1:** DP-SGD with BoundedRELU

To ensure that BoundedRELU works within its optimal performance we need to understand how to improve it, and as mentioned previously its role within DPSGD is to constrain the number of clipping gradients to ensure the minimum information loss. In our case, we clipped the gradient if its value exceed a certain threshold  $C$  given that  $x_i < a$  so we need to make sure that the chosen value of a minimize the number of values that its gradient exceeds a certain threshold  $C$  given that  $x_i < a$ , after the bound is tuned DPSGD then is executed as always without any changes with its steps. By doing so we soften the effect of clipping gradient in DPSGD and therefore we ensure

Model for MNIST	Acc	$\epsilon$
Genetic algorithm with RELU	73.745	1.75
<b>Genetic algorithm with BoundedRELU (ours)</b>	<b>82.99</b>	1.74

TABLE 4: Comparison of genetic algorithm convergence with RELU vs BoudedRELU for MNIST.

maximum gain in accuracy for the same level of privacy compared to unbounded RELU.

After trying to evaluate its impact we aimed to integrate BoundedRELU in state-of-the-art works. Table 4 and Table 5 show the results for the genetic optimization algorithm explained in section 4, to run this optimization algorithm we put population strength=10, mutation rate=0.25, the segments=10 and number-of-generation=10, where the search space is  $(lr, a)$  and  $a$  is the BoundedRELU parameter. The results for RELU are taken from the paper [32] and we repeated the experiments and we got the same results to ensure the fairness of the comparison.

Model for CIFAR-10	Acc	$\epsilon$
Genetic algorithm with RELU	37.999	0.599
<b>Genetic algorithm with BoundedRELU (ours)</b>	37.88	<b>0.514</b>

TABLE 5: Comparison of genetic algorithm convergence with RELU vs BoudedRELU for CIFAR-10.

As Table 6 shows, BoundedRELU outperforms RELU for all three datasets at the same level of privacy, the best improvement is for FashionMnist datasets and by including CIFAR-10 in our experiments we can ensure that our approach works well for more complex datasets. So we can conclude that our new proposed activation function BoundedRELU works well in privacy and is ready to be incorporated in differentially private models rather than unbounded RELU.

Table 6 presents a summary of some results to demonstrate the effectiveness of BoundedRELU in terms of convergence for each model and also shows a comparison with state-of-the-art.

Datasets	Approach	Acc	$\epsilon$
MNIST	DP-SGD with RELU	92.92	1.43
MNIST	<b>DP-SGD with BoundedRELU (ours)</b>	<b>96.02</b>	1.43
CIFAR-10	DP-SGD WITH RELU	42.34	2.20
CIFAR-10	<b>DP-SGD with BoundedRELU (ours)</b>	<b>44.42</b>	2.20
FashinMnist	DP-SGD with RELU	81.9	2.7
FashionMnist	<b>DP-SGD with BoundedRELU (ours)</b>	<b>84.76</b>	2.7

TABLE 6: Results summary.

## 7. Discussion and Concluding Remarks

In this paper, we present a framework analysis for investigating how the hyperparameter space of the model affects the privacy vs. utility trade-off. Our investigation started with an unprotected model, and we used a membership inference attack to assess privacy leakage. We discovered that although some parameters support robustness, others do not. Except in the case of a large model (Alexnet) with a large number of shadows, normalization techniques (Batch

and layer normalization) represent the main improvement. We discovered that dropout helps in all numbers of shadows and lessens the impact of membership inference attacks. Additionally, we discovered that small learning rates enhance model privacy more so than big learning rates. Finally, we address the privacy implications of several activation functions in this case. The impact of hyperparameter space on the privacy vs. utility trade-off in the protected model was then assessed using DPSGD. In this instance, we discovered that bounded activation functions significantly enhanced the model’s output and produced new state-of-the-art results compared to what is commonly used. We highlighted batch normalization’s effects and how they affect the model, as a result, it is clear that this parameter helps while layer normalization does not. We can also see how several parameters might be combined to increase model privacy. Finally, we provide an analysis in the case of a bounded activation function and use the example of BoundedRELU to study this case.

Since the combination of two parameters shows an improvement in our studies, one can propose an optimization method, such as a genetic algorithm, or perform a grid search optimization on the provided set of parameters to build a more efficient model. A different direction is to attempt to incorporate those parameters to produce new, effective privacy-preserving designs. Neural architecture search (NAS) is another topic that requires investigation, one can leverage this work to build more efficient architectures.

## References

- [1] A. Shrestha and A. Mahmood, “Review of deep learning algorithms and architectures,” *IEEE access*, vol. 7, pp. 53 040–53 065, 2019.
- [2] S. Pouyanfar, S. Sadiq, Y. Yan, H. Tian, Y. Tao, M. P. Reyes, M.-L. Shyu, S.-C. Chen, and S. S. Iyengar, “A survey on deep learning: Algorithms, techniques, and applications,” *ACM Computing Surveys (CSUR)*, vol. 51, no. 5, pp. 1–36, 2018.
- [3] F. Wang, L. P. Casalino, and D. Khullar, “Deep learning in medicine—promise, progress, and challenges,” *JAMA internal medicine*, vol. 179, no. 3, pp. 293–294, 2019.
- [4] X. Liu, L. Xie, Y. Wang, J. Zou, J. Xiong, Z. Ying, and A. V. Vasilakos, “Privacy and security issues in deep learning: A survey,” *IEEE Access*, vol. 9, pp. 4566–4593, 2020.
- [5] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in *2017 IEEE symposium on security and privacy (SP)*. IEEE, 2017, pp. 3–18.
- [6] Y. Shen, X. He, Y. Han, and Y. Zhang, “Model stealing attacks against inductive graph neural networks,” in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 1175–1192.
- [7] Z. He, T. Zhang, and R. B. Lee, “Model inversion attacks against collaborative inference,” in *Proceedings of the 35th Annual Computer Security Applications Conference*, 2019, pp. 148–162.
- [8] L. Tran, D. Kong, H. Jin, and J. Liu, “Privacy-cn: A framework to detect photo privacy with convolutional neural network using hierarchical features,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 30, no. 1, 2016.
- [9] J. J. Hathaliya, S. Tanwar, and P. Sharma, “Adversarial learning techniques for security and privacy preservation: A comprehensive review,” *Security and Privacy*, vol. 5, no. 3, p. e209, 2022.

- [10] Z. Xiong, W. Li, Q. Han, and Z. Cai, "Privacy-preserving auto-driving: a gan-based approach to protect vehicular camera data," in *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2019, pp. 668–677.
- [11] C. Dwork, "The differential privacy frontier," in *Theory of Cryptography Conference*. Springer, 2009, pp. 496–502.
- [12] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.
- [13] Z. Bu, H. Wang, and Q. Long, "On the convergence and calibration of deep learning with differential privacy," *arXiv preprint arXiv:2106.07830*, 2021.
- [14] H. Shin, S. Kim, J. Shin, and X. Xiao, "Privacy enhanced matrix factorization for recommendation with local differential privacy," *IEEE Transactions on Knowledge and Data Engineering*, vol. 30, no. 9, pp. 1770–1782, 2018.
- [15] Z. Bu, H. Wang, Q. Long, and W. J. Su, "On the convergence and calibration of deep learning with differential privacy," *arXiv preprint arXiv:2106.07830*, 2021.
- [16] K. Jung, H. Lee, and Y. D. Chung, "Differentially private neural networks with bounded activation function," *IEICE TRANSACTIONS on Information and Systems*, vol. 104, no. 6, pp. 905–908, 2021.
- [17] —, "Differentially private neural networks with bounded activation function," *IEICE TRANSACTIONS on Information and Systems*, vol. 104, no. 6, pp. 905–908, 2021.
- [18] E. Bagdasaryan, O. Poursaeed, and V. Shmatikov, "Differential privacy has disparate impact on model accuracy," *Advances in neural information processing systems*, vol. 32, 2019.
- [19] N. Papernot and T. Steinke, "Hyperparameter tuning with renyi differential privacy," *arXiv preprint arXiv:2110.03620*, 2021.
- [20] X. Xiong, S. Liu, D. Li, Z. Cai, and X. Niu, "A comprehensive survey on local differential privacy," *Security and Communication Networks*, vol. 2020, 2020.
- [21] S. YAGNIK, E. PANCHAL, and P. TRIPATHI, "A novel approach of designing private federated learning model using dpgsd-differentially private stochastic gradient descent optimizer," 2022.
- [22] K. L. van der Veen, R. Seggers, P. Bloem, and G. Patrini, "Three tools for practical differential privacy," *arXiv preprint arXiv:1812.02890*, 2018.
- [23] S. Adhikary, M. Anwar, M. J. M. Chowdhury, I. H. Sarker *et al.*, "Genetic algorithm-based optimal deep neural network for detecting network intrusions," in *Machine Intelligence and Data Science Applications*. Springer, 2022, pp. 145–156.
- [24] Y. LeCun and C. Cortes, "MNIST handwritten digit database," 2010. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [25] A. Krizhevsky, "Learning multiple layers of features from tiny images," Tech. Rep., 2009.
- [26] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *CoRR*, vol. abs/1708.07747, 2017. [Online]. Available: <http://arxiv.org/abs/1708.07747>
- [27] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [28] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: <https://www.tensorflow.org/>
- [29] S. Yeom, I. Giacomelli, M. Fredrikson, and S. Jha, "Privacy risk in machine learning: Analyzing the connection to overfitting," in *2018 IEEE 31st computer security foundations symposium (CSF)*. IEEE, 2018, pp. 268–282.
- [30] C. F. G. D. Santos and J. P. Papa, "Avoiding overfitting: A survey on regularization methods for convolutional neural networks," *ACM Computing Surveys (CSUR)*, vol. 54, no. 10s, pp. 1–25, 2022.
- [31] D. R. Wilson and T. R. Martinez, "The need for small learning rates on large problems," in *IJCNN'01. International Joint Conference on Neural Networks. Proceedings (Cat. No. 01CH37222)*, vol. 1. IEEE, 2001, pp. 115–119.
- [32] A. Priyanshu, R. Naidu, F. Mireshghallah, and M. Malekzadeh, "Efficient hyperparameter optimization for differentially private deep learning," *arXiv preprint arXiv:2108.03888*, 2021.
- [33] N. Papernot, A. Thakurta, S. Song, S. Chien, and Ú. Erlingsson, "Tempered sigmoid activations for deep learning with differential privacy," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 10, 2021, pp. 9312–9321.
- [34] K. Jung, H. Lee, and Y. D. Chung, "Differentially private neural networks with bounded activation function," *IEICE Transactions on Information and Systems*, vol. 104, no. 6, pp. 905–908, 2021.
- [35] A. Priyanshu, R. Naidu, F. Mireshghallah, and M. Malekzadeh, "Efficient hyperparameter optimization for differentially private deep learning," *arXiv preprint arXiv:2108.03888*, 2021.