

# Making Models Shallow Again: Jointly Learning to Reduce Non-Linearity and Depth for Latency-Efficient Private Inference

Souvik Kundu  
Intel Labs  
San Diego, USA

souvik.kundu@intel.com

Yuke Zhang, Dake Chen, Peter A. Beerel  
University of Southern California  
Los Angeles, USA

{yukezhan, dakechen, pabeerel}@usc.edu

## Abstract

Large number of ReLU and MAC operations of Deep neural networks make them ill-suited for latency and compute-efficient private inference. In this paper, we present a model optimization method that allows a model to learn to be shallow. In particular, we leverage the ReLU sensitivity of a convolutional block to remove a ReLU layer and merge its succeeding and preceding convolution layers to a shallow block. Unlike existing ReLU reduction methods, our joint reduction method can yield models with improved reduction of both ReLUs and linear operations by up to  $1.73\times$  and  $1.47\times$ , respectively, evaluated with ResNet18 on CIFAR-100 without any significant accuracy-drop.

## 1. Introduction

Machine learning as a service (MLaaS) helps many users leverage the benefits of artificial intelligence (AI) augmented applications on their private data. However, due to the growing concerns associated with the model IP protection [14], the service providers often prefer to retain the model at its end rather than sharing the black box model IP with the user. Users often do not prefer sharing their personal data due to various data privacy issues. To tackle these concerns, various private inference (PI) methods [7, 18, 20, 21] have been proposed that leverage techniques such as Homomorphic encryption (HE) [1] and secure multi-party computation (MPC) protocols to preserve the privacy of the client’s data as well as the model’s IP. Popular PI frameworks including Gazelle [9], DELPHI [18], CryptoNAS [5], and Cheetah [19] leverage these privacy preserving mechanisms. However, unlike traditional inference, the non-linear ReLU operation latency in PI can increase up to two orders of magnitude. For example, PI methods generally use Yao’s Garbled Circuits (GC) [22] that demand orders of magnitude higher latency and communication than that of linear multiply-accumulate (MAC)




Figure 1. Comparison of the proposed method with the baseline and the existing state-of-the-art (SoTA) [12]. Compared to existing SoTA, our method reduces both the linear and ReLU operations of the model at the cost of in-significant accuracy drop. We normalized each metric with respect to the maximum value of that metric. “Ours-AC” represents the model with AC output. We used ResNet18 on CIFAR-100 for evaluation.

operations.

This has inspired the unique problem of reducing the ReLU operations and thus, the latency overhead associated with PI. In particular, earlier works leveraged neural architecture search [2, 5] to reduce the ReLU layer count. Recent works [3, 12] proposed model linearization via systematic ReLU reduction that can reduce ReLU at the level of layer, channel, and pixel. However, all these works focus on reducing *only* the ReLU operation. Interestingly, the reduced ReLU models may often have comparable total non-linear and total linear MAC operation latency. Moreover, recent improvement of non-linearity operation latency in PI via oblivious transfer [7] has essentially demanded simultaneous reduction of ReLU as well as MAC operations.

**Our contributions.** Towards this goal, we present a model architecture optimization framework that simultaneously learns to reduce both the ReLUs and the MACs of a model. In particular, we systematically replace every convolutional block (example: basic-block layer for ResNet18) having a low ReLU budget [12], with a shallower block having no non-linearity module allowing us to reduce both MACs and ReLUs. We term this method as *gated branching* (GB) as we allow gradual learning of the branch with a shallower block via a gating condition, that changes over epochs (refer to Fig. 3(b)). To further reduce the depth at the

later layers, we present a *auxiliary knowledge distillation* (AKD) to an auxiliary classifier (AC) at a shallower depth of the DNN model. We conduct experiments with ResNet18 on CIFAR-10, CIFAR-100, and WRN-22-8 on CIFAR-100 datasets to evaluate the efficacy of the proposed framework. As Fig. 1 shows, our joint learning method can yield linear and ReLU operation reduction of up to  $1.47\times$  and  $19.23\times$  compared to the baseline models.

## 2. Preliminaries

**Private inference.** We assume a semi-honest client-server PI scenario where a client, holding private data, intends to use inference service from a server having a private model [18]. Here, each party tries to reveal their collaborator’s private data by inspecting the information they received while strictly following the protocol. To protect from data revelation, we assume the PI to happen on encrypted data [5, 17, 18] in an offline-online format [18], that transfers input data independent computations to the offline stage while performing the data-dependent operations online.

**ReLU reduction for efficient PI.** Existing works use model designing with reduced ReLU counts via methods including the search for ReLU-efficient models [2, 5, 17, 18] and manual ReLU-importance-driven non-linearity reduction [8]. More recently, [3] leveraged  $l_1$ -regularization to remove ReLU at the granularity of both pixel and channels to yield improved accuracy vs. non-linearity trade-off. Finally, a contemporary work [12] demonstrated an automated ReLU sensitivity evaluation method and for a given ReLU budget, leveraged that to propose a three-stage training framework, dubbed as SENet, towards linearizing a model via ReLU reduction. In particular, for a given ReLU budget, at stage 1, SENet automatically identifies per-layer ReLU count. During stage 2, it learns the binary mask associated to the ReLU locations in a non-linear layer of the partial ReLU model (PR) model. 1 value in a mask position assigns a ReLU unit for that position while the 0 value assigns an identity unit. Finally, during stage 3, it fine-tunes the PR model with the mask frozen. SENet yields the SoTA accuracy vs. non-linearity trade-off, and can meet a target ReLU budget without any costly hyper-parameter tuning that is necessary for the  $l_1$ -regularization method [3]. For these notable advantages, we adopt SENet’s training method into our optimization framework.

## 3. Motivational Analysis

For a given ReLU budget, we start by evaluating the ReLU sensitivity<sup>1</sup> of the non-linearity layers of a DNN. To perform this analysis, we used a ResNet18 model and

<sup>1</sup>ReLU sensitivity of a layer is measured as the ratio of ReLUs assigned in a layer (mask location value of 1) to the total number of possible ReLU units of that layer.

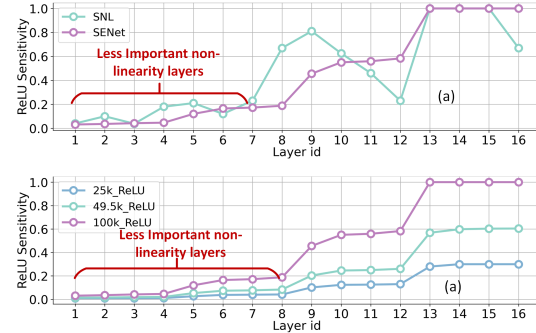


Figure 2. (a) Comparison of layer-wise ReLU sensitivity evaluated via SNL [3] and SENet [12] for a ReLU budget of 100k; (b) Comparison of layer-wise ReLU sensitivity for different ReLU budgets using SENet training method. We used the basic block ReLUs of a ResNet18 trained on CIFAR-100 for both the plots.

trained it on CIFAR-100 using two recent automated ReLU reduction frameworks, SNL [3] and SENet [12]. We keep a target ReLU budget of 100k for both. As demonstrated in Fig. 2(a), both the frameworks yield ReLU distribution sensitivity with the initial layers having less sensitivity than the later layers. As we understand, high ReLU sensitivity may be co-related to high non-linearity importance [12], the earlier non-linearity layers may play a less significant role in retaining the classification performance as compared to the later non-linearity layers.

We further conducted ReLU sensitivity analysis for three different ReLU budgets of 25k, 49.5k, and 100k with the ResNet18 on CIFAR-100. As demonstrated in Fig. 2(b), the earlier non-linearity layers consistently show less importance compared to the later ones, for all the three ReLU budgets. Inspired by these observations, we pose the following intriguing question.

*Is it possible to design a shallow DNN by removing the less important ReLU layers while retaining the classification performance?*

## 4. Methodology

We now describe our training framework that can jointly reduce the ReLU and the depth of a DNN. While we follow [12], to train a partial ReLU model via a three-stage optimization framework, we replace the stage 3 of [12] with our novel fine-tuning stage. In the first stage, we identify the per-layer ReLU count for a given ReLU budget by computing the normalized ReLU sensitivity [12]. For a layer  $l$ , the ReLU sensitivity  $\eta_{\alpha^l}$  can be computed as [12],

$$\eta_{\alpha^l} = (1 - \eta_{\theta^l}), \quad (1)$$

Where,  $\eta_{\theta^l}$  represents the pruning sensitivity [11, 13, 15, 16] of the preceding layer represented in terms of the fraction of non-zero weights for a given parameter density. Next, in stage 2, we compute a binary mask  $M^l \in \{0, 1\}^{h^l \times w^l \times c^l}$

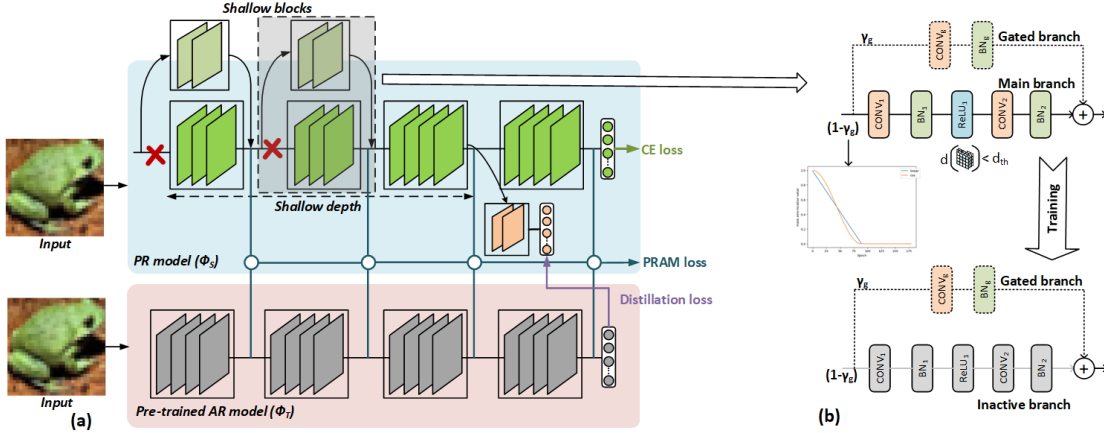


Figure 3. (a) Fine-tuning framework of the partial ReLU (PR) model that uses gated branching to reduce the depth of the earlier blocks, while learning a shallow auxiliary classifier to reduce the depth for later blocks. (b) The gradual training procedure to learn the gated branching path when the non-linearity budget of the main branch is below a certain threshold.

for each non-linearity layer  $l$ , where  $h^l$ ,  $w^l$ , and  $c^l$  represents the height, width, and the number of channels in the activation map. The 1's in  $M^l$  represents the locations with ReLU units while 0's represents identity units in the activation tensor. During the fine-tuning stage (stage 3), we present a training strategy for joint reduction of ReLU operation and depth of the PR model.

In our fine-tuning stage, we first introduce a threshold ReLU sensitivity value denoted by  $d_{th}$ , that represents the minimum ReLU sensitivity for a non-linearity layer to be retained. In particular, at the end of stage 3, for a convolutional layer block (example, basic block for ResNets), we only keep those ReLU layers that have a sensitivity  $d > d_{th}$ . Removal of a non-linear layer effectively connects the preceding and succeeding linear layers (example convolution) one after the other. It is well known that a function represented through a sequence of linear function layers like convolution, can be approximated with a single linear function layer [4]. We thus replace the convolution batch-normalization (CONV-BN) pairs located immediately before and after the dropped ReLU layer, with a single CONV-BN unit.

**Gated branching.** As during fine-tuning, we initialize the PR model weights with that of the best model evaluated during stage 2, modifying the model architecture with shallow layer blocks (as depicted in Fig. 3(a) at the earlier blocks of  $\Phi_S$ ) with untrained weights is not straight forward. In particular, direct replacement of some of the blocks with shallow blocks from the beginning of fine-tuning stage may significantly reduce the PR model's final performance due to sub-optimal training. We thus present a gated branching (GB) method that assigns a shallow branch for the layer blocks having ReLU sensitivity  $< d_{th}$  (see Fig. 3(b)). We assign a gating hyperparameter  $\gamma_g$  to the shallow branch with the main branch having the weight factor of  $(1 - \gamma_g)$ . Thus, for an input tensor  $\mathbf{X}$ , the functional representation of

the combined block is  $\gamma_g(f_{GB}(\mathbf{X})) + (1 - \gamma_g)(f_{MB}(\mathbf{X}))$ , with  $f_{GB}$  and  $f_{MB}$  representing the functions represented by the gated and the main branch, respectively. We start the  $\gamma_g$  from 0 and gradually increase it to 1 with the training epochs using either a cosine or linear increment function. This allows a gradual introduction of the shallow branches helping the PR model to transfer knowledge from the corresponding deeper blocks that finally become inactive at the end of the fine-tuning. In particular, we gradually increase  $\gamma_g$  to 1 till the end of the epoch when the first learning-rate (LR) decay occurs. We keep  $\gamma_g$  fixed to 1 during the remaining part of the fine-tuning allowing the fixed architecture with shallow branches to fine-tune. Note, we keep the first ReLU after the CONV stem, out of the GB method.

**Auxiliary knowledge distillation.** It is note-worthy that reducing the depth via GB, may be applicable for a DNN's initial layers only, due to the less importance of non-linearity in those layers. Later non-linearity layers, on the other hand, play a critical role in retaining the classification performance, and thus ReLU layer removal via the thresholding method may cost a huge accuracy drop. To reduce depth at the later layers, we present an auxiliary classifier (AC) branch located after a shallow depth of the DNN. In particular, inspired by the model architecture used in self-distillation [24], we place the AC before the final convolutional layer block (refer to Fig. 3(a)). This architecture allows the PR model ( $\Phi_S$ ) to learn and predict at two layer depths, one at the original classifier and the other at AC at a shallower depth. To further improve the accuracy of  $\Phi_S$ , we present an *auxiliary knowledge distillation* (AKD) that distills the knowledge from an all-ReLU (AR) baseline model ( $\Phi_T$ ) to the AC of the  $\Phi_S$ .

We thus, combine the two methods of GB and AKD, allowing additional ReLU reduction both at the initial and later layers during the fine-tuning stage, compared to that proposed in SENet [12]. Moreover, the systematic depth

Table 1. Comparison between the proposed method with gated branching (w/ GB) and SoTA reduction method [12] (w/o GB).  $\uparrow$  and  $\downarrow$  indicate the higher the better and the lower the better, respectively.

Model	Baseline Acc (%)	$d_{th}$	#ReLUs (k) $\downarrow$		Accuracy (%) $\uparrow$		Layer depth $\downarrow$		MAC saving $\uparrow$		ReLU ops reduction $\uparrow$	
			w/o GB	w/ GB	w/o GB	w/ GB	w/o GB	w/ GB	w/o GB	w/ GB	w/o GB	w/ GB
Dataset: CIFAR-10												
ResNet18	95.2	0.1	82	<b>76.8</b>	93.05	<b>93.05</b>	16	<b>13</b>	1 $\times$	<b>1.17<math>\times</math></b>	6.8 $\times$	<b>7.3<math>\times</math></b>
Dataset: CIFAR-100												
ResNet18	78.05	0.05	24.9	<b>21.1</b>	<b>70.59</b>	69.10	16	<b>12</b>	1 $\times$	<b>1.3<math>\times</math></b>	21.8 $\times$	<b>26.4<math>\times</math></b>
ResNet18	78.05	0.05	49.6	<b>47.4</b>	<b>75.28</b>	74.62	16	<b>14</b>	1 $\times$	<b>1.16<math>\times</math></b>	11.2 $\times$	<b>11.8<math>\times</math></b>
WRN22-8	80.82	0.1	240	<b>221</b>	<b>79.81</b>	79.62	18	<b>15</b>	1 $\times$	<b>1.15<math>\times</math></b>	5.8 $\times$	<b>6.3<math>\times</math></b>

reduction makes the model yielded by our method have significantly fewer linear operations compared to the SoTA alternatives. let  $\Psi_{pr}^m$  and  $\Psi_{ar}^m$  represent the  $m^{th}$  pair of vectorized post-ReLU activation maps of same layer for  $\Phi_{pr}$  and  $\Phi_{ar}$ , respectively. Our loss function for the fine-tuning phase is given by

$$\mathcal{L} = \underbrace{\lambda \mathcal{L}_{KL} \left( \sigma \left( \frac{z^{ar}}{\rho} \right), \sigma \left( \frac{z^{pr}}{\rho} \right) \right)}_{\text{KL-div. loss}} + \underbrace{(1 - \lambda) \mathcal{L}_{pr}(y, y^{pr})}_{\text{CE loss}} + \underbrace{\frac{\beta}{2} \sum_{m \in I} \left\| \frac{\Psi_{pr}^m}{\|\Psi_{pr}^m\|_2} - \frac{\Psi_{ar}^m}{\|\Psi_{ar}^m\|_2} \right\|_2}_{\text{PRAM loss}}. \quad (2)$$

where  $\sigma$  is the softmax function with  $\rho$  being KL-div. temperature.  $\lambda$  balances the importance between the CE and KL divergence loss components, and  $\beta$  is the weight for the post-ReLU AR-PR activation mismatch (PRAM) loss.

## 5. Experiments

### 5.1. Experimental Setup

To evaluate the efficacy of our method, we performed experiments with ResNet18 [6] and wide residual network 22-8 (WRN22-8) [23]. In particular, we evaluated ResNet18 on two popular datasets, CIFAR-10 and CIFAR-100 [10] and WRN22-8 on CIFAR-100. We used PyTorch API to define and train our models on an Nvidia RTX 2080 Ti GPU.

To perform stage 1 and 2, we followed the same hyperparameters as [12]. To perform joint ReLU and depth reduction during fine-tuning stage, we initialized a PR model with the weights that provided the best accuracy during the mask evaluation stage (stage 2). We then trained the model for 180 epochs with starting LR of 0.01 that decayed by a factor of 10 at epochs 90, 140, and 160. Unless otherwise stated, we used a linear increase of  $\gamma_g$  each epoch by  $\frac{i}{90}$  and kept it to 1 for  $i > 90$ , where  $i$  represents the epoch number. Similar to [12], we weighted the loss of each component by setting  $\lambda = 0.9$ ,  $\beta = 1000$ , and used KD temp.  $\rho = 4$ .

### 5.2. Model Performance Analysis

**Results with GB.** Table 1 demonstrates results of the PR models with GB. Note, for models with GB, similar to [12].

Table 2. Performance of our method with depth reduction both at the earlier layers via GB and later layers via AC output from an intermediate shallow layer. The AC is placed after the 3<sup>rd</sup> basic-block layer of a ResNet18. ‘ $\checkmark$ ’ and ‘ $\times$ ’ in AC output indicates whether output is taken at the AC or the final classifier.

Model	Baseline Acc (%)	With GB	AC output	Accuracy (%) $\uparrow$	ReLU ops reduction $\uparrow$	MAC saving $\uparrow$	Layer depth $\downarrow$
Dataset: CIFAR-10							
ResNet18	95.2	$\checkmark$	$\times$	95.17	3.94 $\times$	1.15 $\times$	14
		$\checkmark$	$\checkmark$	94.51	5.07 $\times$	1.46 $\times$	10
Dataset: CIFAR-100							
ResNet18	78.05	$\checkmark$	$\times$	78.24	5.8 $\times$	1.15 $\times$	14
		$\checkmark$	$\checkmark$	76.31	8.5 $\times$	1.46 $\times$	10

Due to the shallow layer blocks, the models with GB can provide MAC saving and reduce the layer depth by up to 1.3 $\times$  and 1.33 $\times$ , respectively, while yielding reduction of ReLU operation by up to 26.4 $\times$ . To train models with GB, apart from the CE loss of  $\Phi_S$ , we used the PRAM loss and a KL-div. between the final classifiers of  $\Phi_S$  and  $\Phi_T$ .

**Results with GB and AKD.** Table 2 demonstrates the results of the models whose depth are reduced via both GB and AKD with the final classification results taken from the AC of the PR model. In particular, at the cost of modest accuracy drop, the AC can provide an additional depth reduction of 1.4 $\times$  compared to the one using only GB. Further, classification via AC reduces both the ReLU and MAC operations by up to 1.46 $\times$  and 1.27 $\times$ , respectively, than models using only GB.

## 6. Conclusions

We introduced a training method to jointly reduce both the ReLU operations and layer depth of a DNN model suitable for latency-compute-efficient PI. In particular, we present a gated branching method to learn a shallow block in replacing the deep CONV layer blocks having non-linearity layers with low ReLU sensitivity. We further place an auxiliary classifier to allow classification at a shallower depth compared to that with original classifier. Empirical demonstration showed that the shallow-reduced-ReLU models can yield improved reduction of both ReLUs and layer depth by up to 1.73 $\times$  and 1.6 $\times$ , respectively, without any significant accuracy-drop, compared to that yielded via SoTA ReLU reduction methods.

## References

- [1] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. *SIAM Journal on computing*, 43(2):831–871, 2014. [1](#)
- [2] Minsu Cho, Zahra Ghodsi, Brandon Reagen, Siddharth Garg, and Chinmay Hegde. Sphynx: Relu-efficient network design for private inference. *arXiv preprint arXiv:2106.11755*, 2021. [1](#), [2](#)
- [3] Minsu Cho, Ameya Joshi, Brandon Reagen, Siddharth Garg, and Chinmay Hegde. Selective network linearization for efficient private inference. In *International Conference on Machine Learning*, pages 3947–3961. PMLR, 2022. [1](#), [2](#)
- [4] Yonggan Fu, Haichuan Yang, Jiayi Yuan, Meng Li, Cheng Wan, Raghuraman Krishnamoorthi, Vikas Chandra, and Yingyan Lin. Depthshrinker: a new compression paradigm towards boosting real-hardware efficiency of compact neural networks. In *International Conference on Machine Learning*, pages 6849–6862. PMLR, 2022. [3](#)
- [5] Zahra Ghodsi, Akshaj Kumar Veldanda, Brandon Reagen, and Siddharth Garg. Cryptonas: Private inference on a relu budget. *Advances in Neural Information Processing Systems*, 33:16961–16971, 2020. [1](#), [2](#)
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [4](#)
- [7] Zhicong Huang, Wen-jie Lu, Cheng Hong, and Jiansheng Ding. Cheetah: Lean and fast secure two-party deep neural network inference. *IACR Cryptol. ePrint Arch.*, 2022:207, 2022. [1](#)
- [8] Nandan Kumar Jha, Zahra Ghodsi, Siddharth Garg, and Brandon Reagen. Deepreduce: Relu reduction for fast private inference. In *International Conference on Machine Learning*, pages 4839–4849. PMLR, 2021. [2](#)
- [9] 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)*, pages 1651–1669, 2018. [1](#)
- [10] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [4](#)
- [11] Souvik Kundu, Yao Fu, Bill Ye, Peter A Beerel, and Massoud Pedram. Towards adversary aware non-iterative model pruning through dynamic network rewiring of dnns. *ACM Transactions on Embedded Computing Systems (TECS)*, 2022. [2](#)
- [12] Souvik Kundu, Shunlin Lu, Yuke Zhang, Jacqueline Liu, and Peter A Beerel. Learning to linearize deep neural networks for secure and efficient private inference. *ICLR*, 2023. [1](#), [2](#), [3](#), [4](#)
- [13] Souvik Kundu, Mahdi Nazemi, Peter A Beerel, and Massoud Pedram. Dnr: A tunable robust pruning framework through dynamic network rewiring of dnns. In *Proceedings of the 26th Asia and South Pacific Design Automation Conference*, pages 344–350, 2021. [2](#)
- [14] Souvik Kundu, Qirui Sun, Yao Fu, Massoud Pedram, and Peter Beerel. Analyzing the confidentiality of undistillable teachers in knowledge distillation. *Advances in Neural Information Processing Systems*, 34, 2021. [1](#)
- [15] Souvik Kundu and Sairam Sundaresan. Attentionlite: Towards efficient self-attention models for vision. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2225–2229. IEEE, 2021. [2](#)
- [16] Souvik Kundu, Sairam Sundaresan, Sharath Nittur Sridhar, Shunlin Lu, Han Tang, and Peter A Beerel. Sparse mixture once-for-all adversarial training for efficient in-situ trade-off between accuracy and robustness of dnns. *ICASSP*, 2023. [2](#)
- [17] Qian Lou, Yilin Shen, Hongxia Jin, and Lei Jiang. Safenet: A secure, accurate and fast neural network inference. In *International Conference on Learning Representations*, 2021. [2](#)
- [18] Pratyush Mishra, Ryan Lehmkuhl, Akshayaram Srinivasan, Wenting Zheng, and Raluca Ada Popa. Delphi: A cryptographic inference service for neural networks. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 2505–2522, 2020. [1](#), [2](#)
- [19] Brandon Reagen, Woo-Seok Choi, Yeongil Ko, Vincent T Lee, Hsien-Hsin S Lee, Gu-Yeon Wei, and David Brooks. Cheetah: Optimizing and accelerating homomorphic encryption for private inference. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pages 26–39. IEEE, 2021. [1](#)
- [20] Liyan Shen, Ye Dong, Binxiang Fang, Jinqiao Shi, Xuebin Wang, Shengli Pan, and Ruisheng Shi. ABNN2: secure two-party arbitrary-bitwidth quantized neural network predictions. In *Proceedings of the 59th ACM/IEEE Design Automation Conference*, pages 361–366, 2022. [1](#)
- [21] Sijun Tan, Brian Knott, Yuan Tian, and David J Wu. CryptGPU: Fast privacy-preserving machine learning on the GPU. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 1021–1038. IEEE, 2021. [1](#)
- [22] Andrew Chi-Chih Yao. How to generate and exchange secrets. In *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pages 162–167. IEEE, 1986. [1](#)
- [23] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016. [4](#)
- [24] Linfeng Zhang, Jiebo Song, Anni Gao, Jingwei Chen, Chenglong Bao, and Kaisheng Ma. Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3713–3722, 2019. [3](#)