

# MPCViT: Searching for Accurate and Efficient MPC-Friendly Vision Transformer with Heterogeneous Attention

Wenxuan Zeng<sup>1</sup> Meng Li<sup>\*</sup> Wenjie Xiong<sup>2</sup> Tong Tong<sup>1</sup> Wen-jie Lu<sup>3</sup>  
Jin Tan<sup>3</sup> Runsheng Wang<sup>1</sup> Ru Huang<sup>1</sup>  
<sup>1</sup>Peking University <sup>2</sup>Virginia Tech <sup>3</sup>Ant Group

{zwx.andy, tongtong}@stu.pku.edu.cn, {meng.li, wrs, ruhuang}@pku.edu.cn,  
wenjiex@vt.edu, {juhou.lwj, tanjin.tj}@antgroup.com

## Abstract

Secure multi-party computation (MPC) enables computation directly on encrypted data and protects both data and model privacy in deep learning inference. However, existing neural network architectures, including Vision Transformers (ViTs), are not designed or optimized for MPC and incur significant latency overhead. We observe Softmax accounts for the major latency bottleneck due to a high communication complexity, but can be selectively replaced or linearized without compromising the model accuracy. Hence, in this paper, we propose an MPC-friendly ViT, dubbed MPCViT, to enable accurate yet efficient ViT inference in MPC. Based on a systematic latency and accuracy evaluation of the Softmax attention and other attention variants, we propose a heterogeneous attention optimization space. We also develop a simple yet effective MPC-aware neural architecture search algorithm for fast Pareto optimization. To further boost the inference efficiency, we propose MPCViT<sup>+</sup>, to jointly optimize the Softmax attention and other network components, including GeLU, matrix multiplication, etc. With extensive experiments, we demonstrate that MPCViT achieves 1.9%, 1.3% and 3.6% higher accuracy with 6.2×, 2.9× and 1.9× latency reduction compared with baseline ViT, MPCFormer and THE-X on the Tiny-ImageNet dataset, respectively. MPCViT<sup>+</sup> further achieves a better Pareto front compared with MPCViT. The code and models for evaluation are available at <https://github.com/PKU-SEC-Lab/mpcvit>.

## 1. Introduction

As machine learning models are handling increasingly sensitive data and tasks, privacy has become one of the major concerns during the model deployment. Secure multi-party computation (MPC) [14] can protect the privacy of

<sup>\*</sup>Corresponding author.




Figure 1. The latency breakdown of a Transformer block w.r.t (a) # heads and (b) # tokens; and (c) the relative error of Softmax and ReLU Softmax with different input variances.

both data and deep neural network (DNN) models and has gained a lot of attention in recent years [42, 51, 60].

However, existing DNN architectures, especially the recently proposed Vision Transformers (ViTs) [2, 11, 64], are not designed or optimized for MPC (the high-level private ViT inference framework in MPC is shown in Figure 2 and Appendix A.1). Although ViTs have achieved superior performance for various vision tasks [2, 15, 16, 64], they face several realistic limitations when directly deployed in MPC: **1) communication overhead:** in contrast to regular inference on plaintext, operations like Softmax, GeLU, max, etc, require intensive communication in MPC, which usually dominates the total inference latency [42, 60]. For example, Softmax is usually very lightweight in plaintext inference. However, as shown in Figure 1(a) and Figure 1(b), it accounts for the majority of the Transformer inference latency due to the high communication complexity; **2) approximation error:** operations like exponential, tanh, reciprocal, etc, cannot be computed directly and require iterative approximation, limiting the computation accuracy. For instance, as shown in Figure 1(c), the relative error of Softmax significantly increases when the input variance is large due to its narrow dynamic range [60]. In contrast, replacing exponential in Softmax with ReLU reduces the relative error drastically (denoted as ReLU Softmax in §2.2).

Numerous efficient Transformer variants have been pro-

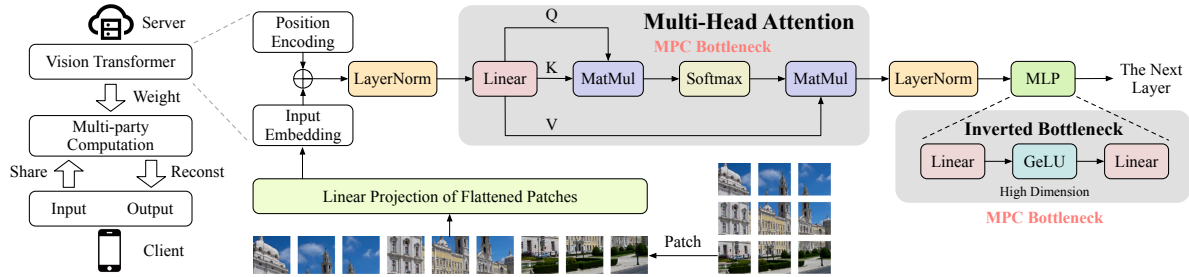


Figure 2. An overall illustration of private Vision Transformer inference in the MPC framework.

posed in recent years [5, 29, 49, 58]. Linear Transformers, including Linformer [58], cosFormer [49], Reformer [29], etc, reduce the quadratic computation complexity of the attentions and significantly accelerate model inference for long sequences. However, these Transformer optimizations have primarily focused on reducing the network computation. Hence, they either retain complex non-linear functions or still require iterative approximation, both resulting in intensive communication overhead in MPC. Recently, MPCFormer [33] and THE-X [6] are proposed to improve the inference efficiency of BERT [28] in MPC. MPCFormer simplifies Softmax by replacing exponential with a more MPC-friendly quadratic operation while THE-X approximates Softmax with a small estimation network. However, their attention variants is not flexible and can still be too expensive for certain latency constraints. Moreover, the network-level optimization suffers from a large accuracy degradation when directly applied to vision tasks.

In this work, we first breakdown Softmax into more atomic operations and analyze their impact on the inference accuracy and efficiency, based on which we propose and compare a set of attention variants comprehensively. The comparison enables us to find MPC-friendly attention variants, which are either highly accurate or highly efficient. We further observe not all attentions are equally important within a ViT. Based on these observations, we propose the first MPC-friendly ViT architecture, dubbed MPCViT. MPCViT features a heterogeneous attention design space to explore the trade-off between accuracy and efficiency, and an MPC-aware differentiable neural architecture search (NAS) algorithm for effective Pareto optimization. Our contributions can be summarized as follows:

- We breakdown Softmax and reveal the impact of atomic operations on the inference accuracy and efficiency, based on which a heterogeneous attention design space is proposed.
- We propose an effective MPC-aware differentiable NAS algorithm to explore the attention design space based on real latency measurements instead of proxies. To the best of our knowledge, we make the first

attempt to introduce the MPC-aware NAS to optimize ViT inference in MPC.

- Our MPC-friendly ViT model family, MPCViT, outperforms prior-art models in MPC in terms of both accuracy and efficiency. MPCViT achieves 1.9% and 1.3% higher accuracy with  $6.2\times$  and  $2.9\times$  latency reduction compared with baseline ViT and MPCFormer on Tiny-ImageNet dataset, respectively.
- We further extend MPCViT to optimize other Transformer components simultaneously with Softmax, named MPCViT<sup>+</sup>, achieving an even better Pareto front compared with MPCViT.

## 2. Preliminaries and Related Works

### 2.1. ViT and Efficient Attention

**ViT architecture.** ViT [11] has demonstrated a great potential to capture the long-range visual dependencies. It takes image patches as input and is composed of an input projection layer, a stack of Transformer layers, and a task-specific multi-layer perceptron (MLP) head. Each layer consists of a multi-head attention (MHA) layer and an MLP block. The core computation is the Softmax attention:

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V,$$

where  $Q, K, V$  are queries, keys and values, respectively.  $d_k$  denotes the embedding dimension of each key, and Softmax normalizes the attention map. The detailed explanation of ViT architecture is shown in Figure 2 and Appendix A.2.

**Linear attention.** Linear attention has been widely studied in previous works [5, 29, 38, 49]. These works aim at reducing the quadratic increase of compute and memory by leveraging special kernel functions. Different kernel functions have been proposed in existing works [4, 49, 53, 58]. UFO-ViT [53] proposes to use  $\ell_2$ -Norm as the kernel function while CosFormer [49] leverages the cosine distance kernel. Linformer [58] approximates self-attention by a low-rank matrix. Hydra Attention [4] proposes to use as many heads as possible to further reduce the computation

Table 1. Comparison among existing network optimizations for the efficient private inference.

Method	Model	Technique	Optimized Component	Granularity	Linear Fusion
Delphi [42]	CNN	ReLU-aware NAS	ReLU	Layer	✗
DeepReDuce [25]	CNN	Manually removal	ReLU	Layer	✗
CryptoNAS [13]	CNN	ReLU-aware NAS	ReLU	Layer	✗
SNL [8]	CNN	ReLU-aware NAS	ReLU	Channel, pixel	✗
Sphinx [7]	CNN	Block-level NAS	ReLU	Block	✗
SENet [31]	CNN	Sensitivity-aware alloc.	ReLU	Layer, pixel	✗
RRNet [48]	CNN	Hardware-aware NAS	ReLU	Layer	✗
DeepReShape [26]	CNN	ReLU-equalization	ReLU	Stage	✗
SAFENet [31]	CNN	Mixed-precision poly. approx.	ReLU	Channel	✗
THE-X [6]	BERT	Polynomial approx.	GeLU, Softmax, LayerNorm	Network	✗
MPCFormer [33]	BERT	Polynomial approx.	Softmax, GeLU	Network	✗
MPCViT (ours)	ViT	MPC-aware NAS Heterogeneous attention	Softmax	Layer, head, row	✗
MPCViT <sup>+</sup> (ours)	ViT	MPC-aware NAS Heterogeneous attention GeLU linearization Linear fusion	Softmax, GeLU, MatMul	Layer, head, row, token	✓

complexity of the linear attention. However, these works primarily focus on reducing the network computation rather than the communication overhead and are not beneficial for ViT inference in MPC.

**Non-local neural networks and Scaling attention.** [59] presents a generic non-local operation to capture long-range dependencies. The formulation is defined as  $\frac{1}{\mathcal{C}(x)} \sum_{\forall j} f(x_i, x_j)g(x_j)$ , where  $x$  denotes the input feature map and  $f(\cdot)$  is a similarity function, e.g., cosine or Euclidean distance.  $g(\cdot)$  computes a certain input representation, and  $\mathcal{C}(x)$  is a normalization factor. By setting  $f(\cdot)$  as the dot-product similarity and  $\mathcal{C}(x) = n$ , where  $n$  is the sequence length, we get an attention variant, named Scaling Attention (ScaleAttn). ScaleAttn is an extremely simple attention and only involves linear operations like multiplication and addition. ScaleAttn can be further re-parameterized to accelerate computation<sup>1</sup> as below:

$$\text{ScaleAttn}(Q, K, V) = \frac{1}{n}(QK^T)V = \frac{Q}{\sqrt{n}}\left(\frac{K^T}{\sqrt{n}}V\right).$$

## 2.2. Multi-Party Computation

MPC [14] is a cryptographic technique that offers a promising solution to protect the privacy of both data and model during deep learning inference. Appendix B depicts the cryptographic primitives for MPC, including Secret Sharing, Oblivious Transfer and Garble Circuit. [14] proposes an efficient MPC protocol for addition and multiplication. [44] proposes efficient two-party computation protocols for various arithmetic operations and approximations for non-linear functions in DNN. For example, [44]

approximates Softmax with ReLU Softmax, i.e.,

$$\text{ReLUSoftmax}(x) = \frac{\text{ReLU}(x_i)}{\sum_i \text{ReLU}(x_i) + \epsilon},$$

where  $\epsilon$  is a very small value to avoid the zero denominator. ABY3 [43] is proposed to enable efficient conversion between different secret sharing schemes in three-party computation. Delphi [42] develops a hybrid MPC protocol and optimizes the network to explore the performance-accuracy trade-off. Iron [18] proposes a set of protocols to reduce the communication overhead of matrix-matrix multiplications (MatMuls), LayerNorm, etc for Transformer-based models.

Although MPC protocols have been improved significantly in recent years, DNN inference in MPC still suffers from intensive communication and latency overhead compared to the plaintext inference. To reduce the latency, different algorithms have been proposed to optimize convolutional NNs (CNNs) with a focus on ReLU. We compare the important characteristics of our methods with existing related work in Table 1. SNL [8] and DeepReDuce [25] propose to remove ReLUs and linearize a subset of neurons selectively. SAFENet [37], Delphi [42], RRNet [48] and PolyMPCNet [47] replace ReLUs with polynomial functions that are more MPC-friendly via NAS for CNNs. Sphinx [7] further develops an MPC-friendly architecture search space for CNNs and leverages NAS to optimize the accuracy given different latency constraints. However, these works mainly focus on the ReLU counts in CNNs, and accelerating ViTs with latency consideration is still a challenge. MPCFormer [33] and THE-X [6] recently propose to replace Softmax with more MPC-friendly operations for BERT, but the network-level method incurs limited latency reduction and a large accuracy degradation. As shown in Table 1, our work proposes an MPC-aware NAS to explore a fined-grained heterogeneous attention optimization space, and further presents a GeLU linearization

<sup>1</sup>In this work, we do not focus on computational overhead in MPC since communication dominates the overhead.

Table 2. Top-1 accuracy and latency comparison of different attention variants on CIFAR-10. The rows in grey are attention variants with the lowest latency and the highest accuracy, respectively. We also show the properties and remaining operations of each variant.

Attention Variant	Monotonicity	Non-negativity	Normalization	Exponential	Max	Reciprocal	Acc. (%)	Lat. (s)
Softmax Attention [11]	✓	✓	✓	✓	✓	✓	92.69	6.82
Linformer Attention [58]	✓	✓	✓	✓	✓	✓	90.85	5.44
ReLU Attention	✓	✓	✗	✗	✓	✗	fail	fail
ReLU6 Attention	✓	✓	✗	✗	✓	✗	90.50	3.02
Sparsemax Attention [40]	✓	✓	✗	✗	✓	✓	91.23	3.23
XNorm Attention [53]	✓	✗	✗	✗	✗	✗	91.24	13.25
Square Attention	✗	✓	✗	✗	✗	✗	91.27	0.72
2Quad Attention [33]	✗	✓	✓	✗	✗	✓	91.86	4.22
Scaling Attention (ScaleAttn) [59]	✓	✗	✗	✗	✗	✗	91.52	<b>0.66</b>
ReLU Softmax Attention (RSAttn) [44]	✓	✓	✓	✗	✓	✓	<b>92.31</b>	5.32




Figure 3. (a) Latency breakdown of a 4-head ViT with 256 hidden dimension and 65 tokens on CIFAR-10 dataset; (b) latency breakdown of Softmax. Both (a) and (b) consist of SEMI-2K (outside) and Cheetah (inside) protocols.

technique to reduce the latency of MLP.

### 3. Motivating Inspiration of MPCViT

In this section, we breakdown Softmax for accuracy and efficiency evaluation and compare two ways of building MPC-friendly ViTs. The analysis serves as our motivation of proposing MPCViT, which features heterogeneous attention and NAS-based Pareto optimization.

**Motivation 1: ViT latency bottleneck in MPC.** In Figure 1, we profile the latency of a Transformer block with different number of heads and tokens. We now further dive into Softmax and evaluate its latency breakdown. We profile a 4-head ViT with 65 tokens and 256 hidden dimension assuming the WAN setting as defined in §5.1. We consider two widely used protocols, including SEMI-2K [9] and Cheetah [22]. From both Figure 1 and Figure 3, we make the following observations: 1) Softmax is the major latency bottleneck in most cases. The latency of GeLU and MatMuls also becomes non-negligible when the number of heads or tokens is small; 2) Softmax mainly involves three operations, including exponential, reciprocal and max. All of the operations are important for the latency overhead. Therefore, we are motivated to get rid of these expensive operations as many as possible, and we propose to first focus on optimizing the Softmax and only consider GeLU and MatMul when the number of heads or tokens is small.

### Motivation 2: comparison of different attention variants.

To reduce the latency of Softmax in the attention module, there are two different ways, i.e., 1) reducing the dimension of  $QK^T$  such as Linformer [58], which approximates  $QK^T$  with a low rank matrix; and 2) simplifying Softmax by replacing exponential, reciprocal or max with other operations. For the second method, while the possible activation variants can be abundant, we observe the following three important properties of Softmax, which helps us to select promising candidates: *monotonicity*, *non-negativity*, and *normalized sum to 1*. The first two properties are related to the exponential while the last one is achieved with reciprocal. We propose and evaluate the following variants, including ReLU, ReLU6, Sparsemax [40], XNorm [53], Square, Scaling [59], 2Quad [33], Linformer [58] and ReLU Softmax [44]. We define the attention with ReLU Softmax as RSAttn for short. The formulation of these attention variants is described in Appendix C.

To compare these attention variants, we train each model for 300 epochs on CIFAR-10 and compare the latency, accuracy, remaining operations, and how they satisfy the properties of Softmax in Table 2. We make the following observations: 1) Attention with more expensive operations can be more inefficient. For instance, 2Quad has an extra reciprocal compared with Square such that 2Quad suffers from a higher latency; 2) Linformer suffers from a high accuracy degradation ( $\sim 1.8\%$  compared with Softmax attention) as it significantly reduces the matrix dimensions such that the model learning capability is degraded, and the computation reduction of this optimization cannot benefit latency reduction a lot. Hence, the first method is not preferred; 3) ReLU Softmax and 2Quad achieve higher accuracy among other variants and we hypothesize this is because they both normalize the sum to 1, which is realized with reciprocal and also leads to relatively higher latency. ReLU Softmax is more preferred over Softmax and 2Quad due to its small approximation error and small accuracy degradation; 4) Scaling is equivalent to directly removing the Softmax. Although it is the fastest, directly linearizing all the Softmax leads to a large accuracy degradation; 5) it is interesting to notice that although XNorm attention has linear com-

putation, it incurs the highest latency because computing  $\ell_2$ -norm involves both reciprocal and square root, both of which are expensive in MPC.

**Motivation 3: Not all attentions are equally important.** As observed in [5, 36, 41, 46, 54, 56], not all attentions are equally important and different layer prefers different attention type. For instance, [46] identifies the minor contribution of the early attention layers in recent ViTs, while [34] demonstrates the second head shows higher importance than other heads. Thus, the importance of attentions among different layers or even the attentions within the same layer can be different. It indicates we can safely replace the expensive Softmax or ReLU Softmax with the cheaper Scaling without compromising the network accuracy.

**Remark** The analysis above provide us with the following intuitions when designing an MPC-friendly ViT:

- 1) ReLU Softmax is more preferred over Softmax and more MPC-friendly, and Scaling has the highest inference efficiency;
- 2) by selectively replacing the ReLU Softmax with the cheaper Scaling, it is possible to reduce the inference latency without compromising the inference accuracy;
- 3) when the latency of GeLU and Softmax are comparable, we have an opportunity to further simplify GeLU along with Softmax.

## 4. The Proposed MPCViT Algorithm

In this section, we introduce MPCViT, which for the first time optimizes the ViT architecture with NAS-based algorithm for a fast private inference in MPC.

### 4.1. Overview of Optimization Flow

Based on the motivations described above, we now describe our MPC-aware ViT optimization flow as shown in Figure 4. We propose to replace the Softmax with the accurate yet MPC-friendly ReLU Softmax in the ViT first and then, linearize the relatively “unimportant” ReLU Softmax with efficient Scaling. To find the “unimportant” ReLU Softmax, we first design a heterogeneous search space that includes both ReLU Softmax and Scaling attentions. Architecture parameters are defined to measure the importance of ReLU Softmax and enable the selection between the two attention variants with different structure granularities (§4.2). Then, we propose an MPC-aware differentiable NAS algorithm to learn the architecture parameters and model parameters simultaneously. We select between the attention variants based on the architecture parameters and the latency constraints (§4.3). In the last step, we retrain the ViT with heterogeneous attentions from scratch. We leverage knowledge distillation (KD) to improve the accuracy of the searched networks (§4.4). Appendix D depicts the complete algorithm flow of MPCViT.

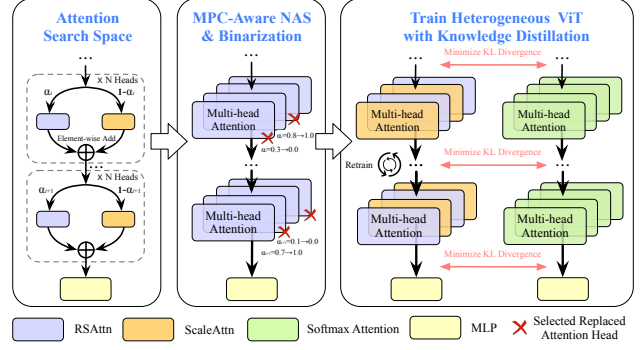


Figure 4. Overview of our proposed MPCViT pipeline.

### 4.2. Heterogeneous Attention Search Space

Our heterogeneous search space combines two candidate attention variants, i.e., MPC-efficient ScaleAttn and high-accuracy RSAttn and has the following three different structural granularities as shown in Figure 5:

- *Layer-wise* search space is coarse-grained, with each layer using either ScaleAttn or RSAttn. The total number of possible architectures in the search space can be limited, especially for shallow ViT models.
- *Head-wise* search space selects ScaleAttn or RSAttn for each attention head of each layer.
- *Row-wise* search space (i.e., token-wise) is the most fine-grained and mixes the two attention variants along each row of the attention map.

Though the search spaces have different granularities, they share very similar NAS formulation. In the rest of the explanation, we will use the head-wise search space as an example to explain our algorithm.

### 4.3. MPC-Aware Differentiable NAS

Many research efforts have been made to NAS including reinforcement learning method (RL) [17, 24, 57], however, RL-based NAS requires significant overhead during searching. Recent years, [10, 27, 55, 61] explore hardware-aware NAS algorithms to make model fit the rigorous hardware requirement. For private ViT inference, we also need a NAS algorithm that takes MPC conditions into consideration.

Given the heterogeneous attention search space, we now introduce our simple yet effective MPC-aware differentiable NAS algorithm.

**Search formulation.** Inspired by [5], we introduce an architecture parameter  $\alpha$  ( $0 \leq \alpha \leq 1$ ) for each head, which is an auxiliary learnable variable that helps to choose between RSAttn and ScaleAttn. Then, each head implements the computation below:

$$\alpha \cdot \text{ReLU Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V + (1-\alpha) \cdot \frac{\text{ScaleAttn}(Q, K, V)}{\sqrt{d_k}}.$$




Figure 5. An illustration of our proposed search space with different structure granularities: (a) layer-wise search space with the size of  $2^{\#\text{layers}}$ ; (b) head-wise search space with the size of  $2^{\#\text{layers} \times \#\text{heads}}$  and; (c) row-wise search space with the size of  $2^{\#\text{layers} \times \#\text{heads} \times \#\text{tokens}}$ . The blocks in blue represent RSAttn and the blocks in grey represent ScaleAttn.

Note that we also add  $\frac{1}{\sqrt{d_k}}$  for ScaleAttn to make the search process more stable and robust.

**Search objective function.** We aim at training ViTs with high accuracy but as few RSAttn as possible. We introduce  $\ell_1$ -penalty into the objective function and formulate the NAS as a one-level optimization problem. To take MPC overhead into consideration, we incorporate realistic inference latency constraints into objective function to construct the MPC-aware NAS algorithm as

$$\text{Cost}(\alpha_{ij}) = \alpha_{ij} \cdot \text{Lat}(\text{ATTN}_{ij}),$$

$$\min_{\theta, \alpha} \ell(f_{\theta, \alpha}(x), y) + \lambda \cdot \left( \sum_{i=1}^L \sum_{j=1}^N \|\text{Cost}(\alpha_{ij})\|_1 \right),$$

where  $x$  and  $y$  are the input-label pairs,  $\ell(\cdot)$  is the loss function,  $\lambda$  is a hyper-parameter to control the weight of inference overhead,  $N$  is the number of heads in one layer, and  $L$  is the number of layers.  $\|\cdot\|_1$  means  $\ell_1$ -norm, and  $\text{ATTN}_{ij}$  means the candidate attention of  $i$ -th layer and  $j$ -th head. We initialize  $\alpha$  for each head to 1.0 and jointly optimize the network parameter  $\theta$  and architecture parameter  $\alpha$  during searching. Note that our differentiable algorithm computes the gradients of  $\alpha$  for all the attention heads across different layers in ViT simultaneously.

**Architecture parameter binarization.** After the network searching, we obtain  $\alpha$  for each attention head in all layers. To select either RSAttn or ScaleAttn for each head, we binarize  $\alpha$  based on the following rule: we use the  $\text{top-}k$  rule according to the ratio of RSAttn budget  $\mu$ , which is defined as  $\#\text{RSAttn}/\#\text{Heads}$ . Specifically, we first find the  $\mu LN$ -th largest  $\alpha$ , denoted as  $\alpha^*$ , and then we binarize  $\alpha$  for each attention head following  $\mathbb{1}(\alpha \geq \alpha^*)$ . Our proposed search method does not need to train candidate networks, but only needs to train once and then select, which accelerates our search process. Now, we obtain a heterogeneous ViT, and by changing  $\mu$ , we can obtain a family of




Figure 6. MPCViT<sup>+</sup>: The optimization of our proposed GeLU linearization and linear fusion for the MLP block.

ViTs with different accuracy and efficiency trade-off.

#### 4.4. Train the Searched ViT for Better Performance

After binarization, RSAttn with a small  $\alpha$  are replaced with ScaleAttn. We find that directly training a heterogeneous ViT results in an accuracy drop. So, *how to effectively train the heterogeneous ViT and restore the model accuracy* is an important question for us.

**Token-wise feature-based KD.** KD [21] is a powerful tool for a student network to learn from a teacher network, which is usually larger and more complex with a higher learning capacity. To improve the accuracy of MPCViT, we leverage token-wise feature-based self-distillation along with a logits-based KD, which uses the baseline ViT as the teacher network while other larger networks can be easily plugged in and used in our framework. For the feature-based KD, we take the output features of the last ViT layer from both the student and the teacher networks and compute the  $\ell_2$ -distance, denoted as  $\mathcal{L}_{feature}$ . We also use KL-divergence loss  $\mathcal{L}_{KL}$  for the logits-based KD, denoted as  $\mathcal{L}_{logits}$ . Combining distillation loss with the task-specific cross-entropy (CE) loss  $\mathcal{L}_{CE}$ , we obtain the final training objective under a distillation temperature  $\mathcal{T}$  as follows:

$$\mathcal{L}_{train} = \chi \mathcal{L}_{CE} + \beta \mathcal{L}_{logits} + \gamma \mathcal{L}_{feature}$$

$$= \chi \mathcal{L}_{CE} + \beta \mathcal{L}_{KL}(\sigma(\frac{z_s}{\mathcal{T}}), \sigma(\frac{z_t}{\mathcal{T}})) + \gamma \|z_s - z_t\|_2,$$

where  $\chi$ ,  $\beta$  and  $\gamma$  are training hyper-parameters to balance different loss functions in the training objective,  $z_s$  and  $z_t$  are feature maps from student and teacher, and  $\sigma$  is Softmax. This self-distillation requires no extra model and only introduces negligible training cost, and has no influence on inference latency or communication overhead in MPC.

#### 4.5. MPCViT<sup>+</sup>: Support for GeLU Linearization

The above design of MPCViT leverages NAS algorithm to build the efficient yet effective attention mechanism. Besides attention, in MLP blocks, non-linear GeLU also impacts the communication cost in some protocols, e.g., Cheetah when the number of heads or tokens is small. And two linear layers near GeLU are also costly because

the inverted-bottleneck structure of MLP leads to high-dimensional MatMul, which is shown in Figure 1 and Figure 3(a). To further improve the efficiency of ViTs, we propose an extended version called MPCViT<sup>+</sup>. Specifically, MPCViT<sup>+</sup> searches and removes relatively “unimportant” GeLUs. Given architecture parameter  $\beta$  for GeLU, the NAS algorithm is similar to the attention optimization:

$$(\beta \cdot \text{GeLU}(x) + (1 - \beta) \cdot x) \cdot \mathcal{W}.$$

The architecture parameters  $\alpha$  and  $\beta$  are jointly learnt during the search, thus the NAS loss is defined along with the classification loss  $\min_{\theta, \alpha, \beta} \ell(f_{\theta, \alpha, \beta}(x), y)$ :

$$\begin{aligned} \text{Cost}(\beta_{ij}) &= \beta_{ij} \cdot \text{Lat}(\text{Act}_{ij}), \\ \lambda \cdot \left( \sum_{i=1}^L \sum_{j=1}^N \|\text{Cost}(\alpha_{ij})\|_1 \right) &+ \eta \cdot \left( \sum_{i=1}^L \sum_{j=1}^T \|\text{Cost}(\beta_{ij})\|_1 \right), \end{aligned}$$

here, we take token-wise linearization as an example, and  $\text{Lat}(\text{Act}_{ij})$  denotes the latency of either GeLU or identity of  $i$ -th layer and  $j$ -th token,  $T$  denotes the number of tokens in each layer. We can consider the balance between  $\lambda$  and  $\eta$  through the private inference latency:  $\frac{\lambda}{\eta} = \frac{\text{Lat}(\text{ATTN})}{\text{Lat}(\text{GeLU})}$ .

**Fuse two linear layers together.** After GeLU linearization, we observe an opportunity for further latency reduction. As shown in Figure 6, we fuse the linear layer before and after GeLU of a certain token, which reduces the communication associated with the MLP. Denote the input of a MLP block of  $i$ -th layer as  $X^{(i)} \in \mathbb{R}^{B \times H \times W \times C}$ , the linear layers as  $\mathcal{W}_1^{(i)} \in \mathbb{R}^{C \times K}$  and  $\mathcal{W}_2^{(i)} \in \mathbb{R}^{K \times C}$ , the output of MLP as  $X^{(i+1)} \in \mathbb{R}^{B \times H \times W \times C}$ . The input dimension of GeLU is  $B \times H \times W \times K$ , and  $C, K$  are channel numbers of linear layers. Note that we generally set the ratio of MLP (i.e.,  $\frac{K}{C}$ ) to 2 or 4 in ViT models. The linearized MLP can be formulated as  $X^{(i+1)} = (X^{(i)} \mathcal{W}_1^{(i)}) \mathcal{W}_2^{(i)} = X^{(i)} (\mathcal{W}_1^{(i)} \mathcal{W}_2^{(i)})$ . In this way,  $\mathcal{W}_1^{(i)} \mathcal{W}_2^{(i)}$  is fused into  $\mathcal{W}_f^{(i)} \in \mathbb{R}^{C \times C}$ , which optimizes two high-dimensional MatMuls to only a single low-dimensional MatMul, further reducing the latency of MLP.

**Adding additional activations.** After GeLU removal and linear fusion, there is no non-linearity in these tokens. To boost the accuracy, we additionally add a non-linear function after the fused linear positions. Here, we choose to add ReLU which incurs a cheaper cost than GeLU.

## 5. Experiments

### 5.1. Experimental Setup

**Experimental and MPC settings.** In our experiments, we leverage the SecretFlow (SPU [39] V0.3.1b0)<sup>2</sup> framework for private inference, which is popular for privacy-

<sup>2</sup><https://github.com/secretflow/secretflow>

preserving deep learning. We adopt the SEMI-2K [9] protocol, which is a semi-honest two-party computation protocol<sup>3</sup>. We follow Cheetah [22] and use the WAN mode for communication. Specifically, the communication bandwidth between the cloud instances is set to 44 MBps and the round-trip time is set to 40ms. Our experiments are evaluated on an Intel Xeon CPU@2.40 GHz with 62 GB RAM.

**Model architectures and datasets.** Many researches have studied training ViTs on small datasets [12, 19, 23, 32, 35]. We consider two ViT architectures on three commonly used datasets following [19]. For the CIFAR-10/100 dataset, we set the ViT depth, # heads, hidden dimension, and patch size to 7, 4, 256, and 4, respectively. For Tiny-ImageNet dataset, we set the ViT depth, # heads, hidden dimension, and patch size to 9, 12, 192, and 4, respectively.

**Searching settings.** We run the NAS algorithm for 300 epochs with AdamW optimizer and a cosine learning rate across three datasets. We set  $\epsilon = 10^{-8}$  to avoid the zero denominator in RSAttn, and  $\lambda = 10^{-5}$ .

**Training settings.** Following [19], we train the searched heterogeneous ViTs for 600 epochs on CIFAR-10/100 and 300 epochs on Tiny-ImageNet. We leverage data augmentations as [19]. For KD, we set the temperature  $\mathcal{T}$  to 1, and set  $\chi, \beta, \gamma$  to 1 as well. To explore the trade-off between inference accuracy and latency, we set  $\mu$  to {0.1, 0.3, 0.5, 0.7}.

### 5.2. Comparison with Prior-Art Efficient Attentions

**Baselines.** We compare MPCViT with prior-art models on the inference accuracy and efficiency. We compare with MPCFormer [33], Linformer [58, 60], THE-X [6] and ViT variants in Table 2. Linformer is studied in [60] as an efficient Transformer variant for MPC and the compression rate of attention dimension is also defined as  $\mu$ . For networks compared with MPCViT, we only modify attention layer and keep MLP unchanged as the baseline ViT. For MPCViT<sup>+</sup>, we optimize attention and MLP simultaneously.

**Results and analysis.** The main results are shown in Figure 7, where we report top-1 accuracy across three datasets and the network inference latency. The main findings are as follows:

- 1) MPCViT outperforms prior-art methods. Without KD, on Tiny-ImageNet, MPCViT with  $\mu = 0.1$  outperforms Linformer with  $\mu = 0.7$  by 0.57% better accuracy with  $4.9\times$  latency reduction; Compared with THE-X, MPCViT achieve 3.6% accuracy improvement with  $1.9\times$  latency reduction;
- 2) with proper KD, MPCViT even achieves 1.9% and 1.3% better accuracy with  $6.2\times$  and  $2.9\times$  latency reduction with  $\mu = 0.1$  and 0.3, compared with baseline ViT and MPCFormer on Tiny-ImageNet, respectively;

<sup>3</sup>Although the actual latency of each operator may be different for different MPC protocols, our proposed algorithm can be generally applied.




Figure 7. Top-1 Accuracy and inference latency comparison with baseline ViTs across three datasets.




Figure 8. Latency and communication cost comparison of different layers in the 9-layer 12-head ViT with  $\mu = 0.5$ .

- 3) we mix Softmax attention and ScaleAttn. Although it has a slightly higher accuracy in plaintext inference, it incurs at least  $1.4\times$  and  $1.2\times$  latency on CIFAR-10/100 and Tiny-ImageNet compared with MPCViT.

We visualize the latency and communication cost comparison of Softmax ViT, ReLU Softmax ViT, and MPCViT in Figure 8. We can observe the latency changes proportionally with the communication. Also, we observe for MPCViT, RSAtn in middle layers (e.g., layer 2 to layer 5) tend to be preserved while RSAtn in early and final layers are likely to be replaced with ScaleAttn.

### 5.3. Comparison on Structure Granularities

The choice of search space is important for MPCViT, so we compare different structure granularities of search space proposed in §4.2. We use the CIFAR-10 dataset for the comparison. As shown in Table 3, the head-wise search space achieves the best accuracy. We hypothesize although row-wise structure leads to the largest search space, such a fine-grained search space may introduce difficulties in NAS, and hence, leads to network architecture with inferior accuracy. Head-wise search space achieves the right balance between flexibility and optimizability.

### 5.4. Ablation Study of MPCViT

**Contribution of KD.** One of our key techniques for training heterogeneous ViT is KD. We enumerate different

Table 3. Top-1 accuracy (%) comparison with different structure granularities of attention search space with different  $\mu$ .

Granularity	Layer-wise	Row-wise	Head-wise
$\mu = 0.5$	93.01	93.16	93.21
$\mu = 0.7$	93.32	93.13	93.38

Table 4. Top-1 Accuracy (%) comparison of different combinations of training loss function with  $\mu = 0.7$ .

CE	Logits-based KD	Feature-based KD	CIFAR-10	Tiny-ImageNet
✓			93.38	59.02
✓	✓		94.18	62.39
✓		✓	94.12	62.26
✓	✓	✓	94.14	61.80
✓	✓	✓	<b>94.27</b>	<b>63.03</b>

combinations of KD and the results are shown in Table 4. We find that: 1) both logits-based and feature-based KD significantly improve the performance of ViT compared to using cross-entropy loss only; and 2) combining the two KD losses further improves the accuracy, especially on a larger dataset. The above findings indicate the importance of KD and the indispensability of three parts of  $\mathcal{L}_{train}$ .

**Consistency and scalability of NAS algorithm.** We hope our NAS algorithm to be robust against hyper-parameter choices and datasets. To analyze this consistency of our algorithm, we adjust the coefficient  $\lambda$  to three different values, i.e.,  $10^{-3}$ ,  $10^{-4}$ ,  $10^{-5}$ , and train a 7-layer ViT with 4 heads. We visualize the distribution of the architecture parameter  $\alpha$  in Figure 9, and we find that  $\alpha$  in each layer has a quite similar trend under different  $\lambda$ 's. Furthermore, when we change the number of heads to 8, or even train the same architecture on different datasets like CIFAR-100 and ImageNet, the distribution still shows a similar trend, empirically proving the consistency and scalability of our proposed NAS algorithm. Appendix F shows more distribution trends under other settings.

**Comparison with per-layer search.** We also compare MPCViT with per-layer search which replaces RSAtn heads in each layer uniformly given a certain  $\mu$  across three



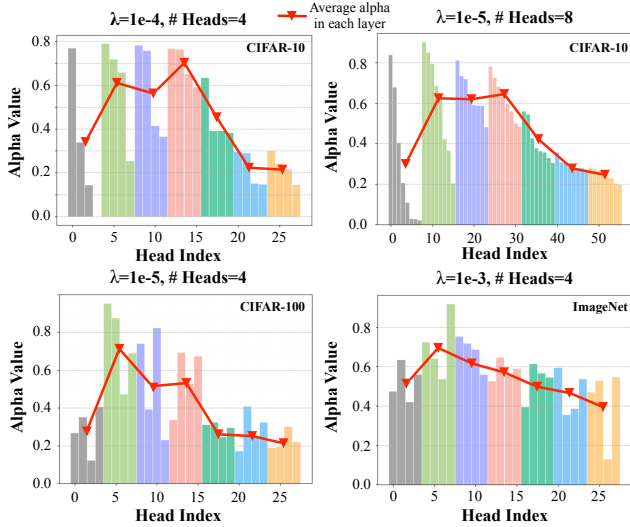


Figure 9. The distribution of architecture parameter  $\alpha$  for each head under different  $\lambda$  and different # heads on different datasets. We draw the average  $\alpha$  in each layer to show the similar trend.

datasets. As shown in Figure 7, MPCViT achieves a better Pareto front for accuracy and efficiency compared to the per-layer NAS. The results demonstrate the importance to search and replace RSAtn heads across different layers.

### 5.5. MPCViT+ Evaluation

As shown in Figure 7, when  $\mu$  decreases to 0.1, the major latency bottleneck comes from GeLU in MLP. We evaluate MPCViT+ for GeLU reduction on CIFAR-10/100 datasets with KD, which correspond to the ViT models with significant GeLU latency. We take token-wise optimization as an example (exploration for other granularity, i.e., layer-wise is shown in Appendix E), and control the GeLU linearization rate using a threshold  $\sigma$  such that  $\beta \leftarrow 0 (\beta \leq \sigma)$  while  $\beta \leftarrow 1 (\beta > \sigma)$ . As shown in Figure 10, GeLU linearization leads to a little accuracy degradation, indicating that GeLU plays an important role in ViTs. MPCViT+ further reduces the latency on the basis of attention optimization, which proves the effectiveness of MPCViT+.

**Ablation study of additional ReLU.** To analyze the help of adding additional ReLU after the fused linear layer, we fix  $\sigma = 0.75$  and perform the ablation study for MPCViT+ under different latency constraints on CIFAR-10 dataset in Table 5. The results show that additional ReLU improves the representation ability of MPCViT+.

Table 5. Ablation study of adding additional ReLU after the fused linear layer under different latency constraints.

$\mu$	0.3		0.7	
	Accuracy (%)	Latency (s)	Accuracy (%)	Latency (s)
MPCViT+	93.92	43.88	94.27	54.56
w/o ReLU	93.72	41.37	93.81	52.11

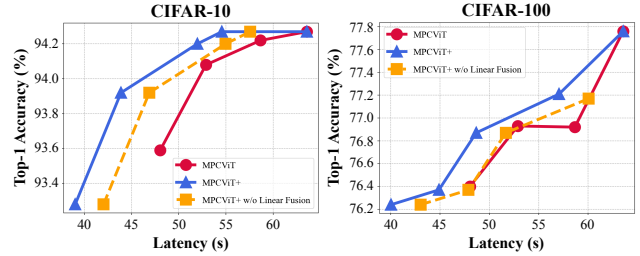


Figure 10. Comparison between MPCViT and MPCViT+.

## 6. Conclusion

In this paper, we present an MPC-friendly ViT family, dubbed MPCViT, to enable accurate yet efficient ViT inference in MPC. We propose a heterogeneous attention optimization space, and design an MPC-aware differentiable NAS algorithm to explore the search space. To further boost the inference efficiency, we propose MPCViT+ to jointly optimize attention and MLP. Extensive experiments demonstrate that MPCViT consistently outperforms prior-art ViT architectures with both lower latency and higher accuracy. Our method can be further applied to large language models (LLMs) including BERT and GPT.

## Acknowledgement

This work is supported in part by the NSFC (62125401) and the 111 Project (B18001).

## References

- [1] Saba Akram and Quarrat Ul Ann. Newton raphson method. *International Journal of Scientific & Engineering Research*, 6(7):1748–1752, 2015. 13
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6836–6846, 2021. 1
- [3] Gilad Asharov, Yehuda Lindell, Thomas Schneider, and Michael Zohner. More efficient oblivious transfer and extensions for faster secure computation. In *ACM SIGSAC conference on Computer & Communications Security (CCS)*, pages 535–548, 2013. 14
- [4] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, and Judy Hoffman. Hydra attention: Efficient attention with many heads. In *European Conference on Computer Vision*, pages 35–49. Springer, 2022. 2
- [5] Jason Ross Brown, Yiren Zhao, Iliia Shumailov, and Robert D Mullins. Dartformer: Finding the best type of attention. *arXiv preprint arXiv:2210.00641*, 2022. 2, 5
- [6] Tianyu Chen, Hangbo Bao, Shaohan Huang, Li Dong, Bin-xing Jiao, Daxin Jiang, Haoyi Zhou, Jianxin Li, and Furu Wei. THE-X: Privacy-preserving transformer inference with homomorphic encryption. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3510–3520,

- Dublin, Ireland, May 2022. Association for Computational Linguistics. 2, 3, 7
- [7] 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. 3
- [8] 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. 3, 15
- [9] Ronald Cramer, Ivan Damgård, Daniel Escudero, Peter Scholl, and Chaoping Xing.  $\text{Spd}\mathbb{Z}_{2^k}$ : Efficient mpc mod  $2^k$  for dishonest majority. In *Annual International Cryptology Conference*. Springer, pages 769–798, 2018. 4, 7, 13
- [10] Xiaoliang Dai, Peizhao Zhang, Bichen Wu, Hongxu Yin, Fei Sun, Yanghan Wang, Marat Dukhan, Yunqing Hu, Yiming Wu, Yangqing Jia, et al. Chamnet: Towards efficient network design through platform-aware model adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11398–11407, 2019. 5
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 1, 2, 4, 13
- [12] Hanan Gani, Muzammal Naseer, and Mohammad Yaqub. How to train vision transformer on small-scale datasets? *arXiv preprint arXiv:2210.07240*, 2022. 7
- [13] 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. 3
- [14] Oded Goldreich. Secure multi-party computation. 1998. 1, 3
- [15] Chengyue Gong, Dilin Wang, Meng Li, Xinlei Chen, Zhicheng Yan, Yuandong Tian, Vikas Chandra, et al. Nasvit: Neural architecture search for efficient vision transformers with gradient conflict aware supernet training. In *International Conference on Learning Representations*, 2021. 1
- [16] Jiaqi Gu, Hyoukjun Kwon, Dilin Wang, Wei Ye, Meng Li, Yu-Hsin Chen, Liangzhen Lai, Vikas Chandra, and David Z Pan. Multi-scale high-resolution vision transformer for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12094–12103, 2022. 1
- [17] Minghao Guo, Zhao Zhong, Wei Wu, Dahua Lin, and Junjie Yan. Irlas: Inverse reinforcement learning for architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9021–9029, 2019. 5
- [18] Meng Hao, Hongwei Li, Hanxiao Chen, Pengzhi Xing, Guowen Xu, and Tianwei Zhang. Iron: Private inference on transformers. *Advances in neural information processing systems*, 2022. 3
- [19] Ali Hassani, Steven Walton, Nikhil Shah, Abulikemu Abuduweili, Jiachen Li, and Humphrey Shi. Escaping the big data paradigm with compact transformers. *arXiv preprint arXiv:2104.05704*, 2021. 7
- [20] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 13
- [21] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. 6
- [22] Zhicong Huang, Wen jie Lu, Cheng Hong, and Jiansheng Ding. Cheetah: Lean and fast secure Two-Party deep neural network inference. In *USENIX Security Symposium (USENIX Security)*, pages 809–826, Boston, MA, Aug. 2022. 4, 7, 13
- [23] Ethan Huynh. Vision transformers in 2022: An update on tiny imagenet. *arXiv preprint arXiv:2205.10660*, 2022. 7
- [24] Yesmina Jaafra, Jean Luc Laurent, Aline Deruyver, and Mohamed Saber Naceur. Reinforcement learning for neural architecture search: A review. *Image and Vision Computing*, 89:57–66, 2019. 5
- [25] 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. 3
- [26] Nandan Kumar Jha and Brandon Reagen. Deepreshape: Redesigning neural networks for efficient private inference. *arXiv preprint arXiv:2304.10593*, 2023. 3
- [27] Weiwen Jiang, Lei Yang, Edwin Hsing-Mean Sha, Qingfeng Zhuge, Shouzhen Gu, Sakyasingha Dasgupta, Yiyu Shi, and Jingtong Hu. Hardware/software co-exploration of neural architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 39(12):4805–4815, 2020. 5
- [28] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, volume 1, page 2, 2019. 2
- [29] Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. 2
- [30] Brian Knott, Shobha Venkataraman, Awni Hannun, Shubho Sengupta, Mark Ibrahim, and Laurens van der Maaten. Crypten: Secure multi-party computation meets machine learning. *Advances in Neural Information Processing Systems*, 34:4961–4973, 2021. 13
- [31] Souvik Kundu, Shunlin Lu, Yuke Zhang, Jacqueline Tiffany Liu, and Peter A. Beerel. Learning to linearize deep neural networks for secure and efficient private inference. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. 3
- [32] Seunghoon Lee, Seunghyun Lee, and Byung Cheol Song. Improving vision transformers to learn small-size dataset from scratch. *IEEE Access*, 10:123212–123224, 2022. 7
- [33] Dacheng Li, Hongyi Wang, Rulin Shao, Han Guo, Eric P. Xing, and Hao Zhang. MPCFORMER: fast, performant and provate transformer inference with MPC. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net,

2023. [2](#), [3](#), [4](#), [7](#), [14](#)
- [34] Zhixin Li, Tong Yang, Peisong Wang, and Jian Cheng. Q-vit: Fully differentiable quantization for vision transformer. *arXiv preprint arXiv:2201.07703*, 2022. [5](#)
- [35] Yahui Liu, Enver Sangineto, Wei Bi, Nicu Sebe, Bruno Lepri, and Marco Nadai. Efficient training of visual transformers with small datasets. *Advances in Neural Information Processing Systems*, 34:23818–23830, 2021. [7](#)
- [36] Zexiang Liu, Dong Li, Kaiyue Lu, Zhen Qin, Weixuan Sun, Jiacheng Xu, and Yiran Zhong. Neural architecture search on efficient transformers and beyond. *arXiv preprint arXiv:2207.13955*, 2022. [5](#)
- [37] Qian Lou, Yilin Shen, Hongxia Jin, and Lei Jiang. Safenet: A secure, accurate and fast neural network inference. In *International Conference on Learning Representations*, 2020. [3](#)
- [38] Jiachen Lu, Jinghan Yao, Junge Zhang, Xiatian Zhu, Hang Xu, Weiguo Gao, Chunjing Xu, Tao Xiang, and Li Zhang. Soft: Softmax-free transformer with linear complexity. *Advances in Neural Information Processing Systems*, 34:21297–21309, 2021. [2](#)
- [39] Junming Ma, Yancheng Zheng, Jun Feng, Derun Zhao, Haoqi Wu, Wenjing Fang, Jin Tan, Chaofan Yu, Benyu Zhang, and Lei Wang. SecretFlow-SPU: A performant and User-Friendly framework for Privacy-Preserving machine learning. In *2023 USENIX Annual Technical Conference (USENIX ATC 23)*, pages 17–33, Boston, MA, July 2023. USENIX Association. [7](#)
- [40] Andre Martins and Ramon Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *International conference on machine learning*, pages 1614–1623. PMLR, 2016. [4](#), [14](#)
- [41] Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one? *Advances in neural information processing systems*, 32, 2019. [5](#), [15](#)
- [42] Pratyush Mishra, Ryan Lehmkuhl, Akshayaram Srinivasan, Wenting Zheng, and Raluca Ada Popa. Delphi: A cryptographic inference service for neural networks. In *USENIX Security Symposium*, pages 2505–2522, 2020. [1](#), [3](#)
- [43] Payman Mohassel and Peter Rindal. Aby3: A mixed protocol framework for machine learning. In *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, pages 35–52, 2018. [3](#)
- [44] Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *IEEE symposium on security and privacy (S&P)*, pages 19–38. IEEE, 2017. [3](#), [4](#)
- [45] Moni Naor and Benny Pinkas. Computationally secure oblivious transfer. *Journal of Cryptology*, 18(1):1–35, 2005. [14](#)
- [46] Zizheng Pan, Bohan Zhuang, Haoyu He, Jing Liu, and Jianfei Cai. Less is more: Pay less attention in vision transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2035–2043, 2022. [5](#)
- [47] Hongwu Peng, Shanglin Zhou, Yukui Luo, Shijin Duan, Nuo Xu, Ran Ran, Shaoyi Huang, Chenghong Wang, Tong Geng, Ang Li, et al. Polymcnet: Towards relu-free neural architecture search in two-party computation based private inference. *arXiv preprint arXiv:2209.09424*, 2022. [3](#)
- [48] Hongwu Peng, Shanglin Zhou, Yukui Luo, Nuo Xu, Shijin Duan, Ran Ran, Jiahui Zhao, Shaoyi Huang, Xi Xie, Chenghong Wang, et al. Rrnet: Towards relu-reduced neural network for two-party computation based private inference. *arXiv preprint arXiv:2302.02292*, 2023. [3](#)
- [49] Zhen Qin, Weixuan Sun, Hui Deng, Dongxu Li, Yunshen Wei, Baohong Lv, Junjie Yan, Lingpeng Kong, and Yiran Zhong. cosformer: Rethinking softmax in attention. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. [2](#)
- [50] Deevashwer Rathee, Mayank Rathee, Rahul Kranti Kiran Goli, Divya Gupta, Rahul Sharma, Nishanth Chandran, and Aseem Rastogi. Sirnn: A math library for secure rnn inference. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 1003–1020. IEEE, 2021. [14](#)
- [51] Deevashwer Rathee, Mayank Rathee, Nishant Kumar, Nishanth Chandran, Divya Gupta, Aseem Rastogi, and Rahul Sharma. Cryptflow2: Practical 2-party secure inference. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 325–342, 2020. [1](#)
- [52] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979. [13](#)
- [53] Jeong-geun Song. Ufo-vit: High performance linear vision transformer without softmax. *arXiv preprint arXiv:2109.14382*, 2021. [2](#), [4](#), [14](#)
- [54] Xiu Su, Shan You, Jiyang Xie, Mingkai Zheng, Fei Wang, Chen Qian, Changshui Zhang, Xiaogang Wang, and Chang Xu. Vitas: Vision transformer architecture search. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXI*, pages 139–157. Springer, 2022. [5](#)
- [55] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2820–2828, 2019. [5](#)
- [56] Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In Anna Korhonen, David R. Traum, and Lluís Màrquez, editors, *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 5797–5808. Association for Computational Linguistics, 2019. [5](#), [15](#)
- [57] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8612–8620, 2019. [5](#)
- [58] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. [2](#), [4](#), [7](#), [14](#)
- [59] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. [3](#), [4](#)
- [60] Yongqin Wang, G Edward Suh, Wenjie Xiong, Benjamin

- Lefaudeux, Brian Knott, Murali Annavaram, and Hsien-Hsin S Lee. Characterization of mpc-based private inference for transformer-based models. In *2022 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 187–197. IEEE, 2022. [1](#), [7](#), [13](#), [14](#)
- [61] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10734–10742, 2019. [5](#)
- [62] Mengzhou Xia, Zexuan Zhong, and Danqi Chen. Structured pruning learns compact and accurate models. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 1513–1528. Association for Computational Linguistics, 2022. [15](#)
- [63] Andrew C Yao. Protocols for secure computations. In *23rd annual symposium on foundations of computer science (sfcs 1982)*, pages 160–164. IEEE, 1982. [14](#)
- [64] Pengchuan Zhang, Xiyang Dai, Jianwei Yang, Bin Xiao, Lu Yuan, Lei Zhang, and Jianfeng Gao. Multi-scale vision long-former: A new vision transformer for high-resolution image encoding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2998–3008, 2021. [1](#)