

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.

ShiftAddAug: Augment Multiplication-Free Tiny Neural Network with Hybrid Computation

Yipin Guo, Zihao Li, Yilin Lang, Qinyuan Ren College of Control Science and Engineering, Zhejiang University

{guoyipin, lzh_jeong, langyilin, renqinyuan}@zju.edu.cn

Abstract

Operators devoid of multiplication, such as Shift and Add, have gained prominence for their compatibility with hardware. However, neural networks (NNs) employing these operators typically exhibit lower accuracy compared to conventional NNs with identical structures. ShiftAddAug uses costly multiplication to augment efficient but less powerful multiplication-free operators, improving performance without any inference overhead. It puts a ShiftAdd tiny NN into a large multiplicative model and encourages it to be trained as a sub-model to obtain additional supervision. In order to solve the weight discrepancy problem between hybrid operators, a new weight sharing method is proposed. Additionally, a novel two stage neural architecture search is used to obtain better augmentation effects for smaller but stronger multiplication-free tiny neural networks. The superiority of ShiftAddAug is validated through experiments in image classification and semantic segmentation, consistently delivering noteworthy enhancements. Remarkably, it secures up to a 4.95% increase in accuracy on the CIFAR100 compared to its directly trained counterparts, even surpassing the performance of multiplicative NNs.

1. Introduction

The application of deep neural networks (DNNs) on resource-constrained platforms is still limited due to their huge energy requirements and computational costs. To obtain a small model deployed on edge devices, the commonly used techniques are pruning[15, 28], quantization[14, 38], and knowledge distillation[16].

However, the NNs designed by the above works are all based on multiplication, which is not hardware-efficient. The common hardware design practice in digital signal processing tells that multiplication can be replaced by bit-wise shifts and additions[13, 42] to achieve faster speed and lower energy consumption. Introducing this idea into NNs design, DeepShift[11] and AdderNet[5] proposed ShiftConv operator and AddConv operator respectively.

This paper takes one step further along the direction of

multiplication-free neural networks, proposing a method to augment tiny multiplication-free NNs by hybrid computation, which significantly improves accuracy without any inference overhead. Considering that multiplication-free operators cannot restore all the information from the original operator, tiny NNs employing ShiftAdd calculations exhibit pronounced under-fitting. Drawing inspiration from NetAug[4], ShiftAddAug chooses to build a larger hybrid computing NN for training and sets the multiplication-free part as the target model used in inference and deployment. The stronger multiplicative part as augmentation is expected to push the target multiplication-free model to a better condition.

In augmented training, the hybrid operators share weights, but because different operators have varying weight distributions, effective weights for multiplication may not suit shift or add operations. This led us to develop a strategy for heterogeneous weight sharing in augmentation.

Furthermore, since NetAug limits augmentation to width, ShiftAddAug aims to extend this by exploring depth and operator change levels. Thus, we adopt a two-step neural architecture search strategy to find highly efficient, multiplicationfree tiny neural networks.

ShiftAddAug is evaluated on MCU-level tiny models. Compared to the multiplicative NNs, directly trained multiplication-free NNs can obtain considerable speed improvement $(2.94 \times \text{ to } 3.09 \times)$ and energy saving $(\downarrow 67.75\% \sim 69.09\%)$ at the cost of reduced accuracy. ShiftAddAug consistently enhances accuracy($\uparrow 1.08\% \sim 4.95\%$) while maintaining hardware efficiency. Our contributions can be summarized as follows:

- For the multiplication-free tiny neural network, we propose hybrid computing augmentation that leverages multiplicative operators to enhance the target multiplication-free network. While maintaining the same model structure, yields a more expressive and ultra-efficient network.
- A new weight sharing strategy is introduced for hybrid computing augmentation, which solves the weight discrepancy in heterogeneous (e.g., Gaussian vs. Laplacian) weight sharing during the augmentation process.

• Based on the idea of augmentation, a two-stage architecture search approach is adopted. Initially extract an augmented large network from the search space, followed by searching for deployable tiny networks within this augmented network.

2. Related Works

Multiplication-Free NNs. To mitigate the high energy and time costs associated with multiplication, a key strategy involves employing hardware-friendly operators instead of multiplications. ShiftNet[7, 40] proposes a zero-parameter, zero-flop convolution. DeepShift[11] retains the calculation method of original convolution, but replaces the multiplication with bit-shift and bit-reversal. BNNs[9, 25, 32] binarize the weight or activation to build DNNs consisting of sign changes. AdderNet [5, 35] chooses to replace multiplicative convolution with less expensive addition, and design an efficient hardware implementation[39]. ShiftAddNet[43] combines bit-shift and add, getting up to $196 \times$ energy savings on hardware as shown in Tab. 1. ShiftAddVit[45] puts this idea into the vision transformer and performs hybrid computing through mixture of experts.

Table 1. Energy cost (pJ) under 45nm CMOS.

OPs	FP32	FP16	INT32	INT8
MULT.	3.7	0.9	3.1	0.2
ADD	1.1	0.4	0.1	0.03
Shift	-	-	0.13	0.024

Network Augmentation. Research on tiny neural networks are developing rapidly. Networks and optimization techniques designed for MCU have already appeared at present[22, 23]. Once-for-all[3] proposes the Progressive Shrinking and finds that the accuracy of the obtained model is better than the counterpart trained directly. Inspired by this result, NetAug[4] raises a point that tiny neural networks need more capacity rather than regularization in training. Therefore, they chose a scheme that is the opposite of regularization methods like Dropout[36]: expand the model width and let the large model lead the small model to achieve better accuracy.

Neural Architecture Search. NAS remarkably succeeded in automating the creation of efficient NN architectures [26, 27], enhancing accuracy while incorporating hardware considerations like latency [37, 41] and memory [22] usage into the design process. NAS also extends its utility to exploring faster operator implementations [6] and integrating network structures for optimization [24, 34], bringing designs closer to hardware requirements. ShiftAddNAS [44] pioneered a search space that includes both multiplicative and multiplication-free operators.

3. ShiftAddAug

3.1. Preliminaries

Shift. The calculation for the shift operator parallels that of standard linear or convolution operators using weight W, except that W is rounded to the nearest power of 2. Bit-shift and bit-reversal techniques are employed to achieve calculation results equivalent to those obtained through traditional methods as Equ. 1. Inputs are quantized before computation and dequantized upon obtaining the output.

$$\begin{cases} S = \operatorname{sign}(W) \\ P = \operatorname{round}(\log_2(|W|)) \\ Y = X \tilde{W_q}^T = X (S \cdot 2^P)^T, \quad train. \\ Y = \sum_{i,j} \sum_k \pm (X_{i,k} << P_{k,j}), \quad eval. \end{cases}$$
(1)

Add. Add operator replaces multiplication with subtractions and ℓ_1 distance since subtractions can be easily reduced to additions by using complement code.

$$Y_{m,n,t} = -\sum_{i=0}^{d} \sum_{j=0}^{d} \sum_{k=0}^{c_{in}} |X_{m+i,n+j,k} - F_{i,j,k,t}|.$$
 (2)

NetAug. Network Augmentation encourages the target tiny multiplicative NNs to work as a sub-model of a large model expanded in width. The target tiny NN and the augmented large model are jointly trained. The training loss and parameter updates are as follows:

$$\mathcal{L}_{aug} = \mathcal{L}(W_t) + \alpha \mathcal{L}(W_a), W_t \in W_a$$

$$W_t^{n+1} = W_t^n - \eta \left(\frac{\partial \mathcal{L}(W_t^n)}{\partial W_t^n} + \alpha \frac{\partial \mathcal{L}(W_a^n)}{\partial W_t^n}\right).$$
(3)

where \mathcal{L} is the loss function, W_t is the weight of the target tiny NN, W_a is the weight of the augmented NN, and W_t is a subset of W_a .

3.2. Hybrid Computing Augment

ShiftAddAug goes a step further based on NetAug[4], using strong operators to augment weak operators.

Taking an *n*-channel DepthWise Conv as an example, NetAug expands it by a factor of α , resulting in a convolution weight for αn channels. During calculation, only the first *n* channels are used for the target model, while the augmented model employs all αn channels. After trained as Equ. 3, the important weights within αn channels are reordered into the first *n*, and only these *n* channels are exported for deployment.

As shown in Fig. 1, ShiftAddAug uses [0, n) channels (target part) and $[n, \alpha n)$ channels (augmented part) for different calculation methods. The target part will use multiplication-free convolution (MFConv, ShiftConv[11] or AddConv[5] can be chosen) while multiplicative convolution (MConv, i.e. original Conv) is used as the augmented part.

Because the channels of Conv are widened, the *input* of each convolution is also widened and can be conceptually



Figure 1. ShiftAddAug augments weak operators with stronger ones. In this framework, pink indicates the multiplication-free kernels designated for the target model (either Shift or Add), while orange represents the multiplicative kernels of the augmented portion. During training, a wide weight is maintained, with the initial n channels processed in a multiplication-free manner and the subsequent channels utilizing multiplicative operations. The weights of both are updated but only the multiplication-free part is exported for deployment. Therefore, the important weights in the wide weight will be reordered into the multiplication-free part. Obtained tiny models have higher accuracy (up to 4.95%) than their directly trained counterparts.

split into the target part X_t and the augmented part X_a , so does the *output* Y_t , Y_a . In ShiftAddAug, X_t and Y_t mainly carry information of MFConv, while X_A and Y_A are obtained by original Conv.

Here three types of operators commonly used to build tiny NNs are discussed: Convolution (Conv), DepthWise Convolution (DWConv), and Fully Connected (FC) layer. The hybrid computing augmentation for DWConv is the most intuitive: split the input into X_t and X_a , then use MFConv and MConv to calculate respectively and connect the obtained Y_t and Y_a in the channel dimension. For Conv, We use all input X to get Y_a through MConv. But to get Y_t , we still need to split the input and calculate it separately, and finally add the results. Since the FC layer is only used as the classification head, its output does not require augmentation. We divide the input and use *Linear* and *ShiftLinear* to calculate respectively, and add the results. If bias is used, it will be preferentially bounded to multiplication-free operators.

$$DWConv: \begin{cases} Y_t = \mathsf{MFConv}(X_t) \\ Y_a = \mathsf{MConv}(X_a) \\ Y = \mathsf{cat}(Y_t, Y_a) \end{cases}$$
$$Conv: \begin{cases} Y_t = \mathsf{MFConv}(X_t) + \mathsf{MConv}(X_a) \\ Y_a = \mathsf{MConv}(X) \\ Y = \mathsf{cat}(Y_t, Y_a) \end{cases} \tag{4}$$
$$FC: \begin{cases} Y_t = \mathsf{ShiftLinear}(X_t) \\ Y_a = \mathsf{Linear}(X_a) \\ Y = Y_t + Y_a \end{cases}$$

3.3. Heterogeneous Weight Sharing



Figure 2. Weight distribution of different convolution operators in MobileNetV2-w0.35. Inconsistent weight distribution leads to discrepancy, making weight sharing difficult.

Dilemma. The important weights will be reordered to the target part at each end of the training epoch (ℓ_1 norm for importance). This is a process of weight sharing and is the key to effective augmentation.

However, the weight distribution of the multiplicationfree operator is inconsistent with the original Conv. It causes the weight discrepancy, i.e. *good weights in original Conv may not be good in MFConv*. As shown in Fig. 2, the weight in the original Conv conforms to Gaussian distribution, while ShiftConv has spikes at some special values. The weight in AddConv conforms to the Laplace distribution. The weight in ShiftConv is the one of the original Conv plus a Laplace distribution with a small variance.

ShiftAddNas[44] adds a penalty term to the loss function, guides the weight to conform to the same distribution. It affects the network to achieve its maximum performance. The Transformation Kernel they proposed also doesn't work on our approach since the loss diverges as Tab. 8. We argue that their approach makes training unstable. This dilemma motivated us to propose our heterogeneous weight sharing strategy.



Figure 3. Weight remapping strategy creates mappings between different weight distributions, making weight sharing workable. cpf_g for cumulative probability function (CPF) for Gaussian distribution and cpf_l CPF for Laplace distribution

Solution: heterogeneous weight sharing. To solve the dilemma above, we propose a new heterogeneous weight sharing strategy for the shift and add operators. This method is based on original Conv and remap parameters to weights of different distribution through a mapping function $\mathcal{R}(\cdot)$. In this way, all weights in memory will be shared under the Gaussian distribution, but will be remapped to an appropriate state for calculation.

When mapping the Gaussian distribution to the Laplace distribution, we hope that the cumulative probability of the original value and mapping result is the same. Firstly, calculate the cumulative probability of the original weight in Gaussian. Then put the result in the percent point function of Laplacian. The workflow is shown in Fig. 3. The mean and standard deviation of the Gaussian can be calculated through the weights, but for the Laplace, these two values need to be determined through prior knowledge.

$$W_{l} = \mathcal{R}(W_{g}) = r(\text{FC}(W_{g}))$$

$$r(\cdot) = \text{ppf}_{l}(\text{cpf}_{g}(\cdot))$$

$$\text{cpf}_{g}(x) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{x} e^{\left(-\frac{(x-u)^{2}}{2\sigma^{2}}\right)} dx$$

$$\text{ppf}_{l}(x) = u - b * \text{sign}(x - \frac{1}{2}) * \ln(1 - 2\left|x - \frac{1}{2}\right|)$$
(5)

Where W_g is the weight in original Conv that conforms to the Gaussian distribution, and W_l is the weight obtained by mapping that conforms to the Laplace distribution. FC is a fully connected layer, which is previously trained and frozen in augmented training. We need this because the weights don't fit the distribution perfectly. $cpf_g(\cdot)$ is the cumulative probability function of Gaussian, $ppf_l(\cdot)$ is the percentage point function of Laplace.

Table 2. Search space for Augmentation.

Ιτεμ	OPTIONAL
SUBNET AS AUG. MODEL	SUPERNET
TARGERNET	SUBNET
BLOCK TYPES	[Mult. , Shift, Add]
WIDTH AUG. MULTIPLES	[2.2, 2.4, 2.8, 3.2]
EXPAND AUG. MULTIPLES	[2.2, 2.4, 2.8, 3.2]
TIME NODE: MUTATION START	[5%, 10%, 15%, 20%]
TIME NODE: MUTATION STOP	[50%, 60%, 70%]

3.4. Neural Architecture Search

To obtain SOTA multiplication-free model at the tiny model size, a two-stage NAS is proposed.

Based on the idea of augmentation, ShiftAddAug starts from a multiplicative SuperNet and cuts a deep SubNet from it as depth-augmented NN. Then select some layers on the SubNet to form the tiny TargetNet for final use. The Target-Net should meet the pre-set hardware limitation. This setup allows the TargetNet to be a part of the SubNet, facilitating joint training with weight sharing as Equ. 3. Layers within the SubNet that are not selected for the TargetNet serve as a form of depth augmentation. Moreover, the layers used for deep augmentation are initially selected but gradually phased out from the target network in training progresses.

A new block mutation training is also proposed, which tends to gradually transform multiplication operators into multiplication-free states during training to make the training process more stable. The training process starts with all multiplications, and the layers of the target network become multiplication-free one by one from shallow to deep. At the end of training, a completely multiplication-free TargetNet can be obtained.

While ShiftAddNas[44] directly uses hybrid computing to train their SuperNet and directly cut SubNets that meet the hardware requirements, we start from the multiplicative SuperNet and split the search process into two steps. The middle step is used for augmented training, which is the unique part of ShiftAddAug.

Combining the Width Augmentation and Expand Augmentation we used in section 3.2, we construct our search space for the augmentation part according to Tab. 2.

We follow tinyNAS[22] to build SuperNet and cut SubNet. Then use evolutionary search to search for subsequent steps.

4. Experiments

4.1. Setup

Datasets. We conduct experiments on several image classification datasets, including CIFAR10[19], CIFAR100[19], ImageNet-1K[10], Food101[1], Flowers102 [29], Cars[18], Pets[31]. We also evaluated our method on VOC[12] and OpenEDS[30] for segmentation task.



Figure 4. Light orange block for MConv; pink block for MFConv; deep orange block for Depth Augmentation. (a).Search process: two-stage search. Find a SubNet as depth augmented one and then further cut out a tiny TargetNet on it for development. Start with multiplication and gradually convert the TargetNet to multiplication-free during training. (b).Methods to augment. *Augment Width:* use MConv to widen the MFConv channel; *Augment Expand:* increase expand ratio of depthwise separable convolution; *Augment Depth:* Augmented blocks for depth. Only participate in training and will not be exported; *Block Mutation:* start with MConv, mutate the block into MFConv during training.

Training Details. The NNs are trained with batch size 128 using 2 GPUs. The SGD optimizer is used with Nesterov momentum 0.9 and weight decay 4e-5. By default, the Mult. and Shift models are trained for 250 epochs and Add models are trained for 300 epochs. The initial learning rate is 0.05 and gradually decreases to 0 following the cosine schedule. Label smoothing is used with a factor of 0.1. For ShiftConv and ShiftLinear, the weight is quantized to 5bit Power-of-2 and the activation is quantized to 16bit. AddConv is calculated under 32bit.

Hardware Performance. Since many works have verified the efficiency of shift and add on proprietary hardware[39, 43–45], we follow their evaluation metrics. Hardware energy and latency are measured based on a simulator of Eyeriss-like hardware accelerator[8, 46], which calculates not only computational but also data movement energy.

4.2. ShiftAddAug vs. Baseline

We validate our method on MobileNetV2[33], MobileNetV3[17], MCUNet[22], ProxylessNAS[2], MobileNetV2-Tiny[22]. ShiftAddAug provides consistent accuracy improvements (average $\uparrow 2.82\%$) for ShiftConv augmentation over the multiplicative baselines. For AddConv augmentation, it improves the accuracy compared with direct training (average $\uparrow 1.59\%$). The resulting model will be faster ($3.0 \times$ for Shift) and more energy-efficient (\downarrow 68.58% for Shift and \downarrow 52.02% for Add) due to the use of hardware-friendly operators. As shown in Tab. 3, these multiplication-free operators usually hurt the performance of the network. Changing all operators to Shift will cause \downarrow 0.82% accuracy drop on average compared to the multiplication baseline. But after using our method, the accuracy increased by \uparrow 3.63% on average under the same energy cost. ¹

In addition, our method achieves higher results than multiplicative NetAug on some models (MobileNetV3: \uparrow 1.17%, MCUNet: \uparrow 1.44%, ProxylessNAS: \uparrow 1.54%). This means that our method enables the multiplication-free operator to be stronger than those of the original operator.

To verify the generality of our method, we also conduct experiments on more datasets. As shown in Tab. 4, our method can achieve $\uparrow 0.89\%$ to $\uparrow 4.28\%$ accuracy improvements on different datasets. Hybrid computing augmentation works better on smaller models and datasets with less classification. On Flower102, MobileNetV2-w0.35 has $\uparrow 3.83\%$ accuracy improvements with our method, while MCUNet has only $\uparrow 1.47\%$. This shows that a smaller model capacity can achieve a better effect on this dataset. The larger the model, the smaller the gain brought by augmentation. The same phenomenon also occurs in CIFAR10. For bigger datasets such as ImageNet, even if it is augmented, the capac-

¹Loss becomes NaN when we use AddConv on MobileNetV3, both direct training and augmented training.

Table 3. ShiftAddAug vs. Multiplicative and directly trained Multiplication-free Baseline in terms of accuracy and efficiency on CIFAR100 classification tasks. The results obtained by our method are bolded. The accuracy on the left is directly trained and the one on the right is obtained with augmentation. *Base* for Mult.; *Shift* for ShiftConv in DeepShift and *Add* for AddConv in AdderNet.

Model	Method	Params (M)	Mult (M)	Shift (M)	Add (M)	Accuracy(%)	Energy (mj)	Latency (ms)
MobileNetV2	Base / NetAug	0.52	29.72	0	29.72	70.59 / 71.98	2.345	0.73
widdlienet v 2	Shift / AugShift	0.52	0	29.72	29.72	69.25 / 71.83 (^ 2.58)	0.74	0.246
w0.35	Add / AugAdd	0.52	4.52	0	56.88	67.85 / 69.38 (†1.5)	1.091	0.753
MobileNetV3	Base / NetAug	0.96	18.35	0	18.35	69.32 / 72.2	1.726	0.485
widdheitet v 5	Shift / AugShift	0.96	0	18.35	18.35	68.42 / 73.37 (†4.95)	0.536	0.16
w0.35	Add / AugAdd	0.96	3.5	0	34.34	- / -	0.699	0.512
	Base / NetAug	0.59	65.72	0	65.72	71.38 / 73.15	4.28	1.682
MCUNet - in3	Shift / AugShift	0.59	0	65.72	65.72	70.87 / 74.59 (†3.72)	1.323	0.545
	Add / AugAdd	0.59	20.91	0	113.09	70.25 / 72.72 (†2.47)	2.345	1.72
ProvuleseNAS	Base / NetAug	0.63	34.56	0	34.56	70.86 / 72.32	2.471	0.883
110XyIC3514A5	Shift / AugShift	0.63	0	34.56	34.56	70.54 / 73.86 (†3.32)	0.774	0.294
w0.35	Add / AugAdd	0.63	8.81	0	61.97	68.87 / 70.18 (†1.31)	1.281	0.881
MobileNetV2	Base / NetAug	0.35	27.31	0	27.31	69.3 / 71.62	2.161	0.67
Tiny	Shift / AugShift	0.35	0	27.31	27.31	68.29 / 71.89 († 3.6)	0.697	0.228
- 1 Iny	Add / AugAdd	0.35	4.43	0	52.09	66.57 / 67.65 (†1.08)	0.999	0.693

Table 4. Accuracy of MobileNetV2 (w0.35) and MCUNet-in3 on more datasets. ShiftAddAug can improve performance without any inference overhead on fine-grained classification tasks. Please refer to Tab.12 for results on more models

Model	Methods	CIFAR10	ImageNet	Food101	Flower102	Cars	Pets
MobileNetV2 - w0.35	Shift	88.59	51.92	72.99	92.25	72.83	75.4
	AugShift	92.51	53.86	74.67	96.08	74.47	79.59
	Shift	90.61	56.45	78.46	95.59	80.51	79.67
MCUNet - in3	AugShift	93.08	57.34	79.96	97.06	83.29	83.95

ity of the model is still not enough. It achieves $\uparrow 1.94\%$ for MobileNetV2-w0.35 and $\uparrow 0.89\%$ for MCUNet on ImageNet. But this goes beyond the common work scenarios of IOT devices.

4.3. ShiftAddAug vs. SOTA Mult.-Free Models

We compare ShiftAddAug over SOTA multiplication-free models, which are designed manually for tiny computing devices, on CIFAR-10/100 to evaluate its effectiveness. As shown in Fig. 5, the model structures we use are smaller and have better energy performance. With ShiftAddAug, the accuracy still exceeds existing work. For DeepShift and Adder-Net, our method boosts $\uparrow 0.67\%$ and $\uparrow 1.95\%$ accuracy on CIFAR100 with $\downarrow 84.17\%$ and $\downarrow 91.7\%$ energy saving. Compared with the SOTA shift quantization method APoT[21], we achieve an improved accuracy of $\uparrow 3.8\%$. With the same accuracy on CIFAR10, our model saves $\downarrow 84.2\%$ of the energy compared with Deepshift, and $\downarrow 56.45\%$ of the energy compared with AdderNet.

4.4. ShiftAddAug on Segmentation Task

To further verify the effectiveness of our method, we conduct experiments on the semantic segmentation task. We



Figure 5. Accuracy and energy cost of ShiftAddAug over SOTA manually designed multiplication-free model and tiny multiplicative models. Tested on CIFAR-100/10.

choose VOC[12] for general performance and OpenEDS[30] for specific tasks on MCU level devices. As shown in Tab. 5, ShiftAddAug has greater improvement on OpenEDS, which

D-44	Madha d	MobileNetV2	ProxylessNAS
Dataset	Niethod	w0.35	w0.35
VOC	Base/NetAug	66.23 / 69.72	66.63 / 68.09
VOC	Shift/AugShift	66.02 / 68.78	65.8 / 68.98
OpenEDS	Base/NetAug	88.68 / 92.41	86.27 / 92.84
	Shift/AugShift	88.01 / 94.52	86.01 / 95.12
	Add/AugAdd	83.94 / 91.20	82.01 / 90.45

Table 5. ShiftAddAug on semantic segmentation task (mIoU %)

illustrates the application potential of our method.

From the results shown in Fig.6, we can see that the model trained with augmentation will have fewer abnormal segmentation areas.



Figure 6. ShiftAddAug on OpenEDS.

4.5. ShiftAddAug on Neural Architecture Search

Based on hybrid computing augmentation, we introduce neural architecture search into our method to get stronger tiny neural networks. We conduct our experiments on Cifar-10/100 and compare them with the results of ShiftAddNAS[44] which is better than multiplication-based FBNet[41]. As shown in Tab. 6, the multiplication-free model we obtained achieved higher accuracy (\uparrow 3.61%) than ShiftAddNas. For hybrid-computed models, we have to use a smaller input resolution (96 instead of 160) and larger models. While the input resolution of ShiftAddNas is 32, this would give us 9× the number of calculations if we have the same model architecture. Even so, we can still save 37.1% of calculations on average with similar accuracy.

4.6. Ablation Study

Hybrid Computing Augment. In order to prove that hybrid computing works better, we add an experiment using only multiplication-free operators for augmentation. We exert experiments based on NetAug, and replace all the original operators with Shift operators. The difference from our method is that the Shift operator is also used in the augmentation part, while our method uses multiplication in it. As shown in Tab. 7, it yields an average accuracy improvement of $\uparrow 1.02\%$.

Then without using the heterogeneous weight sharing (HWS) method, just augmenting tiny NNs with the multiplicative operator will cause $\downarrow 0.9\%$ accuracy drop on average due to the weight tearing problem. However, the situation changed after we applied for HWS. Compared with using the shift operator for augmentation, the accuracy increased by $\uparrow 2.2\%$.

Heterogeneous Weight Sharing. Since the help of HWS on training results has been discussed above, here we visualize the weight distributions of Conv layers in tiny NNs under three scenarios, (a) w/o weight sharing; (b) heterogeneous weight sharing; (c) weight after remapped, as shown in Fig. 7. We consistently observe that the three operators exhibit different weight distributions without weight sharing. With our HWS, the saved weights are based on the original Conv and conform to Gaussian distribution. After remapping, the weights can show different distribution states in Shift/Add Conv calculations.

Different methods for solving weight tearing are also compared. As shown in Tab. 8, both the ShiftAddNas method and direct linear remapping will cause loss to diverge. The additional penalty has a larger impact on tiny NNs, causing accuracy to decrease.

Additionally, Tab. 9 shows why weight sharing is important, which also explains why we have to face the problem of heterogeneous weights. Tab. 10 tells that HWS is not a parameterized trick directly improving accuracy, but just a compensation in augmentation.

Neural Architecture Search. Our neural architecture search approach is dependent on our proposed hybrid computing augmentation. And it can help the multiplication-free operator to be as strong as the original operator. Block augmentation and block mutation help us further improve the performance of multiplication-free tiny NNs. As shown in Tab. 11, under similar energy consumption and latency, block augmentation improves accuracy by $\uparrow 0.5\%$, and block mutation improves by $\uparrow 0.79\%$. Combining all methods, the accuracy of the target model is increased by $\uparrow 1.22\%$.

4.7. Limitation

Since we use additional multiplication structures for augmentation, training consumes more resources. This is consistent with NetAug.

In terms of memory usage, which depends on the size of the augmented NN, ShiftAddAug usually doubles or triples it compared with direct training. This makes ShiftAddAug not suitable for on-device training.

Table 6. ShiftAddAug vs. SOTA NAS method for hybrid operators in term of accuracy and efficiency on Cifar-10/100 classification tasks. ShiftAddAug can further improve the performance of the multiplication-less model.

Detect	Mathad	Dec	Calculation Amount			A course ou (07-)	MACa Compose
Dataset	Method	Kes.	Mult (M)	Shift (M)	Add (M)	Accuracy(%)	MACS Compare
	ACGhostNet-B[20]	32	28	0	61	95.2	1.62×
CIFAR10	ShiftAddNas	32	17	19	58	95.83	$1.71 \times$
	ShiftAddAug		12.3	14.5	28	95.92	1.0 imes
	ShiftAddNas (Mult-free)	32	3	35	48	71.0	$1.27 \times$
	ShiftAddNas	32	22	21	62	78.6	$1.44 \times$
CIFAR100	ShiftAddAug (Mult-free)	160	0.13	33.5	33.6	74.61	1.0 imes
	ShiftAddAug		16.2	20.2	36.4	78.72	1.0 imes
Mult. C	onv Shift Conv Add (Conv	Shift C	onv Ad	ld Conv	Shift Conv	Add Conv
	(a). w/o weight sharing			(b). HWS		(c). HWS	after remap

Figure 7. The weight distribution of original Conv / ShiftConv / AddConv layers

Table 7. The ablation study of hybrid computing augmentation and heterogeneous weight sharing in terms of accuracy on CIFAR100.

Mathad	MobileNetV2	MCUNet	ProxylessNAS
Method	w0.35	in3	w0.35
Mult. baseline	70.59	71.38	70.86
To shift op	69.25	70.87	70.54
Aug. with Shift	70.12	72.68	70.91
Aug. with Hybrid Computation	69.41	71.02	70.6
Aug. with HWS	71.83	74.59	73.86

 Table 8. The ablation study of different method for HWS

Method	MobileNetV2 w0.35	MCUNet	ProxylessNAS w0.35
ShiftAddNAS	Nan	Nan	Nan
Linear remap	Nan	Nan	Nan
KL-loss only	69.84	71.12	70.52
Linear + skip connect + freeze	71.02	74.44	72.99
Ours	71.83	74.59	73.86

Table 9. Ablation study on weight sharing. Augmented with Shift op. Tells why weight sharing is important.

Method	MobileNetV2 w0.35	MCUNet	ProxylessNAS w0.35
Direct trained	69.25	70.87	70.54
w/o. weight sharing	69.42	71.09	70.53
w/. weight sharing	70.12	72.68	70.91

Table 10. HWS is not a parameterization trick that can directly improve target model

Method	MobileNetV2 w0.35	MCUNet	ProxylessNAS w0.35	
w/o. aug. , w/o. HWS	69.25	70.87	70.54	
w/o. aug. , with HWS	68.32	71.13	69.88	
with aug., with HWS	71.83	74.59	73.86	

Table 11. The ablation study of depth augmentation and block mutation in terms of accuracy on CIFAR100. Results are obtained under similar hardware cost.

Method	Accuracy(%)	Energy (mj)	Latency (ms)
Aug. Width & Expand	75.13	1.52	0.57
Aug. Width & Expand & Depth	75.63	1.42	0.66
Aug. Width & Expand, Mutation	75.92	1.4	0.67
All	76.35	1.632	0.554

Table 12. Accuracy of MobileNetV3-w0.35 / ProxylessNAS-w0.35 / MobileNetV2-Tiny on more datasets.

Model	Methods	CIFAR10	ImageNet	Food101	Flower102
MobileNetV3	Shift	88.85	54.19	73.3	93.82
w0.35	AugShift	92.83	56.07	75.1	96.28
ProxylessNAS	Shift	87.71	55.4	75.66	94.31
w0.35	AugShift	92.37	56.3	77.02	96.76
MobileNetV2	Shift	88.02	48.92	72.91	94.63
Tiny	AugShift	91.93	50.0	74.58	96.56

5. Conclusion

In this paper, ShiftAddAug is proposed for training multiplication-free tiny neural networks, which can improve accuracy without expanding the model size. It's achieved by putting the target multiplication-free tiny NN into a larger multiplicative NN to get auxiliary supervision. To relocate important weights into the target model, a novel heterogeneous weight sharing strategy is used to approach the weight discrepancy caused by inconsistent weight distribution. Based on the work above, a two stage neural architecture search is utilized to design more powerful models. Extensive experiments on image classification and semantic segmentation task consistently demonstrate the effectiveness of ShiftAdAug.

References

- Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative components with random forests. In *European Conference on Computer Vision*, pages 446–461, Cham, 2014. Springer International Publishing. 4
- Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019.
 5
- [3] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. In *International Conference on Learning Representations*, 2020. 2
- [4] Han Cai, Chuang Gan, Ji Lin, and song han. Network augmentation for tiny deep learning. In *International Conference* on Learning Representations, 2022. 1, 2
- [5] Hanting Chen, Yunhe Wang, Chunjing Xu, Boxin Shi, Chao Xu, Qi Tian, and Chang Xu. Addernet: Do we really need multiplications in deep learning? In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 1465–1474, 2020. 1, 2
- [6] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Meghan Cowan, Haichen Shen, Leyuan Wang, Yuwei Hu, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. Tvm: an automated end-to-end optimizing compiler for deep learning. In *Proceedings of the 13th USENIX Conference on Operating Systems Design and Implementation*, page 579–594, USA, 2018. USENIX Association. 2
- [7] Weijie Chen, Di Xie, Yuan Zhang, and Shiliang Pu. All you need is a few shifts: Designing efficient convolutional neural networks for image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [8] Yu-Hsin Chen, Tushar Krishna, Joel S. Emer, and Vivienne Sze. Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks. *IEEE Journal of Solid-State Circuits*, 52(1):127–138, 2017. 5
- [9] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1, 2016. 2
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 248–255, 2009. 4
- [11] Mostafa Elhoushi, Zihao Chen, Farhan Shafiq, Ye Henry Tian, and Joey Yiwei Li. Deepshift: Towards multiplication-less neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2359–2368, 2021. 1, 2
- [12] Mark Everingham, Luc van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 4, 6
- [13] Bah-Hwee Gwee, Joseph S. Chang, Yiqiong Shi, Chien-Chung Chua, and Kwen-Siong Chong. A low-voltage micropower asynchronous multiplier with shift-add multiplica-

tion approach. *IEEE Transactions on Circuits and Systems I:* Regular Papers, 56(7):1349–1359, 2009. 1

- [14] Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. *Computer Vision and Pattern Recognition*, 2015. 1
- [15] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In Advances in Neural Information Processing Systems. Curran Associates, Inc., 2015. 1
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015. 1
- [17] Andrew Howard, Mark Sandler, Bo Chen, Weijun Wang, Liang-Chieh Chen, Mingxing Tan, Grace Chu, Vijay Vasudevan, Yukun Zhu, Ruoming Pang, Hartwig Adam, and Quoc Le. Searching for mobilenetv3. In 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pages 1314–1324, 2019. 5
- [18] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In 2013 IEEE International Conference on Computer Vision Workshops, pages 554–561, 2013. 4
- [19] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 4
- [20] Wenshuo Li, Xinghao Chen, Jinyu Bai, Xuefei Ning, and Yunhe Wang. Searching for energy-efficient hybrid adderconvolution neural networks. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 1942–1951, 2022. 8
- [21] Yuhang Li, Xin Dong, and Wei Wang. Additive powers-oftwo quantization: An efficient non-uniform discretization for neural networks. In *International Conference on Learning Representations*, 2020. 6
- [22] Ji Lin, Wei-Ming Chen, Yujun Lin, john cohn, Chuang Gan, and Song Han. Mcunet: Tiny deep learning on iot devices. In Advances in Neural Information Processing Systems, pages 11711–11722. Curran Associates, Inc., 2020. 2, 4, 5
- [23] Ji Lin, Wei-Ming Chen, Han Cai, Chuang Gan, and Song Han. Memory-efficient patch-based inference for tiny deep learning. In Advances in Neural Information Processing Systems, pages 2346–2358. Curran Associates, Inc., 2021. 2
- [24] Yujun Lin, Mengtian Yang, and Song Han. Naas: Neural accelerator architecture search. In 2021 58th ACM/IEEE Design Automation Conference (DAC), pages 1051–1056, 2021. 2
- [25] Zhouhan Lin, Matthieu Courbariaux, Roland Memisevic, and Yoshua Bengio. Neural networks with few multiplications, 2016. 2
- [26] Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan L. Yuille, and Li Fei-Fei. Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 82–92, 2019.
- [27] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *International Conference* on *Learning Representations*, 2019. 2

- [28] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. In *International Conference on Learning Representations*, 2017. 1
- [29] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing, pages 722–729, 2008. 4
- [30] Cristina Palmero, Abhishek Sharma, Karsten Behrendt, Kapil Krishnakumar, Oleg V. Komogortsev, and Sachin S. Talathi. Openeds2020: Open eyes dataset, 2020. 4, 6
- [31] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In 2012 IEEE Conference on Computer Vision and Pattern Recognition, pages 3498–3505, 2012. 4
- [32] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks, 2016. 2
- [33] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR), pages 4510–4520, 2018. 5
- [34] Huihong Shi, Haoran You, Yang Zhao, Zhongfeng Wang, and Yingyan Lin. Nasa: Neural architecture search and acceleration for hardware inspired hybrid networks. In 2022 IEEE/ACM International Conference On Computer Aided Design (ICCAD), pages 1–9, 2022. 2
- [35] Dehua Song, Yunhe Wang, Hanting Chen, Chang Xu, Chunjing Xu, and Dacheng Tao. Addersr: Towards energy efficient image super-resolution. In 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR), pages 15643–15652, 2021. 2
- [36] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958, 2014. 2
- [37] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 2815–2823, 2019. 2
- [38] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision, 2019. 1
- [39] Yunhe Wang, Mingqiang Huang, Kai Han, Hanting Chen, Wei Zhang, Chunjing Xu, and Dacheng Tao. Addernet and its minimalist hardware design for energy-efficient artificial intelligence, 2021. 2, 5
- [40] Bichen Wu, Alvin Wan, Xiangyu Yue, Peter Jin, Sicheng Zhao, Noah Golmant, Amir Gholaminejad, Joseph Gonzalez, and Kurt Keutzer. Shift: A zero flop, zero parameter alternative to spatial convolutions. In 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition(CVPR), pages 9127–9135, 2018. 2
- [41] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing

Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 10726–10734, 2019. 2, 7

- [42] Ping Xue and Bede Liu. Adaptive equalizer using finite-bit power-of-two quantizer. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 34(6):1603–1611, 1986. 1
- [43] Haoran You, Xiaohan Chen, Yongan Zhang, Chaojian Li, Sicheng Li, Zihao Liu, Zhangyang Wang, and Yingyan Lin. Shiftaddnet: A hardware-inspired deep network. In Advances in Neural Information Processing Systems, pages 2771–2783. Curran Associates, Inc., 2020. 2, 5
- [44] Haoran You, Baopu Li, Huihong Shi, Yonggan Fu, and Yingyan Lin. Shiftaddnas: Hardware-inspired search for more accurate and efficient neural networks, 2022. 2, 3, 4, 7
- [45] Haoran You, Huihong Shi, Yipin Guo, Yingyan, and Lin. Shiftaddvit: Mixture of multiplication primitives towards efficient vision transformer, 2023. 2, 5
- [46] Yang Zhao, Chaojian Li, Yue Wang, Pengfei Xu, Yongan Zhang, and Yingyan Lin. Dnn-chip predictor: An analytical performance predictor for dnn accelerators with various dataflows and hardware architectures. In ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 1593–1597, 2020. 5