

All Rivers Run to the Sea: Private Learning with Asymmetric Flows

Yue Niu¹ Ramy E. Ali^{2*} Saurav Prakash³ Salman Avestimehr¹

¹ University of Southern California ² Samsung ³ University of Illinois Urbana-Champaign

yueniu@usc.edu ramy.ali@samsung.com sauravp2@illinois.edu avestime.usc.edu

Abstract

*Data privacy is of great concern in cloud machine-learning service platforms, when sensitive data are exposed to service providers. While private computing environments (e.g., secure enclaves), and cryptographic approaches (e.g., homomorphic encryption) provide strong privacy protection, their computing performance still falls short compared to cloud GPUs. To achieve privacy protection with high computing performance, we propose *Delta*, a new private training and inference framework, with comparable model performance as non-private centralized training. *Delta* features two asymmetric data flows: the main information-sensitive flow and the residual flow. The main part flows into a small model while the residuals are offloaded to a large model. Specifically, *Delta* embeds the information-sensitive representations into a low-dimensional space while pushing the information-insensitive part into high-dimension residuals. To ensure privacy protection, the low-dimensional information-sensitive part is secured and fed to a small model in a private environment. On the other hand, the residual part is sent to fast cloud GPUs, and processed by a large model. To further enhance privacy and reduce the communication cost, *Delta* applies a random binary quantization technique along with a DP-based technique to the residuals before sharing them with the public platform. We theoretically show that *Delta* guarantees differential privacy in the public environment and greatly reduces the complexity in the private environment. We conduct empirical analyses on CIFAR-10, CIFAR-100 and ImageNet datasets and ResNet-18 and ResNet-34, showing that *Delta* achieves strong privacy protection, fast training, and inference without significantly compromising the model utility.*

1. Introduction

In the current machine learning (ML) era, cloud ML services with high-end GPUs have become indispensable. On the other hand, ensuring data privacy is one of the most critical

challenges in the ML platforms. During training, privacy breaches may occur if training data is exposed to ML service providers, increasing vulnerability to potential attacks. Additionally, users' inference queries can also be susceptible to attacks when accessing ML services with sensitive data [36, 44]. In particular, an untrusted ML platform can cache, learn, and leak queries without users' awareness [45].

Related Works Overview. While prior privacy-preserving machine learning (PPML) frameworks [1, 7, 35, 39, 42, 51] mitigate privacy concerns in training and inference, they also come with different tradeoffs. Differential privacy (DP) based methods perturb the data before outsourcing to an untrusted cloud [38, 61] to ensure privacy, but they usually result in degraded model utility even under moderate privacy constraints [8]. On the other hand, crypto-based techniques [39, 51], that provide data protection with encryption schemes, have not yet proved efficient and scalable to large models due to their prohibitive complexities.

PPML with private environments (e.g., trusted execution environments (TEEs), local environments) presents a promising solution by physically isolating the running computing environments. Such private environments are, however, usually resource-constrained compared to public cloud services with high-end GPUs, resulting in lower computing performance [19, 47]. However, the prior works based on leveraging these private environments along with the public GPUs also incur high complexity [20, 40, 41, 55]. One reason for these methods being inefficient is the heavy communication between the private and the public environments.

Proposed Solution Overview. In this paper, we consider a private training and inference setting where users can access both private and public environments, as shown in Figure 1a. We propose a new private training and inference framework, *Delta*, that achieves strong privacy protection with high model utility and low complexity. The core idea of *Delta* originates from an observation that the intermediate representations (IRs) in ML models exhibit an *asymmetric* structure. Specifically, the primary sensitive information is usually encoded in a low-dimensional space, while the high-dimensional residuals contain very little information. Inspired by such an observation, we design a two-way train-

*R.E. Ali was with the University of Southern California.

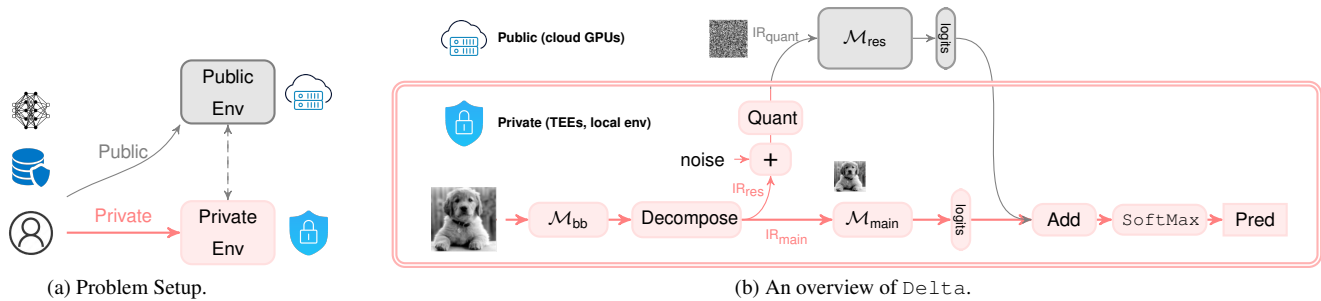


Figure 1. Overview of Δ : the backbone \mathcal{M}_{bb} acts as a feature extractor. The features are decomposed into low-dimensional (information-sensitive) and high-dimensional (residual) parts: IR_{main} and IR_{res} . IR_{main} is fed to a small model \mathcal{M}_{main} , while IR_{res} are outsourced to a large model \mathcal{M}_{res} . \mathcal{M}_{bb} and \mathcal{M}_{main} run in a resource-constrained private environment, whereas \mathcal{M}_{res} is offloaded to a public environment while ensuring privacy through a DP scheme. While only the forward pass is shown, backpropagation is also private (See Sec 4.4).

ing framework that respectively learns low-dimensional but information-sensitive features IR_{main} with a *small* model \mathcal{M}_{main} , and learns the high-dimensional residuals IR_{res} with a *large* model \mathcal{M}_{res} , as illustrated in Figure 1b. Given a model, Δ selects a few front layers as a backbone (\mathcal{M}_{bb}) for extracting features and generating intermediate representations. Δ uses singular value decomposition (SVD) and discrete cosine transformation (DCT) to extract the low-dimensional information-sensitive representation IR_{main} and the residuals IR_{res} . We design a new low-dimensional model (\mathcal{M}_{main}) to learn the information-sensitive IR_{main} . While \mathcal{M}_{bb} and \mathcal{M}_{main} run in a private environment, the rest of the model (\mathcal{M}_{res}) learns residuals in a public environment. Δ further applies DP perturbation and binary quantization on the residuals, leading to further privacy protection and a communication reduction. Owing to the asymmetric structure in IR, Δ guarantees differential privacy on IR_{res} with much smaller noise added compared to the naive scheme that directly adds noise to IR (naive-DP), leading to much-improved model utility. In both training and inference phases, Δ ensures that only the residual part is public, while the information-sensitive part is secured in private environments.

Δ is a generic PPML solution that can be flexibly deployed in several scenarios. For instance, in a cloud ML platform with private TEEs and public GPUs [5, 6], \mathcal{M}_{bb} and \mathcal{M}_{main} can run inside TEEs to preserve privacy, while \mathcal{M}_{res} runs in GPUs to speed up the computations. In general distributed settings, Δ can let clients train \mathcal{M}_{bb} and \mathcal{M}_{main} locally, while a cloud server performs side training with \mathcal{M}_{res} . In summary, our contributions are as follows.

1. We propose a PPML framework for training and inference with a much-improved privacy-utility trade-off compared to the naive-DP methods, a low computing complexity in the private environments and a low communication cost.

2. We design an asymmetric decomposition layer that extracts the low-dimensional information-sensitive representations using SVD and DCT, and design a low-complexity model for resource-constrained computing environments.

3. We provide a formal differential privacy analysis for the proposed framework. In addition, we show empirically that Δ provides strong privacy protection against model inversion and membership inference attacks.

4. We conduct comprehensive evaluations, showing that Δ leads to a better privacy-utility trade-off than the naive DP-based method that directly adds noise to IR. Specifically, under the same privacy budget, Δ improves the accuracy by up to 31%. Moreover, Δ greatly reduces the running time compared to other PPML solutions.

2. System Model

We start by describing the problem setting, the threat model, and our notations.

Problem Setting. We consider a setting where users have private resource-constrained environments (e.g., cloud TEEs, local CPUs/GPUs), but they can also access public, untrusted, and high-end services (e.g., cloud GPUs) to accelerate training and inference. The goal is to protect users’ training and inference data, while maintaining computing performance and high model utility. Note that this setting differs from DP-SGD [1, 30, 60], where the goal is to protect the training data from attacks on gradients or models.

Threat Model. Our threat model is similar to the model considered in prior works leveraging private environments [41, 47, 55]. Specifically, we assume that the private environment is protected against all untrusted, unauthorized access to data and models inside. However, denial-of-service and side-channel attacks are out of our scope. On the other hand, the untrusted, public environment is semi-honest (honest-but-curious). That is, the untrusted public server follows the training and inference protocol faithfully but may attempt to learn as much as possible from what it receives.

Remark 1. While we only consider the semi-honest model in our work, verifiable computing techniques [3, 18, 52] can be incorporated to enhance Δ against malicious parties.

Notations. We denote tensors by capital bold letters as \mathbf{X} , matrices by capital letters as X , and vectors by small

bold letters as \mathbf{x} . X^i denotes a tensor slice $\mathbf{X}(i, :)$. $\|\cdot\|_F$ denotes the Frobenius norm or in general the square root of the sum of squares of elements in a tensor. “ \cdot ” denotes matrix multiplication or in general batch matrix multiplication. “ \otimes ” denotes a convolution operation. X^* indicates a transpose.

3. Asymmetric Structure in IRs

We first observe that intermediate representations (IRs) in neural networks (NNs) exhibit highly asymmetric structures in multiple dimensions. These asymmetric structures are essential for an asymmetric decomposition that embeds the privacy-sensitive information into low-dimensional space.

3.1. Asymmetric Structure in Channel Dimension

In NNs such as convolutional neural networks (CNNs), each layer’s input and output consists of multiple channels, denoted as $\mathbf{X} \in \mathbb{R}^{c \times h \times w}$, where c is the number of channels and h and w denote the height and the width.

We analyze the channel correlation by first flattening \mathbf{X} as $X \in \mathbb{R}^{c \times hw}$ and then performing singular value decomposition (SVD) on the flattened tensor X as

$$X = \sum_{i=1}^c s_i \cdot \mathbf{u}_i \cdot \mathbf{v}_i^*,$$

where s_i is the i -th singular value, $\mathbf{u}_i \in \mathbb{R}^c$ and $\mathbf{v}_i \in \mathbb{R}^{hw}$ are the i -th left and right singular vectors, respectively. We reshape \mathbf{v}_i to the original dimensions as a *principal channel* $V^i \in \mathbb{R}^{h \times w}$, and \mathbf{u}_i as a tensor $U^i \in \mathbb{R}^{c \times 1 \times 1}$.

We then obtain a low-rank representation of \mathbf{X} as follows

$$\mathbf{X}_{\text{lr}} = \sum_{i=1}^r s_i \cdot U^i \cdot V^i,$$

where r denotes the number of principal channels in $\mathbf{X}_{\text{lr}} \in \mathbb{R}^{c \times h \times w}$. Figure 2a shows the normalized error $\frac{\|\mathbf{X} - \mathbf{X}_{\text{lr}}\|}{\|\mathbf{X}\|}$ versus r after the first convolutional layer in ResNet-18 (based on ImageNet). We observe that \mathbf{X}_{lr} with a small r is sufficient to approximate \mathbf{X} . That is, most information in \mathbf{X} can be embedded into \mathbf{X}_{lr} in a low-dimensional space. We notice that 3LegRace [41] also investigated such a property. However, unlike 3LegRace, we leverage the low-rank property to design a new low-complexity model (See Sec. 4.3).

3.2. Asymmetric Structure in Spatial Dimension

The asymmetric structure of the IR also exists over the spatial dimension due to correlations among pixels in each channel $X^i \in \mathbb{R}^{h \times w}$. We use discrete cosine transform (DCT) to analyze the spatial correlation. Specifically, we obtain frequency components using $t \times t$ block-wise DCT [25] as,


$$C^i = \text{DCT}(X^i, T) = \{ C_{k,j}^i = T \cdot X_{k,j}^i \cdot T^* \}_{k,j=1}^{h/t, w/t},$$

where $C_{k,j}^i \in \mathbb{R}^{t \times t}$, $X_{k,j}^i = X^i(kt - t : kt, jt - t : jt)$ and $T \in \mathbb{R}^{t \times t}$ is the DCT transformation matrix. $C^i \in$


$\mathbb{R}^{h \times w}$ is obtained by simply concatenating $C_{k,j}^i$ for $k \in 1, \dots, h/t, j \in 1, \dots, w/t$. Then, we obtain a low-frequency representation, \mathbf{X}_{lf}^i , using inverse $t' \times t'$ block-wise DCT as,

$$\mathbf{X}_{\text{lf}}^i = \text{IDCT}(C_{\text{lf}}^i, T_{\text{lf}}) = \{ X_{\text{lf},k,j}^i = T_{\text{lf}}^* \cdot C_{\text{lf},k,j}^i \cdot T_{\text{lf}} \}_{k,j=1}^{h/t, w/t},$$

where $t' < t$, $X_{\text{lf}}^i \in \mathbb{R}^{\frac{h}{t'} t' \times \frac{w}{t'} t'}$, $C_{\text{lf},k,j}^i = C_{k,j}^i(0 : t', 0 : t')$, and $T_{\text{lf}} \in \mathbb{R}^{t' \times t'}$ is the DCT matrix. Note that t'^2 represents the number of low-frequency components in $X_{\text{lf},k,j}^i$.



(a) Error of the low-rank approximation vs. $r/c \times 100\%$.



(b) Error of the low-frequency approximation vs. $(t'/t)^2 \times 100\%$.

Figure 2. Asymmetric structures along channel and spatial dimension (based on ResNet-18 on ImageNet). Most information in \mathbf{X} can be embedded into low-rank and low-frequency representations.

Owing to the spatial correlation, \mathbf{X}_{lf} can sufficiently approximate \mathbf{X} using a few low-frequency components. Figure 2b shows the error $\frac{\|\mathbf{X} - \mathbf{X}_{\text{lf}}\|}{\|\mathbf{X}\|}$ versus the number of low-frequency components in \mathbf{X}_{lf} after the first convolutional layer in ResNet-18. We can also observe that most information in \mathbf{X} can be embedded into \mathbf{X}_{lf} with low-frequency components.

The asymmetric structure also exists in other tasks, such as language models (See Appendix 10.1). Observing the asymmetric structures in different dimensions, we aim to design an *asymmetric learning* framework that learns privacy-sensitive low-dimensional IR with a low-complexity model in a private environment, while sending residuals to a larger model trained on an untrusted platform.

4. Delta: Private Asymmetric Learning

Overview. At a high level, as shown in Fig. 1, given a model, Delta keeps the first few layers in a private environment as a backbone model \mathcal{M}_{bb} and a main model $\mathcal{M}_{\text{main}}$. Delta then offloads all remaining layers to a public environment as residual model \mathcal{M}_{res} . Specifically, Delta uses the backbone model \mathcal{M}_{bb} as a feature extractor to compute intermediate representations (IRs). Delta then *asymmetrically* decomposes the IRs, obtaining a low-dimensional information-sensitive part (IR_{main}) and residuals IR_{res} . Delta designs a new low-dimensional model $\mathcal{M}_{\text{main}}$ for IR_{main} to reduce computation complexity in the private environment. On the other hand, the residuals IR_{res} are

perturbed and quantized before being outsourced to the untrusted public environment that hosts \mathcal{M}_{res} . At last, output logits from $\mathcal{M}_{\text{main}}$ and \mathcal{M}_{res} are added in the private environment, leading to the final predictions. These final predictions are not disclosed to the public environment. Hence, during training and inference, Δ allows only minimal residual information to be leaked to the public environment.

4.1. Asymmetric IR Decomposition

As observed in Section 3, IRs after \mathcal{M}_{bb} exhibit asymmetric structures in multiple dimensions. Hence, we can decompose IRs such that most information is encoded in the low-dimensional IR denoted as IR_{main} .

Singular Value Decomposition (SVD). Given an IR $\mathbf{X} \in \mathbb{R}^{c \times h \times w}$ obtained from the backbone model \mathcal{M}_{bb} , we first apply SVD as explained in Sec. 3.1 to obtain the principal channels $\{V^i \in \mathbb{R}^{h \times w}\}_{i=1}^c$ and the corresponding coefficients $\{U^i \in \mathbb{R}^{c \times 1 \times 1}\}_{i=1}^c$. We then select the r most principal channels as a low-rank representation of \mathbf{X} as in Sec. 3.1. The rest of the channels are saved as SVD residuals as

$$\mathbf{X}_{\text{SVD res}} = \mathbf{X} - \mathbf{X}_{\text{lr}} = \sum_{i=r+1}^c s_i \cdot U^i \cdot V^i. \quad (1)$$

Discrete Cosine Transform (DCT). After the decomposition along channels, we further decompose V^i over the spatial dimension using DCT.

Specifically, for each principal channel V^i , we first obtain the frequency-domain coefficients as $\mathbf{C}^i = \text{DCT}(V^i, T)$. Then, we only use the low-frequency component to reconstruct a representation as $V_{\text{lf}}^i = \text{IDCT}(\mathbf{C}_{\text{lf}}^i, T_{\text{lf}}) \in \mathbb{R}^{\frac{h}{2} \times \frac{w}{2}}$ as in Sec 3.2. \mathbf{C}_{lf}^i has a reduced dimension and keeps only top-left low-frequency coefficients in \mathbf{C}^i (as shown in Figure 3). $T_{\text{lf}} \in \mathbb{R}^{t' \times t'}$ corresponds to DCT transformation matrix with reduced spatial dimension. The rest of the high-frequency components are saved as DCT residuals as

$$V_{\text{DCT res}}^i = \text{IDCT}(\mathbf{C}_{\text{res}}^i, T), \quad (2)$$

where $\mathbf{C}_{\text{res}}^i$ denotes \mathbf{C}^i with zeros on the top-left corner.

After SVD and DCT, we obtain the privacy-sensitive low-dimensional features as

$$\text{IR}_{\text{main}} = \sum_{i=1}^r s_i \cdot U^i \cdot V_{\text{lf}}^i. \quad (3)$$

On the other hand, the residuals to be offloaded to the untrusted public environment are given as

$$\begin{aligned} \text{IR}_{\text{res}} &= \mathbf{X} - \text{IR}_{\text{main}} = \mathbf{X}_{\text{SVD res}} + \mathbf{X}_{\text{DCT res}} \\ &= \sum_{i=r+1}^c s_i \cdot U^i \cdot V^i + \sum_{i=1}^r s_i \cdot U^i \cdot V_{\text{DCT res}}^i. \end{aligned} \quad (4)$$

IR_{res} is further normalized as

$$\text{IR}_{\text{res, norm}} = \text{IR}_{\text{res}} / \max(1, \|\text{IR}_{\text{res}}\|_2 / C), \quad (5)$$

where C is a scaling parameter for ℓ_2 normalization. The normalization is necessary to bound *sensitivity* for DP. For simplification, we denote the normalized residuals as IR_{res} .

Hence, IR_{main} has fewer principal channels and smaller spatial dimensions but contains most information in IR. The non-principal channels and high-frequency components, on the other hand, are saved in IR_{res} . IR_{main} and IR_{res} are then respectively fed into a small model $\mathcal{M}_{\text{main}}$ in a private environment and a large model \mathcal{M}_{res} in a public environment.

4.2. Residuals Perturbation and Quantization

While the low-dimensional representation IR_{main} has the most important information in the IR, IR_{res} might still contain some information such as a few high-frequency components. Furthermore, as the communication from a private environment (e.g., TEEs) to a public environment (e.g., GPUs) is usually slow, sending floating-point high-dimensional residuals can significantly increase the communication overhead and prolong the total running time.

In this section, we perturb IR_{res} with a Gaussian mechanism and then apply binary quantization on the perturbed IR_{res} to reduce the inter-environment communication cost. Given a DP budget ϵ , we consider the Gaussian mechanism, and compute the corresponding noise parameter σ (See Section 5). For each tensor IR_{res} , we generate a noise tensor $\mathbf{N} \in \mathbb{R}^{c \times h \times w} \sim \mathcal{N}(0, \sigma^2 \cdot \mathbf{I})$, and add it to IR_{res} in the private environment. With noisy residuals, $\text{IR}_{\text{noisy}} = \text{IR}_{\text{res}} + \mathbf{N}$, we apply a binary quantizer as follows

$$\text{IR}_{\text{quant}} = \text{BinQuant}(\text{IR}_{\text{noisy}}) = \begin{cases} 0 & \text{IR}_{\text{noisy}} < 0, \\ 1 & \text{IR}_{\text{noisy}} \geq 0. \end{cases} \quad (6)$$

As a result, the tensor to be offloaded to the public environment is a binary representation of the residuals. Compared to floating-point values, such a binary representation reduces communication by $32\times$. Owing to the asymmetric decomposition, the values in IR_{res} are usually close to zero. Hence, a small noise is sufficient to achieve strong privacy protection (See formal analysis in Section 5). Further ablation studies can be found in Section 9.3.

4.3. Model Design for Low-Dimensional IR_{main}

Knowing that IR_{main} has a low-rank as given in Equation (3), in this section, we show that developing an efficient $\mathcal{M}_{\text{main}}$ with low computation complexity is attainable.

In linear layers such as convolutional layer, low-rank inputs lead to low-rank outputs [41]. Assuming inputs and outputs have rank r and q , we split the convolution layer into two sub-layers as shown in Figure 4 (right). The first layer has q ($k \times k$) kernels to learn the principal features,




Figure 3. The asymmetric IR decomposition is shown (See Figure 1 for the whole pipeline). We use SVD and DCT to encode channel and spatial information into a low-dimensional representation, and offload the residuals to public environments. The low-dimensional IR_{main} has fewer channels and smaller sizes but still encode most sensitive information. The residuals IR_{res} have the same dimension as the original IR.




Figure 4. Model design for the low-dimensional IR_{main} . Knowing the rank in data, the number of channels in convolution layers can be reduced, leading to a reduction in computation complexity.

whereas the second layer has c (1×1) kernels to combine the principal channels. We further add a kernel orthogonality regularization [59] to the first layer to enhance the orthogonality of the output channels.

Such a design comes with fundamental reasoning. Theorem 1 shows that, with knowledge of the ranks of inputs and outputs in a convolutional layer, it is possible to optimize a low-dimensional layer as in Figure 4 (right) such that it results in the same output as the original layer.

Theorem 1. For a convolution layer with weight $\mathbf{W} \in \mathbb{R}^{n \times c \times k \times k}$ with an input \mathbf{X} with rank r and output \mathbf{Y} with rank q , there exists an optimal $\mathbf{W}^{(1)} \in \mathbb{R}^{q \times c \times k \times k}$, $\mathbf{W}^{(2)} \in \mathbb{R}^{n \times q \times 1 \times 1}$ in the low-dimensional layer such that the output of this layer denoted as \mathbf{Y}' satisfies

$$\min_{\mathbf{W}^{(1)}, \mathbf{W}^{(2)}} \left\| \mathbf{Y} - \mathbf{Y}' \right\| = 0. \quad (7)$$

The proof is deferred to Appendix 8.1.

Remark 2. While the low-dimensional layer in Figure 4 (right) shares similar architecture as low-rank compression methods [23, 24, 63], low-rank compression inevitably incurs information loss in outputs. However, the layer in Figure 4 (right) can theoretically preserve all information given low-rank inputs and outputs.

4.4. Private Backpropagation

While the asymmetric IR decomposition together with the randomized quantization mechanism ensures privacy in forward passes, information of IR_{main} can still be leaked through backpropagation on logits. In this section, we further pro-

pose a private backpropagation that removes the logits of $\mathcal{M}_{\text{main}}$ from the gradients to \mathcal{M}_{res} .

In detail, the gradients to $\mathcal{M}_{\text{main}}$ are calculated through a standard backpropagation algorithm that uses logits from both $\mathcal{M}_{\text{main}}$ and \mathcal{M}_{res} . While gradients to \mathcal{M}_{res} are calculated solely using \mathcal{M}_{res} logits. Hence, the logits from $\mathcal{M}_{\text{main}}$ will not be revealed to the outside. Specifically, given the output logits from $\mathcal{M}_{\text{main}}$ and \mathcal{M}_{res} : \mathbf{z}_{main} and \mathbf{z}_{res} , we compute Softmax for i -th label $\mathcal{M}_{\text{main}}$ and \mathcal{M}_{res} as

$$\begin{aligned} \mathcal{M}_{\text{main}} : \mathbf{o}_{\text{tot}}(i) &= e^{\mathbf{z}_{\text{main}}(i) + \mathbf{z}_{\text{res}}(i)} / \sum_{j=1}^L e^{\mathbf{z}_{\text{main}}(j) + \mathbf{z}_{\text{res}}(j)}, \\ \mathcal{M}_{\text{res}} : \mathbf{o}_{\text{res}}(i) &= e^{\mathbf{z}_{\text{res}}(i)} / \sum_{j=1}^L e^{\mathbf{z}_{\text{res}}(j)}, \end{aligned}$$

where L denotes the number of labels in the current task. Following the backpropagation in Softmax , we compute gradients to $\mathcal{M}_{\text{main}}$ and \mathcal{M}_{res} : \mathbf{g}_{main} , \mathbf{g}_{res} , as

$$\mathbf{g}_{\text{main}} = \mathbf{o}_{\text{tot}} - \mathbf{y}, \quad \mathbf{g}_{\text{res}} = \mathbf{o}_{\text{res}} - \mathbf{y},$$

where \mathbf{y} denotes one-hot encoding for labels.

With the separate backpropagation on gradients, $\mathcal{M}_{\text{main}}$ avoids revealing its logit to \mathcal{M}_{res} , while still using \mathcal{M}_{res} 's logits for its own backpropagation.

4.5. Training Procedure

The model training using Δ consists of two stages.

Stage 1. We train the backbone, \mathcal{M}_{bb} , and the main model, $\mathcal{M}_{\text{main}}$, with IR_{main} only. IR_{res} is ignored and not shared with the public. Therefore, there is no privacy leakage at this stage, as all data and model parameters are kept in the private environment. After stage 1, we cache all residual data IR_{res} from SVD and DCT decomposition, and apply the randomized quantization mechanism once on IR_{res} .

Stage 2. We freeze the backbone model, \mathcal{M}_{bb} , and continue to train the main model, $\mathcal{M}_{\text{main}}$, and the residual model, \mathcal{M}_{res} . As \mathcal{M}_{bb} is frozen, we directly sample residual inputs for \mathcal{M}_{res} from the cached residual data. While for \mathcal{M}_{bb} and $\mathcal{M}_{\text{main}}$, we fetch data from the raw datasets, apply SVD and DCT decomposition, and send IR_{main} to $\mathcal{M}_{\text{main}}$. The training procedure is provided in algorithm 1.

5. Privacy Analysis

This section provides the differential privacy analysis for Δ . As the backbone model \mathcal{M}_{bb} , the low-dimensional

Algorithm 1 Delta training procedure

Require: ep_1 : #epochs at stage 1, ep_2 : #epochs at stage 2

Require: ϵ : privacy constraint

- 1: Initialize \mathcal{M}_{bb} , $\mathcal{M}_{\text{main}}$, \mathcal{M}_{res}
 - 2: **for** $t = 1, \dots, \text{ep}_1$ **do**
 - 3: **for** a batch from the dataset **do**
 - 4: Train \mathcal{M}_{bb} and $\mathcal{M}_{\text{main}}$.
 - 5: Freeze \mathcal{M}_{bb} and cache all residual data IR_{res} .
 - 6: Apply the randomized quantization based on Eq (6).
 - 7: **for** $t = 1, \dots, \text{ep}_2$ **do**
 - 8: **for** a batch from the dataset, cached residuals **do**
 - 9: Train $\mathcal{M}_{\text{main}}$ and \mathcal{M}_{res} .
-

model $\mathcal{M}_{\text{main}}$ and final predictions remain private, the only public information are the residuals and the residual model \mathcal{M}_{res} at stage 2 during training. Therefore, we analyze the privacy leakage of the perturbed residuals at stage 2.

Given neighboring datasets $\mathcal{D} = \{X^1, \dots, X^i, \dots, X^N\}$ and $\mathcal{D}' = \{X^1, \dots, 0, \dots, X^N\}$ with i -th record removed, the global ℓ_2 -sensitivity is defined as $\Delta_2 = \sup_{\mathbf{x}_{\text{res}}, \mathbf{x}'_{\text{res}}} \|\mathbf{x}_{\text{res}} - \mathbf{x}'_{\text{res}}\|_F \leq C$, where $\mathbf{x}_{\text{res}}, \mathbf{x}'_{\text{res}}$ denote residuals obtained from \mathcal{D} and \mathcal{D}' .

We provide our DP guarantee in Theorem 2.

Theorem 2. *Delta ensures that the perturbed residuals and operations in the public environment satisfy (ϵ, δ) -DP given noise $\mathbf{N} \sim \mathcal{N}(0, 2C^2 \cdot \log(2/\delta')/\epsilon' \cdot \mathbf{I})$ given sampling probability p , and $\epsilon = \log(1 + p(e^{\epsilon'} - 1))$, $\delta = p\delta'$.*

The proof of Theorem 2 relies on the analysis of Gaussian mechanism and the post-processing rule of DP [17, 28], and is provided in Appendix 8.2 for completeness.

Remark 3. (Training with Cached Residual Improves DP). Sampling an image multiple times does not incur additional privacy leakage. As described in Sec. 4.5, the perturbation is only performed once before training the residual model \mathcal{M}_{res} . When training the residual model \mathcal{M}_{res} , Delta directly samples input for \mathcal{M}_{res} from the perturbed cached residuals, with no need to perform perturbation on the fly [48].

6. Empirical Evaluations

In this section, we evaluate Delta in terms of model accuracy, running time, and resilience against attacks. We conduct further ablation studies on different ways of merging logits and effects of perturbation in Appendix 9.

Datasets and Models. We use CIFAR10/100 [29] and ImageNet [15]. For models, we choose ResNet-18 and ResNet-34 [21]. Hyperparameters are provided in Appendix 11.2

Model Configuration. For ResNet models, \mathcal{M}_{bb} consists of the first convolution layer, while all *ResBlocks* [21] and the fully-connected layer are offloaded in \mathcal{M}_{res} . On the other

hand, $\mathcal{M}_{\text{main}}$'s details are deferred to Appendix 11.1.

Asymmetric Decomposition. For the SVD, based on [41], we keep $r = 8$ principal channels in IR_{main} for CIFAR-10/100 and 12 for ImageNet to keep $> 95\%$ information in the private environment. For the DCT, we set $t, t' = 16, 8$ on CIFAR, and 14, 7 on ImageNet to avoid noticeable information leakage in the residuals. Note that while decomposition with larger r and t' leads to less information leakage, it also incurs more computations in the private environment. Therefore, we choose the parameters to maintain a reasonable trade-off between privacy leakage and complexity.

6.1. Model Accuracy

In this section, we evaluate model accuracy. We train ResNet-18 on CIFAR-10/100, and ResNet-18/34 on ImageNet. We first compare the accuracy of the following schemes.

- $\mathcal{M}_{\text{main}}$: Train $\mathcal{M}_{\text{bb}} + \mathcal{M}_{\text{main}}$ with no residuals.
- $\mathcal{M}_{\text{main}} + \mathcal{M}_{\text{res}}$: Train $\mathcal{M}_{\text{bb}} + (\mathcal{M}_{\text{main}}, \mathcal{M}_{\text{res}})$ with DP.
- Orig: Train the original model without Delta.

Figure 5 and 6 show the final accuracy on CIFAR-10/100 and ImageNet. First, owing to the effective asymmetric IR decomposition and the low-dimensional model, $\mathcal{M}_{\text{main}}$ already gives an accuracy that is close to the original model on both CIFAR-10/100 and ImageNet datasets. With adding \mathcal{M}_{res} , Delta achieves a comparable accuracy as the original model. With the residual information further protected by Gaussian noise, we observe that Delta strikes a much improved privacy-utility trade-off, with slight performance degradation under a low privacy budget (small ϵ).




Figure 5. Val acc of ResNet-18 on CIFAR-10, CIFAR-100. $\mathcal{M}_{\text{main}}$ gives accuracy close to the original model. With adding \mathcal{M}_{res} , Delta achieves comparable accuracy as the original model. By adding noise to IR_{res} , Delta achieves strong DP while still preserving the model performance.

To further study the privacy-utility trade-off, we compare Delta with a scheme that directly adds noise to IR of the original model (naive-DP) under the same privacy guarantee. As in Equation (5), we perform normalization on IR before perturbation. As shown in Table 1, training with noise added to IR incurs significant performance degradation given the same ϵ . With the asymmetric IR decomposition, Delta offers a much better privacy-utility trade-off.




Figure 6. Val acc of ResNet-18 and ResNet-34 on ImageNet.

Table 1. Model accuracy of ResNet-18 by perturbing IR and IR_{res} under the same privacy budget ($\epsilon = 1.4$).

| Dataset | Delta: perturb IR _{res} | naive-DP: perturb IR |
|-----------|----------------------------------|------------------------------|
| CIFAR-10 | 92.4% | 69.6% ($\downarrow -22.8$) |
| CIFAR-100 | 71.4% | 48.3% ($\downarrow -23.1$) |
| ImageNet | 65.9% | 34.4% ($\downarrow -31.5$) |

6.2. Running Time Analysis

In this section, we compare running time using Delta with a private training framework 3LegRace [41] and private inference framework Slalom [55]. As a reference, we also include the time of running the original model solely in a private environment (Priv-only).

Theoretical Complexity. Table 2 lists the theoretical computational complexities by MACs (i.e., multiply-accumulate) of ResNet-18 on CIFAR-10/100, and ResNet-34 on ImageNet during the inference phase with a batch of size 1. For SVD, we use an approximation algorithm [41] that only computes the first r principal channels. First, we can see that SVD and DCT only account for a very small fraction of the total computations, which aligns with the real running time in Table 4. On the other hand, compared to MACs in \mathcal{M}_{res} , the computation complexity of $\mathcal{M}_{bb} + \mathcal{M}_{main}$ is much smaller, only accounting for $\sim 10\%$ of MACs in \mathcal{M}_{res} . This shows that, with the asymmetric decomposition that embeds most information into a low-dimensional representation, the low-dimensional model effectively reduces the computation cost of the resource-constrained private environments.

Table 2. MACs of ResNet-18 and ResNet-34 during forward passes with batch size 1. Compared to the residual model, the complexity of the backbone and the main model is much smaller.

| Model | $\mathcal{M}_{bb} + \mathcal{M}_{main}$ | SVD | DCT | \mathcal{M}_{res} |
|-----------|---|--------|--------|---------------------|
| ResNet-18 | 48.3 M | 0.52 M | 0.26 M | 547M |
| ResNet-34 | 437 M | 1.6 M | 0.7 M | 3.5G |

Real Running Time. We test ResNet-18 on CIFAR-100 with batch size 32 and average per-iteration time across one epoch. We use Intel SGX [14] as a private environment, and Nvidia RTX 5000 as an untrusted public environment.

Table 3 shows the per-iteration training and inference time

Table 3. Running time of Delta and other baselines (ResNet-18/CIFAR-100, $b = 32$). Delta achieves significant speedup compared to 3LegRace, and Slalom. Each cell denotes (time (ms), speedup)

| Task | Priv-only | 3LegRace | Slalom | Delta |
|--------|-----------|-------------------|------------------|-------------------|
| Train | 1372 | 237 (6 \times) | - | 62 (22 \times) |
| Infer. | 510 | 95 (5 \times) | 84 (6 \times) | 20 (25 \times) |

Table 4. Time breakdown on ResNet-18/CIFAR-100. Delta significantly shrinks the time gap between the private and public environments.

| | \mathcal{M}_{bb} | Decomp. | Parallel $\mathcal{M}_{main}/\mathcal{M}_{res}$ |
|----------|--------------------|---------|---|
| Forward | 1 ms | 3 ms | 5/16 ms |
| Backward | 2 ms | 1 ms | 11/39 ms |

for Delta and the baselines. In both phases, owing to the effective asymmetric decomposition, the required computation in the private environment is greatly reduced. Hence, Delta achieves significant speedups compared to running a model solely in the private environment (Priv-only). Moreover, unlike the existing private inference method Slalom and training method 3LegRace, Delta obviates the need for frequent and costly inter-environment communication. Hence, Delta offers much faster training and inference. Table 4 further lists the running time breakdown for Delta. We observe that the time on \mathcal{M}_{bb} and IR decomposition is marginal compared to \mathcal{M}_{main} and \mathcal{M}_{res} . While the theoretical computation complexity in \mathcal{M}_{main} is much lower than \mathcal{M}_{res} , the real running time of \mathcal{M}_{main} still dominates the forward and backward passes. Nevertheless, compared to Slalom and 3LegRace, the time gap between the private and public environments is significantly shrunk.

6.3. Protection against Attacks

This section evaluates Delta against two privacy attacks: a model inversion attack called SecretRevealer [62], and a membership inference attack called ML-Leaks [49].

Model Inversion Attacks. SecretRevealer can leverage prior knowledge, such as blurred images, when reconstructing training samples. In our case, we allow the attack to use the quantized residuals as the prior knowledge when reconstructing images. Other model inversion attacks [4, 12, 27, 50] need confidence scores/predicted labels from the whole model, hence they do not apply to our setting.

Following the attack protocol, we first take quantized residuals as a prior knowledge and train a generative model G . Then, we optimize the latent inputs z to minimize the loss on the residual model and the discriminator used to penalize unrealistic images (L_D): $z^* = \arg \min_z L_D(G(z)) + \lambda L_{\mathcal{M}_{res}}(G(z))$, where λ controls the weights of $L_{\mathcal{M}_{res}}$.

Table 5 shows Delta’s performance against the model inversion attack on ResNet-18/CIFAR-100. We use the structural similarity index (SSIM) to measure the similarity be-

Table 5. Delta’s performance against SecretRevealer on ResNet-18/CIFAR-100. Without DP, the attacker gains some prior information. With noise added, the quality of reconstruction degrades significantly.

| No DP | | DP with $\epsilon = 1.4$ | | No residuals | |
|-------|---------------------|--------------------------|---------------------|--------------|---------------------|
| SSIM | $Acc_{\mathcal{M}}$ | SSIM | $Acc_{\mathcal{M}}$ | SSIM | $Acc_{\mathcal{M}}$ |
| 0.18 | 6.75% | 0.09 | 2.13% | 0.08 | 1% |

tween the reconstructed and the original images [58]. We also measure the target model’s accuracy $Acc_{\mathcal{M}}$ given the reconstructed images as inputs. High $Acc_{\mathcal{M}}$ indicates the target model regarding the reconstructed images by the attacker close to the original training samples. First, we observe that given quantized residuals without noise added, the attack can generate images that have slightly high accuracy on the target model. This indicates that the attacker can leverage the residual information during attacks. The observation is also visually reflected in reconstructed samples in Figure 7. Compared to the original training samples, some reconstructed images reveal the outline of objects (e.g., row 1, col 3). However, with DP, the attacker fails to use the residual information to generate images similar to the training samples. In particular, the accuracy of generated images on the target model is significantly reduced and also SSIM between the reconstructed and the original samples. As a result, the attacker behaves like the case that generates images without prior knowledge of the residuals (col 3 in Table 5).



Figure 7. Visualization of reconstructed images compared to the original training samples. With no DP, the model inversion attack recovers certain features ((b), row 1, col 3). However, when adding noise to the residual, the generated images are drastically affected.

Membership Inference Attacks. We choose ML-Leaks, a strong membership inference attack that infers membership based on confidence scores [49]. In our case, we allow the attack to use confidence scores from \mathcal{M}_{res} . Other membership inference attacks either required labels [13, 33, 34] or susceptible to noise [10]. As the predicted labels are secure with Delta and residuals are perturbed, these methods do not apply to our setting. Following the procedure in [49], we feed private and public samples to the shadow model, and obtain confidence score vectors from the residual model \mathcal{M}_{res} . We then use the vectors from public and private datasets to train an attacker model, which learns to classify whether the confidence score vector comes from public or private datasets.

Table 6 lists the attack performance on CIFAR-100 with ResNet-18. To align with ML-Leaks’s protocol, we use $5k$ and $10k$ samples from the training dataset as a public dataset. The rest of the samples are used as a private dataset to train

Table 6. Membership inference attack on Delta (ResNet-18, CIFAR-100). The DP mechanism is essential to provide further protection for the residual.

| | Attack w. $5k$ samples | | Attack w. $10k$ samples | |
|-----|------------------------|------------------|-------------------------|------------------|
| | No DP | $\epsilon = 1.4$ | $\epsilon = \infty$ | $\epsilon = 1.4$ |
| Acc | 0.56 | 0.52 | 0.60 | 0.55 |
| F1 | 0.68 | 0.57 | 0.73 | 0.68 |

the target model (using Delta). We train the attack model for 50 epochs, with an initial learning rate of 0.1.

We observe that with perturbed residuals, attacks through \mathcal{M}_{res} ’s outputs result in a poor performance compared to training without noise added. It indicates the DP mechanism provides further protection for the residuals, and prevents attackers from inferring membership. Furthermore, as the number of public samples reduces, the attack performance degrades further, implying the attack also heavily depends on prior knowledge via accessing a subset of the target dataset. Such an observation reveals one critical limitation of the attacks. That is, they need to get access to a subset of the target data to obtain a good estimate of the target data distributions. Otherwise, the attack performance collapses. However, in real scenarios, private data can be completely out of attackers’ reach, rendering the target distribution’s estimation impossible. As a result, the attacks can easily fail.

7. Conclusion

We proposed a generic private training and inference framework, Delta, with strong privacy protection, high model accuracy, and low complexity. Delta decomposes the intermediate representations into asymmetric flows: information-sensitive and residual flows. We design a new low-dimensional model to learn the information-sensitive part in a private environment, while outsourcing the residual part to a large model in a public environment. A DP mechanism and binary quantization scheme further protect residuals and improve inter-environment communication efficiency. Our evaluations show that Delta achieves strong privacy protection while maintaining model accuracy and computing performance. While we evaluate Delta in a TEE-GPU environment, Delta can be generalized to other setups such as federated settings, with resource-constrained client side as a private environment, and the server as a public environment.

Acknowledgments

This material is based upon work supported by ONR grant N00014-23-1-2191, ARO grant W911NF-22-1-0165, Defense Advanced Research Projects Agency (DARPA) under Contract No. FASTNICS HR001120C0088 and HR001120C0160. The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

- [1] Martin Abadi, Andy Chu, and et al. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, Vienna Austria, 2016. ACM. 1, 2
- [2] Ahmad Al Badawi and et al. Demystifying bootstrapping in fully homomorphic encryption. *Cryptology ePrint Archive*, 2023, 2023. 4
- [3] Ramy E Ali, Jinhyun So, and A Salman Avestimehr. On polynomial approximations for privacy-preserving and verifiable relu networks. *arXiv preprint arXiv:2011.05530*, 2020. 2
- [4] Shengwei An, Guan hong Tao, and et al. Mirror: Model inversion for deep learning network with high fidelity. In *Proceedings of the 29th Network and Distributed System Security Symposium*, 2022. 7
- [5] Amazon AWS. Aws aws nitro enclaves. <https://aws.amazon.com/ec2/nitro/nitro-enclaves/>, Accessed: 2023-11-09. 2
- [6] Microsoft Azure. Azure confidential computing enclaves. <https://learn.microsoft.com/en-us/azure/confidential-computing/confidential-computing-enclaves>, Accessed: 2023-11-09. 2
- [7] Sara Babakniya, Souvik Kundu, Saurav Prakash, Yue Niu, and Salman Avestimehr. Revisiting sparsity hunting in federated learning: Why does sparsity consensus matter? *Transactions on Machine Learning Research*, 2023. 1
- [8] Eugene Bagdasaryan, Omid Poursaeed, and Vitaly Shmatikov. Differential privacy has disparate impact on model accuracy. In *Advances in Neural Information Processing Systems*, Vancouver, Canada, 2019. Curran Associates, Inc. 1
- [9] Borja Balle and et al. Privacy amplification by subsampling: Tight analyses via couplings and divergences. *Advances in Neural Information Processing Systems*, 31, 2018. 1
- [10] Nicholas Carlini and et al. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1897–1914. IEEE, 2022. 8
- [11] Kumar Chellapilla, Sidd Puri, and et al. High performance convolutional neural networks for document processing. In *Tenth international workshop on frontiers in handwriting recognition*. Suvisoft, 2006. 1
- [12] Si Chen, Mostafa Kahla, and et al. Knowledge-enriched distributional model inversion attacks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16178–16187, Virtual, 2021. IEEE. 7
- [13] Christopher A Choquette-Choo and et al. Label-only membership inference attacks. In *International conference on machine learning*, pages 1964–1974, Virtual, 2021. PMLR. 8
- [14] Victor Costan and Srinivas Devadas. Intel sgx explained. *Cryptology ePrint Archive*, 2016. 7, 4
- [15] Jia Deng, Wei Dong, and et al. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255, Florida, US, 2009. IEEE. 6
- [16] Jacob Devlin, Ming-Wei Chang, and et al. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 3
- [17] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014. 6, 1
- [18] Zahra Ghodsi, Tianyu Gu, and Siddharth Garg. Safetynets: Verifiable execution of deep neural networks on an untrusted cloud. *Advances in Neural Information Processing Systems*, 30, 2017. 2
- [19] Lucjan Hanzlik, Yang Zhang, and et al. Mlcapsule: Guarded offline deployment of machine learning as a service. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3300–3309, 2021. 1, 4
- [20] Hanieh Hashemi, Yongqin Wang, and Murali Annavaram. Darknight: An accelerated framework for privacy and integrity preserving deep learning using trusted hardware. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 212–224, 2021. 1
- [21] Kaiming He, Xiangyu Zhang, and et al. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778. IEEE, 2016. 6
- [22] Ehsan Hesamifard, Hassan Takabi, and et al. Privacy-preserving machine learning as a service. *Proc. Priv. Enhancing Technol.*, 2018(3):123–142, 2018. 4
- [23] Yani Ioannou and et al. Training CNNs with low-rank filters for efficient image classification. *International Conference on Learning Representations (ICLR)*, 2016. 5
- [24] Max Jaderberg, Andrea Vedaldi, and et al. Speeding up convolutional neural networks with low rank expansions. *British Machine Vision Conference (BMVC)*, 2014. 5
- [25] Anil K Jain. *Fundamentals of digital image processing*. Prentice-Hall, Inc., 1989. 3
- [26] Patrick Jauernig, Ahmad-Reza Sadeghi, and Emmanuel Stempf. Trusted execution environments: properties, applications, and challenges. *IEEE Security & Privacy*, 18(2):56–60, 2020. 4
- [27] Mostafa Kahla, Si Chen, Hoang Anh Just, and Ruoxi Jia. Label-only model inversion attacks via boundary repulsion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15045–15053, 2022. 7
- [28] Gautam Kamath. Approximate differential privacy. <http://www.gautamkamath.com/CS860notes/lec5.pdf>, Accessed: 2024-03-29. 6, 1
- [29] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images, 2009. 6
- [30] Jaewoo Lee and et al. Scaling up differentially private deep learning with fast per-example gradient clipping. *Proceedings on Privacy Enhancing Technologies*, 2021(1), 2021. 2
- [31] Joon-Woo Lee and et al. Privacy-preserving machine learning with fully homomorphic encryption for deep neural network. *IEEE Access*, 10:30039–30054, 2022. 4
- [32] Ping Li, Jin Li, Zhengan Huang, and et al. Multi-key privacy-preserving deep learning in cloud computing. *Future Generation Computer Systems*, 74:76–85, 2017. 4
- [33] Zheng Li and Yang Zhang. Membership leakage in label-only exposures. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 880–895, 2021. 8

- [34] Zhaobo Lu, Hai Liang, and et al. Label-only membership inference attacks on machine unlearning without dependence of posteriors. *International Journal of Intelligent Systems*, 37(11):9424–9441, 2022. 8
- [35] Brendan McMahan, Eider Moore, and et al. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. 1
- [36] Midjourney. Midjourney. <https://www.midjourney.com/home>, Accessed: 2023-04-21. 1
- [37] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013. 3
- [38] Fatemehsadat Miresghallah, Mohammadkazem Taram, and et al. Shredder: Learning noise distributions to protect inference privacy. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 3–18, 2020. 1
- [39] Fan Mo, Ali Shahin Shamsabadi, and et al. Darknetz: towards model privacy at the edge using trusted execution environments. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, pages 161–174, 2020. 1
- [40] Krishna Giri Narra, Zhifeng Lin, and et al. Origami inference: Private inference using hardware enclaves. In *2021 IEEE 14th International Conference on Cloud Computing (CLOUD)*, pages 78–84. IEEE, 2021. 1
- [41] Yue Niu, Ramy E Ali, and Salman Avestimehr. 3legrace: Privacy-preserving dnn training over tees and gpus. *Proceedings on Privacy Enhancing Technologies*, 4:183–203, 2022. 1, 2, 3, 4, 6, 7
- [42] Yue Niu, Saurav Prakash, Souvik Kundu, Sunwoo Lee, and Salman Avestimehr. Overcoming resource constraints in federated learning: Large models can be trained with only weak clients. *Transactions on Machine Learning Research*, 2023. 1
- [43] Nvidia. Nvidia confidential computing. <https://www.nvidia.com/en-us/data-center/solutions/confidential-computing/>, Accessed: 2023-04-21. 4
- [44] OpenAI. Introducing chatgpt. <https://openai.com/blog/chatgpt>, Accessed: 2023-04-21. 1
- [45] OpenAI. March 20 chatgpt outage: Here’s what happened. <https://openai.com/blog/march-20-chatgpt-outage>, Accessed: 2023-04-21. 1
- [46] Saurav Prakash, Hanieh Hashemi, and et al. Secure and fault tolerant decentralized learning. *arXiv preprint arXiv:2010.07541*, 2020. 4
- [47] Do Le Quoc, Franz Gregor, and et al. Securetf: A secure tensorflow framework. In *Proceedings of the 21st International Middleware Conference*, pages 44–59, 2020. 1, 2, 4
- [48] Sina Sajadmanesh, Ali Shahin Shamsabadi, Aurélien Bellet, and Daniel Gatica-Perez. Gap: Differentially private graph neural networks with aggregation perturbation. In *USENIX Security 2023-32nd USENIX Security Symposium*, 2023. 6
- [49] Ahmed Salem, Yang Zhang, and et al. MI-leaks: Model and data independent membership inference attacks and defenses on machine learning models. In *Network and Distributed Systems Security (NDSS) Symposium 2019*, 2019. 7, 8
- [50] Lukas Struppek, Dominik Hintersdorf, and et al. Plug and play attacks: Towards robust and flexible model inversion attacks. In *International Conference on Machine Learning*, pages 20522–20545. PMLR, 2022. 7
- [51] Hassan Takabi, Ehsan Hesamifard, and Mehdi Ghasemi. Privacy preserving multi-party machine learning with homomorphic encryption. In *29th Annual Conference on Neural Information Processing Systems (NIPS)*, 2016. 1, 4
- [52] Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In *Annual Cryptology Conference*, pages 71–89. Springer, 2013. 2
- [53] Chandra Thapa, Mahawaga Arachchige Pathum Chamikara, and Seyit A Camtepe. Advancements of federated learning towards privacy preservation: from federated learning to split learning. *Federated Learning Systems: Towards Next-Generation AI*, pages 79–109, 2021. 4
- [54] Chandra Thapa, Pathum Chamikara Mahawaga Arachchige, and et al. Splitfed: When federated learning meets split learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8485–8493, 2022. 4
- [55] Florian Tramèr and Dan Boneh. Slalom: Fast, verifiable and private execution of neural networks in trusted hardware. In *International Conference on Learning Representations*, 2019. 1, 2, 7
- [56] Aravind Vasudevan, Andrew Anderson, and David Gregg. Parallel multi channel convolution using general matrix multiplication. In *2017 IEEE 28th international conference on application-specific systems, architectures and processors (ASAP)*, pages 19–24. IEEE, 2017. 1
- [57] Praneeth Vepakomma, Otkrist Gupta, and et al. Split learning for health: Distributed deep learning without sharing raw patient data. *arXiv preprint arXiv:1812.00564*, 2018. 4
- [58] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 8
- [59] Di Xie, Jiang Xiong, and Shiliang Pu. All you need is beyond a good init: Exploring better solution for training extremely deep convolutional neural networks with orthonormality and modulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6176–6185, 2017. 5
- [60] Da Yu, Huishuai Zhang, and et al. Do not let privacy overbill utility: Gradient embedding perturbation for private learning. *arXiv preprint arXiv:2102.12677*, 2021. 2
- [61] Tianwei Zhang, Zecheng He, and Ruby B Lee. Privacy-preserving machine learning through data obfuscation. *arXiv preprint arXiv:1807.01860*, 2018. 1
- [62] Yuheng Zhang, Ruoxi Jia, Hengzhi Pei, Wenxiao Wang, Bo Li, and Dawn Song. The secret revealer: Generative model-inversion attacks against deep neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 253–261, 2020. 7
- [63] Wei Zhou, Yue Niu, and et al. Sensitivity-oriented layer-wise acceleration and compression for convolutional neural network. *IEEE Access*, 7:38264–38272, 2019. 5