

CondenseNet V2: Sparse Feature Reactivation for Deep Networks

Le Yang^{1*} Haojun Jiang^{1*} Ruojin Cai² Yulin Wang¹ Shiji Song¹ Gao Huang^{1†} Qi Tian³

¹Department of Automation, Tsinghua University, Beijing, China

Beijing National Research Center for Information Science and Technology (BNRist),

²Cornell University ³Huawei Cloud & AI

{yangle15, jhj20, wang-y119}@mails.tsinghua.edu.cn, rc844@cornell.edu,

{shijis, gaohuang}@tsinghua.edu.cn, tian.qil@huawei.com

Abstract

Reusing features in deep networks through dense connectivity is an effective way to achieve high computational efficiency. The recent proposed CondenseNet [14] has shown that this mechanism can be further improved if redundant features are removed. In this paper, we propose an alternative approach named sparse feature reactivation (SFR), aiming at actively increasing the utility of features for reusing. In the proposed network, named CondenseNetV2, each layer can simultaneously learn to 1) selectively reuse a set of most important features from preceding layers; and 2) actively update a set of preceding features to increase their utility for later layers. Our experiments show that the proposed models achieve promising performance on image classification (ImageNet and CIFAR) and object detection (MS COCO) in terms of both theoretical efficiency and practical speed.

1. Introduction

Deep convolutional neural networks (CNNs) have achieved remarkable success in the past few years [30, 7, 15]. However, their state-of-the-art performance is usually fueled with sufficient computational resources, which hinders deploying deep models on low-compute platforms, e.g., mobile phones and Internet of Things (IoT) products. This issue has motivated a number of researchers on designing efficient CNN architectures [15, 3, 11, 40, 24, 31]. Among these efforts, DenseNet [15] is a promising architecture that improves the computational efficiency by reusing early features with dense connections.

Recently, it has been shown that dense connectivity may introduce a large number of redundancies when the network becomes deeper [14]. In a dense network, the output of a layer will never be modified once it is produced. Given that




Figure 1. Different feature reuse patterns in (a) DenseNet [15], (b) CondenseNet [14] and (c) Ours.

shallow features will be repeatedly processed by their following layers, directly exploiting them in deep layers might be inefficient or even redundant. CondenseNet [14] alleviates this problem via strategically pruning less important connections in DenseNet. ShuffleNetV2 [24] shares a similar spirit, where early features are dropped according to layer-distance, leading to an exponentially decaying of long-distance feature re-usage. Although both models show their effectiveness, we hypothesize that straightforwardly abandoning long connections is overly aggressive. These early features which are considered to be “obsolete” at deeper layers may contain useful information, which can benefit network generalization ability, and potentially contribute to a more efficient model if properly utilized.

In this paper, instead of directly discarding obsolete features, we are interested in whether we can revive them to make obsolete features useful again. To this end, we develop a novel module to conduct *feature reactivation*, which learns to update shallow features and enables them to be more efficiently reused by deep layers. Our main idea is illustrated in Figure 1. Compared to DenseNet [15] and CondenseNet [14], where earlier features keep unchanged

*Equal contribution.

†Corresponding author.

throughout the whole feed-forward process, we propose to allow the outputs of a layer to be reactivated by later layers. Such a way keeps feature maps always “fresh” at each dense layer, and therefore the redundancy in dense connections can be largely reduced.

Although the *feature reactivation* procedure effectively reduces the redundancy in dense connections, naively reactivating all features will introduce excessive extra computation, which still hurts the overall efficiency. In fact, it is unnecessary to reactivate all features since a large number of them can be already effectively reused without any change in dense connections, resulting in that only *sparse feature reactivation* (SFR) is required. For this purpose, we develop a cost-efficient SFR module which actively and selectively *reactivates* early features at each layer, using the increments learned from the newly produced feature maps. Importantly, both the features to be updated and the updating formulas are determined *automatically* via learning. During the training process, we first assume all previous features require reactivating, and then gradually remove the reactivation that have less effect on feature re-usage. Moreover, the resulting SFR modules can be converted to efficient group convolutions at test time. As a consequence, the proposed method involves minimal extra computational cost or latency and keeps early features “fresh” even through very deep layers, which leads to a significant efficiency gain.

We implement SFR on the basis of the efficient CondenseNet [14], where the SFR along with the learned group convolutions (LGCs) [14] can be learned compatibly to improve the efficiency of dense networks. The resulting network, CondenseNetV2, are empirically evaluated on image classification benchmarks (ImageNet and CIFAR) and the COCO object detection task. The results demonstrate that SFR significantly boosts the performance by encouraging long-distance feature reusing, and that CondenseNetV2 compares favorably with even state-of-the-art light-weighted deep models. We also show that SFR can be plugged into any CNNs that adopt the concatenation based feature reusing mechanism to further improve their efficiency, such as ShuffleNetV2 [24].

2. Related Work

Efficient network architectures. Designing better network architectures is an effective way to improve the computation efficiency of deep networks. Efficient building units are introduced for light-weighted CNN architectures. For instance, MobileNets [11, 31, 10] propose inverted residuals and linear bottlenecks to build network architectures. Sandglass blocks, which flip the inverted residuals, are developed in MobileNeXts [41]. Cheap operations are developed for generating features in GhostNet [5]. In addition, shuffle layer and learned group convolution (LGC) are employed by ShuffleNets [40, 24] and CondenseNet [14], respectively.

Recent study also shows that developing dynamic neural networks [6] can obviously improve the efficiency of deep models, such as [13, 37, 34]. In this paper, we follow the first line of the research and propose a novel efficient unit named SFR module. The proposed deep models with SFR module retains the simplicity of CondenseNet while significantly improves its accuracy on image classification and detection tasks for mobile applications.

Densely connected neural network. Compared to ResNet [7] and its variants [36, 39], DenseNet architectures [15, 14, 13, 37] can achieve a higher computational efficiency by encouraging feature reuse. However, superfluous re-usage may introduce redundant connections. To address this problem, existing work mainly proposes to remove dense connections to feature maps that are less useful [14, 33], or to discard long-range connections according to a predefined probability [24]. However, as these seemingly redundant connections may have large potential in deep layers if properly utilized, we propose to conduct sparse feature reactivation to deal with the redundant connections rather than pruning them.

Filter pruning. Although the sparsifying procedure in SFR module is related to filter pruning methods [35, 9, 18, 8, 22], our method differs greatly from filter pruning methods in the way dealing with the redundant connections. Instead of removing connections between layers where the feature re-usage is superfluous, our approach aims at building reactivation connections to revive obsolete features to increase their utility. Notably, our reactivation idea is orthogonal to filter pruning, and both are utilized for building CondenseNetV2. Additionally, compared with recent work [26], which proposes to graft new weights to the unimportant filters, the proposed SFR module reactivates obsolete features to improve efficiency.

3. Method

A recently confirmed inefficiency in DenseNet [15] architecture lies in the presence of long-distance connections [14], where the deeper layers seem to consider the early features as “obsolete” ones and ignore them during learning new representations. CondenseNet [14] and ShuffleNetV2 [24] alleviate this inefficiency through strategically pruning redundant connections and exponentially discarding cross-layer connections, respectively. In this paper, we postulate in this paper that directly abandoning shallow features can be an overly aggressive design. To be specific, we find that by involving a learnable *sparse feature reactivation* (SFR) module with a negligible computational cost at each layer, the originally “obsolete” features can be “reactivated” and hence effectively exploited by the later layers. In this section, we first describe the details of the proposed SFR module, and then implement it to build our light-weighted networks.

3.1. Sparse Feature Reactivation

Feature reuse mechanism. We first formulate the feature reusing mechanism introduced in [15]. Assume that a standard network block of L layers produces L feature maps $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_L$, where \mathbf{x}_ℓ is the output of the ℓ -th layer and \mathbf{x}_0 denotes the input feature. Since all previous layers are connected to the ℓ -th layer via the dense connections, the composite function $H_\ell(\cdot)$ of the ℓ -th layer will take all of $\mathbf{x}_0, \dots, \mathbf{x}_{\ell-1}$ as inputs:

$$\mathbf{x}_\ell = H_\ell([\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{\ell-1}]). \quad (1)$$

In CondenseNet [14], learned group convolutions (LGC) are employed in $H_\ell(\cdot)$ to automatically learn the input groupings and remove unimportant connections, while in ShuffleNetV2 [24], inputs of H_ℓ will be dropped according to their distance to layer ℓ . Consequently, both of them remove the superfluous long-distance connections between layers, which are proven effective in term of efficiency. However, given that the output \mathbf{x}_ℓ will never change once it is produced, a side effect is that these seemingly less useful features from shallow layers tend to be permanently discarded by deeper layers. This *static* design may impede exploring more efficient feature reusing mechanisms. To this end, we propose a cost-efficient *sparse feature reactivation* module, enabling obsolete features to be cheaply revived.

Reactivating obsolete features. We start by describing the details of feature reactivation. For ℓ -th layer, we introduce a reactivation module denoted by $G_\ell(\cdot)$. The module takes \mathbf{x}_ℓ as input, and its output \mathbf{y}_ℓ is used to reactive features from preceding layers. In this paper, we define the reactivation operation, $U(\cdot, \cdot)$, as adding¹ the increment \mathbf{y}_ℓ . A dense layer with feature reactivation can be written as

$$\mathbf{x}_\ell^{\text{in}} \leftarrow [\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{\ell-1}], \quad \mathbf{x}_\ell = H_\ell(\mathbf{x}_\ell^{\text{in}}), \quad (2)$$

$$\mathbf{y}_\ell = G_\ell(\mathbf{x}_\ell), \quad \mathbf{x}_\ell^{\text{out}} = U(\mathbf{x}_\ell^{\text{in}}, \mathbf{y}_\ell), \quad (3)$$

$$[\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{\ell-1}] \leftarrow \mathbf{x}_\ell^{\text{out}}, \quad (4)$$

where $\mathbf{x}_\ell^{\text{out}}$ is the reactivated output feature. With $H_\ell(\cdot)$, the ℓ -th layer learns to produce new feature \mathbf{x}_ℓ . Additionally, previous representations ($\mathbf{x}_i, i = 1, \dots, \ell-1$) will be reactivated to increase their utility.

Obviously, it is unnecessary to reactivate all features since a large number of them can be effectively reused without any change (shown in DenseNet [15]). We also empirically observe that dense reactivation will introduce much computation and degrade the overall efficiency of the network. Therefore, we seek to automatically find the features required to be reactivated and merely refresh them. In the following,

¹Other reactivation schemes may also be considered, such as applying channel-wise attention or spatial attention. However, in this paper, we note that a straightforward sum has already achieved good performance.

this aim is formulated by a pruning based approach and can be achieved gradually during training (shown in Figure 2). The resulting architecture is named as sparse feature reactivation (SFR) module.




Figure 2. Sparsify the feature reactivation.

Sparse feature reactivation (SFR). WLOG, we assume that the reactivation module $G_\ell(\cdot)$ consists of a regular 1×1 convolutional layer followed by batch normalization (BN) [16] and a rectified linear unit (ReLU) [27]. The size of the filter weight matrix \mathbf{F} in $G_\ell(\cdot)$ is represented as (O, I) , where O and I denote the number of output and input channels². In each $G_\ell(\cdot)$, we divide \mathbf{x}_ℓ into G groups, and the \mathbf{F} is split correspondingly among the input channel dimension to obtain G groups $\mathbf{F}^1, \dots, \mathbf{F}^G$, and each of them has the size of $(O, I/G)$. To sparsify the reactivation connections, we further define a sparse factor S , which may differ from G , and allow each group to only select $\frac{O}{S}$ output channels to reactivate after training.

During training, in each $G_\ell(\cdot)$, the connection pattern is controlled by G binary masks. Therefore, we aim at learning these binary masks $\mathbf{M}^g \in \{0, 1\}^{O \times \frac{I}{G}}, g = 1, \dots, G$ to screen out unnecessary connections in \mathbf{F}^g by zeroing the corresponding values. In other words, the weight of g -th group can be obtained by $\mathbf{M}^g \odot \mathbf{F}^g$, where \odot denotes the element-wise multiplication.

We then introduce how to train a network with SFR modules in an end-to-end manner. Inspired by [14], the whole training process consists of $S - 1$ sparsification stages followed by an optimization stage. Assuming that E denotes the total number of training epochs, we set the training epochs of each sparsification stage to $\frac{E}{2(S-1)}$ and optimization stage to $\frac{E}{2}$. During training, the SFR module first reactivates all features, and then gradually removes the superfluous connections. Therefore, at the beginning of training, we set all \mathbf{M}^g to all-ones matrices, thus all input feature maps in g -th group are connected with all output features. During sparsification, the importance of reactivating i -th output within g -th group is measured by the L1-norm of the corresponding weights $\sum_{j=1}^{I/G} |\mathbf{F}_{i,j}^g|$. At the end of each sparsification stage, $\frac{O}{S}$ output features whose L1-norms are smaller than others are pruned in g -th group, and $\mathbf{M}_{i,j}^g$ is set to zero for all j in g -th group for each pruned output feature map i . Note that, if i -th output feature map is pruned from every input group,

²We can apply max-pooling on the absolute value of the 4D weights, $\mathbf{F} \in \mathbb{R}^{O \times I \times k \times k}$, to generate the matrix with the size of (O, I) when dealing with larger convolutional kernels.

the i -th feature map in y will equal to 0 , which implies that the i -th feature map from previous layers do not need to reactivate. Therefore, after training, each input group will only update the outputs with a portion of $1/S$, which means that for final M^g , we have $\sum_{i=1}^O M_{i,j}^g = \frac{O}{S}$. The higher the value of S , the sparser the connection pattern is.

Convert to standard group convolution. At test time, our SFR model can be implemented using a standard group convolution and an index layer, allowing for efficient computation in practice. This is illustrated in Figure 3: the converted group convolution contains G groups with output and input channels as $(\frac{OG}{S}, I)$. After generating the intermediate features with the group convolution, the index layer is applied to rearrange the features by their indices to obtain the y_ℓ . Note that the intermediate features with the same index will first be summed and then arranged according to the index.




Figure 3. Convert to standard group convolution ($S=3$ and $G=3$).

3.2. Architecture Design

Architecture of CondenseNetV2. Based on the proposed SFR module, the new dense layer of CondenseNetV2 is shown in Figure 4 (right), which is designed on the basis of CondenseNet [14]. In the proposed architecture, the LGC first selects important connections and the new representations x_ℓ are generated based on these selected features using Eq. (2). Then, the SFR module takes x_ℓ as input and learns to reactivate the obsolete representations. The refreshed features can be derived by Eq. (3). Following [14], we shuffle the output channels of each convolutional layer to ensure communication between different groups. It is worth knowing that the CondenseNetV2 is essentially different from CondenseNet [14]: the outputs of each layer in CondenseNet [14] will never change once it is produced. Therefore, the potential re-usage of previous features can be blocked. In contrast, old features can be reactivated in each layer of CondenseNetV2, resulting in a more efficient and effective feature reuse mechanism.

The architecture of CondenseNetV2 follows the exponentially increasing growth rate and fully dense connectivity design principle [14]. Based on the newly designed SFR-DenseLayer, we develop our CondenseNetV2 as presented in Table 1. The squeeze and excite (SE) module [12] and hard-swish nonlinearity function (HS) are also applied following [10]. The presented architecture provides a basic design for reference, further hyper-parameters tuning or network architecture searching can further boost the performance.




Figure 4. A dense layer in CondenseNet (left), and CondenseNetV2 (right). (LGC: learned group convolution; GC: group convolution)

Table 1. Network architecture of CondenseNetV2. The number of layers and the growth rate of i -th dense block are denoted by d_i and k_i , respectively. SE and HS denote whether using SE and HS module in this block.

Input	Operator	d	k	SE	HS
224×224	Conv2d 3×3 (stride 2)	-	-	-	-
112×112	SFR-DenseLayer	d_1	k_1	-	-
112×112	AvgPool 2×2 (stride 2)	-	-	-	-
56×56	SFR-DenseLayer	d_2	k_2	-	-
56×56	AvgPool 2×2 (stride 2)	-	-	-	-
28×28	SFR-DenseLayer	d_3	k_3	-	1
28×28	AvgPool 2×2 (stride 2)	-	-	-	-
14×14	SFR-DenseLayer	d_4	k_4	1	1
14×14	AvgPool 2×2 (stride 2)	-	-	-	-
7×7	SFR-DenseLayer	d_5	k_5	1	1
1×1	AvgPool 7×7	-	-	-	-
1×1	Conv2d 1×1	-	-	1	-
1×1	FC	-	-	-	-

3.3. Sparse Feature Reactivation in ShuffleNetV2

As SFR is able to be plugged into any CNNs with the feature reusing mechanism, we claim that ShuffleNetV2 [24] can also benefit from the proposed SFR module. The implementation details are illustrated in Figure 5. We refer to the modified ShuffleNetV2 with sparse feature reactivation as SFR-ShuffleNetV2. Note that in SFR-ShuffleNetV2, only basic units conduct the feature reactivation, and the units for spatial down sampling keep unchanged. The detailed architecture is provided in the Appendix.

4. Experiments

We empirically demonstrate the effectiveness of the proposed SFR module and CondenseNetV2 on image classification and object detection tasks, and compare with state-of-the-art light-weighted CNN architectures. Code is available at <https://github.com/jianghaojun/CondenseNetV2>.

Dataset. Experiments are conducted on several benchmark visual datasets, including CIFAR-10 and CIFAR-100 [17],




Figure 5. A building unit in ShuffleNetV2 (left), and the ShuffleNetV2 implemented with SFR module (right). (DWC: depth-wise convolution)

ImageNet (ILSVRC2012 [4]), and MS COCO object detection benchmark [21].

The CIFAR-10 and CIFAR-100 datasets consist of 32×32 RGB images, with 10 and 100 classes of natural scene objects, respectively. Both CIFAR datasets contain 50,000 training images and 10,000 test images. On the two CIFAR datasets, following [15], we apply a set of transformations to augment the training set. The ImageNet (ILSVRC2012 [4]) classification dataset contains 1.2 million training images and 50,000 validation images, with 1000 classes. Data augmentation schemes are applied at training and we adopt a 224×224 center crop at test time. On MS COCO dataset, following [29, 19], we use the *trainval35k* split as training data and report the results in mean Average Precision (mAP) on *minival split* with 5000 images.

4.1. Efficiency of Sparse Feature Reactivation

In this subsection, we conduct a series of experiments on densely connected networks with sparse feature reactivation to verify its effectiveness.

Reactivated features. As the proposed SFR module is designed for reactivating the redundant features to improve the network efficiency, a natural question is whether the utility of these obsolete features indeed shows great importance at later layers after being reactivated. To investigate this question, Figure 6 and 7 visualize the learned weights for CondenseNet [14] and CondenseNetV2 to verify whether those reactivation will indeed encourage feature reuse.

Figure 6 shows detailed weight strength (averaged absolute value of non-pruned weights) between a filter group of a certain layer (corresponding to a column in the figure) and an input feature map (corresponding to a row in the figure). For each layer, there are four filter groups (consecutive columns). Red dots are connections with significant contributions and




Figure 6. Norm of weights between layers per filter group of CondenseNet and CondenseNetV2 trained on CIFAR-10.

white dots are connections that have been pruned. One can observe that connections in CondenseNet are more concentrated in neighbor layers, while long-distance connections appear more frequently in CondenseNetV2 (shown by dense colored dots in the top-right corner). This implies that later layers make more use of feature maps produced by early layers in CondenseNetV2 than in CondenseNet.

Figure 7 shows the overall connection strength between two layers in the CondenseNet and CondenseNetV2. The color shows the L1-norm of weights between layers and red means large weight. We notice that the top-right parts of figures for CondenseNetV2 are more brilliant than these corresponding areas of figures for CondenseNet, which implies that the utility of early features largely increases in CondenseNetV2. As the performance of CondenseNetV2 is shown to be superior to CondenseNet, we can conclude that the dense network will benefit from sparse feature reactivation. This validates our hypothesize: although the features produced by early layers seem unimportant at deep layers in dense networks, they may have potential after being reactivated. Moreover, from the results in Figure 7 in (b) and (d), we further observe that early features are more frequently utilized at deep layers in the model trained on ImageNet than the model trained on CIFAR-10, from which we can infer that the reactivated early features show more importance in complicate tasks.

Necessity of sparse feature reactivation. Although the dense feature reactivation can improve the network performance, the involved extra computation is much, which degrades the overall efficiency of the network. Therefore, it is necessary to make the reactivation sparse. We conduct experiments on CIFAR-10 using CondenseNetV2 with $S = 1$ and $S > 1$. The experimental results are shown in Figure 8 (b). From the results, we observe that the dense feature reactivation ($S = 1$) becomes inefficient due to the heavy extra computational overheads. On the contrary, conducting sparse feature reactivation ($S > 1$) can deal with the aforementioned problems effectively, and therefore, can boost the network performance with minor extra computation.

The hyper-parameters in SFR module. We conduct the experiments with CondenseNetV2 on CIFAR-10 for evalua-




Figure 7. Norm of weights between layers per filter block of CondenseNet ((a) and (c)) and CondenseNetV2 ((b) and (d)). The models are trained on CIFAR-10 ((a) and (b)) and ImageNet ((c) and (d)).




Figure 8. (a): SFR module with different group number (G) in CondenseNetV2. (b): CondenseNetV2 with different sparse factor (S) for SFR module.

tion. We fix the group number and the condense factor for LGC layers, and we only compare different settings for SFR module. In Figure 8 (a), the sparse factor is fixed to 4 and we show the effect of group number G which actually does not affect the FLOPs of the network. One can observe that as G increases, the performance improves gradually. This is due to that a finer-grained sparsification, which corresponds to a large G , is usually able to achieve higher efficiency. In Figure 8 (b), we compare CondenseNetV2s with varying sparse factors in SFR module. A network with a sparse factor S of 1 means that all reactivation connections are preserved (dense reactivation). If S is increased to 4, then each layer only keeps a quarter of reactivation connections. The results show that $S = 4$ outperforms other settings. Also, all the settings with $S > 1$ perform better than the setting with $S = 1$, indicates that removing a set of unnecessary connections is indeed important for building efficient CondenseNetV2.

4.2. Experiments on ImageNet

Implementation details. Following the common practice [11, 40, 24], the proposed network has three levels of computational complexity. The CondenseNetV2 with different sizes are summarized in Table 2, where d and k are number of layers and growth rate of each dense block, respectively. In all experiments, we use the same condense factor (C), group number (G), and sparse factor (S) for all LGCs and SFR modules in the network. CondenseNetV2 are trained using the stochastic gradient descent (SGD) optimizer with an initial learning rate of 0.4, the cosine learning rate [23], and The batch size of 1024. To compare with SOTA base-

Table 2. Network configurations for ImageNet models.

Setting	$\{d\}$	$\{k\}$	C, S, G
CondenseNetV2-A	1-1-4-6-8	8-8-16-32-64	8
CondenseNetV2-B	2-4-6-8-6	6-12-24-48-96	6
CondenseNetV2-C	4-6-8-10-8	8-16-32-64-128	8

Table 3. Top-1 and Top-5 classification error rate (%) on ImageNet.

Model	FLOPs	Top-1 err.	Top-5 err.
CondenseNet-A [14]	56M	43.5	20.2
CondenseNetV2-A	46M	35.6	15.8
CondenseNet-B [14]	132M	33.9	13.1
CondenseNetV2-B	146M	28.1	9.7
ShuffleNetV2 1.0 \times [24]	146M	30.6	11.1
SFR-ShuffleNetV2 1.0 \times	150M	29.9	10.9
ShuffleNetV2 1.5 \times [24]	299M	27.4	9.4
SFR-ShuffleNetV2 1.5 \times	306M	26.5	8.6

lines, we implement an Augmented Setting differing from the original setting in [14]. More details can be found in Appendix C.

Results on ImageNet. We conduct experiments on ImageNet to evaluate the effectiveness of the proposed methods. As the SFR module can be deployed in both CondenseNet and ShuffleNetV2, we first compare the original model and the network with SFR module in Table 3. From the results, we can see that the latter clearly exceeds the former. The computational cost of CondenseNetV2-A is 18% lower than CondenseNet-A (46M v.s. 56M). Using SFR on ShuffleNet can also boost the efficiency with minor extra FLOPs (fewer than 3%). Moreover, we conduct ablation studies on ImageNet to show how each additional design benefit the original CondenseNet. The results are shown in Table 4, from which we observe that implementing the proposed SFR can boost the performance of CondenseNet by a large margin.

We further compare our networks with several efficient network architectures designed by handcraft, including MobileNetV2 [31], CondenseNet [14], MobileNeXt [41] and ShuffleNetV2[24]. The results are shown in Figure 9 (a). One can observe that the proposed CondenseNetV2 outperforms these models in terms of the computational efficiency.




Figure 9. (a) Top-1 err. v.s. FLOPs on ImageNet. (b) Inference speed on the ARM processor. (c) Inference speed on iPhone XS Max.

Table 4. Ablations of CondenseNetV2-A on ImageNet.

Model	SFR	SE/HS	Augmented Setting	Top-1 err.
CondenseNet [14]				43.5
CondenseNetV2	✓			39.8(↓3.7)
CondenseNetV2	✓	✓		37.5(↓6.0)
CondenseNetV2	✓	✓	✓	35.6(↓7.9)

In addition to the deep models in Figure 9, networks based on neural architecture search (NAS) are further compared with the proposed networks, including MobileNetV3 [10], RegNetX [28], ProxylessNAS[1] and MnasNet [32]. The results are summarized in Table 5. We group different models according to their computational costs. Importantly, the proposed CondenseNetV2 does not leverage any technique of NAS, however, it can outperform most of the competitive baselines with similar FLOPs.

To show how our method can benefit from NAS, we further deploy our SFR module on ShuffleNetV2+³, a strengthened version of ShuffleNetV2 obtained by one-shot NAS based on ShuffleNet Units. The proposed SFR-ShuffleNetV2+ outperforms both MobileNet V3 and the original ShuffleNetV2+, which confirms the effectiveness of the proposed SFR module. More experiments on ImageNet are provided in Appendix B.

Actual inference time. Since the proposed CondenseNetV2 is designed for edge devices, we further measure the actual inference speed of CondenseNetV2 on an ARM processor⁴ and an iPhone XS Max (with Apple A12 Bionic). The single-thread mode with batch size 1 is used following [10] and we use a 224×224 input image.

On the ARM processor, all models are implemented in PyTorch1.6.0. From Figure 9 (b), one can observe that the proposed CondenseNetV2 achieves faster runtime under the same error compared with other light-weighted deep models.

³<https://github.com/megvii-model/ShuffleNet-Series/tree/master/ShuffleNetV2%2B>

⁴Quad-Core ARM Cortex-A57 MPCore combined with Dual-Core NVIDIA Denver 2 64-Bit CPU.

Table 5. Comparison of Top-1 and Top-5 classification error rate (%) with state-of-the-art efficient deep learning models on ImageNet. (L., M., and S. represent Large, Medium, Small, respectively.)

Model	FLOPs	Params	Top-1 err.	Top5 err.
ShuffleNetV2 0.5× [24]	41M	1.4M	38.9	17.4
0.4 MobileNetV2 [31]	43M	–	43.4	–
MobileNetV2-0.35 [41]	80M	1.8M	35.3	–
CondenseNetV2-A	46M	2.0M	35.6	15.8
ShuffleNetV2 1.0× [24]	146M	2.3M	30.6	11.1
0.75 MobileNetV2 [31]	145M	–	32.1	–
MobileNetV3 L. 0.75× [10]	155M	4.0M	26.7	–
RegNetX [28]	200M	2.7M	31.1	–
MobileNetV2-0.75 [41]	210M	2.5M	28.0	–
ShuffleNetV2+ S.	156M	5.1M	25.9	8.3
CondenseNetV2-B	146M	3.6M	28.1	9.7
SFR-ShuffleNetV2+ S.	161M	5.2M	25.5	8.2
ShuffleNetV2 1.5× [24]	299M	–	27.4	9.4
1.0 MobileNetV2 [31]	300M	3.4M	28.0	9.0
MobileNetV3 L. 1.0× [10]	219M	5.4M	24.8	–
RegNetX [28]	400M	5.2M	27.3	–
MobileNetV2-1.00 [41]	300M	3.4M	26.0	–
MnasNet-A1 [32]	312M	3.9M	24.8	7.5
FE-Net 1.0× [2]	301M	3.7M	27.1	–
ESPNetV2 [25]	284M	3.5M	27.9	–
ProxylessNAS[1]	320M	4.1M	25.4	7.8
ShuffleNetV2+ M.	222M	5.6M	24.3	7.4
CondenseNetV2-C	309M	6.1M	24.1	7.3
SFR-ShuffleNetV2+ M.	229M	5.7M	23.9	7.3

Specifically, our model obtains about 0.5% lower top-1 error than MobileNetV3 with slightly lower latency. It is noteworthy that the power of the tested processor is lower than most smart mobile phones. We believe that such a speed test is necessary: although mobile phones and processors with high performance have been widely deployed and popularized nowadays, the computational resources of most edge devices, such as IoT products, are still highly limited. The proposed CondenseNetV2 outperforms most light-weighted networks in such a resource-limited scenario.

We further test the inference time on an iPhone XS Max, which can be considered as a high performance edge device.

Table 6. Comparison of error rate (%) with other state-of-the-art efficient models on CIFAR-10 and CIFAR-100.

Model	FLOPs	Params	C-10	C-100
ResNet-based				
CP [9]	62M	-	8.20	-
PFEC [18]	90M	0.73M	6.94	-
LECN [22]	124M	1.21M	5.27	23.91
NISP [38]	142M	0.96M	6.88	-
FPGM [8]	121M	-	6.24	-
DenseNet-based				
LECN [22]	190M	0.66M	5.19	25.28
CondenseNet [14]	65M	0.52M	5.00	23.64
CondenseNetV2-110	41M	0.48M	4.65	23.94
CondenseNetV2-146	62M	0.78M	4.35	22.52

Our implementation is based on the Pytorch Mobile⁵. The results are presented in Figure 9 (c), from which we see that the SFR-ShuffleNetV2 outperforms other competitors. Here, only SFR-ShuffleNetV2 are tested because we found that the Pytorch Mobile might have poor support for group convolution operator: although CondenseNetV2-A only has 46M FLOPs, its latency on iPhone is still up to 34.6ms.

4.3. Experiments on CIFAR

Implementation details. We apply SGD to train all the models with similar hyper-parameters setting as in [14]. We use the cosine learning rate annealing with an initial learning rate of 0.1. The training process lasts for 300 epochs with a mini-batch size of 64. Other training settings are the same as the experiments on ImageNet. CondenseNetV2s on CIFAR follow the configuration listed below. The network consists of three dense blocks with the same number of layers, and the resolutions of feature maps are 32×32 , 16×16 , and 8×8 , respectively. The growth rates are set to 8, 16, 32 for each block. The C , S , and G are all set to 4. We modify the number of blocks d in each stage to change the computational complexity of CondenseNetV2. Moreover, we do not implement SE and HS in CondenseNetV2 for CIFAR models, and the last Conv2d 1×1 is also removed.

Results on CIFAR. We show the comparison results of CondenseNetV2s and other competitive baselines in Table 6. The baselines include several recently proposed network pruning algorithms. It can be observed that the CondenseNetV2s outperform all other approaches with lower error rates and less computational costs — indicating that the effectiveness of the proposed feature reuse mechanisms.

4.4. Experiments on MS COCO

MS COCO [21] is used for evaluating the generalization ability of our networks. Following [29], we use the

⁵<https://pytorch.org/mobile/home/>.

Table 7. Results on the MS COCO dataset.

Detection Framework	Backbone	Backbone FLOPs	mAP
Faster R-CNN	ShuffleNetV2 $0.5 \times$ [24]	41M	22.1
	CondenseNetV2-A	46M	23.5
	ShuffleNetV2 $0.5 \times$ [24]	146M	27.4
	CondenseNetV2-B	146M	27.9
RetinaNet	MobileNetV2 $1.0 \times$ [31]	300M	30.6
	ShuffleNetV2 $1.5 \times$ [24]	299M	30.2
	SFR-ShuffleNetV2 $1.5 \times$	306M	30.7
	CondenseNetV2-C	309M	31.4
RetinaNet	MobileNetV2 [31]	300M	29.7
	ShuffleNetV2 $1.5 \times$ [24]	299M	29.1
	CondenseNetV2-C	305M	31.7

trainval35k split as training data and report the results in mean Average Precision (mAP) on *minival split*. Faster R-CNN [29] with Feature Pyramid Networks (FPN) [19] and RetinaNet [20] are implemented as detection frameworks. Only backbone networks are replaced during experiments. Models are pretrained on ImageNet and then finetuned on the detection task. During finetuning, we train all models using SGD for 12 epochs. The input images are resized to a short side of 800 and a long side not exceed 1333. The backbone FLOPs are calculated with 224×224 input size following [24]. The detection results are shown in Table 7. As we can see, with comparable computational cost, our CondenseNetV2-C achieves higher mAP compared with ShuffleNetV2 and MobileNetV2, both on RetinaNet and Faster R-CNN frameworks.

5. Conclusion

In this paper, we proposed a novel *sparse feature reactivation* module, which can strategically reactivate a set of previous features to increase their utility for later layers. Importantly, the features to be reactivated are not *pre-defined*, but *learned* automatically during training. Due to the sparsity of the feature reactivation, this procedure can be highly computational-efficient. Therefore, the resulting model, CondenseNetV2, based on the proposed SFR module can achieve high efficiency during inference. Encouraging results have been obtained on the image classification tasks (ImageNet and CIFAR) and the COCO object detection task, without resorting to neural architecture search.

Acknowledgement

This work is supported in part by the National Key R&D Program of China (2020AAA0105200), the National Natural Science Foundation of China (61906106, 62022048), the Institute for Guo Qiang of Tsinghua University and Beijing Academy of Artificial Intelligence.

References

- [1] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *ICLR*, 2019.
- [2] 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 *CVPR*, 2019.
- [3] François Chollet. Xception: Deep learning with depthwise separable convolutions. *arXiv preprint arXiv:1610.02357*, 2016.
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [5] Kai Han, Yunhe Wang, Qi Tian, Jianyuan Guo, Chunjing Xu, and Chang Xu. Ghostnet: More features from cheap operations. In *CVPR*, 2020.
- [6] Yizeng Han, Gao Huang, Shiji Song, Le Yang, Honghui Wang, and Yulin Wang. Dynamic neural networks: A survey. *arXiv preprint arXiv:2102.04906*, 2021.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [8] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *CVPR*, 2019.
- [9] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *ICCV*, 2017.
- [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 mobilenet3. In *ICCV*, 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 preprint arXiv:1704.04861*, 2017.
- [12] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.
- [13] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Q Weinberger. Multi-scale dense networks for resource efficient image classification. In *ICLR*, 2018.
- [14] Gao Huang, Shichen Liu, Laurens Van der Maaten, and Kilian Q Weinberger. Condensenet: An efficient densenet using learned group convolutions. In *CVPR*, 2018.
- [15] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.
- [16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [17] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. In *Tech Report*, 2009.
- [18] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *ICLR*, 2017.
- [19] T. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [20] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014.
- [22] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *ICCV*, 2017.
- [23] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *ICLR*, 2017.
- [24] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. ShuffleNet v2: Practical guidelines for efficient cnn architecture design. In *ECCV*, 2018.
- [25] Sachin Mehta, Mohammad Rastegari, Linda Shapiro, and Hannaneh Hajishirzi. Espnetv2: A light-weight, power efficient, and general purpose convolutional neural network. In *CVPR*, 2019.
- [26] Fanxu Meng, Hao Cheng, Ke Li, Zhixin Xu, Rongrong Ji, Xing Sun, and Guangming Lu. Filter grafting for deep neural networks. In *CVPR*, June 2020.
- [27] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010.
- [28] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *CVPR*, 2020.
- [29] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.
- [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] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation. In *CVPR*, 2018.
- [32] 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 *CVPR*, 2019.
- [33] Xijun Wang, Meina Kan, Shiguang Shan, and Xilin Chen. Fully learnable group convolution for acceleration of deep neural networks. In *CVPR*, 2019.
- [34] Yulin Wang, Kangchen Lv, Rui Huang, Shiji Song, Le Yang, and Gao Huang. Glance and focus: a dynamic approach to reducing spatial redundancy in image classification. In *NeurIPS*, 2020.
- [35] Yulong Wang, Xiaolu Zhang, Lingxi Xie, Jun Zhou, Hang Su, Bo Zhang, and Xiaolin Hu. Pruning from scratch. In *AAAI*, 2019.

- [36] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. *arXiv preprint arXiv:1611.05431*, 2016.
- [37] Le Yang, Yizeng Han, Xi Chen, Shiji Song, Jifeng Dai, and Gao Huang. Resolution adaptive networks for efficient inference. In *CVPR*, 2020.
- [38] Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, and Larry S Davis. Nisp: Pruning networks using neuron importance score propagation. In *CVPR*, 2018.
- [39] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In *BMVC*, 2016.
- [40] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, 2018.
- [41] Daquan Zhou, Qibin Hou, Yunpeng Chen, Jiashi Feng, and Shuicheng Yan. Rethinking bottleneck structure for efficient mobile network design. In *ECCV*, 2020.