

# CONVOLUTIONAL NEURAL NETWORKS WITH LAYER REUSE

Okan Köpüklü, Maryam Babae, Stefan Hörmann, Gerhard Rigoll

Technical University of Munich  
Institute for Human-Machine Communication  
{okan.kopuklu, maryam.babae, s.hoermann, rigoll}@tum.de

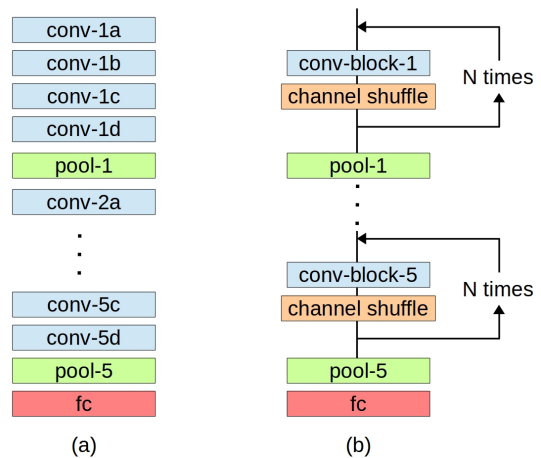
## ABSTRACT

A convolutional layer in a Convolutional Neural Network (CNN) consists of many filters which apply convolution operation to the input, capture some special patterns and pass the result to the next layer. If the same patterns also occur at the deeper layers of the network, why wouldn't the same convolutional filters be used also in those layers? In this paper, we propose a CNN architecture, Layer Reuse Network (LruNet), where the convolutional layers are used repeatedly without the need of introducing new layers to get a better performance. This approach introduces several advantages: (i) Considerable amount of parameters are saved since we are reusing the layers instead of introducing new layers, (ii) the Memory Access Cost (MAC) can be reduced since reused layer parameters can be fetched only once, (iii) the number of nonlinearities increases with layer reuse, and (iv) reused layers get gradient updates from multiple parts of the network. The proposed approach is evaluated on CIFAR-10, CIFAR-100 and Fashion-MNIST datasets for image classification task, and layer reuse improves the performance by 5.14%, 5.85% and 2.29%, respectively. The source code and pretrained models are publicly available <sup>1</sup>.

**Index Terms**— layer reuse, convolutional neural networks, inference routing

## 1. INTRODUCTION

The conventional way of designing a Convolutional Neural Network (CNN) is to put convolutional, pooling and batch normalization layers one after another with some nonlinearity functions in between. With the invention of residual connections [4], deeper CNN architectures started to be trained without overfitting and they achieved better results compared to shallow ones. Consequently, the primary trend for solving major visual recognition tasks has become building deeper and larger CNNs [10, 4, 15]. The best performing CNNs usually have hundreds of layers and millions of parameters. However, deeper architectures also mean more parameters which makes the designed architecture not appropriate




**Fig. 1:** (a) Conventional design of CNNs, (b) CNN design with layer reuse. Instead of stacking convolutional layers and feeding one layer's output as input to another layer, we feed the output of a convolutional block as input to itself for  $N$ -times before passing it to the next block.

for embedded and mobile devices. But the question always comes to mind: "Do we really need all these layers and parameters to achieve a better performance?". Fig. 2 shows the learned convolutional kernels in the first convolutional layer of AlexNet [10]. It is obvious to notice that some of the kernels are very similar, hence redundant. Therefore, instead of introducing all these similar kernels separately, they can simply be reused.

In this paper, we propose a network architecture, Layer Reuse Network (LruNet), where we have reused some convolutional layers repeatedly. Instead of stacking convolutional layers one after another, as in Fig. 1 (a), we feed the output of the convolutional blocks to itself for a given  $N$  times before passing the output to the next layer, as in Fig. 1 (b). While doing this, we apply channel shuffling in order to ensure feeding the outputs of convolution filters as inputs to other filters in the same block. Layer Reuse (LRU) brings several advantages to the system: (i) The number of parameters in the de-

<sup>1</sup><https://github.com/okankop/CNN-layer-reuse>



**Fig. 2:** Learned convolutional kernels in the first convolutional layer of AlexNet. Some of the convolutional kernels are very similar, which are highlighted with the same color. Better view in color.

signed architecture drops considerably since we are reusing layers instead of inserting additional new ones, (ii) the Memory Access Cost (MAC) can be reduced since the computing device can load the reused layer’s parameters only once, (iii) convolutional filters get gradient update from all reuse operations, and (iv) the number of nonlinearities increases as the number of LRU increases.

The LruNet architecture is constructed following the guidelines of recent resource efficient CNN architectures [7, 13, 16, 5, 14, 19, 12, 3]. These architectures are built mostly using group convolutions and depthwise separable convolutions. Group convolutions are first introduced in AlexNet [10] and effectively used in ResNeXt [18]. Depthwise separable convolutions are proposed in Xception [1] and they are the main building blocks for recent lightweight architectures.

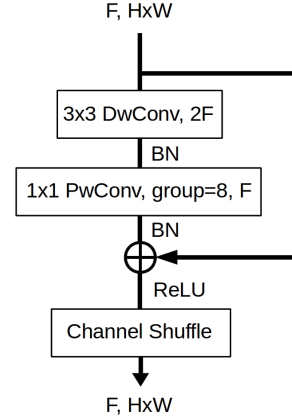
Some other architectures are constructed facilitating better gradient update mechanism to get improved performance [2, 6, 12]. Those works concentrate on how to feed the output of a layer as input to the next layers as efficiently as possible. In contrast, we reuse the layers multiple times, and as we increase the number-of-reuse ( $N$ ), convolutional filters also get more gradient updates.

## 2. APPROACH

### 2.1. Layer Reuse

Layer Reuse, referred as LRU, is the concept of using a convolutional layer or block multiple times at multiple places of a CNN architecture. However, the idea of parameter reuse in CNNs should not be limited at layer level. Convolutional layers can be split into smaller chunks up to filters, and reuse can be applied on these chunks throughout the whole network. In this case, it should be addressed as filter reuse (FRU) or kernel reuse (KRU). We believe parameter reuse concept will open a new era for deep learning practitioners for designing new CNN architectures.

LRU block used in the LruNet architecture is depicted in Fig. 3 where we use group convolutions and depthwise separable



**Fig. 3:** LRU block architecture. We represent DwConv and PwConv as depthwise and pointwise convolutions, respectively.  $F$  is the number of feature maps of the convolutional layers and *group* is the number of groups for group convolutions. *BN* and *ReLU* refer to Batch Normalization and Rectified Linear Unit nonlinearity, respectively.  $\oplus$  represents summation for the shortcut connection.

convolutions in order to keep the number of network parameters as small as possible. We first use depthwise convolutions increasing the channel number from  $F$  to  $2F$ . Then, we apply pointwise group convolutions with 8 groups. We have tried different group numbers here, but 8 groups experimentally proved to be the best in terms of accuracy and number of parameters. We have applied Batch Normalization (BN) [8] after these two layers, and ReLU after the shortcut connection. Finally, channel shuffle is applied at the end of the LRU block in order to feed different channels to different filters at each reuse. Channel shuffle is implemented as in [19], with little modification. In the original implementation, the first and the last channel of the feature volume always remain same. So, we switched the first half of the input volume with the second half before applying shuffling. We have not applied channel shuffle at the very last reuse since there is no need.

### 2.2. Network Architecture

The complete LruNet architecture is given in Table 1. We have built different LruNet architectures varying the number-of-reuse  $N$ . Similar to [5], width multiplier  $\alpha$  can also be applied to scale the number of filters. The network architecture in Table 1 is denoted as LruNet-1x. We use  $N$ -LruNet- $\alpha$ x to denote network architecture with  $N$  LRU and the number of filters in LruNet-1x are scaled by  $\alpha$ .

Although we reuse the convolutional layers repeatedly, we need to use a new BN layer for every reuse since the output feature volume has a different data distribution. Therefore, the number of parameters increases slightly for increased

Layer / Stride	Kernel Size	Layer Reuse	Output size
Conv/s2	3x3		64x16x16
LRU Block, F:64		N	64x16x16
MaxPool/s2	3x3		64x8x8
Concatenate			128x8x8
LRU Block, F:128		N	128x8x8
MaxPool/s2	3x3		128x4x4
Concatenate			256x4x4
LRU Block, F:256		N	256x4x4
Concatenate			512x4x4
LRU Block, F:512		N	512x4x4
MaxPool/s2	3x3		512x2x2
PwConv/s1, group=8	1x1		256x2x2
PwConv/s1	1x1		<i>NumCls</i> x2x2
AvgPool/s1	2x2		<i>NumCls</i>

**Table 1:** LruNet architecture. N is the number of times LRU block is reused. PwConv refers to pointwise convolutions. Concatenation of output feature map with itself is applied at the end of each LRU Block for channel expansion, except for the last one.

LRU due to the newly introduced BN layers. It must be also noted that the complexity and inference time of the network linearly depends on the number-of-reuse  $N$ .

### 2.3. Training Details

All the models are trained from scratch. We use stochastic gradient descent (SGD) with mini-batch of 256, and apply categorical cross-entropy loss. For the momentum and weight decay, 0.9 and  $5 \times 10^{-4}$  are used, respectively.

For regularization, several techniques are used to reduce overfitting. Firstly, weight decay ( $\gamma=5 \times 10^{-4}$ ) is applied on all the parameters of the network. Secondly, we used dropout before the last pointwise convolution with dropout ratio of 0.5. Lastly, we applied several data augmentation techniques: (a) Random cropping (padding=4), (b) random spatial rotation ( $\pm 10$ ), and (c) random horizontal flipping.

We have trained the networks with learning rate of 0.1 for 200 epochs, and 50 more epochs with learning rates of 0.01 and 0.001. Our approach is implemented in PyTorch with a single Nvidia Titan Xp GPU.

## 3. EXPERIMENTS

The proposed approach is evaluated for image classification task on three publicly available datasets: CIFAR-10, CIFAR-100 and Fashion-MNIST datasets. Each experiment is repeated 5 times in order to obtain more robust results due to the random initialization of the network parameters.

Model	Params	MFLOPs	Acc.(%)
1-LruNet-1x	131k	3.47	84.20
2-LruNet-1x	137k	6.30	84.95
4-LruNet-1x	149k	11.97	86.87
6-LruNet-1x	160k	17.63	87.91
8-LruNet-1x	172k	23.29	88.45
10-LruNet-1x	183k	28.95	88.66
12-LruNet-1x	195k	34.61	88.73
14-LruNet-1x	206k	40.27	<b>89.34</b>
16-LruNet-1x	218k	45.93	88.45

**Table 2:** Results for different Layer Reuse (LRU) on the validation set of CIFAR-10. All networks contain the same number of convolution parameters, which is 125k.

Model	Acc.(%)
8-LruNet-1x (without shuffling)	86.53
8-LruNet-1x (with shuffling)	<b>88.45</b>
14-LruNet-1x (without shuffling)	86.74
14-LruNet-1x (with shuffling)	<b>89.34</b>

**Table 3:** Results for Layer Reuse (LRU) with/without shuffle on the validation set of CIFAR-10.

### 3.1. Results Using CIFAR-10 Dataset

The CIFAR-10 [9] is a fundamental dataset in computer vision containing 50k training and 10k testing images in 10 classes with image resolution of 32x32. Initially we investigate the effect of LRU on the performance. Results in Table 2 show that acquired accuracy increases as we increase the applied LRU until 14-LRU. Afterwards, the performance does not improve with further reusing. As it has been mentioned earlier, the computational complexity (floating point operations - FLOPs) depends on the number-of-reuse  $N$  linearly. 14-LruNet-1x achieves 5.14% better classification accuracy than 1-LruNet-1x.

Secondly, we investigate the effect of channel shuffle. Results in Table 3 show that channel shuffle plays an important role in layer reuse. For both 8-LruNet-1x and 14-LruNet-1x cases, networks with channel shuffle perform better compared to the networks without channel shuffle. It is also interesting to note that even without channel shuffle, networks perform better than 1-LruNet-1x.

Model	Params	Acc.(%)
67-depth Network ( $\sim$ 8-LruNet-1x)	<b>172k</b>	88.45
67-depth Network (with new layers)	902k	<b>90.27</b>
115-depth Network ( $\sim$ 14-LruNet-1x)	<b>206k</b>	89.34
115-depth Network (with new layers)	1562k	<b>90.93</b>

**Table 4:** Comparison of LruNet with the networks having same depth containing all newly introduced layers.

Model	Params	MFLOPs	Acc.(%)
1-LruNet-2x	514k	10.44	63.02
2-LruNet-2x	525k	19.77	64.77
4-LruNet-2x	549k	38.44	65.50
6-LruNet-2x	572k	57.10	66.02
8-LruNet-2x	595k	75.76	67.57
10-LruNet-2x	618k	94.42	67.98
12-LruNet-2x	641k	113.08	68.30
14-LruNet-2x	664k	131.74	<b>68.87</b>
16-LruNet-2x	687k	150.40	68.02

**Table 5:** Results for different Layer Reuse (LRU) on the validation set of CIFAR-100. All networks contain the same number of convolution parameters, which is 501k.

Model	Params	MFLOPs	Acc.(%)
1-LruNet-1x	130k	2.85	91.17
2-LruNet-1x	136k	5.40	92.28
4-LruNet-1x	148k	10.50	92.69
6-LruNet-1x	159k	15.61	93.13
8-LruNet-1x	171k	20.71	93.27
10-LruNet-1x	182k	25.81	<b>93.46</b>
12-LruNet-1x	194k	30.92	93.34

**Table 6:** Results for different Layer Reuse (LRU) on the validation set of Fashion-MNIST. All networks contain the same number of convolution parameters, which is 124k.

Lastly, Table 4 compares the LruNet with the networks having same depth containing all newly introduced layers. Although networks with newly introduced layers achieves comparatively better results, networks with LRU have 5.24 and 7.58 times less parameters for 67-depth and 115-depth networks, respectively.

### 3.2. Results Using CIFAR-100 Dataset

CIFAR-100 [9] is very similar to the CIFAR-10, except it has 100 classes containing 600 images each. Since it is a more challenging task, we have increased the dropout rate to 0.7 and used width multiplier  $\alpha$  of 2. We again analyzed the effect of LRU on the performance, and 14-LruNet-2x achieves 5.14% better classification accuracy than 1-LruNet-2x, as depicted by Table 5.

### 3.3. Results Using Fashion-MNIST Dataset

Fashion-MNIST [17] is very similar to MNIST [11], but contains images of various articles of clothing and accessories. There are 50k training and 10k testing images in grayscale for 10 classes with image resolution of 28x28. 10-LruNet-1x achieves 2.29% better classification accuracy than 1-LruNet-1x, as depicted by Table 6. This performance improvement is relatively small compared to the performance improvements using CIFAR-10 and CIFAR-100. This might be due

	Model	Total Params	Conv. Params	Acc.(%)
CIFAR100	ShuffleNet-0.5x(g=3)	229k	217k	<b>90.32</b>
	ShuffleNetV2-0.25x	211k	205k	89.31
	MobileNet-0.25x	216k	210k	84.72
	MobileNetV2-0.25x	249k	239k	89.57
	14-LruNet-1x	206k	<b>125k</b>	89.34
CIFAR100	ShuffleNet-0.75x(g=3)	560k	542k	<b>69.97</b>
	ShuffleNetV2-0.75x	568k	559k	69.15
	MobileNet-0.4x	567k	558k	60.54
	MobileNetV2-0.4x	603k	588k	69.95
	14-LruNet-2x	664k	<b>501k</b>	68.87
Fash.-MNIST	ShuffleNet-0.5x(g=3)	228k	216k	<b>94.11</b>
	ShuffleNetV2-0.25x	211k	205k	93.42
	MobileNet-0.25x	216k	210k	90.63
	MobileNetV2-0.25x	249k	239k	93.43
	10-LruNet-1x	182k	<b>124k</b>	93.46

**Table 7:** Comparison of LruNet with state-of-the-art results.

to the relative simpleness of the Fashion-MNIST dataset compared to the CIFAR-10 and CIFAR-100. 1-LruNet-1x already achieves 91.17% classification accuracy on Fashion-MNIST.

Table 7 shows the comparison of LruNet with the state-of-the-art results. For ShuffleNetv2, channel numbers are adjusted accordingly for width multipliers of 0.25 and 0.75. Results in Table 7 show that LruNet achieves comparatively similar results, although it has much smaller number of convolutional parameters.

## 4. CONCLUSION

This paper proposes a parameter reuse strategy, Layer Reuse (LRU), where convolutional layers of a CNN architecture (LruNet) are used repeatedly. We evaluated the LRU on several publicly available datasets and achieved improved classification performance. LRU especially boosts the CNNs with small number of parameters which is of utmost importance for embedded applications. We believe this work will open up a new research direction for deep learning practitioners for designing novel CNN architectures.

As a future work, we would like to analyze different parameter reuse strategies. It must be noted that parameter reuse is not only restricted to layer level, but we can also apply Filter Reuse (FRU) (or Kernel Reuse (KRU)) for efficient CNN architecture designs.

## Acknowledgements

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

## 5. REFERENCES

- [1] F. Chollet. Xception: Deep learning with depthwise separable convolutions. *arXiv preprint*, pages 1610–02357, 2017.
- [2] C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slowfast networks for video recognition. <https://arxiv.org/abs/1812.03982>, 2018.
- [3] I. Freeman, L. Roese-Koerner, and A. Kummert. Effnet: An efficient structure for convolutional neural networks. *arXiv preprint arXiv:1801.06434*, 2018.
- [4] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [5] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [6] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *CVPR*, volume 1, page 3, 2017.
- [7] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.
- [8] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [9] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [10] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [11] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [12] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. *arXiv preprint arXiv:1807.11164*, 5, 2018.
- [13] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European Conference on Computer Vision*, pages 525–542. Springer, 2016.
- [14] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4510–4520. IEEE, 2018.
- [15] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [16] J. Wu, C. Leng, Y. Wang, Q. Hu, and J. Cheng. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4820–4828, 2016.
- [17] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [18] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5987–5995. IEEE, 2017.
- [19] X. Zhang, X. Zhou, M. Lin, and J. Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6848–6856. IEEE, 2018.