PRIVATE SPLIT INFERENCE OF DEEP NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Splitting network computations between the edge device and the cloud server is a promising approach for enabling low edge-compute and private inference of neural networks. Current methods for providing the privacy train the model to minimize information leakage for a given set of private attributes. In practice, however, the test queries might contain private attributes that are not foreseen during training. We propose an alternative solution, in which, instead of obfuscating the information corresponding to a set of attributes, the edge device discards the information irrelevant to the main task. To this end, the edge device runs the model up to a split layer determined based on its computational capacity and then removes the activation content that is in the null space of the next layer of the model before sending it to the server. It can further remove the low-energy components of the remaining signal to improve the privacy at the cost of reducing the accuracy. The experimental results show that our methods provide privacy while maintaining the accuracy and introducing only a small computational overhead.

1 INTRODUCTION

The surge in cloud computing and machine learning in recent years has led to the emergence of Machine Learning as a Service (MLaaS), where the compute capacity of the cloud is used to analyze the data that lives on edge devices. One shortcoming of the MLaaS framework is the leakage of the clients' private data to the cloud server. Several cryptography-based solutions have been proposed to address this privacy leakage issue (Juvekar et al. (2018); Riazi et al. (2019)); however, the large communication and computational overhead of cryptographic tools hinders their wide-scale adoption in resource-constrained edge devices. Alternatively, one can manually hide private information on the edge device; For instance, sensitive information in an image can be blurred on the edge device before sending it to the service provider (Vishwamitra et al. (2017)). This approach, however, is task-specific and may not be viable for generic applications.

The objective of split inference, as shown in Figure 1, is to provide a generic and computationally efficient data obfuscation scheme (Kang et al. (2017); Chi et al. (2018)). The service provider trains the model and splits it into M_1 and M_2, where M_1 contains the first few layers of the whole model and M_2 contains the rest. The client runs M_1 on the edge device and sends the resulting feature vector $z = M_1(x)$ to the server, which computes the public label as $y^{\text{pub}} = M_2(z)$. To preserve the privacy, the client desires z to only contain information related to the underlying task. For instance, when sending facial features for cell-phone authentication, the client does not want to disclose other information such as the type of makeup he/she is wearing. To quantify the privacy leakage, Figure 1 also depicts an adversary that trains the model M_3 to extract private label y^{pri} from feature vector z.

Current methods of private split inference aim to censor information corresponding to a black-list of known private attributes. Feutry et al. (2018) utilize adversarial training to minimize the accuracy of M_3 on the private task. Osia et al. (2018) minimize the mutual information between the query z and the private label y^{pri} at training time. Mireshghallah et al. (2020) propose to add structured noise to the query z and send the noisy features to the server. The above works commonly assume knowledge of private labels at training time. Such private information or attributes, however, can vary from one query to another. Hence, it is not feasible to foresee all information that could be considered private for a specific MLaaS application. Moreover, gathering labels and training models for all such private attributes significantly increase the cost of deploying ML models on the cloud.

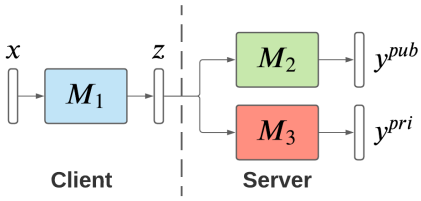

Figure 1: Split inference setup. Client runs M_1 locally and sends the features $z = M_1(x)$ to the server. The server predicts the intended attribute as $y^{\text{pub}} = M_2(z)$. An adversary trains a separate model M_3 to predict secret attributes as $y^{\text{pri}} = M_3(z)$.

In this paper, we propose an alternative solution where, instead of censoring information that is utilized to extract known private labels, we hide information that is not used by the main model. Our contributions are summarized in the following.

- We characterize the information that is not relevant to the prediction of the public attribute as the content in the feature vector z that will be removed by the server-side model. We then define the null content of the feature vector, $z_{\mathcal{N}}$, as the content in z that is in the null-space of the following linear layer. The remaining content is called signal content denoted by $z_{\mathcal{S}}$. We have $M_2(z) = M_2(z_{\mathcal{S}} + z_{\mathcal{N}}) = M_2(z_{\mathcal{S}})$.
- We show that removing $z_{\mathcal{N}}$ from features does not affect the accuracy of the main model (M_2) but reduces the accuracy of the adversary (M_3). The null content, however, does not necessarily capture all the private information in z as some information might still be leaked through $z_{\mathcal{S}}$; We show that, by removing the low-energy components of $z_{\mathcal{S}}$, we can achieve higher privacy (lower accuracy of M_3) at the cost of a small reduction of utility (accuracy of M_2).
- We show the inherent tradeoff between edge computation-efficiency, privacy, and accuracy. By computing more layers on the edge-device, the client discards more private information from features, which results in a less accurate adversary model. Also, for the same edge computation (a given split layer), removing part of the signal content provides a tradeoff between accuracy and privacy.
- We perform extensive experiments on several datasets used for split inference privacy. We show our method provides a better tradeoff between accuracy and privacy compared to existing approaches such as adversarial training, despite assuming that the private attribute is not known at training or inference times. We show that our methods provide privacy as the private accuracy can be lowered down close to the random classifier, while maintaining the performance on the main task.

2 Background and Related Work

We consider the supervised learning setting of Figure 1, where the model $M_2 \circ M_1$ is trained with a set of examples $\{x_i\}_{i=1}^N$ and their corresponding public labels $\{y_i^{\text{pub}}\}_{i=1}^N$. The client runs M_1 on their data and sends the intermediate feature vector $z = M_1(x)$ to the server. The goal of private inference is to ensure that z does not contain information about private attributes.

Measuring privacy. Several methods have been proposed to measure the privacy leakage of the feature vector. One approach is computing the mutual information between the query x and the feature vector z (Kraskov et al. (2004)). In practice, measuring the mutual information is not tractable for high-dimensional random variables, unless certain assumptions are made about the probability distribution of the random variables of interest. A more practical approach defines privacy based on reconstruction error (Mahendran & Vedaldi (2015)); given a reconstruction algorithm that takes z as input and outputs an estimate $\widetilde{x}$, the error $||\widetilde{x} - x||$ quantifies privacy. Finally, attribute privacy is defined based on the accuracy of an adversary model trained with z on a private label.

In this paper, we use the attribute privacy notion as it applies to a wide range of applications. Assume each example $\{x_i\}_{i=1}^N$ has one or multiple private labels $\{y_i^{\text{pri}}\}_{i=1}^N$. The adversary trains a separate model M_3 on z, the output of the pre-trained M_1, and a private label. The split learning framework should satisfy *utility*, i.e., the server should be able to infer the public attribute from z accurately, while providing *privacy*, i.e., z should not contain information about y^{pri}. We refer to the accuracy of $M_2 \circ M_1$ on y^{pub} as public accuracy and the accuracy of $M_3 \circ M_1$ on y^{pri} as private accuracy.

Earlier work has shown that the representations learned by deep neural networks can be used to extract private information (Song & Shmatikov (2019)) or even reconstruct the raw data (Mahendran & Vedaldi (2015)). Current methods for private inference can be categorized as follows.

Noise Injection. A line of work suggests obfuscating private attributes by adding noise to the features, i.e., instead of z, the client sends $z + \mu$ to the server, with the noise designed to maintain public accuracy while reducing private accuracy (Mireshghallah et al. (2020)). While noise addition can improve privacy, it has been to shown to reduce the public accuracy significantly (Liu et al. (2019)).

Information Bottlenecks. The notion of mutual information can be utilized as a proxy to provide privacy (Tishby & Zaslavsky (2015)). Let $I(a, b)$ denote the mutual information between random variables a and b. The idea is to train M_1 that maximizes $I(z, y^{\text{pub}})$ while minimizing $I(z, y^{\text{pri}})$ (Osia et al. (2018); Moyer et al. (2018)). The optimization is formulated as follows:

$$\max_{M_1} \quad \mathbb{E}_{x,y^{\text{pub}},y^{\text{pri}}}[I(M_1(x), y^{\text{pub}}) - \gamma I(M_1(x), y^{\text{pri}}) - \beta I(M_1(x), x)]. \tag{1}$$

The use of mutual information for privacy, however, has been challenged by practical attacks that extract secret information even if $I(z, y^{\text{pri}})$ is minimized (Song & Shmatikov (2019)).

Adversarial Training. This defense solves the following min-max optimization problem:

$$\max_{M_1,M_2} \min_{M_3} \mathbb{E}_{x,y^{\text{pub}},y^{\text{pri}}}[\gamma \mathcal{L}(y^{\text{pri}}, M_3 \circ M_1(x)) - \mathcal{L}(y^{\text{pub}}, M_1 \circ M_2(x))], \tag{2}$$

where $\mathcal{L}$ denotes the cross-entropy loss and γ is a scalar. The above objective can be achieved through adversarial training (Edwards & Storkey (2015); Hamm (2017); Xie et al. (2017); Li et al. (2018); Feutry et al. (2018); Li et al. (2019)). At convergence, the trained M_1 generates z such that $M_3(z)$ is not an accurate estimation of y^{pri} while $M_2(z)$ accurately describes y^{pub}.

Existing methods for private split inference have several limitations. First, the underlying assumption in all above defenses is that a set of private attributes along with the public label are provided at training time. In practice, however, it might not be feasible to foresee and identify all possible private attributes and annotate the training data accordingly. It also contradicts deployment at-scale since whenever a new private attribute is identified, the model M_1 needs to be retrained and re-distributed to all edge devices that use the service. Second, current approaches for private inference often provide a poor tradeoff between accuracy and privacy. In addition, the tradeoff of accuracy and privacy with the client-side computation is not well studied in the split learning framework. In this paper, we characterize this tradeoff and propose an alternative approach that, instead of obfuscating the information related to the private attributes, the edge device removes the feature content that is irrelevant to the public task. We empirically show our method successfully reduces the accuracy on private attributes at a small or no cost to public accuracy.

3 Proposed Methods

3.1 Signal and Null contents of Feature Vector

Let $z \in \mathbb{R}^n$ be a feature vector and $W \in \mathbb{R}^{m \times n}$ be a matrix. The operation of fully-connected and convolutionals layers can be modeled by matrix-vector and matrix-matrix multiplication. Herein, we let z represent the vector in fully-connected layer and a column of the second matrix in convolution layer. Let the singular value decomposition (SVD) of W be $W = U \cdot S \cdot V$. Since the rows of V form an orthonormal basis, we can write the feature vector as

$$z = \sum_{i=1}^{n} \alpha_i v_i^T, \qquad \alpha_i =< v_i^T, z >, \tag{3}$$

where v_i is the i-th row of V and the $< \cdot, \cdot >$ operator denotes inner-product.

Definition 1. *The signal content of z with respect to matrix W, or simply the signal content of z, denoted by z_S is defined as*

$$z_{\mathcal{S}} = \arg\min_h \|h\|_2, \quad \textit{s.t.}, W \cdot (z - h) = 0. \tag{4}$$

The null content of z is also defined as $z_{\mathcal{N}} = z - z_{\mathcal{S}}$.

Lemma 1. *We have*

$$z_{\mathcal{S}} = \sum_{i=1}^{m} \alpha_i v_i^T \quad \textit{and} \quad z_{\mathcal{N}} = \sum_{i=m+1}^{n} \alpha_i v_i^T. \tag{5}$$

Proof. We write h as the composition of orthonormal vectors v_i's as $h = \sum_{i=1}^{n} \beta_i v_i^T$. We have

$$W(z-h) = \sum_{i=1}^{n} (\alpha_i - \beta_i) W v_i^T = \sum_{i=1}^{n} (\alpha_i - \beta_i) US \underbrace{V v_i^T}_{q_i \in \mathbb{R}^n} \tag{6}$$

Since the rows of V are orthonormal, then $q_i = V v_i^T$ is a one-hot vector with its i-th element equal to 1. By substituting q_i in (6) we obtain

$$W(z-h) = \sum_{i=1}^{n} (\alpha_i - \beta_i) US_{[:,i]} = \sum_{i=1}^{m} (\alpha_i - \beta_i) US_{[:,i]} = \sum_{i=1}^{m} s_i (\alpha_i - \beta_i) U_{[:,i]} \tag{7}$$

where $S_{[:,i]}$ and $U_{[:,i]}$ are the i-th columns of S and U, respectively. Note that, since S is a diagonal matrix, we have $S_{[:,i]} = 0, \forall i \in \{m+1, \cdots, n\}$, thus reducing the summation from n to m components. Also, for $i \in \{1, \cdots, m\}$, $S_{[:,i]}$ is a column vector with only one non-zero element, s_i, at the i-th dimension.

As a result, to obtain $W(z-h) = 0$, we must have $\beta_i = \alpha_i, \forall i \in \{1, \cdots, m\}$. Since v_i's are orthonormal, we have $\|h\|_2 = \sqrt{\sum_{i=1}^{n} \beta_i^2}$. Hence, to minimize $\|h\|_2$, we set $\beta_i = 0, \forall i \in \{m+1, \cdots, n\}$. Therefore, $z_{\mathcal{S}} = \sum_{i=1}^{m} \alpha_i v_i^T$. The null content $z_{\mathcal{N}}$ can be then computed as $z_{\mathcal{N}} = z - z_{\mathcal{S}} = \sum_{i=m+1}^{n} \alpha_i v_i^T$. □

Definition 2. *The normalized signal and null contents of z are defined as $C_{\mathcal{S}}(z) = \frac{||z_{\mathcal{S}}||_2^2}{||z||_2^2}$ and $C_{\mathcal{N}}(z) = \frac{||z_{\mathcal{N}}||_2^2}{||z||_2^2}$, respectively. We have $C_{\mathcal{S}}(z) + C_{\mathcal{N}}(z) = \frac{\sum_{i=1}^{m} \alpha_i^2}{\sum_{i=1}^{n} \alpha_i^2} + \frac{\sum_{i=m+1}^{n} \alpha_i^2}{\sum_{i=1}^{n} \alpha_i^2} = 1$.*

3.2 Defense 1: Obfuscating Null content of Feature Vector

We propose to remove all the content in feature vector that is irrelevant to the public attribute. To do so, given a feature vector z, we find a minimum-norm vector z' that generates the same prediction for the public attribute as z, i.e., $M_2(z') = M_2(z)$. Formally,

$$z' = \arg\min_{h} \|h\|_2, \quad \text{s.t., } M_2(h) = M_2(z). \tag{8}$$

Due to the complex (nonlinear) nature of deep networks, finding such a vector would require doing multiple backpropagations on M_2 for each given z. This is, however, not feasible for resource-constrained edge devices. In order to reduce the edge computation, we relax the constraint of (8) to $Wz' = Wz$, where W is the weight matrix of the first layer of M_2 (server-side model). As discussed in Section 3.1, the solution to this relaxed optimization problem is $z_{\mathcal{S}}$, the signal content of z. Removing or obfuscating the null content of z does not change the model prediction on public attribute. It, however, might harm the private accuracy since part of the null content of $z_{\mathcal{N}}$ might fall into the signal content of the first linear layer of M_3. The method is described in Figure 2 (left).

At inference time, to obfuscate $z_{\mathcal{N}}$, the client constructs z_o using either of the following methods.

- Client constructs $\mu = \sum_{i=m+1}^{n} \eta_i v_i^T$, with η_i coefficients chosen at random, and sends $z_o = z + \mu$ to the server. The adversary can recover $z_{\mathcal{S}} = V_{1:m}^T \cdot V_{1:m} \cdot z_o$ but cannot recover $z_{\mathcal{N}}$.
- Client computes the signal content of z and sends $z_o = z_{\mathcal{S}} = \sum_{i=1}^{m} \alpha_i v_i^T$ to the server.

For the first case, since μ is independent of z, the client can compute it offline, e.g., when the edge device is plugged in and not in use, and store it for later use. The second approach does not require storage on the edge device but extra computation has to be done during inference to extract $z_{\mathcal{S}}$.

Removing null content	Removing signal content
Training: **Server** trains $M_2 \circ M_1$ with $(x_i, y_i^{\text{pub}}), i \in [N]$ **Adversary** computes $z_i = M_1(x_i), i \in [N]$ **Adversary** computes $z_i \rightarrow z_{i\mathcal{S}}, i \in [N]$ **Adversary** trains M_3 with $(z_{i\mathcal{S}}, y_i^{\text{pri}}), i \in [N]$ **Inference:** **Client** computes $z = M_1(x)$ **Client** computes $z \rightarrow z_o$ **Client** sends z_o to server **Server** computes $\hat{y}^{\text{pub}} = M_2(z_o) = M_2(z_{\mathcal{S}})$ **Adversary** computes $z_o \rightarrow z_{\mathcal{S}}$ **Adversary** computes $\hat{y}^{\text{pri}} = M_3(z_{\mathcal{S}})$	**Training:** **Server** trains $M_2 \circ M_1$ using $(x_i, y_i^{\text{pub}}), i \in [N]$ **Server** augments $M_1 \rightarrow \widetilde{M}_1$ **Server** freezes $\widetilde{M}_1$ and fine-tunes $M_2 \circ \widetilde{M}_1$ **Adversary** computes $\widetilde{z}_i = \widetilde{M}_1(x_i), i \in [N]$ **Adversary** trains M_3 with $(\widetilde{z}_i, y_i^{\text{pri}}), i \in [N]$ **Inference:** **Client** computes $\widetilde{z} = \widetilde{M}_1(x)$ **Client** sends $\widetilde{z}$ to server **Server** computes $\hat{y}^{\text{pub}} = M_2(\widetilde{z})$ **Adversary** computes $\hat{y}^{\text{pri}} = M_3(\widetilde{z})$

Figure 2: (left): Training and inference for defense 1 (client removes null content of features). (right): Training and inference for defense 2 (client removes low-energy content of z). $\widetilde{M}_1$ is constructed by augmenting M_1 with the module that removes the signal content.

3.3 Defense 2: Discarding Low-Energy Signal Content of Feature Vector

In the first defense method, we proposed to discard the content of feature vector which would have been removed by the first layer of M_2. The following layers of M_2 will further remove more content from feature vector. Hence, we can potentially discard more content from z and still get similar prediction as the original feature vector. One candidate for discarding is the low-energy component of $z_{\mathcal{S}}$. For a linear layer, following the same process in Section 3.1, the output is computed as:

$$W \cdot z = W \cdot z_{\mathcal{S}} = \sum_{i=1}^{m} \alpha_i U \cdot S \cdot q_i = \sum_{i=1}^{m} \alpha_i U \cdot S_{[:,i]} = \sum_{i=1}^{m} s_i \alpha_i U_{[:,i]} \tag{9}$$

where s_i is the i-th eigenvalue in S, α_i is defined in (3), and $U_{[:,i]}$ denotes the i-th column of U. From (9) we observe that components with larger $s_i \alpha_i$ are contributing more to the layer output since $||U_{[:,i]}||_2 = 1$ for all columns of U. As such, we approximate $z \rightarrow \widetilde{z}$ by only keeping $m' < m$ components of the right-hand-side of (9) that have the largest coefficients. To improve public accuracy when removing signal content of features, the server fine-tunes M_2 on $\widetilde{z}$. The method is described in Figure 2 (right).

Note that s_i and $U_{[:,i]}$ are fixed at the inference time, thus, the client only needs to send the selected α_i values along with their indices to the server; the server knows s_i and U and can reconstruct $\widetilde{z}$ accordingly. Hence, choosing $m' \ll m$ also reduces the communication cost of sending the feature vector to the server. Unlike null content filtering, this approach will affect the public attribute accuracy since it changes the main network output, but can further reduce the private accuracy.

4 Experiments

4.1 Setup

Datasets. We perform our experiments on four visual datasets listed below.

- **EMNIST** (Cohen et al. (2017)) is an extended version of the MNIST dataset where the labels are augmented with writer identities. To have a uniform distribution of writer identities, we selected 13000 samples from EMNIST written by 100 writers with 130 examples per writer. We then split this dataset into 10000, 1500 and 1500 training, validation and test sets. We use the digit label and writer ID as the public and private attributes, respectively.
- **FaceScrub** (Ng & Winkler (2014); facescrub (2020)) is a dataset of celebrity faces labeled with gender and identity. We use gender and identity as the public and private attributes, respectively. In experiments, we cropped the face region using the bounding boxes in the image annotations and resized images to 50×50.

- **UTKFace** (Zhang et al. (2017)) is a dataset of face images labeled with gender and race. We use gender and race as the public and private attributes, respectively. We cropped the face region using the bounding boxes in the image annotations and resized images to 50×50.
- **CelebA** (Liu et al. (2015)) is a dataset of celebrity images. Each image is labeled with 40 binary attributes. Out of these, we select "Smiling" as the public attribute and {Male, Heavy_Makeup, High_Cheekbones, Mouth_Slightly_Open, Wearing_Lipstick, Attractive} as private attributes. These attributes have near balanced distribution of positive and negative examples. In experiments, we cropped the face region using the bounding boxes in the image annotations and resized images to 73×60.

Model architecture. We present the experimental results on a model used in prior work (Song & Shmatikov (2019)). The model architecture and baseline test accuracies are summarized in Table 1 and Table 2 in Appendix.

Adversary capabilities. We use the same architecture for adversary's model M_3 as the server model M_2. The model M_3 is trained using the features extracted by M_1 and the associated private labels. We also assume that the adversary knows the parameters of M_1 and M_2.

Training settings. We use Adam optimizer with an initial learning rate of 0.001. We drop the learning rate by a factor of 10 after 20 and 40 epochs. All models including the adversary and main models are trained for 50 epochs unless stated otherwise.

4.2 Evaluations

We start our analysis by computing the null and signal contents in every layer of $M = M_2 \circ M_1$. Figure 3 (left) shows the content of the input remained at each layer; for the i-th layer, this content is computed as $\prod_{j=1}^{i} C_{\mathcal{S}}(z_j)$ where z_j denotes the activation vector at the j-th layer and $C_{\mathcal{S}}(\cdot)$ is defined in Section 3.1. As the feature vector propagates through network layers, more content is gradually removed from z until the model outputs the prediction on the public task. We also split the network at different layers and measure the private accuracy of M_3 trained with the feature vector. As seen in Figure 3 (right), the private accuracy also decreases as we get closer to the output layer, indicating that the discarded content contained information relevant to the private attribute.

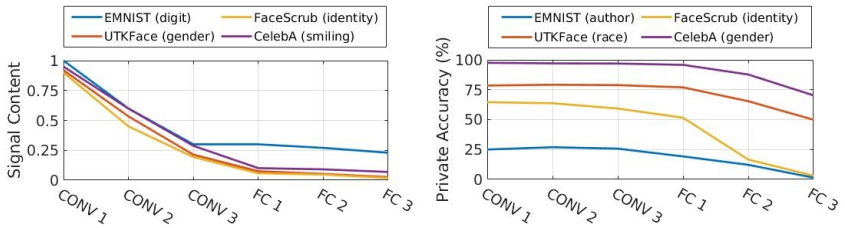

Figure 3: (left) Remaining signal content at the output of each layer of the main model, (right) private accuracy for different split layers. The content of the input significantly decreases as we get closer to the output layer. Similarly, the private accuracy decreases, indicating that the discarded content contained information relevant to the private attribute.

To reduce privacy leakage, we proposed to filter out the null content and part of the signal content of the feature vector. Figure 4 shows public and private accuracy for the two defense methods and for different split layers. As can be seen, removing the null content reduces the private accuracy without affecting the public accuracy. Also, removing part of the signal content further reduces the private accuracy, in some cases to the level of random guess, while slightly reducing the public accuracy. Figure 4 shows that the information leakage can be controlled using the following factors factors:

- **Number of signal components sent to the server:** For the same edge-computation (a given split layer), the number of preserved features (m') can be tuned so as to achieve a desired tradeoff between utility (higher public accuracy) and privacy (lower private accuracy).
- **The split layer:** Better tradeoffs can be achieved when the network is split from deeper layers, i.e., the edge device performs more computation.

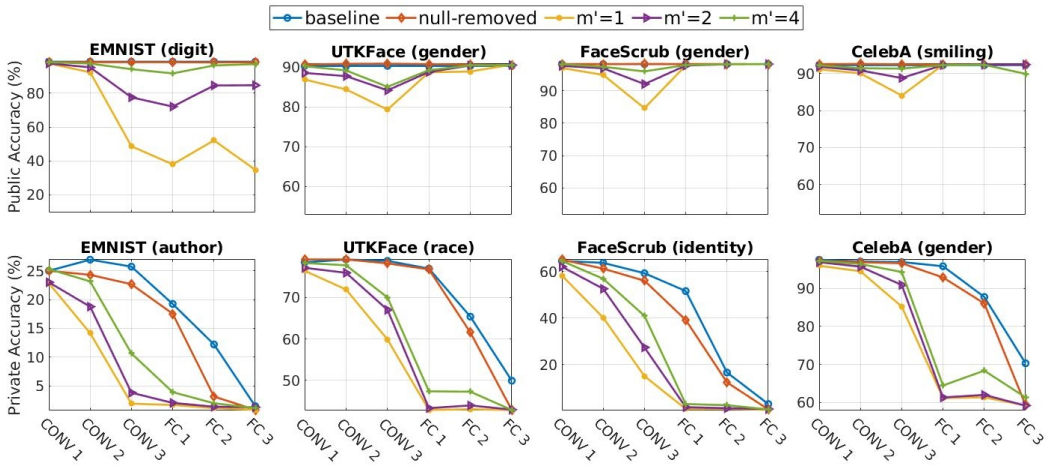

Figure 4: The effect of the split layer and our proposed methods for filtering features vector on (top) public and (bottom) private accuracy. The Y-axis in each figure covers the range between random guess and the baseline accuracy. For a given split layer, the number of preserved features (m') can be tuned so as to achieve a desired tradeoff between utility (higher public accuracy) and privacy (lower private accuracy). Also, better tradeoffs can be achieved when the network is split from deeper layers.

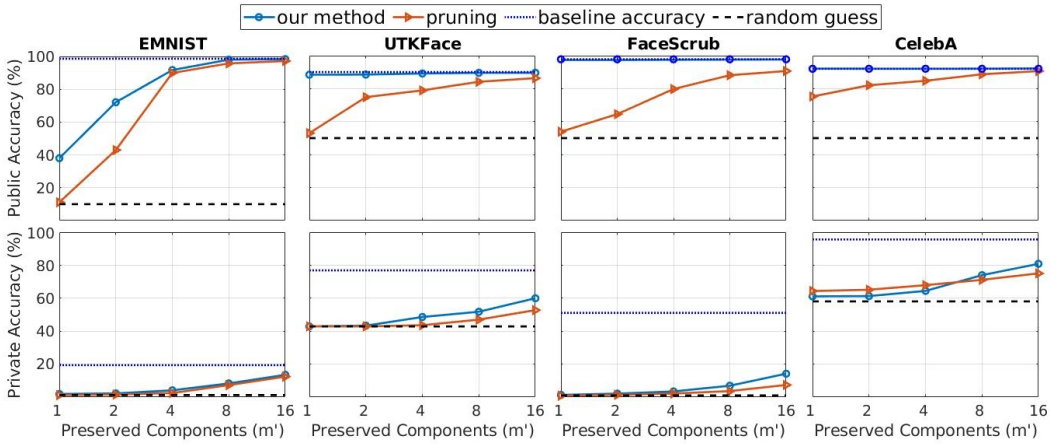

Figure 5: Comparison between our method and feature pruning. Both methods reduce the private accuracy. Pruning, however, significantly reduces the public accuracy.

Relation to Pruning. Similar to our approach, pruning network layers can eliminate features that do not contribute to the public attributed. In the following, we compare our method with pruning in terms of public and private accuracy. We split the network from the middle layer, i.e., at the input of the FC-1 layer. For our method, we keep the top m' components of z from its signal content and filter out the rest. For pruning, we keep m' elements in z and set the rest to zero. We adopt the pruning algorithm proposed by (Li et al. (2016)) which selects the (non-pruned) elements based on the L_1 norm of the columns of the following layer's weight matrix. After pruning, we fine-tune the M_2 model to enhance the public accuracy. Figure 5 shows the public and private accuracy for each dataset. As seen, with small m', both our method and pruning achieve a low private accuracy. Pruning, however, significantly reduces the public accuracy as well. For example, for the UTKFace dataset, with $m' = 1$, both methods result in a private accuracy close to the random guess (42.84%). However, pruning reduces the public accuracy from 92.25% to 53% (also close to random guess), whereas our method keeps the public accuracy at 88.63%.

Comparison with adversarial training. We implement adversarial training framework proposed by (Feutry et al. (2018)) and present the utility-privacy tradeoff in Figure 6. To achieve the best tradeoff for adversarial training, we train the models in multiple settings with different γ parameters (Eq. 2) in range of $[0.1, 1]$. Note that, unlike our method, adversarial training assumes that the private

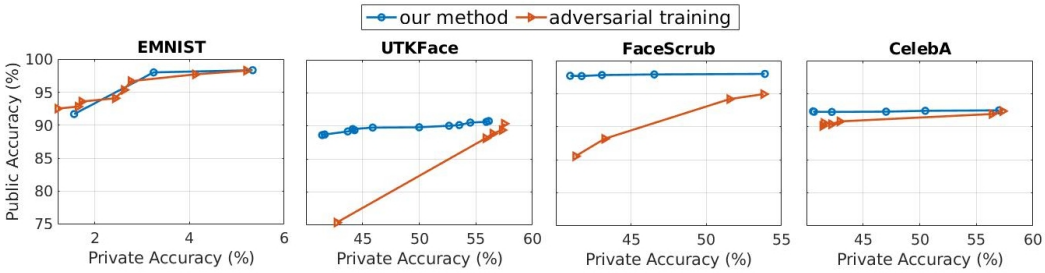

Figure 6: Comparison between our method and adversarial training. For a given privacy level, our method provides higher utility (higher public accuracy) compared to adversarial training method.

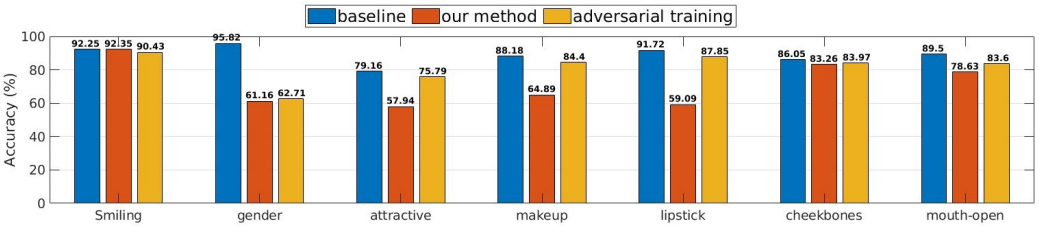

Figure 7: Accuracy on public ("smiling") and private attributes. Our method obfuscates the feature vector without the knowledge of private attribute at training or inference times. Adversarial training method maximizes the accuracy on "smiling," while minimizing the accuracy on "gender." As can be seen, our method reduces accuracy on all private attributes. The adversarially trained model successfully reduces accuracy on "gender" attribute, but fails to remove information of other attributes. This highlights the applicability of our method in practical setting as a generic obfuscator compared to specialized techniques such as adversarial training.

attribute is known at training time. Despite this, Figure 6 shows that our method achieves a better utility-accuracy tradeoff than adversarial training.

We also do experiments for the case with multiple (unseen) private labels. Specifically, we consider the CelebA model trained to detect "smiling" and evaluate two methods, 1) our method: we keep only $m' = 1$ component from the signal content of feature vector and then train one adversary model per private attribute, and 2) adversarial training: we first adversarially train an M_1 model to obfuscate "gender," and then train one model for each private attribute.

For both of the above methods, the network is split at the input of the FC-1 layer. Figure 7 shows the results. Our method outperforms adversarial training method on both public and private accuracy. In our method, the accuracy on all private attributes are significantly lower than the baseline private accuracy. The only exceptions are "high cheekbones" and "mouth open" attributes, which have correlations with public attribute, that is, a smiling person likely has high cheekbones and their mouth open. The correlation between public and private attributes causes the signal content of server and adversary's models to have large overlaps and, hence, results in high private accuracy. The adversarially trained model successfully hides the information that it has been trained to obfuscate (the "gender" attribute). Such a model, however, fails to remove information of other attributes such as "makeup" or "lipstick". This highlights the applicability of our method in practical setting as a generic obfuscator compared to specialized techniques such as adversarial training.

5 Conclusion

We proposed a private inference framework, in which edge devices run several layers locally and obfuscate the intermediate feature vector before sending it to the server to execute the rest of the model. For obfuscation, we proposed to remove information that is not relevant to the main task or does not significantly change the predictions. In doing so, we developed two methods of removing the content of the feature vector in the null space of the following linear layer and also removing the low-energy content of the remaining signal. We show that, unlike existing methods, our methods improve privacy without requiring the knowledge of private attributes at training or inference times.

REFERENCES

Jianfeng Chi, Emmanuel Owusu, Xuwang Yin, Tong Yu, William Chan, Patrick Tague, and Yuan Tian. Privacy partitioning: Protecting user data during the deep learning inference phase. *arXiv preprint arXiv:1812.02863*, 2018.

Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist to handwritten letters. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 2921–2926. IEEE, 2017.

Harrison Edwards and Amos Storkey. Censoring representations with an adversary. *arXiv preprint arXiv:1511.05897*, 2015.

facescrub. *the facescrub dataset*, 2020. `http://engineering.purdue.edu/~mark/puthesis`, (accessed July 3, 2020).

Clément Feutry, Pablo Piantanida, Yoshua Bengio, and Pierre Duhamel. Learning anonymized representations with adversarial neural networks. *arXiv preprint arXiv:1802.09386*, 2018.

Jihun Hamm. Minimax filter: learning to preserve privacy from inference attacks. *The Journal of Machine Learning Research*, 18(1):4704–4734, 2017.

Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. {GAZELLE}: A low latency framework for secure neural network inference. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pp. 1651–1669, 2018.

Yiping Kang, Johann Hauswald, Cao Gao, Austin Rovinski, Trevor Mudge, Jason Mars, and Lingjia Tang. Neurosurgeon: Collaborative intelligence between the cloud and mobile edge. *ACM SIGARCH Computer Architecture News*, 45(1):615–629, 2017.

Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical review E*, 69(6):066138, 2004.

Ang Li, Jiayi Guo, Huanrui Yang, and Yiran Chen. Deepobfuscator: Adversarial training framework for privacy-preserving image classification. *arXiv preprint arXiv:1909.04126*, 2019.

Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.

Yitong Li, Timothy Baldwin, and Trevor Cohn. Towards robust and privacy-preserving text representations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 25–30, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-2005. URL `https://www.aclweb.org/anthology/P18-2005`.

Sicong Liu, Anshumali Shrivastava, Junzhao Du, and Lin Zhong. Better accuracy with quantified privacy: representations learned via reconstructive adversarial network. *arXiv preprint arXiv:1901.08730*, 2019.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5188–5196, 2015.

Fatemehsadat Mireshghallah, Mohammadkazem Taram, Prakash Ramrakhyani, Ali Jalali, Dean Tullsen, and Hadi Esmaeilzadeh. 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*, pp. 3–18, 2020.

Daniel Moyer, Shuyang Gao, Rob Brekelmans, Aram Galstyan, and Greg Ver Steeg. Invariant representations without adversarial training. In *Advances in Neural Information Processing Systems*, pp. 9084–9093, 2018.

Hong-Wei Ng and Stefan Winkler. A data-driven approach to cleaning large face datasets. In *2014 IEEE international conference on image processing (ICIP)*, pp. 343–347. IEEE, 2014.

Seyed Ali Osia, Ali Taheri, Ali Shahin Shamsabadi, Kleomenis Katevas, Hamed Haddadi, and Hamid R Rabiee. Deep private-feature extraction. *IEEE Transactions on Knowledge and Data Engineering*, 32(1):54–66, 2018.

M Sadegh Riazi, Mohammad Samragh, Hao Chen, Kim Laine, Kristin Lauter, and Farinaz Koushanfar. {XONN}: Xnor-based oblivious deep neural network inference. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, pp. 1501–1518, 2019.

Congzheng Song and Vitaly Shmatikov. Overlearning reveals sensitive attributes. In *International Conference on Learning Representations*, 2019.

Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pp. 1–5. IEEE, 2015.

Nishant Vishwamitra, Bart Knijnenburg, Hongxin Hu, Yifang P Kelly Caine, et al. Blur vs. block: Investigating the effectiveness of privacy-enhancing obfuscation for images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 39–47, 2017.

Qizhe Xie, Zihang Dai, Yulun Du, Eduard Hovy, and Graham Neubig. Controllable invariance through adversarial feature learning. In *Advances in Neural Information Processing Systems*, pp. 585–596, 2017.

Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017.

A APPENDIX

Table 1: Network Architecture. Each row shows a split layer, e.g., at layer 1, the raw data and, at layer 6, the input to the last fully-connected layer is sent to the server.

Split Layer				
1	CONV $(3, 16)$	ReLU	Max-Pooling (2×2)	Batch-Normalization
2	CONV$(3, 32)$	ReLU	Max-Pooling (2×2)	Batch-Normalization
3	CONV $(3, 64)$	ReLU	Max-Pooling (2×2)	Batch-Normalization
4	FC(128)	ReLU	-	Batch-Normalization
5	FC(64)	ReLU	-	Batch-Normalization
6	FC(n_classes)	Softmax	-	-

Table 2: Model accuracy for public and private attributes of different datasets.

Dataset	MNIST	UTKFace	FaceScrub	CelebA					
number of classes	10	2	2	2					
public attribute	digit	gender	gender	smiling					
public accuracy (%)	98.60	90.25	97.90	92.25					
number of classes	100	5	530	2	2	2	2	2	2
private attribute	writer	race	identity	gender	makeup	cheekbones	mouth-open	lipstick	attractive
private accuracy (%)	26.93	79.18	65.52	97.53	90.00	86.29	92.94	93.40	80.46