

This CVPR Workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Block Selective Reprogramming for On-device Training of Vision Transformers

Sreetama Sarkar¹ Souvik Kundu² Kai Zheng¹ Peter A. Beerel¹ ¹Universiy of Southern California, Los Angeles, USA ²Intel Labs, San Diego, USA

{sreetama, kzheng44, pabeerel}@usc.edu

Abstract

The ubiquity of vision transformers (ViTs) for various edge applications, including personalized learning, has created the demand for on-device fine-tuning. However, training with the limited memory and computation power of edge devices remains a significant challenge. In particular, the memory required for training is much higher than that needed for inference, primarily due to the need to store activations across all layers in order to compute the gradients needed for weight updates. Previous works have explored reducing this memory requirement via frozen-weight training as well storing the activations in a compressed format. However, these methods are deemed inefficient due to their inability to provide training or inference speedup. In this paper, we first investigate the limitations of existing on-device training methods aimed at reducing memory and compute requirements. We then present block selective reprogramming (BSR) in which we fine-tune only a fraction of total blocks of a pre-trained model and selectively drop tokens based on self-attention scores of the frozen layers. To show the efficacy of BSR, we present extensive evaluations on ViT-B and DeiT-S with five different datasets. Compared to the existing alternatives, our approach simultaneously reduces training memory by up to $1.4 \times$ and compute cost by up to $2 \times$ while maintaining similar accuracy. We also showcase results for Mixture-of-Expert (MoE) models, demonstrating the effectiveness of our approach in multitask learning scenarios. Code will be available at: https://github.com/sreetamasarkar/BSR.

1. Introduction

Over the past several years, there has been an unprecedented growth in the deployment of deep learning applications on edge devices. Mutitask learning (MTL) is gaining momentum in edge applications due to their ability to dynamically adapt to different tasks with minimum overhead. Currently, edge devices primarily handle inference, while the training or fine-tuning processes take place in the cloud. This not only involves substantial communication overhead for





Figure 1. Test accuracy, training time, and memory comparison for a ViT-B on CIFAR-10 with a batch size of 32. In particular, we achieve a $5.47 \times$ memory reduction and $2.43 \times$ training speedup while yielding similar test accuracy. Our benefits are even more significant for higher batch sizes.

transferring both the model and data, but also raises concerns about data privacy [28]. Although on-device training is preferred, memory limitations of edge devices pose challenges. Naive solutions include directly using pre-trained models or fine-tuning only the last layer. However, these approaches can lead to a notable drop in accuracy if the distribution of the new data differs significantly from the pre-training data.

Previous works on efficient on-device learning [3, 22] have pointed out that the training memory is primarily dominated by activations, rather than parameters. In order to reduce activation memory, training residual modules while keeping the original backbone network frozen, has been widely explored in convolutional neural networks (CNNs) [3, 44]. However, this approach is not effective in vision transformers (ViTs) due to the presence of non-linear layers such as self-attention, softmax, and GELU, that require storing the input activation for gradient computation, even with frozen weights. Another line of research has explored reducing the memory cost of ViTs by introducing irregular sparsity in activation tensor [22]. However, such approaches might not yield significant memory savings due to the overhead of sparse representations and do not provide any potential speedup.

In contrast, this work presents a memory and parameterefficient fine-tuning approach called block selective reprogramming (BSR) that fine-tunes only a fraction of the pretrained weight tensors such that the activation memory requirement to propagate the gradients is minimized.

Our Contributions We first investigate and identify the limitations of existing SoTA methods applied to CNNs in the context of ViTs, as well as the shortcomings of existing on-device ViT fine-tuning approaches. In particular, we observe that residual fine-tuning does not necessarily reduce activation memory and can suffer from increased computational overhead.

Based on our observation, we propose BSR where we selectively fine-tune a small fraction of blocks of a pre-trained network. Additionally, to reduce the compute cost, we employ a token-dropping method that drops tokens primarily based on the self-attention score of the frozen layers. This saves activation memory and yields training speedup, unlike activation sparsification [22].

We present a detailed empirical evaluation of the impact of the block selection choice and token drop locations on accuracy. We demonstrate the efficacy of BSR on two different transformer models with five different datasets. Specifically, compared to the existing alternatives, our approach yields additional training memory saving of up to $1.4\times$ while saving compute cost by up to $2\times$. Compared to fully fine-tuning an ImageNet pre-trained ViT-B [13] model on CIFAR-10 (Figure 1), our approach provides $5.47\times$ training memory reduction along with $2.43\times$ reduction in training time for a batch size of 32.

Finally, we consider multi-task learning (MTL) as a target for our fine-tuning approach because of the benefits of maintaining a single model across different tasks in edge applications. In particular, mixture-of-expert (MoE) models [8, 31] are becoming popular in MTL due to their ability to disentangle the parameter space yielding improved MTL performance. We achieve close to baseline accuracy for fine-tuning MoE using BSR while reducing training memory and FLOPs by $2.3 \times$ and $1.5 \times$ respectively, demonstrating the effectiveness of our approach in MTL scenarios.

2. Related Work

On-device Training TinyTL [3] was one of the first papers to show that storing activation, not parameters, is the primary bottleneck for on-device training. Therefore, model compression techniques like weight pruning [15, 17, 19, 20, 24, 25, 34, 40], quantization [11, 18, 27], and knowledge distillation [26, 41] do not significantly reduce the cost of on-device training. Parameter efficient training methods, like Bitfit [47] or training only Batch Normalization parameters [21], reduces parameter memory, which is only a small fraction of the total train memory [3, 22]. Recomputing discarded activations during back-propagation [6, 16] reduces memory cost at the expense of large computational overhead, making them unsuitable for edge devices. TinyTL proposes training only biases with frozen weights. To compensate for reduced learning capacity, they introduce resid-

ual blocks, which are much smaller in width and employ group convolution to reduce memory footprints. Rep-Net [44] improves on this residual learning approach by designing six small residual modules which exchange features intermittently with a frozen pre-trained network, achieving a lower memory footprint and improved accuracy. However, introducing residual blocks leads to an increase in parameters, as well as both training and inference compute. Other methods for activation memory reduction include activation quantization [14, 27, 36] and pruning [4, 22, 25]. Back-Razor [22] proposes pruning the activation stored for backpropagation after the forward pass. They sparsify activations up to 95%, yielding a $5.5 \times$ reduction in training memory for ViTs. Their training memory reduction is limited by the memory required to save a binary mask of the sparse activations and hence, cannot achieve memory reduction proportional to activation sparsity. Notably, none of the current on-device learning methods provide speedup or computational efficiency during inference.

Efficient Vision Transformers Design for efficient ViTs is an active area of research and can be partitioned into two broad categories, namely architecture design for efficient ViTs and optimization of the operations for ViTs. In the first category, researchers have explored different self-attention alternatives including Linformer [43], SAL-ViT [48], and Performer [10], that targeted reducing its quadratic compute complexity. Other works presented resource and latency friendly efficient architectures including Mobile-former [9] and Efficientformer [30]. However, these works do not target on-device fine-tuning. For the second category, researchers have developed various compression [29, 45, 46], token dropping [33], and token merging [1] schemes to yield compute efficiency which are designed to provide inference speedup. However, despite reducing compute and often latency, such methods alone cannot yield the desirable reduction in activation and gradient storage necessary for on-device fine-tuning. In this paper, we leverage token dropping for on-device fine-tuning with resource-constrained memory and compute budget.

3. Motivational Analysis

3.1. Preliminaries

ViT models partition the input image into N patches, also called tokens, and embed each token into an embedding vector of length L. An extra classification token cls_token is added to the set of image tokens, creating N + 1 tokens. The cls_token is responsible for aggregating global image information. A trainable positional embedding is added to each of the token embeddings, which are then passed through a series of transformer encoder blocks, each consisting of a multi-head self-attention (MHSA) layer fol-

lowed by a feed forward network (FFN). The tokens are mapped into Query (Q), Key (K), and Value (V) matrices, each having a dimension of [N + 1, d], where d = L/H and H is the number of heads in the MHSA. Each head performs the self-attention operation

Attention
$$(\boldsymbol{Q}, \boldsymbol{K}, \boldsymbol{V}) = \text{Softmax}(\frac{\boldsymbol{Q}\boldsymbol{K}^T}{\sqrt{d}})\boldsymbol{V}.$$
 (1)

The FFN is usually a 2-layer network with GELU activation and a hidden dimension denoted D_{ffn} .

3.2. Analyzing Memory Cost

We perform a detailed analysis of the activation memory required for gradient computation in a ViT block. The gradient for back-propagation for linear and non-linear layers are given by Equation 2 and Equation 3 respectively, as follows.

$$\frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i}} = \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+1}} \frac{\partial \mathbf{a}_{i+1}}{\partial \mathbf{a}_{i}} = \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+1}} \mathbf{w}_{i}^{T} \text{ for linear layers}$$
(2)
$$= \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+1}} \mathbf{h} (\mathbf{a}_{i}) \text{ for non-linear layers}$$
(3)

Here, the input and output activations for the i^{th} layer are denoted by \mathbf{a}_i and \mathbf{a}_{i+1} . For linear layers, $\mathbf{a}_{i+1} = \mathbf{w}_i \mathbf{a}_i + b$ and the gradient $\frac{\partial \mathbf{a}_{i+1}}{\partial \mathbf{a}_i}$ is independent of \mathbf{a}_i , whereas, for non-linear layers, the gradient is a function of \mathbf{a}_i , given by $\mathbf{h}(\mathbf{a}_i)$. The gradient for the weights are given by Equations 4 and 5 where g denotes the gradient function with respect to weights for non-linear layers, as follows.

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}_{i}} = \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+1}} \frac{\partial \mathbf{a}_{i+1}}{\partial \mathbf{w}_{i}} = \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+1}} \mathbf{a}_{i}^{T} \text{ for linear layers} \qquad (4)$$
$$= \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+1}} \mathbf{g}(\mathbf{a}_{i}) \text{ for non-linear layers} \qquad (5)$$

The gradients of weights are dependent on input activations, whether the layer is linear or non-linear. Thus, trainable blocks must store activations both for linear layers, including LayerNorm or fully-connected layers, and non-linear layers, including Softmax, GELU, and Attention. In contrast, frozen blocks situated in the gradient flow path only need to store activations for non-linear layers. The memory cost for the stored activations for non-linear layers in DeiT-S [39] are presented in Table 1. A detailed analysis of the gradient activation memory required by each block reveals that a frozen block stores ~2.9 MB whereas a fully trainable block stores ~5.5 MB, indicating that training with frozen weights reduces memory requirements by ~ 2×. From Table 1, it is clear that activation memory is largely dependent on the number of input tokens. Softmax input exhibits

| Module | Stored Activation | Dimension | Memory Cost (MB) |
|---------------|---|------------------------------|---------------------|
| Self-Attentio | on $oldsymbol{Q},oldsymbol{K},oldsymbol{V}$ | [B, H, N+1, L/H] | 0.87 |
| Softmax | $oldsymbol{Q}oldsymbol{K}^T$ | [B, H, N+1, N+1] | 0.89 |
| GELU | | $[B, N+1, L \times D_{ffn}]$ | 1.15 |

Table 1. Memory cost for DeiT-S ($patch_size = 16, N = 196, L = 384, H = 6, D_{ffn} = 4$) model with a batch size B=1 and input image dimension of 224×224

quadratic dependence, while all saved activations exhibit linear dependence. With this motivation, we employ an efficient token dropping mechanism, where we selectively drop uninformative tokens, achieving over linear reduction in activation memory without sacrificing accuracy.

4. Proposed Approach

4.1. Block Reprogramming in ViTs

To improve the SoTA in fine-tuning ViTs, we propose Block Selective Reprogramming (BSR) which selectively trains a small fraction of blocks of a pre-trained model coupled with token dropping. Token-dropping approaches have been widely studied in ViTs for latency and energy improvement [1, 33]. However, token dropping has not been used in the context of activation memory reduction for on-device training. In this work, we couple token dropping with frozenweight training, obtaining activation memory reduction up to $6\times$. Our token-dropping method is inspired by EViT [32]. We calculate token importance using self-attention scores and fuse low importance tokens based on a token drop rate. Thus, fewer tokens are passed on to the trainable blocks, significantly reducing activation memory.

The token importance is calculated by the MHSA module of a block using the attention from the classification token to all other tokens, given by a score S_{token} as follows

$$S_{token} = rac{q_{class}K^T}{\sqrt{d}}.$$
 (6)

The first row of the attention map QK^T is the dot product between the query vector obtained from classification token (q_{class}) and key matrix. We use this dot product to calculate token importance because tokens in the value matrix are linearly combined according to these scores (S_{token}) to predict the output class. The top-K tokens according to the importance score S_{token} are preserved, K being determined by the token drop rate, while the unimportant ones are fused and passed onto the subsequent FFN layer. The tokens are incrementally discarded across the length of the network through blocks termed token drop locations. Since the classification token shows an exactly identical distribution for a pre-trained and fully fine-tuned model 4, we leverage the pre-trained classification token for importance calculation.



Figure 2. Block selective reprogramming framework for a ViT model with 12 layers. The token drop locations are set at 4^{th} , 7^{th} , and 10^{th} blocks, where 50% of incoming tokens are dropped based on token importance calculated by the MHSA module. The 4^{th} , 8^{th} , and 12^{th} blocks along with the classification head are trainable. The gradient propagation does not occur beyond the last trainable block.

This alleviates the need to transmit gradients all the way back to the start of the network, thereby saving memory.

4.2. Understanding the Gradient Flow

In this section, we derive the gradient flow for residual learning and BSR, and justify why we choose to finetune parts of the main network instead of inserting residual blocks. We further argue that activation memory reduction cannot be achieved due to an additive relation between trainable and frozen blocks in the context of ViTs, contrary to Rep-Net [44].



Figure 3. *Left:* Gradient flow in a residual learning architecture like [44] *Right:* Gradient flow in our block selective reprogramming approach

In Figure 3, on the left, we consider a residual learning network like Rep-net, where the frozen i^{th} main block, modeled as f_{M_i} exchanges features with a trainable side block f_{S_i} . The input and output activations of the i^{th} blocks are denoted by \mathbf{a}_i and \mathbf{a}_{i+1} . The individual outputs of the main and side blocks are denoted by \mathbf{a}_{M_i} and \mathbf{a}_{S_i} , where $\mathbf{a}_{M_i} = f_{M_i}(\mathbf{a}_i)$ and $\mathbf{a}_{S_i} = f_{S_i}(\mathbf{a}_i)$. The derivative of block f_k is given by f'_k . The derivative of the loss \mathcal{L} with respect to the trainable weights \mathbf{w}_{S_i} of the i^{th} side block S_i is given as follows,

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}_{S_{i}}} = \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+2}} \frac{\partial \mathbf{a}_{i+2}}{\partial \mathbf{w}_{S_{i}}}$$

$$= \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+2}} \frac{\partial}{\partial \mathbf{w}_{S_{i}}} \left(\mathbf{a}_{i+1}^{M} + \mathbf{a}_{i+1}^{S} \right)$$

$$= \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+2}} \frac{\partial}{\partial \mathbf{w}_{S_{i}}} \left(f_{M_{i+1}} \left(\mathbf{a}_{i+1} \right) + f_{S_{i+1}} \left(\mathbf{a}_{i+1} \right) \right)$$

$$= \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+2}} \frac{\partial}{\partial \mathbf{w}_{S_{i}}} \left(f_{M_{i+1}} \left(\mathbf{a}_{i}^{M} + \mathbf{a}_{i}^{S} \right) + f_{S_{i+1}} \left(\mathbf{a}_{i}^{M} + \mathbf{a}_{i}^{S} \right) \right)$$

$$= \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+2}} \frac{\partial \mathbf{a}_{i}^{S}}{\partial \mathbf{w}_{S_{i}}} \left(f'_{M_{i+1}} \left(\mathbf{a}_{i}^{M} + \mathbf{a}_{i}^{S} \right) + f'_{S_{i+1}} \left(\mathbf{a}_{i}^{M} + \mathbf{a}_{i}^{S} \right) \right)$$
(7)

where the last step of the derivation utilizes the fact that \mathbf{a}_i^M is independent of \mathbf{w}_{S_i} .

Equation 7 clearly shows that the gradient of weights of the trainable side block depends on the gradient through the subsequent main block $f'_{M_{i+1}}(\mathbf{a}_{i+1})$. Because the derivative of the loss depends on the input activation for the non-linear components, we still need to store these input activations along the main path.

The advantage of freezing weights in CNNs stems from the fact that the derivative of the loss for linear layers like convolution need only weights, and not their input activation, as given by Equation 2. ReLU, although non-linear, only needs to store a binary mask. Therefore, unlike the claims made by Rep-Net, the additive relationship between frozen and trainable blocks does not play a role in activation memory reduction. The memory reduction is obtained by reducing the activation dimension at the input of trainable blocks feeding downsampled versions of the activation to



Figure 4. Distribution of patch embedding, positional embedding and classification token of two DeiT-S: a pre-trained ImageNet model and a fully fine-tuned model on CIFAR-10

residual blocks. This advantage, however, is lost for ViTs where each encoder block consists of self-attention, softmax and GELU, each of which requires storing the activation input for gradient propagation (Equation 3). Therefore, introducing residual blocks further increases the activation memory, as demonstrated through our experiments detailed in Section 5.1. Moreover, in residual learning the forward pass occurs through both the frozen main blocks and the trainable residual blocks, resulting in an increase in parameters and compute costs, both during training and inference.

In Figure 3, on the right, we show the gradient flow in BSR, where the i^{th} main block is made trainable. The derivative of the loss \mathcal{L} with respect to the trainable weights \mathbf{w}_{M_i} of main block M_i is given as follows

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}_{M_i}} = \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+2}} \frac{\partial \mathbf{a}_{i+2}}{\partial \mathbf{w}_{M_i}} = \frac{\partial \mathcal{L}}{\partial \mathbf{a}_{i+2}} \frac{\partial \mathbf{a}_i^M}{\partial \mathbf{w}_{M_i}} \left(f'_{M_{i+1}}(\mathbf{a}_i^M) \right)$$
(8)

The lack of a residual path simplifies Equation 8 and implies only activations along the main path need to be stored.

4.3. Design Choice Discussion

Discussion 1 Which blocks are more essential during finetuning for transferring to small-scale datasets?

We perform an analysis of intermediate features between an ImageNet pre-trained model and a fully fine-tuned model on CIFAR-10 dataset to observe which layers show the maximum differences. The layer statistics for the three trainable elements before the encoder blocks, patch embedding, positional embedding, and classification token, have been presented in Figure 4. Surprisingly, we observe that these elements have an exactly identical distribution for the pre-trained and fine-tuned models. This is not obvious, particularly for the classification token, which is responsible for aggregating information for final classification. In general, fine-tuning the last layers of a model is more important [5, 12, 38], which is why in transfer learning applications, only the last layers are often fine-tuned. For CNNs, the first few layers or feature extraction layers are mostly similar across tasks. This brings us to our next discussion, namely, is only fine-tuning the last few layers sufficient in terms of accuracy and activation memory?

Discussion 2 Is fine-tuning the last few blocks enough?

Table 2 presents results for training the last two, three, and four blocks, with and without token pruning. We observe that training blocks only towards the end cannot close the accuracy gap, even without token pruning. With token pruning, because only a few tokens are left towards the end, we observe a significant drop in accuracy. This necessitates an intelligent selection of trainable blocks and token drop locations, balancing between accuracy and activation memory to achieve an optimal trade-off.

| Trainable | Token | Train | Reduce | Accuracy |
|-------------------------|----------|--------|------------|----------|
| Blocks | Dropping | Memory | Ratio | |
| Baseline | | 8649 | $1 \times$ | 98.48 |
| (10, 11) | × | 1477 | 5.8× | 96.14 |
| (10, 11) | ✓ | 215 | 40.2× | 93.93 |
| (9, 10, 11) | × | 2187 | 3.9× | 96.35 |
| (9, 10, 11) | ✓ | 333 | 25.9× | 94.81 |
| (8,9,10,11) (8,9,10,11) | × | 2896 | 2.9× | 96.81 |
| | ✓ | 502 | 17.2× | 95.23 |

Table 2. Training last few blocks of a DeiT-S model on CIFAR-10 Dataset with and without token pruning

Discussion 3 Relative Positioning of Token Drop Locations and Trainable Blocks The overall token drop rate of the network depends on the token drop locations and the fraction of tokens dropped in those locations. For example, if a higher token drop rate is chosen, placing token drop locations towards the end of the network will have the same effect as placing token drop locations in the shallower layers with a lower drop rate (Table 8). EViT [32] shows that the token dropping approach causes significant performance degradation when tokens are dropped before the 3^{rd} layer, since transformer models are not able to identify important tokens that early. They further observe that after the 3^{rd} layer, irrespective of drop location and drop rate, the network exhibits a stable performance for the same level of overall token reduction. We verify this observation in Section 5.3. We place token drop locations at the 4^{th} , 7^{th} , and 10^{th} blocks with a token drop rate of 0.5, resulting in an overall token reduction of 50%. With this token drop configuration, the position of trainable blocks is varied to find a suitable trade-off between accuracy and training memory. An ablation for trainable block positions is presented in Table 7. We make several interesting observations.

- Placing trainable blocks before token drop locations improves accuracy, but heavily costs training memory.
- Placing all trainable blocks after token drop locations yields significant memory advantages but is accompanied with a drop in accuracy.
- Training blocks at uniformly distributed depths performs better than solely training the last blocks.

Based on these observations, we keep 3 out of 12, i.e., the 4^{th} , 8^{th} and 12^{th} , blocks as trainable. The 4^{th} block is also a token drop location, where all N tokens are processed by the MHSA for token importance calculation, passing on N/2 tokens to the FFN layer. The 8^{th} and the 12^{th} blocks process only N/4 and N/8 tokens respectively.

5. Experimental Results

Models and Datasets We demonstrate results on DeiT-S [39] and ViT-B [13]. Fine-tuning is performed on models pre-trained on ImageNet1k for DeiT-S, and ImageNet22k for ViT-B. MoE models are constructed from an Imagenet pre-trained ViT-B backbone, with three MoE layers, each consisting of 16 experts. Transfer learning performance is shown on five datasets: CIFAR-10, CIFAR-100 [23], Flowers [35], Pets [37], and Food [2]. The estimated memory results are calculated following [3, 22] and the on-device memory is measured on NVIDIA RTX A6000 GPUs.

Training Hyperparameters ViT-B models are trained in Pytorch following the training settings in [22]. ViT-B models are trained using SGD optimizer for 20k steps with cosine learning rate decay and initial learning rate tuned for each dataset. ViT MoE models are trained using SGD for 8k steps with an initial learning rate of 0.01. DeiT-S models are trained using AdamW optimizer for 50 epochs using cosine learning rate decay. The default image size is set to 224 and default patch size is set as 16.

5.1. Results and Analysis

ViT-B and DeiT-S models for our approach are configured according to the settings described in Discussion 3 (Section 4.3) with an overall token drop rate of 50%. Table 3 presents results for DeiT-S on CIFAR-10 and CIFAR-100. Our approach achieves $6.03 \times$ training memory reduction, while showing accuracy degradation of 0.9% on CIFAR-10



Figure 5. Accuracy and on-device memory for DeiT-S and ViT-B models on CIFAR-100.

and 3.23% on CIFAR-100. We also achieve a FLOPS reduction of $\sim 2\times$.

| Method | Train | Reduce | FLOPS | CIFAR-10 | CIFAR-100 | |
|---------|---------|---------------|---------|----------|-----------|--|
| | Mem. | Ratio | (GMacs) | | | |
| FT-Full | 8649 MB | 1× | 589 | 98.48 | 88.79 | |
| FT-Last | 369 MB | $23.4 \times$ | 589 | 89.9 | 75.3 | |
| Ours | 1433 MB | 6.03× | 295 | 97.5 | 85.56 | |

Table 3. BSR for DeiT-S with batch size 128.

In Table 4, we compare full fine-tuning (FT-Full), finetuning only the last layer or the classification head (FT-Last), Back Razor [22], and our approach. The training memory, on-device memory, and FLOPS are calculated for CIFAR-100. We compare with Back Razor for 80% and 95% activation sparsity. FT-Last provides the highest memory benefits, while performing reasonably well on datasets like Flowers, but suffers substantial degradation for CIFAR-100 and Food. Some of our FT-Full accuracies are marginally lower than baseline accuracies reported in Back Razor. For example, for Flowers-102, our FT-Full accuracy is 99.2% as opposed to their reported accuracy of 99.5%. Both Back Razor and our method suffer 0.1% degradation in Flowers-102. For Pets, the accuracy of our method is at par with Back Razor@95%, although our FT-Full accuracy is lower by 0.4%, implying our method suffers 0.4% less degradation. Our approach outperforms Back Razor@95% for CIFAR-100 and Pets by 1.02% and 0.4%, and performs at par on Flowers dataset. We provide greater training memory reduction than Back Razor@95% while simultaneously achieving $2 \times$ FLOPS reduction.

On-device Memory The accuracy and on-device memory for FT-Full, FT-Last, and BSR on CIFAR-100 are presented in Figure 5. FT-Last provides the highest memory reduction but suffers substantial (\sim 10% on DeiT-S) accuracy degradation. Our approach provides on-device memory reduction of 4.8× for DeiT-S model and 4.4× for ViT-B and outperforms Back Razor by 1.4×.

Residual Learning for CNNs and ViTs We also investigate residual learning inspired from Rep-Net for ViTs.

| Method | Train Mem. (MB)↓ | Reduce . Ratio ↑ | On-device Mem. (MB) ↓ | Reduce , Ratio ↑ | FLOPS↓ (GMacs) | Reduce Ratio ↑ | CIFAR-10 | Accuracy CIFAR-100 | v(%)† Flowers | Pets Food |
|---|---------------------|---------------------------|--------------------------|---------------------|-------------------|---------------------|---------------------|-----------------------|---------------------|---------------------------------------|
| FT-Full FT-Last | 17388 525 | $1 \times$ $33 \times$ | 19098 1629 | 1× 11.7× | 2249 2249 | $1 \times 1 \times$ | 98.98 96.1 | 93.1 84.5 | 99.2 99.02 | 93.7 90.5 92.2 84.1 |
| BackRazor@80 % [22] BackRazor@95% [22] | 4565 3496 | 4.2× 5.5× | 8501 7189 | 2.66× 3.15× | 2249 2249 | $1 \times 1 \times$ | <u>98.9</u> 98.8 | <u>92.9</u> 90.0 | <u>99.4</u> 99.4 | $\frac{93.8}{92.4} \frac{90.5}{88.7}$ |
| Ours | 2938 | 5.92× | 4342 | 4.4 × | 1129 | 2 × | 98.3 | 91.02 | 99.1 | 92.4 88.4 |

Table 4. Comparison of BSR with other transfer learning approaches for ViT-B with batch size 128. The values for Back Razor are taken directly from the paper [22]. The training memory reduction ratios for Back Razor are calculated with respect to their reported baseline. The best accuracy other than FT-Full is underlined. Back Razor@80% gives the best accuracy with limited memory reduction. BSR yields the highest reduction in estimated and measured memory and FLOPS (highlighted in bold).

Three trainable side blocks are introduced while the pretrained backbone is kept frozen. The output activations of main and side blocks are added together, similar to Figure 3. BSR outperforms residual learning in accuracy, training memory and FLOPS, as shown in Table 5.

For testing BSR on CNNs, we train only the last blocks of layers 2, 3, 4 in ResNet-50. We outperform FT-Full and achieve similar performance as Rep-Net with lesser FLOPS and training memory.

| Model | Method | Accuracy ↑ (%) | Activation ↓ Mem.(MB) | FLOPS↓ (GMacs) |
|-----------|---------|----------------|--------------------------|-------------------|
| | FT-Full | 95.13 | 88.4 | 4.1 |
| ResNet-50 | Rep-Net | 96.37 | 9.91 | 6.2 |
| | Ours | 96.21 | 8.02 | 4.1 |
| | FT-Full | 98.48 | 66.9 | 4.6 |
| DeiT-S | Rep-Net | 97.81 | 40.1 | 5.7 |
| | Ours | 98.16 | 34.3 | 4.6 |

Table 5. Residual learning and BSR for CNNs and ViTs on CIFAR-10. The activation memory and FLOPS are reported for a batch size of 1.

5.2. Extending BSR for Mixture-of-Experts

Multi-task learning (MTL) [7, 31] has become pivotal for edge intelligence as it allows to adapt to tasks dynamically with minimum overhead. MTL models learn a shared representation for different tasks, thereby avoiding the overhead of training as well as storing separate models. Mixture-ofexperts (MoE) and its variants [8, 31, 42] constitute a new paradigm in MTL that separate the parameter space by activating parts of the model based on task and input tokens, providing improved MTL performance. In this section, we demonstrate that our approach can be easily extended to fine-tune the MoE models. To the best of our knowledge, no earlier work has demonstrated on-device training for MoE models. **Setup** For ViT MoE, the transformer block in ViT is replaced by a MoE layer. A MoE layer consists of several experts represented as MLPs with a shared MHSA module. A task-dependent router takes the MHSA output into account and sparsely activates a subset of experts for each input token. We use a MoE model with a ViT-B backbone. We only replace the FFN layers for the three blocks (corresponding to the location of trainable blocks in BSR) with MoE MLP layers. In our experiments, we use a single layer MLP as the router for each task for each MoE module, allowing trainability to meaningfully select the top-k experts. The MoE layers have 16 expert candidates, out of which top-4 candidates for each token are selected by the router, adding up the results to generate the output of the MoE layer. The results for BSR fine-tuning with MoE are presented in Table 6. Here, single task learning (STL) consists of separate models for each task. We observe that BSR performs remarkably well for MoE fine-tuning as well, causing negligible accuracy drop of only 0.21%, 0.78%, 0.09% and 0.78% over the MoE baseline for CIFAR-10, CIFAR-100, Flowers, and Food while outperforming the baseline by 1.04% for Pets. We obtain on-device memory reduction by $2.32 \times$ and FLOPs reduction by $1.5 \times$. Notably, our MoE models significantly outperform the STL models in most cases.

5.3. Ablation Studies

Number and Position of Trainable Blocks Table 7 presents an ablation on the number and position of trainable blocks. The indices runs from 0 to 11 corresponding to the model depth of twelve layers. Training with four blocks provides no significant advantages over three blocks. In fact, three blocks often performs better than four blocks, as seen for training blocks [4, 7, 11] vs [4, 7, 10, 11] or [2, 5, 8, 11] vs [2, 5, 8]. Training with two blocks causes much higher degradation. Therefore we choose the number of trainable blocks as three. The position of trainable blocks with respect to the token drop locations provides an accuracy memory trade-off, as discussed in Section 4.3.

| Method | Model | On-device | Reduce | FLOPS ↓ | Reduce | Accuracy(%)↑ | | | | |
|--------|----------|------------------------|------------------|----------------|---------------------|--------------|-----------|---------|-------|-------|
| | | Mem. (MB) \downarrow | Ratio \uparrow | (GMacs) | Ratio \uparrow | CIFAR-10 | CIFAR-100 | Flowers | Pets | Food |
| STL | Baseline | 19098 | 1× | 2249 | $1 \times 2 \times$ | 98.98 | 93.1 | 99.2 | 93.7 | 90.5 |
| STL | Ours | 4342 | 4.4 × | 1129 | | 98.3 | 91.02 | 99.1 | 92.4 | 88.4 |
| MoE | Baseline | 30722 | 1× | 1801 | 1× | 98.93 | 92.77 | 98.33 | 92.21 | 90.12 |
| MoE | Ours | 13244 | 2.32× | 1237 | 1.5× | 98.72 | 91.99 | 98.24 | 93.24 | 89.34 |

Table 6. BSR for MoE models with ViT-B backbone with batch size 128.

| #Trainable Blocks | Indices | Train Mem. (MB) | Reduce Ratio | Accuracy (%) |
|----------------------|----------------|--------------------|-----------------|-----------------|
| | [3, 6, 9, 11] | 1526 | $5.7 \times$ | 97.44 |
| 4 | [2, 5, 8, 11] | 2809 | $3.1 \times$ | 97.82 |
| 4 | [4, 7, 10, 11] | 1132 | $7.6 \times$ | 97.17 |
| | [8, 9, 10, 11] | 502 | $17.2 \times$ | 95.23 |
| | [0, 2, 5, 8] | 3705 | $2.3 \times$ | 97.75 |
| | [3, 7, 11] | 1433 | 6.03× | 97.43 |
| | [4, 7, 11] | 1079 | 8.0 	imes | 97.34 |
| 3 | [2, 7, 11] | 2731 | $3.2 \times$ | 97.81 |
| | [9, 10, 11] | 333 | $25.9 \times$ | 94.81 |
| | [2, 5, 8] | 2785 | $3.1 \times$ | 97.86 |
| | [4, 11] | 986 | 8.8 	imes | 96.80 |
| 2 | [7, 11] | 460 | $18.8 \times$ | 95.84 |
| | [10, 11] | 215 | $40.2 \times$ | 93.93 |

Table 7. Ablation study on the number and position of trainable blocks for DeiT-S on CIFAR-10. The token drop locations are kept frozen at the 4^{th} , 7^{th} , and 11^{th} blocks.

Token Drop Rate and Drop Locations Table 8 presents an ablation study on token drop rates and locations such that the overall drop rate is around 50%. We observe that for similar drops there is minor accuracy variation ($\sim 0.1\%$) while the memory and FLOPs reduction are also similar.

| Drop Rate | Drop Indices | Train Mem. (MB) | FLOPs (GMACs) | Accuracy (%) |
|--------------|-----------------|--------------------|------------------|-----------------|
| 0.3 | [1,3,5,7,9] | 1363 | 270 | 97.54 |
| 0.5 | [3,6,9] | 1433 | 295 | 97.42 |
| 0.7 | [5,7,9] | 1673 | 311 | 97.43 |

Table 8. Ablation study on the token drop rate and locations for DeiT-S on CIFAR-10. The trainable block positions are fixed at the 4^{th} , 8^{th} , and 12^{th} blocks.

6. Conclusions

We propose BSR, an approach for efficient on-device training of ViTs that couples selectively fine-tuning a small fraction of blocks of pre-trained models on downstream tasks with token dropping based on self-attention scores of primarily frozen weights. Unlike existing on-device learning approaches that reduce training memory at the cost of increased computational overhead, BSR reduces compute cost by $2\times$ while reducing training time by $\sim 2.5\times$ and memory up to $6.03\times$. We outperform SoTA approaches in on-device training memory reduction by $1.4\times$ while maintaining similar performance. We further demonstrate the effectiveness of the approach for MoE models in MTL applications. Our approach is orthogonal to activation quantization which can be used to obtain further reductions in activation memory.

References

- Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your vit but faster. *arXiv preprint arXiv:2210.09461*, 2022. 2, 3
- [2] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101-mining discriminative components with random forests. In *ECCV*, 2014. 6
- [3] Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. Tinytl: Reduce memory, not parameters for efficient on-device learning. Advances in Neural Information Processing Systems, 33:11285–11297, 2020. 1, 2, 6
- [4] Ayan Chakrabarti and Benjamin Moseley. Backprop with approximate activations for memory-efficient network training. In Advances in Neural Information Processing Systems. Curran Associates, Inc., 2019. 2
- [5] Ken Chatfield, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Return of the devil in the details: Delving deep into convolutional nets. In *BMVC*, 2014. 5
- [6] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. Training deep nets with sublinear memory cost. arXiv preprint arXiv:1604.06174, 2016. 2
- [7] Tianlong Chen, Xuxi Chen, Xianzhi Du, Abdullah Rashwan, Fan Yang, Huizhong Chen, Zhangyang Wang, and Yeqing Li. Adamv-moe: Adaptive multi-task vision mixture-ofexperts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17346–17357, 2023. 7
- [8] Tianlong Chen, Zhenyu Zhang, AJAY KUMAR JAISWAL, Shiwei Liu, and Zhangyang Wang. Sparse moe as the new dropout: Scaling dense and self-slimmable transformers. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 7

- [9] Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Xiaoyi Dong, Lu Yuan, and Zicheng Liu. Mobileformer: Bridging mobilenet and transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5270–5279, 2022. 2
- [10] Krzysztof Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020. 2
- [11] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *NeurIPS*, 2015. 2
- [12] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, 2014. 5
- [13] 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 *International Conference on Learning Representations*, 2021. 2, 6
- [14] R David Evans and Tor Aamodt. Ac-gc: Lossy activation compression with guaranteed convergence. Advances in Neural Information Processing Systems, 34:27434–27448, 2021. 2
- [15] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR*, 2019. 2
- [16] Audrunas Gruslys, Rémi Munos, Ivo Danihelka, Marc Lanctot, and Alex Graves. Memory-efficient backpropagation through time. In *NeurIPS*, 2016. 2
- [17] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *NeurIPS*, 2015. 2
- [18] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *ICLR*, 2016. 2
- [19] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *ICCV*, 2017.
- [20] Duc Hoang, Souvik Kundu, Shiwei Liu, Zhangyang Wang, et al. Don't just prune by magnitude! your mask topology is a secret weapon. Advances in Neural Information Processing Systems, 36:65056–65068, 2023. 2
- [21] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 2
- [22] Ziyu Jiang, Xuxi Chen, Xueqin Huang, Xianzhi Du, Denny Zhou, and Zhangyang Wang. Back razor: Memoryefficient transfer learning by self-sparsified backpropagation. *Advances in Neural Information Processing Systems*, 35: 29248–29261, 2022. 1, 2, 6, 7
- [23] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009. 6

- [24] Souvik Kundu, Mahdi Nazemi, Massoud Pedram, Keith M Chugg, and Peter A Beerel. Pre-defined sparsity for lowcomplexity convolutional neural networks. *IEEE Transactions on Computers*, 69(7):1045–1058, 2020. 2
- [25] Souvik Kundu, Mahdi Nazemi, Peter A Beerel, and Massoud Pedram. Dnr: A tunable robust pruning framework through dynamic network rewiring of dnns. In *Proceedings of the* 26th Asia and South Pacific Design Automation Conference, pages 344–350, 2021. 2
- [26] Souvik Kundu, Qirui Sun, Yao Fu, Massoud Pedram, and Peter Beerel. Analyzing the confidentiality of undistillable teachers in knowledge distillation. *Advances in Neural Information Processing Systems*, 34:9181–9192, 2021. 2
- [27] Souvik Kundu, Shikai Wang, Qirui Sun, Peter A Beerel, and Massoud Pedram. Bmpq: bit-gradient sensitivity-driven mixed-precision quantization of dnns from scratch. In 2022 Design, Automation & Test in Europe Conference & Exhibition (DATE), pages 588–591. IEEE, 2022. 2
- [28] Souvik Kundu, Shunlin Lu, Yuke Zhang, Jacqueline Liu, and Peter A Beerel. Learning to linearize deep neural networks for secure and efficient private inference. *ICLR*, 2023. 1
- [29] Souvik Kundu, Sharath Sridhar Nittur, Maciej Szankin, and Sairam Sundaresan. Sensi-bert: Towards sensitivity driven fine-tuning for parameter-efficient bert. *ICASSP*, 2024. 2
- [30] Yanyu Li, Geng Yuan, Yang Wen, Ju Hu, Georgios Evangelidis, Sergey Tulyakov, Yanzhi Wang, and Jian Ren. Efficientformer: Vision transformers at mobilenet speed. Advances in Neural Information Processing Systems, 35: 12934–12949, 2022. 2
- [31] Hanxue Liang, Zhiwen Fan, Rishov Sarkar, Ziyu Jiang, Tianlong Chen, Kai Zou, Yu Cheng, Cong Hao, and Zhangyang Wang. M³ViT: Mixture-of-experts vision transformer for efficient multi-task learning with model-accelerator co-design. In Advances in Neural Information Processing Systems, 2022. 2, 7
- [32] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. In *International Conference on Learning Representations*, 2022. 3, 5
- [33] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. arXiv preprint arXiv:2202.07800, 2022. 2, 3
- [34] Pierpaolo Morì, Manoj-Rohit Vemparala, Nael Fasfous, Saptarshi Mitra, Sreetama Sarkar, Alexander Frickenstein, Lukas Frickenstein, Domenik Helms, Naveen Shankar Nagaraja, Walter Stechele, et al. Accelerating and pruning cnns for semantic segmentation on fpga. In *Proceedings of the* 59th ACM/IEEE Design Automation Conference, pages 145– 150, 2022. 2
- [35] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In Sixth Indian Conference on Computer Vision, Graphics & Image Processing, 2008. 6
- [36] Zizheng Pan, Peng Chen, Haoyu He, Jing Liu, Jianfei Cai, and Bohan Zhuang. Mesa: A memory-saving training frame-

work for transformers. *arXiv preprint arXiv:2111.11124*, 2021. 2

- [37] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *CVPR*, 2012. 6
- [38] Ali Sharif Razavian, Hossein Azizpour, Josephine Sullivan, and Stefan Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *CVPR Workshops*, 2014. 5
- [39] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pages 10347–10357. PMLR, 2021. 3, 6
- [40] Manoj-Rohit Vemparala, Nael Fasfous, Alexander Frickenstein, Sreetama Sarkar, Qi Zhao, Sabine Kuhn, Lukas Frickenstein, Anmol Singh, Christian Unger, Naveen-Shankar Nagaraja, Christian Wressnegger, and Walter Stechele. Adversarial robust model compression using in-train pruning. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, pages 66– 75, 2021. 2
- [41] Chaofei Wang, Qisen Yang, Rui Huang, Shiji Song, and Gao Huang. Efficient knowledge distillation from model checkpoints. Advances in Neural Information Processing Systems, 35:607–619, 2022. 2
- [42] Hongyi Wang, Felipe Maia Polo, Yuekai Sun, Souvik Kundu, Eric Xing, and Mikhail Yurochkin. Fusing models with complementary expertise. *arXiv preprint arXiv:2310.01542*, 2023. 7
- [43] 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
- [44] Li Yang, Adnan Siraj Rakin, and Deliang Fan. Rep-net: Efficient on-device learning via feature reprogramming. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12277–12286, 2022. 1, 2, 4
- [45] Lu Yin, Shiwei Liu, Ajay Jaiswal, Souvik Kundu, and Zhangyang Wang. Junk dna hypothesis: A task-centric angle of llm pre-trained weights through sparsity. *arXiv preprint arXiv:2310.02277*, 2023. 2
- [46] Shixing Yu, Tianlong Chen, Jiayi Shen, Huan Yuan, Jianchao Tan, Sen Yang, Ji Liu, and Zhangyang Wang. Unified visual transformer compression. arXiv preprint arXiv:2203.08243, 2022. 2
- [47] Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In ACL, 2022. 2
- [48] Yuke Zhang, Dake Chen, Souvik Kundu, Chenghao Li, and Peter A Beerel. Sal-vit: Towards latency efficient private inference on vit using selective attention search with a learnable softmax approximation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5116–5125, 2023. 2