

Efficient Partitioning Vision Transformer on Edge Devices for Distributed Inference

Xiang Liu*, Yijun Song[†], Xia Li[‡], Yifei Sun[§], Huiying Lan[¶],
Zemin Liu[§], Linshan Jiang^{||}, Jialin Li*

*School of Computing, National University of Singapore

[†]Information and Artificial Intelligence Institute, Zhejiang University of Finance & Economics Dongfang College

[‡]Department of Computer Science, ETH Zurich

[§]College of Computer Science and Technology, Zhejiang University

[¶]Lumia Ltd.

^{||}Institute of Data Science, National University of Singapore

Abstract—Deep learning models are increasingly utilized on resource-constrained edge devices for real-time data analytics. Recently, Vision Transformer and their variants have shown exceptional performance in various computer vision tasks. However, their substantial computational requirements and low inference latency create significant challenges for deploying such models on resource-constrained edge devices. To address this issue, we propose a novel framework, ED-ViT, which is designed to efficiently split and execute complex Vision Transformers across multiple edge devices. Our approach involves partitioning Vision Transformer models into several sub-models, while each dedicated to handling a specific subset of data classes. To further reduce computational overhead and inference latency, we introduce a class-wise pruning technique that decreases the size of each sub-model. Through extensive experiments conducted on five datasets using three model architectures and actual implementation on edge devices, we demonstrate that our method significantly cuts down inference latency on edge devices and achieves a reduction in model size by up to 28.9 times and 34.1 times, respectively, while maintaining test accuracy comparable to the original Vision Transformer. Additionally, we compare ED-ViT with two state-of-the-art methods that deploy CNN and SNN models on edge devices, evaluating metrics such as accuracy, inference time, and overall model size. Our comprehensive evaluation underscores the effectiveness of the proposed ED-ViT framework.

Index Terms—Distributed Inference, Edge Computing, Model Splitting, Vision Transformer

I. INTRODUCTION

In recent years, deep learning models have been increasingly deployed on resource-constrained edge devices to meet the growing demand for real-time data analytics in industrial systems [1]–[3] and have demonstrated remarkable capabilities in various applications such as video image analysis and speech recognition. Convolutional Neural Networks (CNNs) [4] like VGGNet [5] and ResNet [6], as well as Spike Neural Networks (SNNs) [7], have achieved satisfactory performance in many edge computing scenarios. As the field progresses, researchers are exploring the deployment of more complex structured

Xiang Liu and Yijun Song contribute equally to this work. Yijun Song (yijunsong.0377@gmail.com) and Linshan Jiang (linshan@nus.edu.sg) are the corresponding authors. Dr. Jialin Li is supported by the Singapore Ministry of Education Academic Research Fund Tier 1 (T1 251RES2104) and Tier 2 (MOE-T2EP20222-0016).

models on edge devices to further improve performance. Transformer architecture [8], which has revolutionized natural language processing (NLP) tasks, has inspired similar advancements in computer vision. Vision Transformer (ViT) models [9] and their variants have shown outstanding results across various computer vision tasks, including image classification [10], [11], object detection [12]–[14], semantic segmentation [15], [16] and action recognition [17], [18] and audio spectrogram recognition [19]. The success of ViTs has sparked interest in leveraging their capabilities for edge computing applications.

However, the rapid advancement in machine learning technologies has increased the demand for computational resources and memory, given the complexity of these model configurations. Achieving higher accuracy with ViT requires substantial computational power and memory, which poses challenges for deployment on edge devices. For instance, ViT-Base [9] consists of over 86.7 million parameters and requires approximately 330MB of memory. Researchers now face the dilemma of deploying such complex models while dealing with resource-constrained resources.

Previous studies aiming to reduce the deployment overhead primarily focus on compressing Vision Transformer models. These approaches can be classified into three major categories: (1) architecture and hierarchy restructuring [20], [21], (2) encoder block enhancements [22]–[27], and (3) integrated approaches [28], [29]. However, these methods often suffer from either poor inference accuracy or high inference latency as they attempt to compress a large model to fit into a memory-constrained edge device.

To develop a solution that mitigates accuracy drop and enables efficient deployment of the Vision Transformer on resource-constraint edge devices, we aim to utilize the collaboration of multiple edge devices and propose a Vision Transformer splitting framework named **Edge Device Vision Transformer**, shorten as **ED-ViT**. As illustrated in Fig. 1, ED-ViT first partitions the original Vision Transformer into several smaller sub-models, inspired by the concept of split learning (SL). However, unlike traditional SL, which does not consider edge device constraints, each of these small sub-models is

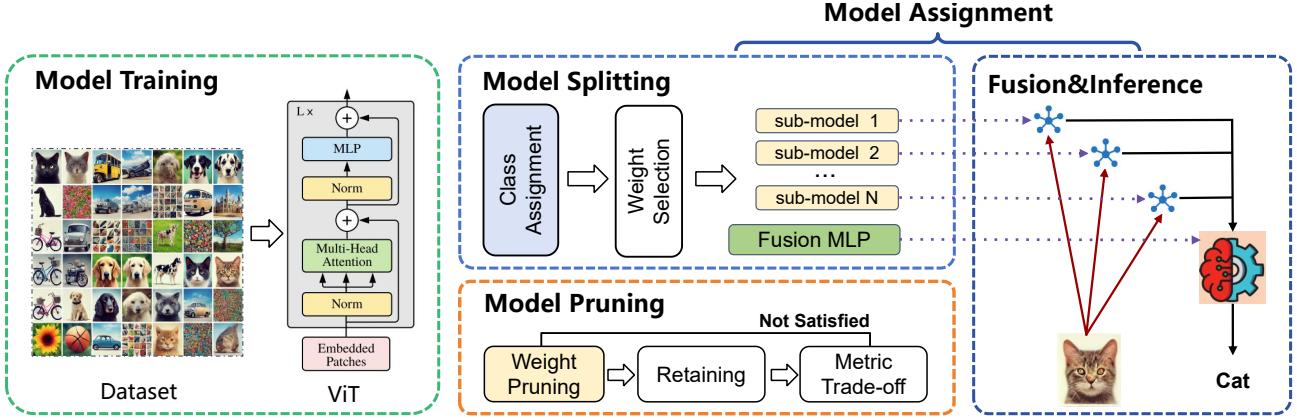


Fig. 1: The overview of ED-ViT, including four steps: Model Splitting, Model Pruning, Model Assignment and Model Fusion.

responsible for detecting a specific subset of the classes and is deployed on resource-constrained edge devices. ED-ViT then employs model pruning techniques to further alleviate the computational load and processing requirement for each sub-model. To optimize model assignment, we design a greedy assignment algorithm that takes into account both the model computational resources and memory resources. Besides integrating the previous steps, ED-ViT uses a multilayer perceptron (MLP) model to fuse the results from all the sub-models. We conduct experiments across five datasets to evaluate the effectiveness of ED-ViT framework on edge devices, particularly in low-power video analytics. The results, measured across three key metrics—accuracy, inference latency, and model size—consistently demonstrate the significant benefits of ED-ViT. Additionally, we compare ED-ViT with methods that split CNN and SNN, highlighting the great potential of deploying Vision Transformer on resource-constrained edge devices to achieve high accuracy while maintaining small model sizes and low inference latency. Our main contributions are summarized as follows:

- This is the first study focusing on combining pruning and splitting, deploying Vision Transformer onto edge devices. We propose a framework that leverages the capabilities of Vision Transformer, allowing for the collaboration of multiple edge devices to achieve distributed inference in practical applications.
- We introduce ED-ViT to address the Vision Transformer splitting problem by decomposing the complex original model into sub-models, and applying pruning techniques to reduce the size of each sub-model. Using a combined greedy method for model assignment, ED-ViT effectively addresses the formulated problem by reducing model sizes, minimizing inference latency, and maintaining high accuracy, achieving a trade-off across these three metrics.
- We conduct extensive experiments with three computer vision datasets and two audio recognition datasets across three ViT structures. Besides, we implement our ED-ViT on Raspberry Pi 4Bs, demonstrating that our framework

significantly reduces inference latency on edge devices and decreases overall memory usage with negligible accuracy loss in various applications.

II. RELATED WORKS

A. General Vision Transformer Compression

Deploying Vision Transformer models in resource-constrained environments poses significant challenges due to their intensive computational and memory demands. These approaches address Vision Transformer resource limitations via pruning, encompassing both local and global strategies as follows.

Local pruning techniques focus on removing redundant components within specific layers of the model. For instance, PVT [30] and its successor PVTv2 [31] introduce a pyramid hierarchical structure to transformer backbones, achieving high accuracy with reduced computation. [32] applies sparsity regularization during training and subsequently prunes the dimensions of linear projections, targeting less significant parameters. [33] prunes multi-head self-attention (MHSA) and feed-forward networks (FFN), which are often redundant components. [34], [35] propose network pruning to eliminate complexity and model sizes by reducing tokens. Other noteworthy contributions include DToP [36], which enables early token exits for semantic segmentation tasks. Conversely, **global pruning techniques** adopt a comprehensive perspective by evaluating and pruning the overall significance of neurons or layers across the entire network. SAViT [37] purposes structure-aware Vision Transformer pruning via collaborative optimization. For instance, CP-ViT [38] systematically assesses the importance of head and attention layers for the purpose of pruning, while Evo-ViT [39] identifies and preserves significant tokens, thereby discarding those of lesser importance. Moreover, the Skip-attention approach [40] facilitates the omission of entire self-attention layers, thereby exemplifying a global pruning methodology. X-pruner [41] employs explainability-aware masks to inform its pruning decisions, thereby advancing a more informed global pruning strategy. In

TABLE I: Performance characteristics for standard Vision Transformer models with their default parametrization at resolution 224×224 : ViT-Small, ViT-Base and ViT-Large, all with patch size of 16×16 . All the experiment results are obtained on Raspberry Pi-4B devices.

Model	Depth	Width	Heads	Params ($\times 10^6$)	Flops ($\times 10^9$)	Latency (ms)	Mem Size (MB)
ViT-Small	12	384	6	22.1	4.25	9628	83
ViT-Base	12	768	12	86.6	16.86	36940	327
ViT-Large	24	1024	16	304.4	59.69	118828	1157

addition, UP-ViT [42] introduces a unified pruning framework that leverages KL divergence to guide the decision-making process for pruning, while LORS [43] optimizes parameter usage by sharing the majority of parameters across stacked modules, thereby necessitating fewer unique parameters.

Among existing pruning methods, UP-ViT [42] has the closest resemblance to our approach. However, it is important to note that these techniques cannot be directly applied to edge devices: they often suffer from poor performance when the pruning ratio is high or incur high computation overhead when the pruning ratio is low, making them unsuitable for resource-constrained edge environments. In contrast, our work introduces a class-based global structured pruning method that addresses these limitations. Our approach is orthogonal to most previous methods and does not involve trainable parameters, contributing to more stable performance.

B. Vision Transformer on Edge Devices

There are several methods focused on deploying Vision Transformer on-edge devices, which can be classified into three major categories.

Architecture and Hierarchy Restructuring: HVT [20] compresses sequential resolutions using hierarchical pooling, reducing computational cost and enhancing model scalability. LeViT [21] is a hybrid model that combines the strengths of CNNs and transformers. For image classification tasks, it utilizes the hierarchical structure of LeNet [4] to optimize the balance between accuracy and efficiency, and uses average pooling in the feature map stage. MobileViTv3 [44] propose changes to the fusion block, which addresses the scaling and simplifies the learning tasks.

Encoder Block Enhancements: ViL [22] introduces a multiscale vision longformer that lessens computational and memory complexity when encoding high-resolution images. Poolformer [23] deliberately replaces the attention module in transformers with a simple pooling layer. LiteViT [24] introduces a compact transformer backbone with two new lightweight self-attention modules (self-attention and recursive atrous self-attention) to mitigate performance loss. Dual-ViT [26] reduces feature map resolution, consisting of two dual-block and two merge-block stages. MaxViT [25] divides attention into local and global components and decomposes it into a sparse form with window and grid attention. Slide-Transformer [27] proposes a slide attention module to address the problem that computational complexity increases quadratically with the attention modules, while EdgeViT [45] enables

attention-based vision models to compete with the best lightweight CNNs when considering the tradeoff between accuracy and on-device efficiency.

Integrated Approaches: Some methods integrate both of the above approaches. CeiT [28] combines Transformer and CNN strengths to overcome the shortcomings of each, incorporating an image-to-tokens module, locally-enhanced feedforward layers, and layer-wise class token attention. CoAt-Net [29] combines depth-wise convolutions and simplifies traditional self-attention by relative attention, enhancing efficiency by stacking convolutions and attention layers. DeViT [46] also decomposes Vision Transformer for collaborative inference. However, DeViT trains a ViT-Large for each sub-model even when splitting ViT-Small and employs model distillation to enhance accuracy, which introduces significant training overhead. In addition, the smallest model size that DeViT provides is larger than 90MB.

However, they never consider linking pruning with specific classes, which limits their methods when both high performance and low memory usage are required.

C. Split Learning

Current works that combine Vision Transformer and split learning primarily focus on federated learning, addressing data privacy and efficient collaboration in multi-client environments [47]–[49], where the inner structure of a large model is split across smaller devices and later fused [50]. However, these approaches do not target the deployment of Vision Transformer on edge devices.

Traditional machine learning model splitting generally involves partitioning a large model into multiple smaller sub-models that can be executed collaboratively on resource-constrained devices, providing a promising technique for deploying models on edge devices. Splitnet [51] clusters classes into groups, partitioning a deep neural network into tree-structured sub-networks. [52] dynamically partitions models based on the communication channel’s state. Nnfacet [1], [2] splits large CNNs into lightweight class-specific sub-models to accommodate device memory and energy constraints, with the sub-models being fused later. [3] follows a similar approach to split deep SNNs across edge devices. Distredge [53] uses deep reinforcement learning to compute the optimal partition for CNN models.

To the best of our knowledge, our work presents the first exploration of Vision Transformer model partitioning for edge deployment, marking a significant contribution to

this field. Drawing inspiration from previous studies [1]–[3], our framework, ED-ViT, introduces an innovative approach to decompose a multi-class ViT model into several class-specific sub-models, each performing a subset of classification. Unlike relying on channel-wise pruning, ED-ViT employs advanced pruning techniques specifically designed for the unique architecture of Vision Transformers.

III. PROBLEM FORMULATION

The structures of three representative Vision Transformer, ViT-Small, ViT-Base, and ViT-Large, are presented in Table I. The number of operations is commonly used to estimate computational energy consumption at the hardware level. In Vision Transformer models, almost all floating-point operations (FLOPs) are multiply-accumulate (MAC) operations.

For Patch Embedding, FFN, and MLP Head, their operation counts are easy to infer as they follow a fully connected (FC) structure, where the MAC count is given by $(2FC_{in} + 1) \times FC_{out}$, where FC_{in} and FC_{out} represent the input and output features, respectively. For MHSA, assuming the number of patches is p , the dimension of each patch is d_p , the embedding dimension is d , and the number of attention heads is h , the MAC for the linear projections to generate the Q , K , and V matrices is $3 \times p \times d^2/h$. The MAC for QK^T is $p^2 \times d/h$, and the MAC for the softmax operation and multiplication with V is $p^2 \times d/h$. For h attention heads, the total MAC is $h \times (3 \times p \times d^2/h + 2 \times p^2 \times d/h) = 3 \times p \times d^2 + 2 \times p^2 \times d$. Based on this analysis, energy consumption can be estimated as being proportional to FLOPs, given that the pruned model follows the same structure.

Thus, we formulate the problem as follows. We assume that we have L inference samples in total to be processed, and the set of N edge devices is represented as D . The available memory and energy (FLOPs of an edge device) [54] for each edge device D_i are denoted as M_i and E_i , respectively. The FLOPs (energy consumption) for each inference sample for the sub-model $Model_j$ from the set of sub-models $\{Models\}$ is represented as e_j , calculated based on the previous energy consumption estimation. To formulate the problem of Vision Transformer partitioning and edge-device-based deployment, we define the objective function as $\max_{\{Model_j\}} \min_{D_i \in D} \{E_i - Le_j\}$, aiming to minimize the maximal inference latency, as inference latency is closely related to the computational power of edge devices. Additionally, the accuracy a_{fus} of the fused results from all N inference samples must be greater than or equal to the required inference accuracy A_{re} ; the total memory sizes of all sub-models should not exceed the memory budget bu .

The optimization problem can be formally formulated as follows, where x_{ie} is a binary decision variable: 1 indicates that the sub-model deployed on edge device D_i is responsible for class e , and 0 otherwise. Each sub-model learns a specific subset of the classes in C . Furthermore, the memory consumption of sub-model j , denoted as $\text{size}(Model_j)$, represented as m_j , must be smaller than the available memory size of the deployed edge device:

$$\left\{ \begin{array}{l} \text{argmax}_{\{Model_j\}} \min_{D_i \in D} \{E_i - Le_j\} \\ \text{s.t. } Le_j \leq E_i, \quad Model_j \text{ deploys on } D_i \\ m_j \leq M_i, \\ a_{fus} \geq A_{re}, \\ \sum_j m_j \leq bu, \\ \sum_{i=1}^{|D|} x_{ie} = 1, \quad \forall e \in C, \forall i \in D \end{array} \right. \quad (1)$$

IV. METHODOLOGY

This section describes the design of the ED-ViT framework proposed to solve the optimization problem outlined in (1). We first explain the main workflow of ED-ViT and then provide detailed descriptions of the four key steps involved.

A. Design Overview

As illustrated in Fig. 1, ED-ViT leverages the unique characteristics of Vision Transformer and the collaboration of multiple edge devices. The framework involves N concurrent edge devices for distributed inference alongside a lightweight MLP aggregation to derive the final classification results. Initially, the original Vision Transformer is trained on the entire dataset to achieve high test accuracy for the classification task. The ED-ViT framework is composed of four main components: model splitting, pruning, assignment, and fusion. During model splitting, the Vision Transformer model is divided into sub-models, each responsible for a subset of classes. To reduce computation overhead, these sub-models are further pruned using model pruning techniques. Subsequently, the sub-models are assigned to the appropriate edge devices, taking the optimization problem into consideration. Finally, the aggregation device fuses the outputs from the edge devices to produce the final inference results. The specific details of each component are provided below.

B. Model Splitting

In the original Vision Transformer, different heads contribute to learning and inferring from the samples. However, for certain classes, maintaining all the connections between the heads can be redundant. As a result, ED-ViT prunes these connections and reconstructs the heads, with more retained heads leading to more parameters and connections being preserved. As illustrated in Algorithm 1, each Vision Transformer sub-model undergoes pruning based on a head number threshold and its associated categories, following a relatively equitable workload distribution. Subsequently, a greedy search mechanism is used to identify the most suitable edge device model assignment plan for deploying a particular sub-model, considering both energy and memory constraints. If the total memory size exceeds the budget or no suitable plan is found, an iterative approach is applied to adjust the number of heads for the sub-model with the biggest memory size to be pruned, repeating the allocation process until all sub-models

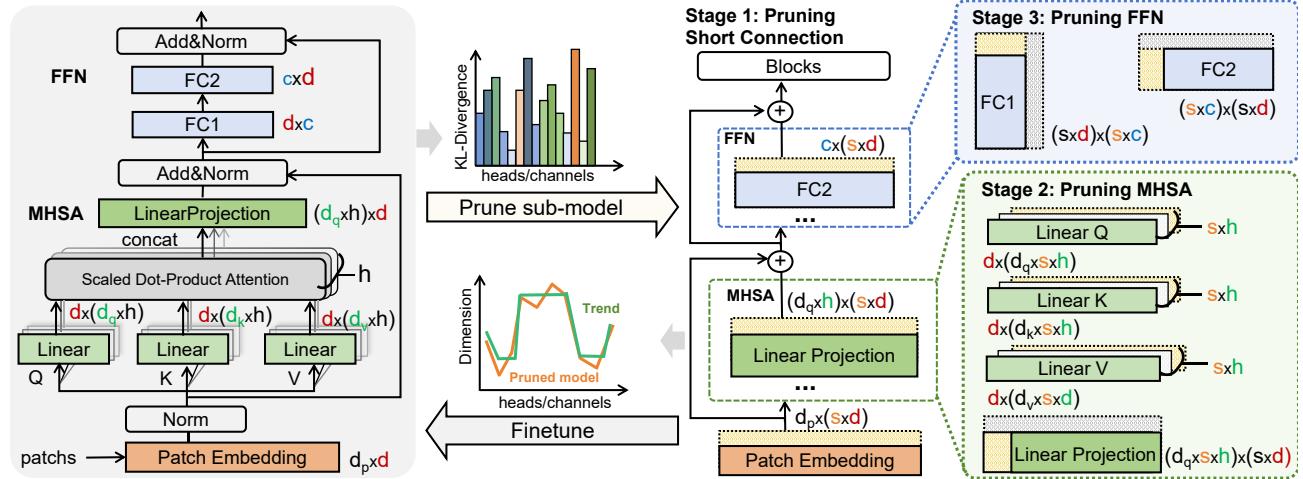


Fig. 2: Structured pruning of a Vision Transformer block. Left: illustration of prunable components in a ViT block. Right: corresponding sequential pruning process. Our approach targets three key components: (1) channels in residual connections (red, denoted as d), (2) the number of heads in the MHSA module (green, denoted as h), and (3) hidden layer channels in the FFN (blue, denoted as c). The pruning process occurs in three stages: residual connection channels, MHSA heads, and FFN hidden dimensions. Yellow regions indicate parameters being pruned in the current stage, while gray regions represent previously pruned parameters.

are successfully assigned to edge devices. The pruning and the greedy assignment methods are shown as Algorithm 2 and Algorithm 3, located in Section IV-C and IV-D, respectively.

C. Model Pruning

We believe that reducing the computational burden of Vision Transformer will significantly contribute to lowering inference latency in distributed edge device settings. We focus on the original ViT architecture [9], chosen for its simplicity and well-defined design space, focusing on redistributing the dimensionality across different blocks to achieve a more balanced tradeoff between computational efficiency and accuracy, as shown in Fig. 2.

Analysis of Prunable Parameters: The main prunable components in a ViT block, as illustrated in Fig. 2.

- Residual Connection Channels (Red, d): The channels across the shortcut connections within the transformer blocks.
- Heads in MHSA (Green, h): The dimensions of the query, key, value projections (d_q, d_k, d_v)¹.
- Feed-Forward Network (FFN) Hidden Dimensions (Blue, c): The dimension c of the hidden layer used for expanding and reducing.

Pruning Process: As illustrated in Fig. 2, The pruning process is carried out in stages, with each stage focusing on one of the prunable components. We compute the KL-Divergence between the output distributions of the original

model and the pruned model to evaluate the importance of each component, as follows:

$$D_{\text{KL}}(P \parallel Q) = \sum_i P(i) \log \frac{P(i)}{Q(i)}$$

where $P(i)$ represents the output distribution of the original model, and $Q(i)$ represents the distribution after pruning.

We focus on pruning the channels of the residual connections (as shown in red) in the first stage. Using KL-Divergence, we identify and prune the channels that contribute the least, reducing the dimensionality from d to $s \times d$, and the pruning factor s controls the degree of reduction in the parameters. We use j -th sub-model as an example: we set $s = (h - hp_j)/h$, effectively controlling the extent of the pruning and parameter reduction. This helps to streamline the flow of information between layers without significantly affecting model performance. Then, instead of directly removing entire heads in the MHSA module, we prune the least important dimensions within the query, key, and value projections (d_q, d_k, d_v) across multiple heads. This process effectively reduces the total number of heads to $s \times h$, without entirely discarding any head, thus maintaining a balanced representation of the attention mechanism while reducing its complexity. The dimensionality of the projections is scaled accordingly to reflect the merging and pruning process, ensuring that the model retains its ability to capture token interactions. The final stage involves pruning the hidden dimension c in the FFN, as shown in blue. By calculating KL-Divergence, we identify the least important neurons and reduce the hidden dimension from c to $s \times c$. Following each pruning stage, the model is finetuned to recover any performance loss that may result from

¹ $d_q = d_k = d_v = d/h$

Algorithm 1 Model Splitting in ED-ViT

Input: The number of edge devices N ; memory budget bu ; initial pruning head number $hp = hp_1, hp_2, \dots, hp_N$ for all sub-models, remaining available memory size M_i and remaining computational resource E_i for device i ; training dataset (\mathbf{X}, \mathbf{y})
Parameter: the classes set C ; trained original Model₀
Output: class-specific sub-models $\{\text{Model}_1, \dots, \text{Model}_N\}$ and a fusion model MLP

```
1: Let  $flag_r = True$ ,  $D = \{\text{device 1, ..., device } N\}$ .
2: Let  $E = \{E_1, \dots, E_N\}$ ,  $M = \{M_1, \dots, M_N\}$ .
3: repeat
4:   Let  $C = \{C_1, C_2, \dots, C_N\}$ , s.t.  $|C| = \sum_{i=1}^N |C_i|$ .
5:    $C_i$  is determined randomly.
6: until  $\forall C_a, C_b \in C, ||C_a| - |C_b|| \leq 1$ .
7: while  $flag_r$  is  $True$  do
8:   for  $i$  in  $N$  do
9:      $\text{Model}_i = \text{prune}(\text{Model}_0, \mathbf{X}, \mathbf{y}, C_i, hp_i)$ .
10:    end for
11:     $MA = \phi$ 
12:    if  $\sum_j m_i \leq bu$  then
13:       $MA = \text{greedySearchAssign}(E, M, \text{Model}_i, D)$ .
14:    end if
15:    if  $MA \neq \phi$  then
16:       $flag_r = False$ .
17:    else
18:       $hp_j = hp_j + 1$  where  $\text{Model}_j$  has the biggest memory size.
19:    end if
20:  end while
21:   $concat_{out} = concat(\text{Model}_1(\mathbf{X}), \dots, \text{Model}_N(\mathbf{X}))$ .
22:   $MLP = \text{train}(concat_{out}, \mathbf{y})$ .
23: return  $MLP, \text{Model}_1, \dots, \text{Model}_N$ .
```

Algorithm 2 Model Pruning in ED-ViT

Input: pruning head number hp_i ; assigned classes subset C_i
Parameter: the raw original Model₀; training dataset (\mathbf{X}, \mathbf{y})
Output: pruned Model_i

```
1:  $\mathbf{X}_i, \mathbf{y}_i = \text{resample}(\mathbf{X}, \mathbf{y}, C_i)$ .
2:  $\text{Model}_i = \text{PruneShortConnection}(\text{Model}_0, hp_i)$ 
3:  $\text{Model}_i = \text{PruneMHS}(\text{Model}_i, hp_i)$ 
4:  $\text{Model}_i = \text{PruneFFN}(\text{Model}_i, hp_i)$ 
5:  $\text{Model}_i = \text{retrain}(\text{Model}_i, X_i, y_i)$ .
6: return  $\text{Model}_i$ .
```

the parameter reduction. This ensures that the pruned model achieves a similar level of accuracy as the original model while requiring fewer computational resources.

In conclusion, the pruning process is outlined in Algorithm 2. An additional advantage of ED-ViT is that, even after pruning, the sub-models still retain the structure of Vision Transformer. This gives our method the potential to

Algorithm 3 Model Assignment in ED-ViT

Input: remaining available memory size set M , remaining computational resource set E , the edge device set D , the sub-model set.

Output: Model assignments MA

```
1:  $\{\text{Models}\} \leftarrow \text{sort}(\{\text{Models}\})$  (Sort the sub-models based on the computation overhead from the highest to lowest).
2: for  $i$  in  $N$  do
3:    $j \leftarrow \text{argmax}_{k \in D} E_k$ .
4:   if  $M_j \geq \text{size}(\text{Model}_i)$  and  $E_j \geq \text{computing}(\text{Model}_i)$  then
5:      $E_j \leftarrow E_j - \text{computing}(\text{Model}_i)$ .
6:      $M_j \leftarrow M_j - \text{size}(\text{Model}_i)$ .
7:   else
8:      $D \leftarrow D - j$ .
9:     if  $D = \phi$  then
10:      return  $\phi$ 
11:    end if
12:  end if
13: end for
14: return  $MA$ 
```

be combined with other horizontal pruning techniques for ViT and its variants and leverage the inherent features of Vision Transformer models to generalize well into downstream tasks.

D. Model Assignment

To address the optimization problem expressed in (1), we propose a greedy search algorithm for assigning Vision Transformer sub-models to edge devices. As shown in Algorithm3, the sub-models are first sorted based on their energy consumption. ED-ViT assigns the most computation-intensive sub-model first based on their model sizes, which is proportional to the computation overhead as in Section III. The algorithm then iteratively assigns the remaining sub-models to maximize the system's available energy. Initially, the device with the highest computational power is selected. If the remaining memory and energy can accommodate the sub-model, we update the device's available memory and energy. Otherwise, if the sub-model exceeds the device's memory capacity, the memory-exhausted device is removed from the set. If no devices remain, it indicates that the current pruning results prevent deployment of all sub-models. In this case, the algorithm terminates, and the ED-ViT framework re-prunes the sub-models based on a new head pruning parameter, as described in Algorithm 1. Finally, the algorithm outputs the model assignment plan MA , representing the mapping of sub-models to edge devices.

As described in Section III, the problem of Vision Transformer sub-model partitioning and assignment can be formulated as a 0-1 knapsack problem, where each edge device has varying available memory and energy. Each sub-model is responsible for a specific set of classes, and multiple sub-models can be deployed on a single device. We perform a col-

laborative optimization of partitioning the Vision Transformer model into multiple sub-models, as shown in Algorithm 1, and deploying these sub-models across edge devices using a greedy search assignment mechanism, as shown in Algorithm 3. This approach provides a relatively optimal solution to the formulated problem. Our extensive experiments demonstrate the effectiveness of our framework design and algorithms.

E. Model Fusion

In the result fusion phase, each sub-model on the edge devices processes inputs and extracts corresponding features. The aggregation edge device aggregates the generated features through concatenation and feeds them into an MLP to produce the final prediction. Notably, the MLP for result fusion requires training only once after all sub-models have been trained.

In our paper, we utilize a tower-structured MLP to process the concatenated tensors received from the various edge devices. Specifically, each transmitted tensor from a device is integrated using a $N \times d \times s \rightarrow \lambda \times N \times d \times s \rightarrow numcls$ MLP structure, where λ is the shrinking hyperparameter and the default value is 0.5, $numcls$ is the number of classes. By utilizing a compact MLP model, we effectively fuse the distributed inference results from the sub-models while consuming only a minimal amount of computational resources.

V. EXPERIMENTS

A. Experiments settings

Datasets. Considering the versatile applicability of the framework across various scenarios, we select three computer vision datasets (i.e., CIFAR-10 [55], MNIST [56], and Caltech256 [57]) and two audio recognition datasets (i.e., GTZAN [58] and Speech Command [59]) to construct the classification tasks for our experiments. For all the computer vision datasets, we resize the sample to $224 \times 224 \times 3$ to support various datasets and downstream tasks via a similar data structure without loss of generality; for the audio recognition datasets, we resize the sample to $224 \times 224 \times 1$ with the same aim.

Implementation Details: All models are implemented using Pytorch [60]. During the training process, we use the Adam optimizer [61] with a decaying learning rate initialized to $1e-4$, and we set the batch size to 256. For the computer vision task, the original Vision Transformer model is pre-trained on the ImageNet dataset [62], followed by fine-tuning the task-specific data for 10 epochs. For the audio recognition task, Vision Transformer is pre-trained on the AudioSet dataset [63] and then fine-tuned on the task data for about 20 epochs. All the experimental results are averaged over five trial runs. We use the server with $8 \times$ NVIDIA A100 GPUs to generate sub-models and the fusion model. Each inference trial is conducted on 1 Raspberry Pi-4B for fusion and 1 to 10 Raspberry Pi-4B devices for sub-models, which serve as the edge devices for evaluating the execution time of processing a single sample on a specific sub-model. The edge devices are all connected with a gigabyte switch S1720-52GWR-PWR-4P as shown in Fig. 3. For bandwidth control, we use the traffic control tool



Fig. 3: Our 5-device example experimental prototype utilizes a switch and Raspberry Pi 4B devices, with one dedicated to the fusion model and the other four allocated to sub-models.

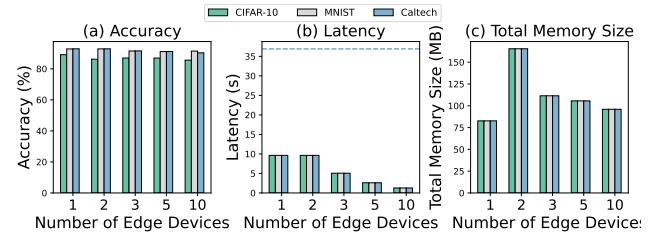


Fig. 4: Performance metrics of Split ViT-Base models on CIFAR-10, MNIST, Caltech dataset. Note that (a) shows the accuracy results; (b) shows the latency results, the dotted lines represent the latency of the original ViT-Base model, and (c) shows the total memory sizes for all the sub-models. All the experiment results are collected on Raspberry Pi-4B.

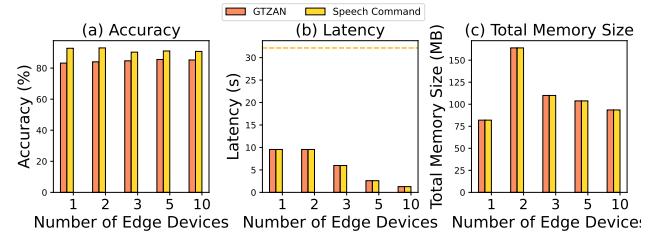


Fig. 5: Performance metrics of Split ViT-Base models on GTZAN and Speech Command dataset.

tc [64], which is able to limit the bandwidth under the setting value. The maximum bandwidth between devices is capped at 2 Mbps to simulate real-world scenarios.

B. Experiments on Computer Vision Datasets

We evaluate our approach using CIFAR-10, MNIST, and Caltech image datasets. The original model size is 327.38 MB. Fig. 4 shows the accuracy, inference latency, and memory usage of the ViT-Base model under the ED-ViT framework,

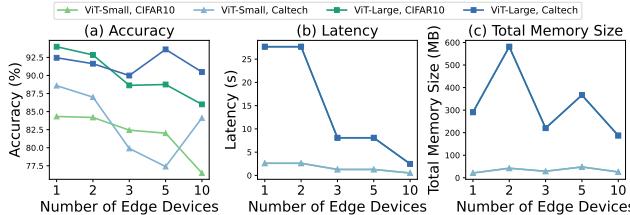


Fig. 6: Performance metrics of Split ViT-Small and ViT-Large models on CIFAR-10, Caltech dataset.

with 1 to 10 edge devices. With only one edge device, we apply model compression by pruning Vision Transformer without decomposition. All experiments are conducted with a total memory budget of 180MB across devices, ensuring fair comparisons.

The results demonstrate that as the number of edge devices increases, the accuracy remains largely consistent and yields strong performance. For CIFAR-10, accuracies are consistently above 85%; for MNIST, they are above 91%; and for Caltech, they exceed 90%. In most cases, the variance in final fusion prediction accuracy is less than one percentage point. The inclusion of more sub-models illustrates the feasibility of deploying larger-scale models without significant accuracy loss. As the number of edge devices increases, the inference latency decreases, as each sub-model is responsible for fewer classes and contains fewer parameters. Notably, the latency for the original model is 36.94 seconds on the CIFAR-10 dataset, which is 28.9 times the smallest latency (1.28s) and 3.84 times the highest latency (9.63s). Our ED-ViT could make multiple edge devices work collaboratively to maintain accuracy while lowering the storage burden and inference time as the number of edge devices increases. The results for other datasets show a similar trend as the model structures are the same.

In terms of total memory usage, ED-ViT provides effective splitting and assignment strategies. Note that as the number of retained heads increases, the memory size grows quadratically. For one edge device, retaining more heads could exceed the budget. However, in a two-device setting, each sub-model retains a similar number of heads, ensuring the total memory usage remains within the budget. This explains the spike in total memory sizes with two edge devices, as shown in Fig. 4. As the number of edge devices increases from 3 to 10, the memory size of each sub-model decreases, reducing computation overhead and demonstrating that many complex model designs and computational operations are redundant for problem-solving. In the 10-edge device setting, the model size on the CIFAR-10 dataset is reduced to just 9.60MB, achieving a size reduction of up to 34.1 times, compared to the original model from ED-ViT.

C. Experiments on Audio Recognition Datasets

We use the GTZAN and Speech Command audio datasets to evaluate the performance of our framework. The original model sizes of Vision Transformer for GTZAN and Speech

TABLE II: The FLOPs for sub-models on different datasets when using ViT-Base.

Dataset	The Number of Edge Devices				
	Original	2	3	5	10
CIFAR-10	16.86G	4.25G	1.90G	1.08G	0.48G
GTZAN	16.79G	4.20G	1.88G	1.059G	0.46G

Command is 325.88MB. Fig. 5 presents the accuracy, inference latency, and total memory size of the ViT-Base model as implemented by the ED-ViT framework, similar to the experiments with the computer vision datasets. We still set the memory budget to 180MB.

The results show that as the number of edge devices increases, ED-ViT is able to maintain the accuracy, delivering robust performance. For GTZAN, accuracies are consistently above 84%, and for the Speech Command dataset, accuracies are above 90%. Similar to the results on the computer vision datasets, the inference latency decreases as the number of edge devices increases. Notably, the latency for the original model is 32.16 seconds on the GTZAN dataset, which is 25.13 times the smallest latency (1.28s) and 3.37 times the highest latency (9.55s). This substantial latency reduction trend is consistent across both datasets. Regarding total memory size, all configurations remain within the set limits. As the number of edge devices increases, the memory size of each sub-model decreases, and the computation overhead is similarly reduced. In the 10-edge devices setting, for each model in the GTZAN dataset, the size is reduced to only 9.35MB, achieving a reduction of up to 34.85 times compared to the original model. Similar results are also observed in the Speech Command dataset from ED-ViT.

D. Overhead of Computation and Communication

We use FLOPs to simulate the energy consumption on the edge devices. Table II presents the FLOPs for each edge device on CIFAR (computer vision/video) and GTZAN (audio recognition) datasets with different numbers of edge devices using the ViT-Base model. The FLOPs for the original model on the CIFAR-10 and GTZAN datasets are 16.86G and 16.79G, respectively. As the number of edge devices increases, the FLOPs decrease across all datasets, and the experimental results are consistent with the parameter counts. These findings demonstrate that ED-ViT significantly reduces computation overhead and saves energy for the edge devices.

When the number of edge devices increases from 1 to 10 using ViT-Base across all the datasets, the size of features for communication on each sub-model decreases from 1536 bytes to 512 bytes. Compared with the original image size (150528 bytes), our method could greatly reduce the communication overhead to 294 times. The maximal communication time for one edge device is 5.86ms, which is acceptable in the practical situation. The results also show that the inferences on sub-models and the fusion model take up most of the latency (order of seconds) in Section V-B and Section V-C.

TABLE III: The accuracy results of splitting CNN and SNN versus ED-ViT with ViT-Base on CIFAR-10 dataset.

Methods	The Number of Edge Devices				
	1	2	3	5	10
Split-CNN	85.05 \pm 0.32	85.11 \pm 0.30	85.17 \pm 0.27	85.33 \pm 0.14	85.31 \pm 0.29
Split-SNN	83.56 \pm 0.01	82.45 \pm 0.12	83.01 \pm 0.79	83.06 \pm 0.50	82.29 \pm 0.32
ED-ViT	89.11 \pm 0.40	86.18 \pm 0.18	86.97 \pm 0.21	86.94 \pm 0.19	85.59 \pm 0.33

E. Experiments on Different Model Structures

We also select two complex datasets (e.g., CIFAR-10 and Caltech) to test different Vision Transformer structures for low-power video analytics tasks. The original model sizes of ViT-Small and ViT-Large are 82.71MB, 1,157MB, respectively. Fig. 6 presents the accuracy, inference latency, and total memory size of the ViT-Small and ViT-Large models as implemented by the ED-ViT framework, similar to Fig. 4. We increase the total memory size limit for ViT-Large to 600MB and decrease the limit for ViT-Small to 50MB.

The results show that as the number of devices increases, the accuracy remains relatively consistent, again showing robust performance. For ViT-Small, the accuracy is over 76.5% on the CIFAR-10 dataset and over 77.39% on Caltech across all settings; for ViT-Large, the accuracy is over 86% on the CIFAR-10 dataset and over 90.48% on Caltech in all settings. In most cases, the accuracy fluctuation for the final fusion prediction remains within a variance of less than one percentage point. The accuracy for ViT-Small is lower than that of ViT-Base, while ViT-Large achieves higher accuracy than ViT-Base, corresponding to the difference in parameter counts. Generally, the more parameters, the better the accuracy. As the number of edge devices increases, the latency decreases for both settings, similar to ViT-Base. The latency for ViT-Small is lower than that of ViT-Base, as ViT-Small requires less computational power, while the latency for ViT-Large is higher due to its larger size. In terms of memory size, in the 10-edge device setting, for each model on the CIFAR-10 dataset, the size for ViT-Small is 2.58MB, achieving a reduction of up to 32.06 times compared to the original model. Similarly, for ViT-Large, the size is 18.73MB, which also achieves a 61.77-fold reduction compared to the original model size.

Note that for the ViT-Small on the CIFAR-10 and CalTech, the input size and the output size are the same; thus, their latency and total memory size on the edge devices are also the same. Similar results are observed across both datasets for ViT-Small and ViT-Large.

F. Comparison with Baseline Methods: Split-CNN and Split-SNN

Vision Transformer achieves better accuracy compared to traditional CNN and SNN models. However, the performance of these models on edge devices has not been directly compared before. Nnfacet [2] proposes a method to split CNNs across multiple edge devices, employing a filter pruning technique [65], which differs from our approach. EC-SNN [3] utilize the convolutional spiking neural network (CSNN) [7]

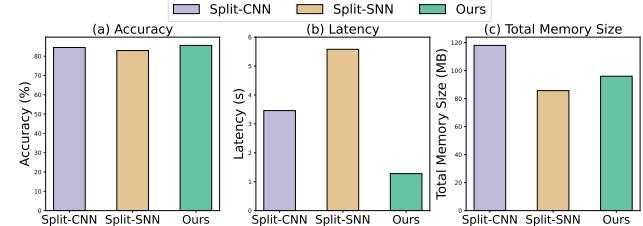


Fig. 7: Performance of splitting method with CNN, SNN, and ViT-Base models on CIFAR-10 dataset with 10 edge devices.

TABLE IV: The impact of retraining for CIFAR-10 dataset on ViT-Base of ED-ViT.

Methods	The Number of Edge Devices				
	1	2	3	5	10
ED-ViT	89.11 \pm 0.02	86.18 \pm 0.34	86.97 \pm 0.70	86.94 \pm 0.10	85.59 \pm 0.24
(w/o) retrain	88.25 \pm 0.30	86.00 \pm 0.45	86.08 \pm 0.03	85.33 \pm 0.11	84.20 \pm 0.14
(w/) entire retrain	89.11 \pm 0.30	92.33 \pm 0.11	91.14 \pm 0.45	89.97 \pm 0.62	90.26 \pm 0.07

to transform CNNs into SNNs, using a similar strategy. Both methods focus on VGGNet [5] backbone networks and are channel-wise methods. In our experiments, the baseline model for these methods is VGGNet-16 in their papers, which also has a memory size similar to ViT-Base and achieves the best original results for comparison. We follow the hyper-parameters in their papers to conduct the experiments.

The accuracy results on the CIFAR-10 dataset are presented in Table III. Based on the results, we observe that CNN outperforms SNN, while our ED-ViT method for ViT-Base yields better accuracy than both CNN and SNN approaches. The original accuracies of ViT-Base, CNN, and SNN are 98.12%, 93.64%, and 93.56%, respectively. Their average accuracy losses are 11.16%, 8.5% and 10.15% across various device numbers. Due to the inherently high accuracy of Vision Transformer, we admit that its accuracy drop is slightly higher than CNN and SNN. This is precisely why we leverage Vision Transformer: to employ its exceptional performance. Our method maintains a comparable accuracy drop with a 28.9 \times size reduction, achieving up to 4.06% and 5.55% higher accuracy than state-of-the-art CNN-based and SNN-based methods shown in Table III.

In addition to the accuracy results, we also compare inference latency and total memory size of the three methods when the number of edge devices is 10. These results are shown in Fig. 7. Based on the results, our ED-ViT method achieves the best accuracy compared to SNN and CNN, while its inference latency is much lower than SNN (4.36 times) and CNN (2.70 times). Furthermore, the total memory size of ED-ViT is significantly lower than CNN and is comparable to SNN, since SNN is known for its small model size. This experiment demonstrates that deploying Vision Transformer onto edge devices can meet latency and memory constraints while delivering superior accuracy results.

G. Experiments on Effects for Retraining

As we quantify model accuracy, we perform an ablation study to assess the impact of retraining. The results are shown in Table IV. The first line shows the results of the original ED-ViT. The second line shows the results from averaging the softmax output of sub-models without the fusion MLP. The third line shows the results based on the retraining of the overall models (sub-models and MLP together) for the fusion stage. When using only one device, the result is the same as the original ED-ViT as the training process remains unchanged in this scenario. Different from the work on splitting SNN and CNN, which are based on channel-wise methods and only get about 0.1% improvement in performance when retaining the overall models, our method is shown to have a great potential to improve performance (up to 6.15%). However, in the practical setting, it may be hard to retrain the sub-models with the fusion MLP.

VI. CONCLUSION

In this study, we are the first to propose a novel framework aimed at deploying Vision Transformer on edge devices, which combines model-partitioning and pruning. The formulation and resolution of the problem offer a viable solution, ED-ViT, which decomposes the Vision Transformer model into smaller sub-models and leverages the state-of-the-art pruning method to streamline the complex network architecture. ED-ViT not only preserves the essential structure of the original model but also enables more efficient inference, maintaining high system accuracy within the memory and energy constraints of edge devices. Extensive experiments and implementations have been conducted on five datasets, three ViT architectures, and two baseline methods, using three evaluation metrics of accuracy, inference latency, and total memory size. The results demonstrate that ED-ViT significantly reduces overall energy consumption and inference latency on edge devices while maintaining high accuracy. Our ED-ViT shows great potential for deployment on edge devices and for future integration with other horizontal methods to achieve better performance.

REFERENCES

- [1] J. Chen, D. Van Le, R. Tan, and D. Ho, “Split convolutional neural networks for distributed inference on concurrent iot sensors,” in *2021 IEEE 27th International Conference on Parallel and Distributed Systems (ICPADS)*. IEEE, 2021, pp. 66–73.
- [2] —, “Nnfacet: Splitting neural network for concurrent smart sensors,” *IEEE Transactions on Mobile Computing*, vol. 23, no. 2, pp. 1627–1640, 2023.
- [3] D. Yu, X. Du, L. Jiang, W. Tong, and S. Deng, “Ec-snn: Splitting deep spiking neural networks for edge devices,” *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 2024. [Online]. Available: <https://api.semanticscholar.org/CorpusID:271507864>
- [4] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” *Advances in neural information processing systems*, vol. 25, 2012.
- [5] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.
- [6] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2015. [Online]. Available: <https://api.semanticscholar.org/CorpusID:206594692>
- [7] S.-W. Deng and S. Gu, “Optimal conversion of conventional artificial neural networks to spiking neural networks,” *ArXiv*, vol. abs/2103.00476, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:232075977>
- [8] A. Vaswani, “Attention is all you need,” *Advances in Neural Information Processing Systems*, 2017.
- [9] D. Alexey, “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv: 2010.11929*, 2020.
- [10] Z. Wu, Z. Liu, J. Lin, Y. Lin, and S. Han, “Lite transformer with long-short range attention,” *arXiv preprint arXiv:2004.11886*, 2020.
- [11] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablayrolles, and H. Jégou, “Training data-efficient image transformers & distillation through attention,” in *International conference on machine learning*. PMLR, 2021, pp. 10 347–10 357.
- [12] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, “End-to-end object detection with transformers,” in *European conference on computer vision*. Springer, 2020, pp. 213–229.
- [13] Z. Dai, B. Cai, Y. Lin, and J. Chen, “Up-detr: Unsupervised pre-training for object detection with transformers,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 1601–1610.
- [14] F. Yang, Q. Zhai, X. Li, R. Huang, A. Luo, H. Cheng, and D.-P. Fan, “Uncertainty-guided transformer reasoning for camouflaged object detection,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 4146–4155.
- [15] Z. Song, F. Wu, X. Liu, J. Ke, N. Jing, and X. Liang, “Vr-dann: Real-time video recognition via decoder-assisted neural network acceleration,” in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 2020, pp. 698–710.
- [16] Y. Wang, Z. Xu, X. Wang, C. Shen, B. Cheng, H. Shen, and H. Xia, “End-to-end video instance segmentation with transformers,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 8741–8750.
- [17] R. Girdhar, J. Carreira, C. Doersch, and A. Zisserman, “Video action transformer network,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 244–253.
- [18] C. Plizzari, M. Cannici, and M. Matteucci, “Spatial temporal transformer network for skeleton-based action recognition,” in *Pattern recognition. ICPR international workshops and challenges: virtual event, January 10–15, 2021, Proceedings, Part III*. Springer, 2021, pp. 694–701.
- [19] Y. Gong, C.-I. Lai, Y.-A. Chung, and J. Glass, “Ssast: Self-supervised audio spectrogram transformer,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 10, 2022, pp. 10 699–10 709.
- [20] Z. Pan, B. Zhuang, J. Liu, H. He, and J. Cai, “Scalable vision transformers with hierarchical pooling,” *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 367–376, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:232290833>
- [21] B. Graham, A. El-Nouby, H. Touvron, P. Stock, A. Joulin, H. J’egou, and M. Douze, “Levit: a vision transformer in convnet’s clothing for faster inference,” *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12 239–12 249, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:233004577>
- [22] P. Zhang, X. Dai, J. Yang, B. Xiao, L. Yuan, L. Zhang, and J. Gao, “Multi-scale vision longformer: A new vision transformer for high-resolution image encoding,” *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 2978–2988, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:232404731>
- [23] W. Yu, M. Luo, P. Zhou, C. Si, Y. Zhou, X. Wang, J. Feng, and S. Yan, “Metaformer is actually what you need for vision,” *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10 809–10 819, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:244478080>
- [24] C. Yang, Y. Wang, J. Zhang, H. Zhang, Z. Wei, Z. L. Lin, and A. L. Yuille, “Lite vision transformer with enhanced self-attention,” *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11 988–11 998, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:245353696>
- [25] Z. Tu, H. Talebi, H. Zhang, F. Yang, P. Milanfar, A. C. Bovik, and Y. Li, “Maxvit: Multi-axis vision transformer,” in *European Conference on Computer Vision*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:247939839>
- [26] T. Yao, Y. Li, Y. Pan, Y. Wang, X. Zhang, and T. Mei, “Dual vision transformer,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*

Intelligence, vol. 45, pp. 10 870–10 882, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:250425982>

[27] X. Pan, T. Ye, Z. Xia, S. Song, and G. Huang, “Slide-transformer: Hierarchical vision transformer with local self-attention,” 2023 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2082–2091, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:258048654>

[28] K. Yuan, S. Guo, Z. Liu, A. Zhou, F. Yu, and W. Wu, “Incorporating convolution designs into visual transformers,” 2021 *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 559–568, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:232307700>

[29] Z. Dai, H. Liu, Q. V. Le, and M. Tan, “Coatnet: Marrying convolution and attention for all data sizes,” *ArXiv*, vol. abs/2106.04803, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:235376986>

[30] W. Wang, E. Xie, X. Li, D.-P. Fan, K. Song, D. Liang, T. Lu, P. Luo, and L. Shao, “Pyramid vision transformer: A versatile backbone for dense prediction without convolutions,” 2021 *IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 548–558, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:232035922>

[31] ———, “Pvt v2: Improved baselines with pyramid vision transformer,” *Computational Visual Media*, vol. 8, pp. 415 – 424, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:235652212>

[32] M. Zhu, Y. Tang, and K. Han, “Vision transformer pruning,” *arXiv preprint arXiv:2104.08500*, 2021.

[33] M. Xia, Z. Zhong, and D. Chen, “Structured pruning learns compact and accurate models,” *arXiv preprint arXiv:2204.00408*, 2022.

[34] Y. Liang, C. Ge, Z. Tong, Y. Song, J. Wang, and P. Xie, “Not all patches are what you need: Expediting vision transformers via token reorganizations,” *arXiv preprint arXiv:2202.07800*, 2022.

[35] Y. Liu, M. Gehrig, N. Messikommer, M. Cannici, and D. Scaramuzza, “Revisiting token pruning for object detection and instance segmentation,” 2024 *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 2646–2656, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:259138783>

[36] Q. Tang, B. Zhang, J. Liu, F. Liu, and Y. Liu, “Dynamic token pruning in plain vision transformers for semantic segmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 777–786.

[37] C. Zheng, Z. Li, K. Zhang, Z. Yang, W. Tan, J. Xiao, Y. Ren, and S. Pu, “Savit: Structure-aware vision transformer pruning via collaborative optimization,” in *Neural Information Processing Systems*, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:258509611>

[38] Z. Song, Y. Xu, Z. He, L. Jiang, N. Jing, and X. Liang, “Cp-vit: Cascade vision transformer pruning via progressive sparsity prediction,” *arXiv preprint arXiv:2203.04570*, 2022.

[39] Y. Xu, Z. Zhang, M. Zhang, K. Sheng, K. Li, W. Dong, L. Zhang, C. Xu, and X. Sun, “Evo-vit: Slow-fast token evolution for dynamic vision transformer,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 3, 2022, pp. 2964–2972.

[40] S. Venkataraman, A. Ghodrati, Y. M. Asano, F. Porikli, and A. Habibian, “Skip-attention: Improving vision transformers by paying less attention,” *arXiv preprint arXiv:2301.02240*, 2023.

[41] L. Yu and W. Xiang, “X-pruner: explainable pruning for vision transformers,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 24 355–24 363.

[42] H. Yu and J. Wu, “A unified pruning framework for vision transformers,” *Science China Information Sciences*, vol. 66, no. 7, p. 179101, 2023.

[43] J. Li, Q. Nie, W. Fu, Y. Lin, G. Tao, Y. Liu, and C. Wang, “Lors: Low-rank residual structure for parameter-efficient network stacking,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15 866–15 876.

[44] S. N. Wadekar and A. Chaurasia, “Mobilevit3: Mobile-friendly vision transformer with simple and effective fusion of local, global and input features,” *arXiv preprint arXiv:2209.15159*, 2022.

[45] J. Pan, A. Bulat, F. Tan, X. Zhu, L. Dudziak, H. Li, G. Tzimiropoulos, and B. Martinez, “Edgevits: Competing light-weight cnns on mobile devices with vision transformers,” in *European Conference on Computer Vision*. Springer, 2022, pp. 294–311.

[46] G. Xu, Z. Hao, Y. Luo, H. Hu, J. An, and S. Mao, “Devit: Decomposing vision transformers for collaborative inference in edge devices,” *IEEE Transactions on Mobile Computing*, 2023.

[47] S. Oh, J. Park, S. Baek, H. Nam, P. Vepakomma, R. Raskar, M. Bennis, and S.-L. Kim, “Differentially private cutmix for split learning with vision transformer,” *ArXiv*, vol. abs/2210.15986, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:253224400>

[48] F. Almalik, N. Alkhunaizi, I. Almakky, and K. Nandakumar, “Fesvibs: Federated split learning of vision transformer with block sampling,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:259251916>

[49] S. Oh, S. Baek, J. Park, H. Nam, P. Vepakomma, R. Raskar, M. Bennis, and S.-L. Kim, “Privacy-preserving split learning with vision transformers using patch-wise random and noisy cutmix,” *ArXiv*, vol. abs/2408.01040, 2024. [Online]. Available: <https://api.semanticscholar.org/CorpusID:271693565>

[50] Z. Su, H. Zhang, J. Chen, L. Pang, C.-W. Ngo, and Y.-G. Jiang, “Adaptive split-fusion transformer,” 2023 *IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1169–1174, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:248392009>

[51] J. Kim, Y. Park, G. Kim, and S. J. Hwang, “Splitnet: Learning to semantically split deep networks for parameter reduction and model parallelization,” in *International Conference on Machine Learning*, 2017. [Online]. Available: <https://api.semanticscholar.org/CorpusID:12078675>

[52] A. Bakhtiarnia, N. Milo, Q. Zhang, D. Bajovi, and A. Iosifidis, “Dynamic split computing for efficient deep edge intelligence,” *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:248986420>

[53] X. Hou, Y. Guan, T. Han, and N. Zhang, “Distredge: Speeding up convolutional neural network inference on distributed edge devices,” 2022 *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pp. 1097–1107, 2022. [Online]. Available: <https://api.semanticscholar.org/CorpusID:246486210>

[54] V. Weaver, “Green machines - energy efficient machines,” https://web.eece.maine.edu/~vweaver/group/green_machines.html, 2024, accessed: 13-Sep-2024.

[55] A. Krizhevsky, “Learning multiple layers of features from tiny images,” 2009. [Online]. Available: <https://api.semanticscholar.org/CorpusID:18268744>

[56] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proc. IEEE*, vol. 86, pp. 2278–2324, 1998. [Online]. Available: <https://api.semanticscholar.org/CorpusID:14542261>

[57] G. Griffin, A. Holub, P. Perona *et al.*, “Caltech-256 object category dataset,” Technical Report 7694, California Institute of Technology Pasadena, Tech. Rep., 2007.

[58] G. Tzanetakis and P. Cook, “Musical genre classification of audio signals,” *IEEE Transactions on speech and audio processing*, vol. 10, no. 5, pp. 293–302, 2002.

[59] P. Warden, “Speech commands: A dataset for limited-vocabulary speech recognition,” *arXiv preprint arXiv:1804.03209*, 2018.

[60] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” *Advances in neural information processing systems*, vol. 32, 2019.

[61] P. K. Diederik, “Adam: A method for stochastic optimization,” (*No Title*), 2014.

[62] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.

[63] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 776–780.

[64] T. L. Foundation, “Tc-show / manipulate traffic control settings,” <https://www.linux.com/tutorials/tc-show-manipulate-traffic-control-settings/>, 2022, [Online; accessed 10-October-2023].

[65] H. Hu, R. Peng, Y.-W. Tai, and C.-K. Tang, “Network trimming: A data-driven neuron pruning approach towards efficient deep architectures,” *ArXiv*, vol. abs/1607.03250, 2016. [Online]. Available: <https://api.semanticscholar.org/CorpusID:2493219>