Finet: Using Fine-grained Batch Normalization to Train Light-weight Neural Networks

Chunjie Luo
Institute of Computing Technology
Chinese Academy of Sciences
`luochunjie@ict.ac.cn`

Jianfeng Zhan
Institute of Computing Technology
Chinese Academy of Sciences
`zhanjianfeng@ict.ac.cn`

Lei Wang
Institute of Computing Technology
Chinese Academy of Sciences
`wanglei_2011@ict.ac.cn`

Wanling Gao
Institute of Computing Technology
Chinese Academy of Sciences
`gaowanling@ict.ac.cn`

Abstract

To build light-weight network, we propose a new normalization, Fine-grained Batch Normalization (FBN). Different from Batch Normalization (BN), which normalizes the final summation of the weighted inputs, FBN normalizes the intermediate state of the summation. We propose a novel light-weight network based on FBN, called Finet. At training time, the convolutional layer with FBN can be seen as an inverted bottleneck mechanism. FBN can be fused into convolution at inference time. After fusion, Finet uses the standard convolution with equal channel width, thus makes the inference more efficient. On ImageNet classification dataset, Finet achieves the state-of-art performance (65.706% accuracy with 43M FLOPs, and 73.786% accuracy with 303M FLOPs), Moreover, experiments show that Finet is more efficient than other state-of-art light-weight networks.

1 Introduction

Since AlexNet [20] won ImageNet Large-Scale Visual Recognition Competition in 2012, deep neural networks have received great successes in many areas of machine intelligence. Modern state-of-art networks [35] [8] [14] [38] [50] become deeper and wider. The requirement of high computational resources hinders their usages on many mobile and embedded applications. As a result, there has been rising interest for the design of light-weight neural networks [11] [33] [10] [39] [49] [33] [13] [48] [44] [37]. The objective of light-weight networks is to decrease the computation complexity with low or no loss of accuracy.

To build light-weight networks, we propose a new normalization, Fine-grained Batch Normalization (FBN). Batch Normalization (BN) [17] has become a standard component in modern neural networks. There are lots of variants of Batch Normalization to meet diverse usage scenarios [16] [5] [41] [43] [22] . FBN aims to improve the training of light-weight networks, while keep the inference efficient. As shown in Figure 1, different from BN, which normalizes the final summation of the weighted inputs, FBN normalizes the intermediate state of the summation. At inference time, the normalization can be fused into the linear transformation. Then there is no need for the intermediate state.

We propose a novel light-weight network based on FBN, called Finet. At training time, the convolutional layer with FBN can be seen as an inverted bottleneck mechanism since the intermediate channels are normalized and then summarized. However, this bottleneck has only one convolutional layer. FBN can be fused into convolution at inference time. After fusion, Finet uses the standard convolution with equal channel width, thus makes the inference more efficient.
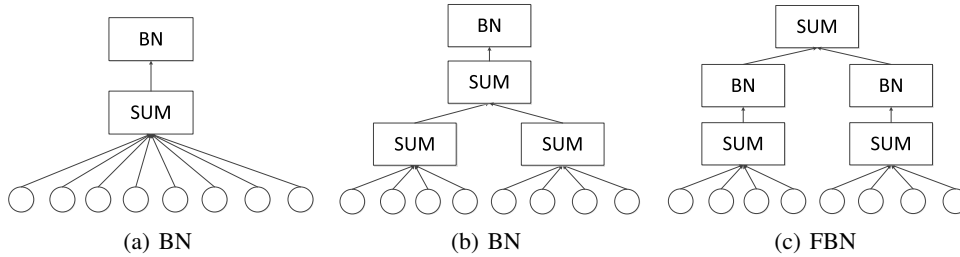
Figure 1: The difference between Batch Normalization (BN) and Fine-grained Batch Normalization (FBN). Each small circle in the figure represents $w_i x_i$. Mathematically, (a) is equal to (b) because of Associative Law of addition. (c) represents FBN which normalizes the intermediate state of the summation. At inference time, the normalization can be fused into the linear transformation $w_i x_i$. After fusion, (a), (b) and (c) are equivalent.

On ImageNet classification dataset, Finet achieves the state-of-art performance. With 43M FLOPs, Finet achieves 65.706% accuracy, outperforming the corresponding model of ShuffleNetV2 (60.3%), MobileNetV2 (58.2%), MobileNetV3 (65.4%). With 303M FLOPs, Finet achieves 73.786% accuracy, outperforming the corresponding model of ShuffleNetV2 (72.6%), MobileNetV2 (72.0%). Moreover, we compare the inference speed of Finet, ShuffleNetV2, MobileNetV2, MobileNetV3, and MnasNet on three different mobile phones. The results show that Finet is more efficient than other state-of-art light-weight networks. We also evaluate Finet on CIFAR-10/100 dataset, and present the influences of different hyper-parameter settings. Finally, we show that FBN also improves the performance of ResNet.

2 Related Work

2.1 Normalization

Batch Normalization [17] performs the normalization for each training minibatch along (N,H,W) dimensions in the case of NCHW format feature. Normalization Propagation [4] uses a data-independent parametric estimate of the mean and standard deviation instead of explicitly calculating from data. Batch Renormalization [16] introduces two extra parameters to correct the fact that the minibatch statistics differ from the population ones. Layer Normalization [5] computes the mean and standard deviation along (C,H,W) dimensions. Instance Normalization [41] computes the mean and standard deviation along (H,W) dimensions. Group Normalization [43] is a intermediate state between layer normalization and instance normalization. Extended Batch Normalization [22] computes the mean along the (N, H, W) dimensions, and computes the standard deviation along the (N, C, H, W) dimensions. Weight Normalization [32] normalizes the filter weights instead of the activations by re-parameterizing the incoming weight vector. Cosine Normalization [23] normalizes both the filter weights and the activations by using cosine similarity or Pearson correlation coefficient instead of dot product in neural networks. Kalman Normalization [42] estimates the mean and standard deviation of a certain layer by considering the distributions of all its preceding layers. Instead of the standard L^2 Batch Normalization, [9] performs the normalization in L^1 and L^∞ spaces. Generalized Batch Normalization [46] investigates a variety of alternative deviation measures for scaling and alternative mean measures for centering. Batch-Instance Normalization [28] uses a learnable gate parameter to combine batch and instance normalization together, and Switchable Normalization [24] uses learnable parameters to combine batch, instance and layer normalization. Virtual Batch Normalization [31] and spectral normalization [27] focus on the normalization in generative adversarial networks. Self-Normalizing [18] focuses on the fully-connected networks. Recurrent Batch Normalization [7] modifies batch normalization to use in recurrent networks. EvalNorm [36] estimates corrected normalization statistics to use for batch normalization during evaluation. [29] provides a unifying view of the different normalization approaches. [34], [25] and [6] try to explain how Batch Normalization works.

2.2 Light-weight Network

MobileNet [11] is a light-weight deep neural network designed and optimized for mobile and embedded vision applications. MobileNet is based on depthwise separable convolutions to reduce the number of parameters and computation FLOPs. MobileNetV2 [33] introduces two optimized mechanisms: 1) inverted residual structure where the shortcut connections are between the thin layers. 2) linear bottlenecks which removes non-linearities in the narrow layers. MnasNet [39] is a neural architecture automated searched for mobile device by using multi-objective optimization and factorized hierarchical search space. MobileNetV3 [10] is also a light-weight network searched by network architecture search algorithm. EfficientNet [40] proposes a new scaling method that uniformly scales all dimensions of depth/width/resolution ShufflfleNet [49] utilizes pointwise group convolution and channel shuffle, to greatly reduce computation cost while maintaining accuracy. ShufflfleNetV2 [33] proposes to evaluate the direct metric on the target platform, beyond only considering FLOPs. Following several practical guidelines, a new efficient architecture based on channel split and shuffle is presented. CondenseNet [13] combines dense connectivity with learned group convolution. IGCV [48] [44] [37] propose interleaved group convolutions to build efficient networks. SqueezeNet [15] aims to decrease the number of parameters while maintaining competitive accuracy.

3 Fine-grained Batch Normalization

3.1 Batch Normalization

Batch normalization (BN) has become a standard technique for training the deep networks. For a neuron, BN is defined as

$$\widehat{x} = (\sum_{i=1}^{N} w_i x_i - \mu)/\sigma \tag{1}$$

where w_i is the input weight, and x_i is the input, N is the size of input, μ is the mean of the neurons along the batch dimension, and σ is the standard deviation. For simplicity, we omit the epsilon in the denominator, and the affine transformation after normalization. The non-linear function $f(\widehat{x})$ is performed after the normalization.

3.2 Fine-grained Batch Normalization

We can divide the inputs and the weights of a neuron into G groups. Because of Associative Law of addition, we have

$$\sum_{i=1}^{N} w_i x_i = \sum_{g=1}^{G} \sum_{j=1}^{N/G} w_{gj} x_{gj} \tag{2}$$

where w_i and w_{gj} are the same weights, x_i and x_{gj} are the same inputs, but using different index notation. Then Equation 1 can be re-written as

$$\widehat{x} = (\sum_{g=1}^{G} \sum_{j=1}^{N/G} w_{gj} x_{gj} - \mu)/\sigma \tag{3}$$

In this paper, we propose Fine-grained Batch Normalization (FBN), which is defined as

$$\widehat{x} = \sum_{g=1}^{G} ((\sum_{j=1}^{N/G} w_{gj} x_{gj} - \mu_g)/\sigma_g) \tag{4}$$

where μ_g and σ_g are the mean and the standard deviation of the intermediate summation along the batch dimension. Different from BN, which normalizes the final summation of the weighted inputs,

FBN normalizes the intermediate state of the summation. When training with BN, neurons can be coordinated in a mini-batch. FBN makes the coordination more fine-grained. Figure 1 shows the difference between BN and FBN. At training time, the intermediate state need to be stored and traced. Thus FBN takes more memory resource for training. Though FBN needs to normalize more channels, the computation overhead is trivial comparing to the convolution operation.

3.3 Inference and Normalization Fusion

At inference time, the mean and the standard deviation are pre-computed from the training data by the moving average. The inference of BN is computed as

$$\widehat{x} = (\sum_{i=1}^{N} w_i x_i - \mu')/\sigma' \tag{5}$$

where μ' refers to the moving average of the mean, σ' refers to the moving average of the standard deviation. Since the mean and the standard deviation are pre-computed and fixed at inference time, the normalization can be fused into the linear transformation, e.g. convolution operation. Re-write Equation 5 as

$$\widehat{x} = \sum_{i=1}^{N} \frac{w_i}{\sigma'} x_i - \frac{\mu'}{\sigma'} \tag{6}$$

Let $w_i' = \frac{w_i}{\sigma'}$, and $b' = \frac{\mu'}{\sigma'}$, then we can fuse the normalization into the linear transformation as

$$\widehat{x} = \sum_{i=1}^{N} w_i' x_i - b' \tag{7}$$

Here w_i' is the new weight of the linear transformation, and b' is the bias. Thus BN is fused into the new linear transformation. Similarly, the inference of FBN is computed as

$$
\begin{aligned}
\widehat{x} &= \sum_{g=1}^{G}((\sum_{j=1}^{N/G} w_{gj} x_{gj} - \mu_g')/\sigma_g') \\
&= \sum_{g=1}^{G}(\sum_{j=1}^{N/G} \frac{w_{gj}}{\sigma_g'} x_{gj} - \frac{\mu_g'}{\sigma_g'}) \\
&= \sum_{g=1}^{G}\sum_{j=1}^{N/G} \frac{w_{gj}}{\sigma_g'} x_{gj} - \sum_{g=1}^{G} \frac{\mu_g'}{\sigma_g'}
\end{aligned}
\tag{8}
$$

where μ_g' and σ_g' are the moving average of the mean and the standard deviation of intermediate summation. Let $w_{gj}' = \frac{w_{gj}}{\sigma_g'}$, and $b' = \sum_{g=1}^{G} \frac{\mu_g'}{\sigma_g'}$. Because of Equation 2, FBN can also be fused into the linear transformation as

$$
\begin{aligned}
\widehat{x} &= \sum_{g=1}^{G}\sum_{j=1}^{N/G} w_{gj}' x_{gj} - b' \\
&= \sum_{i=1}^{N} w_i' x_i - b'
\end{aligned}
\tag{9}
$$

In summary, we can fuse FBN into the linear transformation at inference time. By this way, we do not need to store the intermediate state of the summation. That is to say, there is no computation and memory overhead at inference time.

4   Network Architecture

In this section, we describe the architecture of our light-weight network based on Fine-grained Batch Normalization, called Finet. Figure 2 shows the building blocks of Finet. Similar as other light-weight networks, we use the depthwise convolution to reduce the FLOPs and the parameters. The main difference of our block is that we use FBN in the 1x1 pointwise convolution instead of BN. As pointed in [47] [49] [48], wider layer makes more powerful representation, but brings more FLOPs and parameters. There are many ways to reduce the FLOPs and parameters of wide layer, e.g. bottleneck [8] or inverted bottleneck [33], group convolution [49]. Finet normalizes the intermediate channels, which are G times wider than final channels. Take a convolutional layer as example, as shown in Figure 3, we can use group convolution to implement FBN. The channels are expanded by G times, then the expanded channels are normalized and summarized. It can be seen as an inverted bottleneck mechanism. However, this bottleneck has only one convolutional layer. As pointed in [26], non-equal channel width and group convolution increase memory access cost. Fortunately, FBN can be fused into convolution at inference time. After fusion, Finet uses the standard convolution with equal channel width, thus makes the inference more efficient. That is very helpful for deploy on the mobile and embedded device which has limit computation and memory resource, or online service environment which is sensitive to the latency.
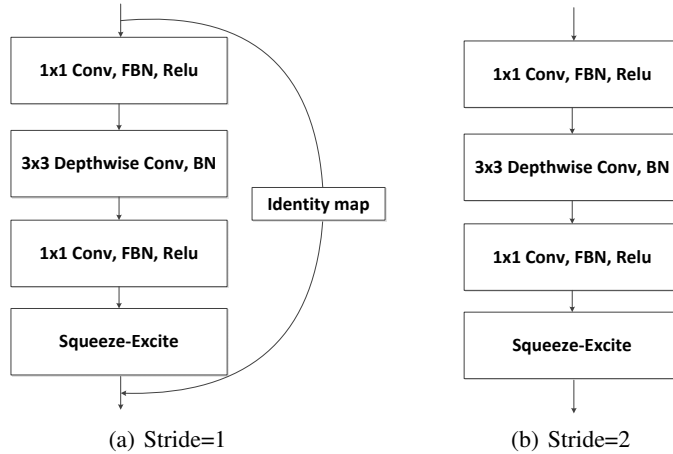

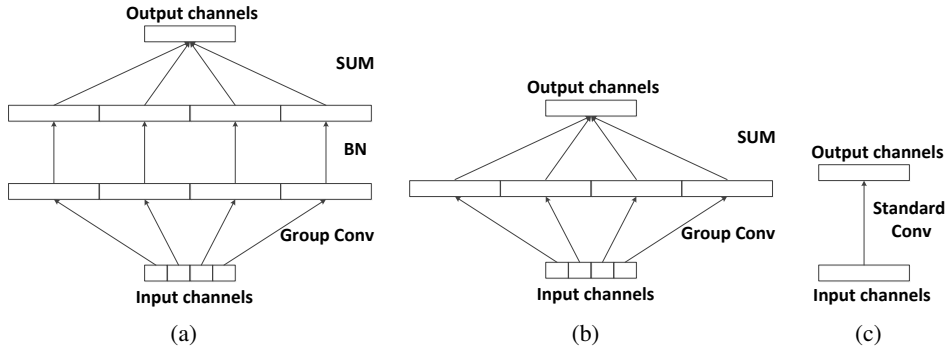Figure 2: The building blocks of Finet

Figure 3: The implementation of Fine-grained Batch Normalization. (a) FBN can be seen as an inverted bottleneck mechanism at training time. (b) Normalization can be fused into convolution at inference time. (c) The standard convolution can be used at inference time. (b) and (c) are equal.

FBN can also be treated as a procedure of splitting, transforming, and aggregating, which is the key design philosophy of ResNeXt [45]. Different from ResNeXt which uses multi-layer network as the transforming, we use normalization as the transforming. Comparing to ResNeXt, we do not need to carry out splitting, transforming, and aggregating at inference time because of normalization fusion.

We can add Squeeze-Excite module [12] which is a light-weight attention mechanism and widely used in other light-weight networks, e.g. Mnasnet [39], MobileNetV3 [10]. For simplicity, we do not use identity map in the block where the stride of depthwise convolution is 2.

Table 1 shows the overall architecture of Finet, for small and large levels of complexities. We adopt the architecture similar with ShufflenetV2. The architectures of the series of MobileNet are more heterogeneous, thus bring difficulties for optimization of performance and memory usage. For Finet, there are two differences from ShufflenetV2: 1) the building blocks shown in Figure 2 are used in each stage, 2) there is an additional fully connected layer before the classifier layer. The fully connected layer is prone to overfitting since it takes much parameters. As a result, modern heavy-weight convolutional networks [21] [8] [14] [38] try to avoid using fully connected layer except the final classifier layer. However, light-weight networks usually suffer from underfitting rather than overfitting. For light-weight networks, fully connected layer brings much parameters but with little computation overhead. The latest light-weight network MobilenetV3 [10] also uses an additional fully connected layer to redesign expensive layers. Squeeze-Excite module also utilizes the fully connected layer to enhance the representative ability with little computation overhead.

Table 1: Overall architecture of Finet, for small and large levels of complexities. The architecture is similar with ShufflenetV2, except 1) the building blocks of Finet are used in each stage, 2) there is an additional fully connected layer before the classifier layer.

Layer	Output size	Ksize	Stride	Repeat	Output channels	
					Small	Large
Image	224x224				3	3
Conv1	112x112	3x3	2	1	24	24
MaxPool	56x56	3x3	2			
Stage2	28x28		2	1	30	100
	28x28		1	3		
Stage3	14x14		2	1	60	200
	14x14		1	7		
Stage4	7x7		2	1	120	400
	7x7		1	3		
Conv5	7x7	1x1	1	1	1024	1024
GlobalPool	1x1	7x7		1	1024	1024
FC1	1x1			1	1024	1024
FC2	1x1			1	1000	1000

5 Experiment

5.1 ImageNet

ImageNet classification dataset [30] has 1.28M training images and 50,000 validation images with 1000 classes. We use Pytorch in our experiments. To augment data, we use the same procedure as the official examples of Pytorch [1]. The training images are cropped with random size of 0.08 to 1.0 of the original size and a random aspect ratio of 3/4 to 4/3 of the original aspect ratio, and then resized to 224x224. Then random horizontal flipping is made. The validation image is resized to 256x256, and then cropped by 224x224 at the center. Each channel of the input is normalized into 0 mean and 1 std globally. We use SGD with 0.9 momentum, and 4e-5 weight decay. Four TITAN Xp GPUs are used to train the networks. The batch size is set to 512. We use linear-decay learning rate policy (decreased from 0.2 to 0). Dropout with 0.2 is used in the last two fully connected layers. We train the networks with 320 epochs. For Squeeze-Excite module, we set the number of hidden unit to 200.

Figure 4 and Table 2 show the results of ImageNet classification. The baseline network uses BN. Actually, BN is equal to FBN when $G = 1$. As G increases, Finet achieves higher accuracy with little overhead of FLOPs and parameters. Moreover, Squeeze-Excite module enhances the accuracy. With Squeeze-Excite and $G = 4$, small Finet achieves 65.706% accuracy with 43M FLOPs, and large Finet achieves 73.786% accuracy with 303M FLOPs. Finet outperforms ShuffleNetV2 (60.3% with 41M FLOPs, and 72.6% with 299M FLOPs) and MobileNetV2 (58.2% with 43M FLOPs, and 72.0% with 300M FLOPs) which are also manual designed architectures. Finet takes more parameters than

ShuffleNetV2 and MobileNetV2. The last two fully connected layers take about 2M parameters. Small version of Finet achieves higher accuracy than the corresponding model of MobileNetV3, while MobileNetV3 and MnasNet achieve higher accuracies for the large version. MobileNetV3 and MnasNet are auto-searched architectures. Neural Architecture Search (NAS) reduces the demand for experienced human experts comparing to hand-drafted design. It can find the optimal combination of existing technique units, but can not invent new techniques. Moreover, the searched architectures are often more heterogeneous, and are less efficient than the homogeneous architectures, as shown in the next Section 5.2.
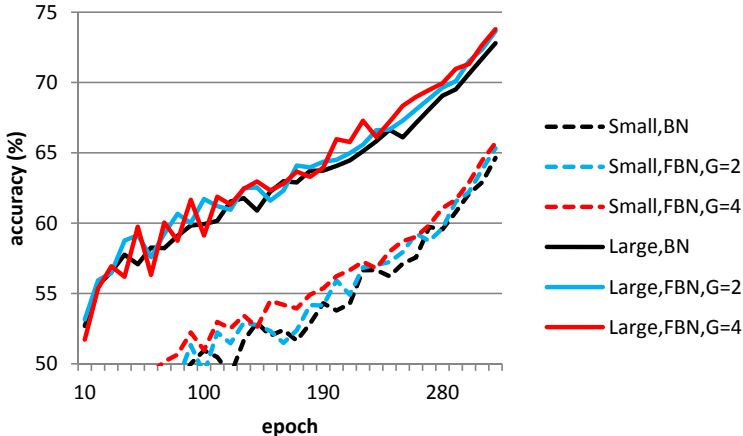

Figure 4: The validation accuracy of Finet with Squeeze-Excite, vs. numbers of training epochs

Table 2: The FLOPs, Parameters (Paras), and Accuracies (Acc) of different networks on ImageNet validation data. BN: Batch Normalization, FBN: Fine-grained Batch Normalization, SE: Squeeze-Excite.

	Small			Large		
	FLOPs (Million)	Paras (Million)	Acc (%)	FLOPs (Million)	Paras (Million)	Acc (%)
baseline (BN)	42.6	2.388	61.710	301.8	4.434	71.256
Finet (FBN,G=2)	42.6	2.392	62.988	301.8	4.448	72.410
Finet (FBN,G=4)	42.6	2.402	63.338	301.8	4.477	72.746
baseline+SE (BN)	43.0	2.809	64.632	303.2	5.812	72.790
Finet+SE (FBN,G=2)	43.0	2.813	65.322	303.2	5.827	73.642
Finet+SE (FBN,G=4)	43.0	2.827	**65.706**	303.2	5.855	**73.786**
ShuffleNetV2 [26]	41	1.4	60.3	299	3.5	72.6
MobileNetV2 [33] [3]	43	1.66	58.2	300	3.4	72.0
MobileNetV3 [10]	44	2.0	65.4	219	5.4	**75.2**
MnasNet [39]	-	-	-	312	3.9	**75.2**

5.2 Inference Speed

We evaluate the inference speed of different light-weight networks on three mobile phones: Samsung Galaxy s10e, Huawei Honor v20, and Vivo x27. Galaxy s10e and Vivo x27 equip the mobile SoC of Qualcomm SnapDragon. SnapDragon uses a heterogeneous computing architecture to accelerate the AI applications. AI Engine of SnapDragon consists of Kryo CPU cores, Adreno GPU and Hexagon DSP. Honor v20 equips HiSilicon Kirin SoC. Different from SnapDragon, Kirin introduces a specialized neural processing unit (NPU) to accelerate the AI applications. The detailed configurations of the devices are shown in Table 3.

We use Pytorch Mobile [2] to deploy the networks on mobile phones. We compare the large versions of Finet, ShufflenetV2, MobileNetV2, MobileNetV3, and MnasNet. Normalization are omitted since it can be fused into convolution operation at inference time. Each model infers 1000 validation

Table 3: The configurations of the measured devices

Device	Soc	CPU	AI Accelerator	RAM
Galaxy s10e	Snapdragon 855	Kryo 485, 2.84 GHz, 7nm	Adreno 640 GPU Hexagon 690 DSP	6GB
Honor v20	Kirin 980	Cortex-A76, 2.6 GHz, 7nm	Cambricon NPU	8GB
Vivo x27	Snapdragon 710	Kryo 360, 2.2 GHz, 10nm	Adreno 616 GPU Hexagon 685 DSP	8GB

images sequentially, and the average throughput is counted. Table 4 shows the results of different models. On all of the three phones, Finet without Squeeze-Excite module is the fastest model, and the second fastest model is Finet with Squeeze-Excite module. The heterogeneous models (MobileNetV2, MobileNetV3, and MnasNet) are slower than homogeneous models (Finet, ShufflenetV2)

Table 4: Inference speed (images per second)

	Galaxy s10e	Honor v20	Vivo x27
Finet, without SE	**7.80**	**6.34**	**2.45**
Finet, with SE	**7.69**	**6.17**	**2.37**
ShufflenetV2	7.35	5.99	2.09
MobileNetV2	3.98	3.65	1.42
MobileNetV3	5.65	4.82	1.79
MnasNet	3.77	3.74	1.48

5.3 CIFAR

CIFAR-10 [19] is a dataset of natural 32x32 RGB images in 10 classes with 50, 000 images for training and 10, 000 for testing. CIFAR-100 is similar with CIFAR-10 but with 100 classes. To augment data, the training images are padded with 0 to 36x36 and then randomly cropped to 32x32 pixels. Then randomly horizontal flipping is made. Each channel of the input is normalized into 0 mean and 1 std globally. Large Finet is evaluated in this section. To adapt Finet to CIFAR datasets, the stride of Conv1 is set to 1, and omit the MaxPool in Table 1. We use SGD with 0.9 momentum, and 5e-4 weight decay. All models are trained on one TITAN Xp GPU. The batch size is set to 128. Learning rate is set to 0.1, and decreased 10 times at epoch 100 and 150. We train the networks with 200 epochs. For Squeeze-Excite module, we set the number of hidden unit to 200.

Table 5 shows the accuracies of Finet on CIFAR-10/100 with different groups. FBN achieves better accuracy than BN (when $G = 1$, FBN is equal to BN). When $G = 8$, Finet achieves the highest accuracy of 93.674% on CIFAR-10, and 77.202% on CIFAR-100. The affine transformation in normalization is define as $y = \gamma \widehat{x} + \beta$, where γ and β are learned parameters for each channel in convolutional networks. Since the intermediate channels of FBN are wider than BN, there are more parameters in the affine transformation. To analyze how FBN improve the performance, we also evaluate FBN without affine transformation. In that case, there is no extra parameter comparing to BN. As shown in Table 5, FBN still achieves better accuracy than BN without affine transformation. Table 6 shows the accuracies of Finet on CIFAR-10/100 by fixing the number of input channels per group. Because the layers have different channel numbers, the group number G changes across layers in this setting. Generally, less input channels per group make more fine-grained training and increase the accuracy.

Table 5: The accuracies (%) of Finet on CIFAR-10/100 with different groups

		G=1	G=2	G=4	G=6	G=8
CIFAR-10	with affine	92.584	93.082	93.532	93.286	**93.674**
	without affine	91.958	93.162	92.478	92.564	**93.430**
CIFAR-100	with affine	75.624	76.752	76.856	76.716	**77.202**
	without affine	74.216	74.382	74.486	**74.772**	74.734

Finally, we evaluate ResNet18 and ResNet50 with FBN on CIFAR datasets. The batch size is set to 64 since training ResNet with FBN consumes a lot of GPU memories. Even decreasing the batch

8

Table 6: The accuracies (%) of Finet on CIFAR-10/100 with different input channels per group

	C/G=20	C/G=50	C/G=100
CIFAR-10	93.470	**93.606**	93.122
CIFAR-100	**76.926**	76.850	75.674

size, there is no enough memory for training ResNet50 with $G = 4$ on single TITAN xp. Thus we only evaluate ResNet50 with $G = 2$. Table 7 shows the accuracies of Resnet18 and ResNet50 with FBN on CIFAR-10/100. The results show that FBN also achieves better accuracy than BN ($G = 1$) for training heavy-weight networks.

Table 7: The accuracies (%) of ResNet with FBN on CIFAR-10/100

		G=1	G=2	G=4
CIFAR-10	ResNet18	94.512	**95.012**	94.962
	ResNet50	94.644	**95.040**	-
CIFAR-100	ResNet18	76.866	**77.754**	77.726
	ResNet50	77.796	**78.122**	-

6 Conclusion

In this paper, we propose a new normalization, Fine-grained Batch Normalization (FBN), and a novel light-weight network based on FBN, called Finet. At training time, the convolutional layer with FBN can be seen as an inverted bottleneck mechanism. At inference time, Finet uses the standard convolution with equal channel width after normalization fusion, thus makes the inference more efficient. We show the effectiveness and efficiency of Finet in our experiments.

References

[1] Pytorch examples. https://github.com/pytorch/examples/tree/master/imagenet.

[2] Pytorch mobile. https://pytorch.org/mobile/home/.

[3] Tensorflow models. https://github.com/tensorflow/models/tree/master/research/slim/nets/mobilenet.

[4] Devansh Arpit, Yingbo Zhou, Bhargava U Kota, and Venu Govindaraju. Normalization propagation: A parametric technique for removing internal covariate shift in deep networks. *arXiv preprint arXiv:1603.01431*, 2016.

[5] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.

[6] Nils Bjorck, Carla P Gomes, Bart Selman, and Kilian Q Weinberger. Understanding batch normalization. In *Advances in Neural Information Processing Systems*, pages 7694–7705, 2018.

[7] Tim Cooijmans, Nicolas Ballas, César Laurent, Çağlar Gülçehre, and Aaron Courville. Recurrent batch normalization. *arXiv preprint arXiv:1603.09025*, 2016.

[8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.

[9] Elad Hoffer, Ron Banner, Itay Golan, and Daniel Soudry. Norm matters: efficient and accurate normalization schemes in deep networks. In *Advances in Neural Information Processing Systems*, pages 2160–2170, 2018.

[10] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for mobilenetv3. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1314–1324, 2019.

[11] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *ArXiv*, abs/1704.04861, 2017.

[12] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7132–7141, 2018.

[13] Gao Huang, Shichen Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Condensenet: An efficient densenet using learned group convolutions. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2752–2761, 2017.

[14] Gao Huang, Zhuang Liu, and Kilian Q. Weinberger. Densely connected convolutional networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269, 2017.

[15] Forrest N. Iandola, Matthew W. Moskewicz, Khalid Ashraf, Song Han, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1mb model size. *ArXiv*, abs/1602.07360, 2017.

[16] Sergey Ioffe. Batch renormalization: Towards reducing minibatch dependence in batch-normalized models. In *Advances in neural information processing systems*, pages 1945–1953, 2017.

[17] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 448–456, 2015.

[18] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. In *Advances in neural information processing systems*, pages 971–980, 2017.

[19] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.

[20] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

[21] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *CoRR*, abs/1312.4400, 2014.

[22] Chunjie Luo, Jianfeng Zhan, Lei Wang, and Wanling Gao. Extended batch normalization. *ArXiv*, abs/2003.05569, 2020.

[23] Chunjie Luo, Jianfeng Zhan, Lei Wang, and Qiang Yang. Cosine normalization: Using cosine similarity instead of dot product in neural networks. *arXiv preprint arXiv:1702.05870*, 2017.

[24] Ping Luo, Jiamin Ren, Zhanglin Peng, Ruimao Zhang, and Jingyu Li. Differentiable learning-to-normalize via switchable normalization. *arXiv preprint arXiv:1806.10779*, 2018.

[25] Ping Luo, Xinjiang Wang, Wenqi Shao, and Zhanglin Peng. Towards understanding regularization in batch normalization. 2018.

[26] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. *ArXiv*, abs/1807.11164, 2018.

[27] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.

[28] Hyeonseob Nam and Hyo-Eun Kim. Batch-instance normalization for adaptively style-invariant neural networks. In *Advances in Neural Information Processing Systems*, pages 2558–2567, 2018.

[29] Mengye Ren, Renjie Liao, Raquel Urtasun, Fabian H Sinz, and Richard S Zemel. Normalizing the normalizers: Comparing and extending network normalization schemes. *arXiv preprint arXiv:1611.04520*, 2016.

[30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

[31] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.

[32] Tim Salimans and Diederik P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in Neural Information Processing Systems*, pages 901–901, 2016.

[33] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018.

[34] Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization? In *Advances in Neural Information Processing Systems*, pages 2483–2493, 2018.

[35] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2015.

[36] Saurabh Singh and Abhinav Shrivastava. Evalnorm: Estimating batch normalization statistics for evaluation. *arXiv preprint arXiv:1904.06031*, 2019.

[37] Ke Sun, Mingjie Li, Dong Liu, and Jingdong Wang. Igcv3: Interleaved low-rank group convolutions for efficient deep neural networks. In *BMVC*, 2018.

[38] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016.

[39] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2815–2823, 2018.

[40] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019.

[41] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.

[42] Guangrun Wang, Ping Luo, Xinjiang Wang, Liang Lin, et al. Kalman normalization: Normalizing internal representations across network layers. In *Advances in Neural Information Processing Systems*, pages 21–31, 2018.

[43] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018.

[44] Guotian Xie, Jingdong Wang, Ting Zhang, Jian-Huang Lai, Richang Hong, and Guo-Jun Qi. Igcv2: Interleaved structured sparse convolutional neural networks. *ArXiv*, abs/1804.06202, 2018.

[45] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5987–5995, 2017.

[46] Xiaoyong Yuan, Zheng Feng, Matthew Norton, and Xiaolin Li. Generalized batch normalization: Towards accelerating deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1682–1689, 2019.

[47] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *ArXiv*, abs/1605.07146, 2016.

[48] Ting Zhang, Guo-Jun Qi, Bin Xiao, and Jingdong Wang. Interleaved group convolutions. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 4383–4392, 2017.

[49] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6848–6856, 2018.

[50] Barret Zoph, V. Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8697–8710, 2018.