

---

# MistNet: Towards Private Neural Network Training with Local Differential Privacy

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 Deep neural networks (DNNs) generally learn parameters from massive amounts  
2 of high-quality training data to provide superb prediction performance. Classical  
3 machine learning approaches centralize decentralized data dispersed across devices  
4 in a common site for effective training but raise serious concerns of data privacy.  
5 In this paper, we design, implement, and evaluate MistNet, a privacy-preserving  
6 model training system that enables the cloud and the edge devices to collaboratively  
7 perform neural network training without revealing users' data as well as model  
8 parameters. MistNet partitions a DNN model into two parts, a lightweight feature  
9 extractor at the edge side to generate meaningful features from the raw training data,  
10 and a classifier including the most model layers at the cloud to be iteratively trained  
11 for specific tasks. Different from prior work, the feature extractor is transferred  
12 from pre-trained models for similar application domains and kept fixed during  
13 training, which eliminates the need to synchronize feature extractors across devices.  
14 Furthermore, MistNet enhances privacy via applying local differential privacy  
15 (LDP) to the intermediate features and assess the privacy leakage with two kinds  
16 of common attacks - membership inference and feature inversion attacks. We  
17 conduct an experimental study on multiple models and datasets, demonstrating that  
18 by choosing an appropriate partition layer and privacy budget, MistNet achieves  
19 acceptable model utility while greatly reducing privacy leakage from the released  
20 intermediate features.

## 21 1 Introduction

22 Deep neural networks (DNN) have been successfully applied to a wide range of areas including  
23 vision, speech, and natural language [1, 2, 3]. The superior prediction ability of DNN models relies on  
24 large amounts of data. Meanwhile, with the proliferation of mobile and IoT technology, tremendous  
25 valuable data are generated by edge devices but live in silos. There is an increasing demand to learn  
26 from the dispersed data, so as to better support machine learning (ML) tasks at the edge. However, in  
27 the most conventional training paradigm, models are placed at a central site. It requires collecting  
28 training data from users, which raises concerns about data privacy and violates data protection laws.  
29 As a result, privacy issues turn to be a barrier for empowering edge intelligence, and collaborative  
30 training without sharing the input training data is highly desired.

31 Recently, two learning paradigms have emerged to address this issue: *federated learning* [4] and  
32 *split learning* [5, 6]. Federated learning pushes the whole model to the edge and *model gradients* are  
33 exchanged across devices to learn a shared model. As DNNs have become deeper and more complex,  
34 model training has incurred prohibitive costs in computational resources, which poses a substantial  
35 challenge to resource-constrained devices. Split learning evolves to train the first few layers of the

36 neural network at the edge and transmit the *intermediate features* to cloud servers with abundant  
37 computational resources to facilitate training the rest of the training.

38 However, the existing split learning method [5, 6] trains the model in a sequential fashion across edge  
39 devices, which slows down the training process - an edge device has to receive updated weights from  
40 the last trained device before training on its local data. There is a need to enable parallel local training  
41 among distributed edge devices. Exposing the intermediate features instead of the training data is  
42 assumed to be safer, but is it sufficient to protect the privacy of the training data? Recent study [7]  
43 devises an inversion attack to recover the inputs from the intermediate features even in the black-box  
44 setting without the need to know the parameters of the model at the edge.

45 In this paper, we propose MistNet<sup>1</sup>, a privacy-preserving collaborative training framework, in which  
46 we divide the neural network into two parts, the first few layers as a feature extractor at the edge and  
47 the rest layers as a classifier at the cloud. We design the feature extractor based on an insight from  
48 transfer learning [8]: the early layer features are general to many datasets and tasks. Take computer  
49 vision tasks for example, we initialize the feature extractor with weights pre-trained on large public  
50 datasets like ImageNet [9]. The feature extractor thus is ready to produce meaningful features without  
51 the need to be further trained with the rest model on cloud servers. It thus supports edge devices to  
52 perform training in parallel by eliminating the need to synchronize updated feature extractors. To  
53 ensure that intermediate features do not reveal sensitive information about any particular training  
54 record, we adopt the rigorous local differential privacy technique - Randomized Response (RR),  
55 which was introduced by Warner et al. [10] for collecting sensitive statistics from survey respondents  
56 and later widely deployed in real systems by Google [11] and Apple [12]. Subject to RR constraints,  
57 we discretize the features of the partition layer by constraining the values to either 1 or 0. Each  
58 feature value will be independently randomized before leaving the edge devices, thus no raw features  
59 are leaked. The probability to preserve the original value is determined by the privacy budget  $\epsilon$ . A  
60 small value of privacy budget  $\epsilon$  strictly guarantees privacy, but also detracts model utility. We have  
61 limited understanding of the range of  $\epsilon$  values for reasonable privacy-accuracy trade-off in practice.  
62 We thus assess the privacy leakage with two common attacks against ML models, model inversion,  
63 and membership inference attack. We summarize our major contributions as follows:

- 64 • We propose MistNet, which uses a fixed-weight pre-trained feature extractor to generate  
65 meaningful features and further apply local differential privacy on features to enhance  
66 privacy.
- 67 • Besides from privacy budget  $\epsilon$ , we use model inversion and membership inference attack to  
68 quantitatively assess the privacy leakage.
- 69 • We experimentally show that MistNet achieves good prediction accuracy while preserving  
70 privacy under different models, datasets, and parameter settings.

## 71 2 Preliminaries

### 72 2.1 Distributed Collaborative Learning

73 To support training on a substantial amount of training data from different sources, distributed collabo-  
74 rative learning emerges to enable multiple parties to contribute to learning a shared model. Depending  
75 on where the model is located, distributed collaborative learning systems can be categorized into the  
76 following paradigms.

77 **Federated Learning.** To build a more privacy-friendly collaborative training approach, Google  
78 proposes federated learning [4] to enable participating devices collaboratively to train a shared model,  
79 while keeping the training data on mobile devices. As illustrated in Figure 1a, each mobile device  
80 trains the model with its local data for several local epochs, and push the local updated parameters to  
81 the central server, where these updates are aggregated to compute a new model with model averaging.  
82 The updated global model is then sent back to edge devices in the next round. During the whole  
83 training procedure, only model parameters are shared, while the training data is kept locally at the  
84 devices.

---

<sup>1</sup>MistNet is a combination of the words **Mist** and Neural **Net**work. This name reflects the fact that preventing users' data from being revealed is like putting a mist around them.

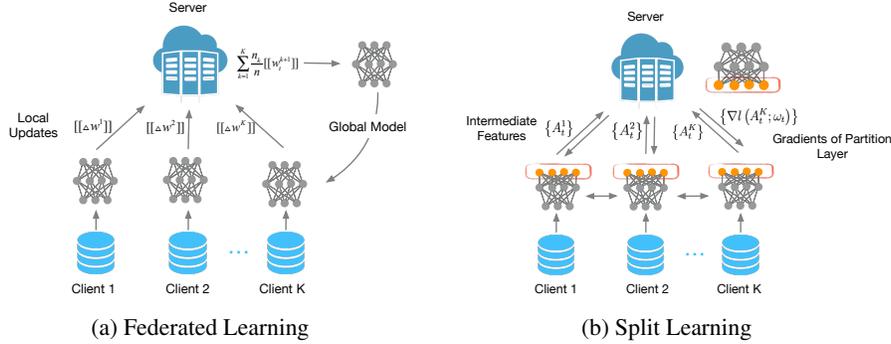


Figure 1: Two distributed collaborative learning paradigms: federated learning and split learning

85 **Split Learning.** Instead of pushing the whole model to the edge, split learning [5, 6] is proposed as  
 86 an alternative collaborative training approach, where a neural network model is partitioned between  
 87 the cloud and edge. As shown in Figure 1b, with the training data, an edge device trains the network  
 88 up to the partition layer and sends the intermediate features of the partition layer to the cloud. Upon  
 89 receiving the features, the cloud takes over training the remaining layers to complete the forward  
 90 pass. In the backward pass, with the gradients of the partition layer backpropagated from the cloud,  
 91 the edge device updates local model parameters. For consistency of local models among devices,  
 92 the edge device then synchronizes the updated model with the next device scheduled to participate  
 93 in training. The sequential training manner results in severe under-utilization of resources on edge  
 94 devices. Only one edge device is active in training with the cloud at any specific point in time. Split  
 95 learning preserves privacy in a way that only the intermediate features are sent out to the cloud while  
 96 the training data are still left to the edge devices. Unfortunately, transmitting features still has the risk  
 97 to leak sensitive information of the input data. Recent work [7] shows the possibility to accurately  
 98 recover the input image from the intermediate features even without access to the edge device.

## 99 2.2 Local Differential Privacy

100 Differential privacy is a statistical definition of privacy that is used to publish aggregate information  
 101 about the entire population while constraining the privacy leakage of each individual. As a kind of  
 102 differential privacy, local differential privacy (LDP) works without assuming a trusted data collector.  
 103 The data owners directly add noise to their data before sharing them with the untrusted data collector,  
 104 which provides a much stronger privacy guarantee [13, 14]. We provide a formal definition of local  
 105 differential privacy below. As the privacy budget  $\epsilon$  measures the extent of privacy leakage of the  
 106 random mechanism  $\pi$ . A lower value of  $\epsilon$  means more privacy.

107 **Definition 2.1** ( $\epsilon$ -LDP [13]). *An random mechanism  $\pi : \mathcal{X} \rightarrow \mathcal{Y}$  satisfies  $\epsilon$ -local differential privacy,*  
 108 *where  $\epsilon \geq 0$ , if and only if for any inputs  $x, x' \in \mathcal{X}$  and  $y \in \mathcal{Y}$ , we have  $Pr[\pi(x) = y] \leq$*   
 109  *$e^\epsilon Pr[\pi(x') = y]$ .*

110 Randomized response (RR) [10] is a typical mechanism to implement LDP. It is initially proposed  
 111 as a survey technique to collect answers to sensitive binary questions. The respondent uses a  
 112 randomization method like a coin flip to randomize the answer "yes" or "no". She answers the  
 113 question truthfully if the coin comes up heads, otherwise returns a false answer. Assume that we  
 114 use a biased coin and the probability to provide a truthful answer (coin flip result is head) is  $p$ . The  
 115 mechanism satisfies  $\epsilon$ -differential privacy with  $p = \frac{e^\epsilon}{1+e^\epsilon}$ .

## 116 2.3 Attacks against ML Models

117 We concentrate on two types of attacks against ML models which are closely related to our work and  
 118 suitable to be used to assess the privacy leakage.

119 **Feature Inversion Attack** As a specific case of model inversion attack [15], feature inversion attack  
 120 [7] is recently devised for the edge-cloud collaborative learning system. The neural network  $f_\theta$  is  
 121 split into two parts:  $f_{\theta_1}$  and  $f_{\theta_2}$  between the edge and the cloud. The adversary aims to recover the  
 122 input data sample  $x_0$  from the intermediate layer features  $f_{\theta_1}(x_0)$  in both white-box and block-box

123 settings. In the white-box setting, the adversary is assumed to compromise with participant and  
124 knows the structure and parameters of model  $f_{\theta_1}(x_0)$ . The adversary performs gradient descent  
125 technique on  $f_{\theta_1}$  to find a generated sample  $x$  whose intermediate  $f_{\theta_1}(x)$  is the most similar to  $f_{\theta_1}(x_0)$   
126 and at the same time following the same distribution as the input data. Recovering inputs is more  
127 challenging in the black-box setting, where the information of  $f_{\theta_1}$  is totally unknown to the adversary.  
128 Suppose the attacker can feed arbitrary inputs to  $f_{\theta_1}$  and receive the corresponding outputs, the  
129 attacker accordingly trains a network to approximate the inversion function of  $f_{\theta_1}$  and then converts  
130 the intermediate output into the input sample with the trained network.

131 **Membership Inference Attack** Membership inference attack aims to find out whether a given  
132 sample is used to train a model or not, which is considered as a direct privacy breach. For example,  
133 knowing a patient record is used to train models for diagnosis (disease presence) reveals that the  
134 patient has the disease. Shokri et al. [16] introduces membership inference in the block-box ML  
135 setting where the model is not accessible by the attacker. The attacker firstly trains multiple "shadow  
136 models" to imitate the behavior of the target model and then trains the binary attack model with the  
137 labeled inputs and outputs of the shadow models. Taking a data sample's prediction output queried  
138 from the target model as input, the binary attack model infers whether the data sample is a member or  
139 non-member of the target model's training dataset.

### 140 3 Related Work

141 The idea of neural network partition in MistNet is inspired by a large body of previous work  
142 [17, 18, 19, 20, 21, 22]. These work mainly focus on optimizing performance (eg., latency and  
143 accuracy) and cost (eg., communication and computation overhead) without considering data privacy.  
144 We break related work on protecting data privacy into the following three categories, injecting noise  
145 on the intermediate features, censoring the intermediate features with private feature extractor, and  
146 leveraging secure computation techniques.

147 **Noise Injection.** Several efforts [23, 24] inject noise to reduce the mutual information between the  
148 input and the intermediate features. They assume that the attacker performs sensitive secondary  
149 inferences. It is unknown whether these noises could successfully defend other types of attacks.  
150 Differential privacy provides strict privacy guarantees in the worst-case scenario without knowing the  
151 types of privacy attacks. There exist different mechanisms to apply differential privacy mechanisms  
152 on the intermediate features, such as the Laplace mechanism [25], the Gaussian mechanism [26] adds  
153 Gaussian noise. MistNet extracts binarized features and perturb them with local differential privacy  
154 technique - randomized response.

155 **Private Feature Extractor.** To defend against attacks while not sacrificing much accuracy, proposals  
156 [27, 28] use adversarial training to find an appropriate feature extractor from two respects - the  
157 number of layers and the strategy to prune output channels of the partition layer. DPFE [24] and  
158 DeepObfuscator [29] train the feature extractor to hide information about sensitive attributes while  
159 keeping useful features for the target task.

160 **Secure Computation.** Secure computation techniques recently are used to deal with privacy-  
161 preserving machine learning. The first line is to use cryptographic protocols such as secure multi-party  
162 computation (MPC) [30, 31] and homomorphic encryption (HE) [32, 33, 34]. Two or more parties  
163 collaboratively train a neural network on encrypted data from clients without the need to decrypt  
164 them. However, the computational and communication cost is usually prohibitively high, which  
165 makes these cryptographic techniques too heavy to be deployed at resource-constrained edge devices.  
166 The second line leverages trusted execution environments (TEE) [35, 36, 37] such as Intel SGX [38]  
167 and ARM TrustZone [39] to protect the training data. Training data is used within an isolated secure  
168 environment which is invisible to unauthorized adversaries. Nevertheless, the potential drawbacks  
169 are the limited scalability of TEEs and the vulnerability to side-channel attacks.

### 170 4 Design

171 In this section, we present MistNet, a framework for privacy-preserving collaborative training between  
172 the cloud and the edge. We first provide an overview of the framework and then the role of every  
173 component of our architecture. Lastly, we show how the intermediate features are perturbed with  
174 local differential privacy to provide a strong privacy guarantee.

175 **4.1 Workflow of MistNet**

176 Figure 2 illustrates the collaborative training process between the edge and the cloud. To protect  
 177 sensitive training data from being abused and support training complex models, we separate a model  
 178 into two parts between the edge and cloud. The lightweight first few layers of the neural network  
 179 are placed at the edge devices as the feature extractor. The rest layers with heavy computation are  
 180 offloaded to the powerful cloud servers as the cloud classifier. We detail the  
 181 workflow of the local feature extractor and cloud classifier below.

182 **Local Feature Extractor.** As explained in §2, maintaining consistency of feature extractors among  
 183 edge devices hinders parallel training, which lowers the efficiency of training among a large scale of  
 184 edge devices. Taking the factor into account, MistNet uses a fixed feature extractor thus eliminate  
 185 the need to perform synchronization among edge devices. To obtain meaningful features, the local  
 186 feature extractor is transferred from pre-trained models that work on a similar application domain via  
 187 transfer learning. Yosinski et al. [8] quantify the transferability of features from different layers in  
 188 deep neural networks. Features from the early layers are more general than that from later layers,  
 189 which show more flexibility to adapt to a wide range of related datasets and tasks. Meanwhile, to  
 190 keep the feature extractor simple and lightweight, the partition point in MistNet is usually set at a  
 191 very early layer in the model. The extracted features thus show high generalization capability and  
 192 are ubiquitous to various tasks. Moreover, the feature extractor is less sensitive to the changes of  
 193 input data, which provides the possibility to apply a fixed pre-trained feature extractor during training.  
 194 With the fixed pre-trained feature extractor, edge devices transform the input training data into feature  
 195 representations in parallel and send them to the cloud for the rest of the training. Thereafter, edge  
 196 devices do not need to repeatedly send the feature representations for the same batch of training  
 197 samples, nor receive backward-propagated feedbacks from the cloud.

198 As indicated in previous work [7], the value of intermediate features has the potential risk to reveal  
 199 sensitive information about the input data. The volume of feature representations of the early layer  
 200 in some models can be even larger than the raw input data, which incurs high communication costs.  
 201 The intermediate representations should be transmitted in a secure and communication-efficient way.  
 202 In MistNet, edge devices binarize each activation value with 1 bit and perturb the binarized feature  
 203 representations conforming to local differential privacy before sending them to the cloud. We further  
 204 explain the perturbing mechanism in §4.2.

205 **Cloud Classifier.** Upon receiving the perturbed feature representations from the edge devices, the  
 206 cloud iteratively trains the rest layers of the network with stochastic gradient descent (SGD) algorithm  
 207 to minimize the loss for a specific task. To reduce the communication cost, intermediate features  
 208 for input samples without any data augmentation effect are transmitted to the cloud. We further  
 209 apply random cropping on the intermediate features of each input sample as the data augmentation  
 210 technique to alleviate overfitting. We have the intermediate features of each input sample reshuffled  
 211 at each epoch during training. As shown in [8, 40], fine-tuning model on new related tasks is faster to  
 212 converge to near optimum than training from scratch. We thus initialize the parameter weights for the  
 213 cloud classifier with the transferred weights from the pre-trained model. During the backward pass,  
 214 the cloud classifier does not need to propagate the loss back to the edge devices and only parameters  
 215 of the cloud classifier are updated.

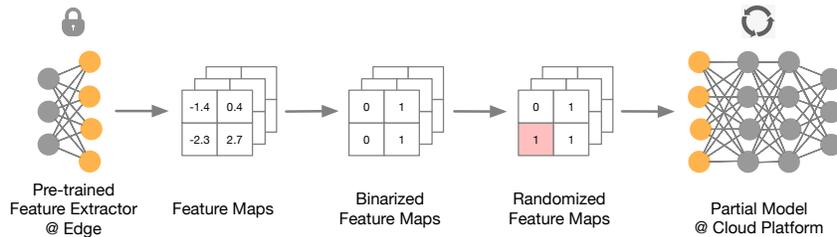


Figure 2: The overview of MistNet architecture. MistNet partitions the model between the edge and the cloud platform. The edge uses a pre-trained feature extractor to transform the local input data into a set of feature maps. Each edge device quantizes each activation into 1-bit and sends the randomized, differentially private version of the binarized feature values to the cloud platform. The partial model at the cloud platform is trained with the perturbed features collected from the edge devices.

## 216 4.2 Differentially Private Feature Representations

217 In this section, we detail how to protect the sensitive information about the training inputs from  
218 being revealed from the feature representations. Without assuming a trusted data collector, we apply  
219 local differential privacy to the intermediate feature representations from edge devices. The LDP  
220 mechanism we apply in MistNet is based on the randomized response method [10], which is widely  
221 deployed in practical systems [11, 12].

222 We denote the local feature extractor at the edge as  $f : \mathbb{R}^d \rightarrow \mathbb{R}^r$ , which transforms local data  
223  $x \in \mathcal{X}$  into feature representations  $A$ , with  $r = \dim(A)$ . With unary encoding [11], we encode each  
224 real-value feature  $A_i$  of the feature representations  $A$  into a bit  $B_i$ , whose value is either 1 or 0. The  
225 binarization function we use is as the following:

$$B_i = \begin{cases} 1 & \text{if } A_i > 0; \\ 0 & \text{if } A_i \leq 0. \end{cases} \quad (1)$$

226 We concatenate these bits as a binary string  $B = (B_1, B_2, B_3, \dots, B_r)$ . Then we apply randomized  
227 response defined in Eq. (2) to perturb each bit  $B_i$  in  $B$  independently and submit the noisy version  $\tilde{B}$   
228 to the cloud. Each bit is preserved as its true value with probability  $p$  or responded with the other  
229 value with probability  $q$ . We have  $p = 1 - q$  in this setting. The privacy budget  $\epsilon$  is calculated with  
230 Theorem 4.1.

$$P(\tilde{B}_i = 1) = \begin{cases} p & \text{if } B_i = 1; \\ q & \text{if } B_i = 0. \end{cases} \quad (2)$$

231 **Theorem 4.1.** *Local feature extractor with randomized response defined in Eq. (2) satisfies  $\epsilon$ -local*  
232 *differential privacy given that  $p \geq q$ , where  $\epsilon = r \cdot \ln \frac{p}{q}$ .*

233 See appendix A for the proof.

## 234 5 Evaluation

235 We evaluate the performance of MistNet on PyTorch and seek to answer the following questions: (1)  
236 How does MistNet perform with different privacy budgets for popular neural network models (§5.2)?  
237 (2) How does the partition layer selection affect the performance (§5.3)? (3) Whether MistNet is  
238 effective to defend model inversion and membership inference attack (§5.4)?

### 239 5.1 Experimental Setup

240 **Datasets and Models.** We evaluate MistNet for image classification on CIFAR-10 [41] and SVHN  
241 [42] dataset. CIFAR-10 has 10 classes and contains 60,000  $32 \times 32$  color pixel images with 3 RGB  
242 channels (50,000 training images and 10,000 testing images). SVHN is a MNIST-like dataset of  
243  $32 \times 32$  images, including 73,257 training digits and 26,032 testing digits. To pre-train neural network  
244 models, we use ImageNet  $32 \times 32$  images extracted from CINIC dataset [43], which downsamples part  
245 of the original ImageNet images from  $224 \times 224$  to  $32 \times 32$  resolution with the Box algorithm from the  
246 Pillow Python library<sup>2</sup>. These Imagenet  $32 \times 32$  images have the same 10 classes as CIFAR-10 (the  
247 number of images for train/validation/test is 70,000/70,000/70,000 respectively) but do not include  
248 any image in neither CIFAR-10 dataset nor SVHN dataset. The models we choose are ResNet-18  
249 [44] and VGG-16 [45].

250 **Compared Schemes.** We compare MistNet with the conventional training paradigm, centralized  
251 learning (CL), which collects data from users in a central site to train the model.

252 **Evaluation Metrics.** The performance of MistNet is evaluated from accuracy and privacy. Specifi-  
253 cally, we detail how membership inference and model inversion attacks are used to assess privacy  
254 risks. For membership inference attack, we adopt two metrics, precision (the fraction of records  
255 inferred as members actually are members of the training dataset) and recall (the fraction of records  
256 which are correctly inferred as training samples over all training samples) as privacy metrics, which

<sup>2</sup><https://python-pillow.org>

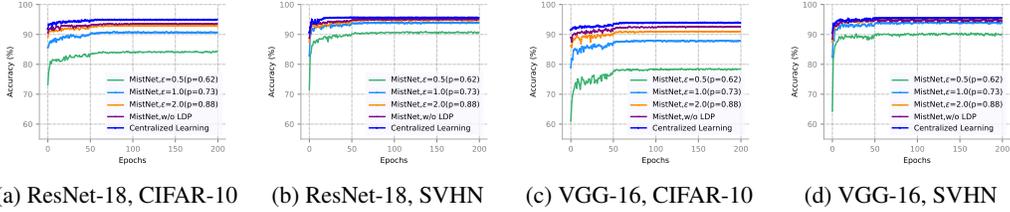


Figure 3: Comparison among MistNet with varying privacy budgets, MistNet without local differential privacy and Centralized Learning for ResNet-18 and VGG-16 on CIFAR-10 and SVHN. ResNet-18 is partitioned at Block 1 and VGG-16 is partitioned at Conv 2.

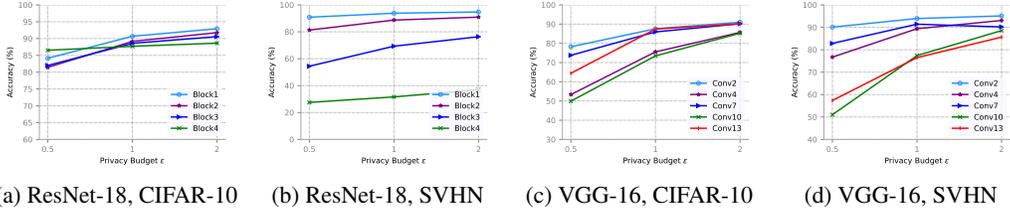


Figure 4: Comparison among MistNet with different partition layers for ResNet-18 and VGG-16 on CIFAR-10 and SVHN.

257 are consistent with previous work [16, 46, 47]. To quantify the quality of images recovered by model  
 258 inversion attack, we use two commonly used image quality metrics, peak signal-to-noise ratio (PSNR)  
 259 as well as the structural similarity index measure (SSIM). PSNR quantifies the pixel-level reconstruction  
 260 quality of the image, which can be expressed as the ratio of the maximum possible value of a  
 261 signal to the cumulative squared error between the reconstructed image and the original image. SSIM  
 262 measures the reconstructed image quality by taking into account the structural information perceived  
 263 by the human vision system including structure, luminance, and contrast.

## 264 5.2 Impact of Privacy Budget $\epsilon$

265 In this section, we vary privacy budgets  $\epsilon$  per feature to investigate its impact on model accuracy. As  
 266 indicated in Sec. 2.2, the value of  $\epsilon$  is proportional to the probability  $p$  of reporting the true value  
 267 of each feature in the intermediate features. The smallest value of  $\epsilon$  can be 0, which is equivalent  
 268 to absolute privacy. The lowest value of  $\epsilon$  we show in our experiments is 0.5, whose corresponding  
 269 probability for each feature to report the true value is 62%. We evaluate MistNet with a range of  
 270 privacy budgets  $\epsilon \in \{0.5, 1, 2\}$  and use MistNet without LDP (i.e., do not randomize the intermediate  
 271 features and is equivalent to  $\epsilon = +\infty$ ) and centralized learning as two baselines. We partition  
 272 ResNet-18 at Block 1 and VGG-16 at Conv 2 (the partition strategy is explained in Sec. 5.3), which  
 273 both are early layers. In Figure 3, we observe that strong privacy is provided at the sacrifice of the  
 274 utility. With a smaller value of  $\epsilon$  (stricter privacy guarantee), the accuracy decreases and MistNet  
 275 converges slower. Particularly, the performance of MistNet without LDP for both datasets achieves  
 276 similar performance as centralized learning, which demonstrates the effectiveness of the pre-trained  
 277 feature extractor.

## 278 5.3 Impact of Partition Layers

279 We explore the robustness of MistNet to different partition layers. More specifically, we partition  
 280 the ResNet-18 model into 4 fused layer blocks (Block 1-4), with each containing 4 convolution and  
 281 batch normalization layers. The VGG-16 model has 13 convolutional layers and is partitioned after  
 282 convolutional layer 2, 4, 7, 10, and 13. As shown in Figure 4, MistNet is robust to most partition  
 283 layers and achieves acceptable utility. In most cases, partitioning at an earlier layer (i.e., Block 1 for  
 284 ResNet-18 and Conv 2 for VGG-16) achieves better performance. The possible reason is twofold.  
 285 First, the transferability of an earlier layer is better and partitioning at an earlier layer leaves more  
 286 space to fine-tune the rest model to adapt to the perturbed features. Second, features from an earlier  
 287 layer generally have more dimensions which contain redundant information. We also notice that

Table 1: Comparison of the quality of reconstructed images from features generated by different schemes with model inversion attack (ResNet-18 and CIFAR-10).

	SSIM					PSNR				
	$\epsilon = 0.5$	$\epsilon = 1$	$\epsilon = 2$	w/o LDP	CL	$\epsilon = 0.5$	$\epsilon = 1$	$\epsilon = 2$	w/o LDP	CL
Block 1	0.354	0.576	0.728	0.775	0.918	12.941	14.126	15.531	16.314	21.028
Block 2	0.211	0.306	0.453	0.515	0.690	12.466	12.918	13.621	14.099	15.724
Block 3	0.165	0.170	0.205	0.206	0.303	12.311	12.422	12.538	12.616	13.010
Block 4	0.155	0.154	0.149	0.169	0.164	12.274	12.296	12.299	12.314	12.366

Table 2: Membership inference attacks on different schemes (ResNet-18 and CIFAR-10).

	Precision					Recall				
	$\epsilon = 0.5$	$\epsilon = 1$	$\epsilon = 2$	w/o LDP	CL	$\epsilon = 0.5$	$\epsilon = 1$	$\epsilon = 2$	w/o LDP	CL
Block 1	0.5194	0.5488	0.5778	0.5973	0.5952	0.5267	0.6021	0.6962	0.7201	0.7551
Block 2	0.4988	0.5275	0.5673	0.5847	—	0.4860	0.5539	0.6595	0.7104	—
Block 3	0.5033	0.5155	0.5427	0.5788	—	0.4924	0.5371	0.5959	0.7101	—
Block 4	0.4997	0.5020	0.4991	0.5070	—	0.4975	0.5038	0.5030	0.5136	—

288 model partitioned at the last convolutional layer achieves reasonable performance for CIFAR-10 while  
 289 performing poorly for SVHN. As described in Sec. 5.1, the chosen ImageNet images for pre-training  
 290 models have the same classes with CIFAR-10. CIFAR-10 dataset is more similar than SVHN dataset.  
 291 Features from later layers still transfer well for CIFAR-10.

## 292 5.4 Effect of Attack Mitigation

293 **Feature Inversion Attack.** We consider feature inversion attacks under the white-box setting  
 294 where the attacker has access to the feature extractor at the edge, since white-box attacks are more  
 295 challenging to defend than block-box attacks. In Table 1, we show the quality of images recovered  
 296 from features generated by different schemes. Images are recovered from raw features for centralized  
 297 learning. As indicated in [7], a recovered image with SSIM value below 0.3 is considered to be  
 298 unrecognizable. We observe that, with ResNet-18 model partitioned at different layers, the value  
 299 of SSIM for MistNet with various privacy budgets is consistently significantly lower than that for  
 300 centralized learning, which demonstrates the effectiveness of applying LDP on the partition layer in  
 301 protecting images from being recovered from the intermediate features. Partitioning ResNet-18 with  
 302  $\epsilon = 0.5$  at Block 2 achieves a good trade-off between accuracy and privacy.

303 **Membership Inference Attack.** We perform membership inference attacks on the model trained  
 304 with different schemes. As membership inference is a binary classification, the precision and recall  
 305 value is always between 0.5 and 1. The value of 0.5 is equivalent to random guessing, which indicates  
 306 that there is no privacy leakage. Table 2 shows that the precision and recall decreases with the value  
 307 of  $\epsilon$  and is reduced to around 0.5 with  $\epsilon = 0.5$ , while MistNet without LDP and centralized learning  
 308 still remain a high precision and recall. An interesting observation is that ResNet-18 partitioned at  
 309 Block 4 mitigates membership inference attacks even with a large  $\epsilon$ . This is because the trainable  
 310 cloud classifier only includes a linear layer, which is not sufficient to fit the training records.

## 311 6 Conclusion

312 In this paper, we presented MistNet as a privacy-preserving collaborative training system for resource-  
 313 constrained edge devices. Our method uses a pre-trained feature extractor to eliminate the need  
 314 to synchronize local weights across edge devices and enhances privacy by applying LDP to the  
 315 intermediate features. We extensively evaluate MistNet with various settings on a variety of models,  
 316 datasets, and attacks. The results show that MistNet partitioned at most layers with privacy budget  
 317  $\epsilon = 0.5$  achieves acceptable utility while effectively reducing privacy leakage.

318 **Broader Impact**

319 Distributed machine learning is a widely used computing paradigm to learn from gigantic amounts of  
320 data generated by edge devices. Our work can be used to enhance data privacy in distributed machine  
321 learning systems, which follows the recent trend to comply with the EU General Data Protection  
322 Regulation (GDPR) law. It is possibly adopted by application developers and service providers as a  
323 tool to collect personal data from users. The possible negative aspect is it will be harder for the service  
324 provider to regulate the submitted training data from users, and detecting data with a significant  
325 detrimental impact on the prediction performance will be challenging. It is not likely to directly raise  
326 any ethical issues.

327 **References**

- 328 [1] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual  
329 networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- 330 [2] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep  
331 recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and*  
332 *signal processing*, pages 6645–6649. IEEE, 2013.
- 333 [3] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. A neural probabilistic  
334 language model. *Journal of machine learning research*, 3(Feb):1137–1155, 2003.
- 335 [4] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh,  
336 and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv*  
337 *preprint arXiv:1610.05492*, 2016.
- 338 [5] Otkrist Gupta and Ramesh Raskar. Distributed learning of deep neural network over multiple  
339 agents. *Journal of Network and Computer Applications*, 116:1–8, 2018.
- 340 [6] Praneeth Vepakomma, Otkrist Gupta, Tristan Swedish, and Ramesh Raskar. Split learn-  
341 ing for health: Distributed deep learning without sharing raw patient data. *arXiv preprint*  
342 *arXiv:1812.00564*, 2018.
- 343 [7] Zecheng He, Tianwei Zhang, and Ruby B. Lee. Model inversion attacks against collaborative  
344 inference. In *Proceedings of the 35th Annual Computer Security Applications Conference*,  
345 *ACSAC '19*, page 148–162, New York, NY, USA, 2019. Association for Computing Machinery.
- 346 [8] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in  
347 deep neural networks? In *Advances in neural information processing systems*, pages 3320–3328,  
348 2014.
- 349 [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-  
350 scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern*  
351 *recognition*, pages 248–255. Ieee, 2009.
- 352 [10] Stanley L Warner. Randomized response: A survey technique for eliminating evasive answer  
353 bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.
- 354 [11] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable  
355 privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on*  
356 *computer and communications security*, pages 1054–1067, 2014.
- 357 [12] Learning with privacy at scale. [https://machinelearning.apple.com/2017/12/06/  
358 learning-with-privacy-at-scale.html](https://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html). Accessed: December 2017.
- 359 [13] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam  
360 Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- 361 [14] Björn Běbensee. Local differential privacy: a tutorial. *arXiv preprint arXiv:1907.11908*, 2019.
- 362 [15] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit  
363 confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC*  
364 *Conference on Computer and Communications Security*, pages 1322–1333, 2015.
- 365 [16] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference  
366 attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy*  
367 *(SP)*, pages 3–18. IEEE, 2017.
- 368 [17] Yiping Kang, Johann Hauswald, Cao Gao, Austin Rovinski, Trevor Mudge, Jason Mars, and  
369 Lingjia Tang. Neurosurgeon: Collaborative intelligence between the cloud and mobile edge.  
370 *ACM SIGARCH Computer Architecture News*, 45(1):615–629, 2017.
- 371 [18] Jiachen Mao, Xiang Chen, Kent W Nixon, Christopher Krieger, and Yiran Chen. Modnn: Local  
372 distributed mobile computing system for deep neural network. In *Design, Automation & Test in*  
373 *Europe Conference & Exhibition (DATE), 2017*, pages 1396–1401. IEEE, 2017.

- 374 [19] Hongshan Li, Chenghao Hu, Jingyan Jiang, Zhi Wang, Yonggang Wen, and Wenwu Zhu. Jalad:  
375 Joint accuracy-and latency-aware deep structure decoupling for edge-cloud execution. In *2018*  
376 *IEEE 24th International Conference on Parallel and Distributed Systems (ICPADS)*, pages  
377 671–678. IEEE, 2018.
- 378 [20] Jong Hwan Ko, Taesik Na, Mohammad Faisal Amir, and Saibal Mukhopadhyay. Edge-host  
379 partitioning of deep neural networks with feature space encoding for resource-constrained  
380 internet-of-things platforms. In *2018 15th IEEE International Conference on Advanced Video*  
381 *and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2018.
- 382 [21] Hyuk-Jin Jeong, Hyeon-Jae Lee, Chang Hyun Shin, and Soo-Mook Moon. Ionn: Incremental  
383 offloading of neural network computations from mobile devices to edge servers. In *Proceedings*  
384 *of the ACM Symposium on Cloud Computing*, pages 401–411, 2018.
- 385 [22] Xiaorui Wu, Hong Xu, Bo Li, and Yongqiang Xiong. Stanza: Layer separation for distributed  
386 training in deep learning. *arXiv preprint arXiv:1812.10624*, 2018.
- 387 [23] Fatemehsadat Mireshghallah, Mohammadkazem Taram, Prakash Ramrakhani, Ali Jalali, Dean  
388 Tullsen, and Hadi Esmaeilzadeh. Shredder: Learning noise distributions to protect inference  
389 privacy. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support*  
390 *for Programming Languages and Operating Systems*, pages 3–18, 2020.
- 391 [24] Seyed Ali Osia, Ali Taheri, Ali Shahin Shamsabadi, Kleomenis Katevas, Hamed Haddadi, and  
392 Hamid R Rabiee. Deep private-feature extraction. *IEEE Transactions on Knowledge and Data*  
393 *Engineering*, 32(1):54–66, 2018.
- 394 [25] Ji Wang, Jianguo Zhang, Weidong Bao, Xiaomin Zhu, Bokai Cao, and Philip S Yu. Not just  
395 privacy: Improving performance of private deep learning in mobile cloud. In *Proceedings of*  
396 *the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*,  
397 pages 2407–2416, 2018.
- 398 [26] Yunlong Mao, Shanhe Yi, Qun Li, Jinghao Feng, Fengyuan Xu, and Sheng Zhong. Learn-  
399 ing from differentially private neural activations with edge computing. In *2018 IEEE/ACM*  
400 *Symposium on Edge Computing (SEC)*, pages 90–102. IEEE, 2018.
- 401 [27] Meng Li, Liangzhen Lai, Naveen Suda, Vikas Chandra, and David Z Pan. Privynet: A  
402 flexible framework for privacy-preserving deep neural network training. *arXiv preprint*  
403 *arXiv:1709.06161*, 2017.
- 404 [28] Shuang Zhang, Liyao Xiang, Congcong Li, Yixuan Wang, Zeyu Liu, Quanshi Zhang, and  
405 Bo Li. Preventing information leakage with neural architecture search. *arXiv preprint*  
406 *arXiv:1912.08421*, 2019.
- 407 [29] Ang Li, Jiayi Guo, Huanrui Yang, and Yiran Chen. Deepobfuscator: Adversarial training  
408 framework for privacy-preserving image classification. *arXiv preprint arXiv:1909.04126*, 2019.
- 409 [30] Andrew Chi-Chih Yao. How to generate and exchange secrets. In *27th Annual Symposium on*  
410 *Foundations of Computer Science (sfcs 1986)*, pages 162–167. IEEE, 1986.
- 411 [31] Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving  
412 machine learning. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 19–38. IEEE,  
413 2017.
- 414 [32] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the*  
415 *forty-first annual ACM symposium on Theory of computing*, pages 169–178, 2009.
- 416 [33] Ran Gilad-Bachrach, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John  
417 Wernsing. Cryptonets: Applying neural networks to encrypted data with high throughput and  
418 accuracy. In *International Conference on Machine Learning*, pages 201–210, 2016.
- 419 [34] Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. {GAZELLE}: A low  
420 latency framework for secure neural network inference. In *27th {USENIX} Security Symposium*  
421 *({USENIX} Security 18)*, pages 1651–1669, 2018.

- 422 [35] Zhongshu Gu, Hani Jamjoom, Dong Su, Heqing Huang, Jialong Zhang, Tengfei Ma, Dimitrios  
423 Pendarakis, and Ian Molloy. Reaching data confidentiality and model accountability on the  
424 caltrain. In *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and*  
425 *Networks (DSN)*, pages 336–348. IEEE, 2019.
- 426 [36] Florian Tramer and Dan Boneh. Slalom: Fast, verifiable and private execution of neural  
427 networks in trusted hardware. *arXiv preprint arXiv:1806.03287*, 2018.
- 428 [37] Tyler Hunt, Congzheng Song, Reza Shokri, Vitaly Shmatikov, and Emmett Witchel. Chiron:  
429 Privacy-preserving machine learning as a service. *arXiv preprint arXiv:1803.05961*, 2018.
- 430 [38] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V Rozas, Hisham Shafi, Vedvyas  
431 Shanbhogue, and Uday R Savagaonkar. Innovative instructions and software model for isolated  
432 execution. *HASP@ ISCA*, 10(1), 2013.
- 433 [39] Tiago Alves. Trustzone: Integrated hardware and software security. *White paper*, 2004.
- 434 [40] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time  
435 object detection with region proposal networks. In *Advances in neural information processing*  
436 *systems*, pages 91–99, 2015.
- 437 [41] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.  
438 2009.
- 439 [42] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng.  
440 Reading digits in natural images with unsupervised feature learning. 2011.
- 441 [43] Luke N Darlow, Elliot J Crowley, Antreas Antoniou, and Amos J Storkey. Cinic-10 is not  
442 imagenet or cifar-10. *arXiv preprint arXiv:1810.03505*, 2018.
- 443 [44] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image  
444 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,  
445 pages 770–778, 2016.
- 446 [45] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale  
447 image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- 448 [46] Bargav Jayaraman and David Evans. Evaluating differentially private machine learning in  
449 practice. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, pages 1895–1912,  
450 2019.
- 451 [47] Daniel Bernau, Philip-William Grassal, Jonas Robl, and Florian Kerschbaum. Assessing differ-  
452 entially private deep learning with membership inference. *arXiv preprint arXiv:1912.11328*,  
453 2019.

454 **A Proof of Theorem 4.1**

455 *Proof.* Given noisy binary vector  $\tilde{B}$ , for any  $x, x' \in \mathcal{X}$  we have:

$$\begin{aligned}
 \frac{P[\tilde{B}|x]}{P[\tilde{B}|x']} &= \frac{P[\tilde{B}|B, x]P[B|x]}{P[\tilde{B}|B', x']P[B'|x']} = \frac{P[\tilde{B}|B]}{P[\tilde{B}|B']} \\
 &= \frac{\prod_{i=1}^r P[\tilde{B}_i|B_i]}{\prod_{i=1}^{r'} P[\tilde{B}_i|B'_i]} = \prod_{i=1}^r \frac{P[\tilde{B}_i|B_i]}{P[\tilde{B}_i|B_i]} \\
 &\leq \prod_i^r \max \left\{ \frac{P[\tilde{B}_i = 1|B_i = 1]}{P[\tilde{B}_i = 1|B'_i = 0]}, \frac{P[\tilde{B}_i = 0|B_i = 0]}{P[\tilde{B}_i = 0|B'_i = 1]} \right\} \\
 &= \prod_{i=1}^r \left\{ \frac{p}{q} \right\} \\
 &= \left\{ \frac{p}{q} \right\}^r
 \end{aligned}$$

456 where the second equality follows from the fact that the mapping from input  $x$  to binary vector  $B$  is  
 457 deterministic, while the fifth inequality is based on the assumption that  $p \geq q$ , under which we need  
 458 not consider another two situations, where

$$\frac{P[\tilde{B}_i = 1|B_i = 0]}{P[\tilde{B}_i = 1|B'_i = 1]} = \frac{q}{p} \quad \text{or} \quad \frac{P[\tilde{B}_i = 0|B_i = 1]}{P[\tilde{B}_i = 0|B'_i = 0]} = \frac{1-p}{1-q}$$

459 .