

---

# Structured Pruning for Efficient ConvNets via Incremental Regularization

---

Huan Wang, Qiming Zhang, Yuehai Wang, Haoji Hu\*  
Zhejiang University  
{huanw, qmzhang, wyuehai, haoji\_hu}@zju.edu.cn

## Abstract

Parameter pruning is a promising approach for CNN compression and acceleration by eliminating redundant model parameters with tolerable performance loss. Despite its effectiveness, existing regularization-based parameter pruning methods usually drive weights towards zero with *large and constant* regularization factors, which neglects the fact that the expressiveness of CNNs is fragile and needs a more gentle way of regularization for the networks to adapt during pruning. To solve this problem, we propose a new regularization-based pruning method (named *IncReg*) to incrementally assign different regularization factors to different weight groups based on their relative importance, whose effectiveness is proved on popular CNNs compared with state-of-the-art methods.

## 1 Introduction

Recently, deep Convolutional Neural Networks (CNNs) have made a remarkable success in computer vision tasks by leveraging large-scale networks learning from big amount of data. However, CNNs usually lead to massive computation and storage consumption, thus hindering their deployment on mobile and embedded devices. To solve this problem, many research works focus on compressing the scale of CNNs. Parameter pruning is a promising approach for CNN compression and acceleration, which aims at eliminating redundant model parameters at tolerable performance loss. To avoid hardware-unfriendly irregular sparsity, structured pruning is proposed for CNN acceleration [1, 23]. In the `im2col` implementation [2, 4] of convolution, weight tensors are expanded into matrices, so there are generally two kinds of structured sparsity, i.e. row sparsity (or filter-wise sparsity) and column sparsity (or shape-wise sparsity) [25, 24].

There are mainly two categories of structured pruning. One is importance-based methods, which prune weights in groups based on some established importance criteria [18, 20, 24]. The other is regularization-based methods, which add group regularization terms to learn structured sparsity [25, 17, 9]. Existing group regularization approaches mainly focus on the regularization form (e.g. Group LASSO [27]) to learn structured sparsity, while ignoring the influence of regularization factor. In particular, they tend to use a *large and constant* regularization factor for all weight groups in the network [25, 17], which has two problems. Firstly, this ‘one-size-fit-all’ regularization scheme has a hidden assumption that all weights in different groups are equally important, which however does not hold true, since weights with larger magnitude tend to be more important than those with smaller magnitude. Secondly, few works have noticed that the expressiveness of CNNs is so fragile [26] during pruning that it cannot withstand a large penalty term from beginning, especially for large pruning ratios and compact networks (like ResNet [8]). AFP [7] was proposed to solve the first problem, while ignored the second one. In this paper, we propose a new regularization-based method named *IncReg* to *incrementally* learn structured sparsity.

---

\*Corresponding author

Apart from pruning, there are also other kinds of methods for CNN acceleration, such as low-rank decomposition [16, 29], quantization [3, 5, 19, 21], knowledge distillation [10] and architecture re-design [12, 11, 28]. Our method is orthogonal to these methods.

## 2 The Proposed Method

Given a conv kernel, modeled by a 4-D tensor  $\mathbf{W}^{(l)} \in \mathbb{R}^{N^{(l)} \times C^{(l)} \times H^{(l)} \times W^{(l)}}$ , where  $N^{(l)}$ ,  $C^{(l)}$ ,  $H^{(l)}$  and  $W^{(l)}$  are the dimension of the  $l$ th ( $1 \leq l \leq L$ ) weight tensor along the axis of filter, channel, height and width, respectively, our proposed objective function for regularization can be formulated as:  $E(\mathbf{W}) = L(\mathbf{W}) + \frac{\lambda}{2}R(\mathbf{W}) + \sum_{l=1}^L \sum_{g=1}^{G^{(l)}} \frac{\lambda_g^{(l)}}{2}R(\mathbf{W}_g^{(l)})$ , where  $\mathbf{W}$  denotes the collection of all weights in the CNN;  $L(\mathbf{W})$  is the loss function for prediction;  $R(\mathbf{W})$  is non-structured regularization on every weight, i.e. weight decay in this paper;  $R(\mathbf{W}_g^{(l)})$  is the structured sparsity regularization term on group  $g$  of layer  $l$  and  $G^{(l)}$  is the number of weight groups in layer  $l$ . In [17, 25], the authors used the same  $\lambda_g$  for all groups and adopted Group LASSO [27] for  $R(\mathbf{W}_g^{(l)})$ . In this work, since we emphasize the key problem of group regularization lies in the regularization factor rather than the regularization form, we use the most common regularization form weight decay, for  $R(\mathbf{W}_g^{(l)})$ , but we vary the regularization factors  $\lambda_g$  for different weight groups and at different iterations.

Our method prunes all the conv layers simultaneously and independently. For simplicity, we omit the layer notation  $l$  for following description. All the  $\lambda_g$ 's are initialized to zero. At each iteration,  $\lambda_g$  is increased by  $\lambda_g^{new} = \max(\lambda_g + \Delta\lambda_g, 0)$ . Like AFP [7], we agree that unimportant weights should be punished more, so we propose a decreasing piece-wise linear punishment function (Eqn.1, Fig.1) to determine  $\Delta\lambda_g$ . Note that the regularization increment is negative (i.e. reward actually) when ranking is above the threshold ranking  $RG$ , since above-the-threshold means these weights are expected to stay in the end. Regularization on these important weights is not only unnecessary but also very harmful via our experiment confirmation.

$$\Delta\lambda_g(r) = \begin{cases} -\frac{A}{RG}r + A & \text{if } r \leq RG \\ -\frac{A}{G(1-R)-1}(r - RG) & \text{if } r > RG \end{cases} \quad (1)$$


where  $R$  is the pre-assigned pruning ratio for a layer,  $G$  is the number of weight groups,  $A$  is a hyper-parameter in our method to describe the maximum penalty increment (set to half of the original weight decay in default),  $r$  is the ranking obtained by sorting in ascending order based on a proposed importance criterion, which is essentially an averaged ranking over time, defined as  $\frac{1}{N} \sum_{n=1}^N r_n$ , where  $r_n$  is the ranking by  $L_1$ -norm at  $n$ th iteration,  $N$  is the number of passed iterations. This averaging is adopted as smoothing for a more stable pruning process.

As training proceeds, the regularization factors of different weight groups increase gradually, which will push the weights towards zero little by little. When the magnitude of a weight group is lower than some threshold ( $10^{-6}$ ), the weights are permanently removed from the network, thus leading to increased structured sparsity. When the sparsity of a layer reaches its pre-assigned pruning ratio  $R$ , that layer automatically stops structured regularization. Finally, when all conv layers reach their pre-assigned pruning ratios, pruning is over, followed by a retraining process to regain accuracy.

## 3 Experiments

### 3.1 Analysis with ConvNet on CIFAR-10

We firstly compare our proposed IncReg with other two group regularization methods, i.e. SSL [25] and AFP [7], with ConvNet [15] on CIFAR-10 [14], where both row sparsity and column sparsity are explored. Caffe [13] is used for all of our experiments. Experimental results are shown in Tab.1. We can see that IncReg consistently achieves higher speedups and accuracies than the other two constant regularization schemes. Notably, even though AFP achieves similar performance as our method under relatively small speedup (about  $4.5\times$ ), when the speedup ratio is large (about  $8\times \sim 10\times$ ), our method outperforms AFP by a large margin. We argue that this is because the incremental way of regularization gives the network more time to adapt during pruning, which is especially important in face of large pruning ratios.



**Figure 1:** The illustration of our proposed punishment function of  $\Delta\lambda_g$ .

| Method          | Row sparsity |             | Column sparsity |             |
|-----------------|--------------|-------------|-----------------|-------------|
|                 | speedup      | accuracy    | speedup         | accuracy    |
| SSL             | 3.6×         | 77.3        | 3.1×            | 78.6        |
| AFP (our impl.) | 4.1×         | 77.7        | 4.5×            | 81.0        |
| Ours            | <b>4.1×</b>  | <b>79.2</b> | <b>4.6×</b>     | <b>81.2</b> |
| SSL             | 9.7×         | 73.0        | 10.0×           | 75.2        |
| AFP (our impl.) | 9.9×         | 73.4        | 8.4×            | 77.5        |
| Ours            | <b>9.9×</b>  | <b>76.0</b> | <b>10.0×</b>    | <b>78.7</b> |

**Table 1:** Comparison of our method with SSL [25] and AFP [7] with ConvNet on CIFAR-10.

### 3.2 VGG-16 and ResNet-50 on ImageNet

We further evaluate our method with VGG-16 [22] (13 conv layers) and ResNet-50 [8] (53 conv layers) on ImageNet [6]. We download the open-sourced caffemodel as our pre-trained model, whose single-view top-5 accuracy on ImageNet validation dataset is 89.6% (VGG-16) and 91.2% (ResNet-50). For VGG-16, following SPP [24], the proportion of remaining ratios of low layers (conv1\_x to conv3\_x), middle layers (conv4\_x) and high layers (conv5\_x) are set to 1 : 1.5 : 2, for easy comparison. For ResNet-50, constant pruning ratio 0.4 is adopted for all conv layers. Pruning is conducted with batch size 64 and fixed learning rate 0.0005, followed by retraining with batch size 256. Experimental results are shown in Tab.2. On VGG-16, our method is slightly better than CP and SPP, and outperforms FP by a significant margin. Notably, since we use the same pruning ratios as SPP does, the only explanation for the performance improvement should be a better pruning process itself, guided by our incremental regularization scheme. On ResNet-50, our method is significantly better than CP and SPP, demonstrating the effectiveness of IncReg when pruning compact networks. Moreover, to confirm the actual speedup, we also evaluate our method with VGG-16 on CPU and GPU. The result is shown in Tab.3.

| Method               | Increased err. (%) |            |      | Method            | Increased err. (%) |
|----------------------|--------------------|------------|------|-------------------|--------------------|
|                      | 2×                 | 4×         | 5×   |                   |                    |
| FP [18] (CP's impl.) | 0.8                | 8.6        | 14.6 | CP (enhanced) [9] | 1.4                |
| CP [9]               | <b>0</b>           | 1.0        | 1.7  | SPP [24]          | 0.8                |
| SPP [24]             | <b>0</b>           | <b>0.8</b> | 2.0  | Ours              | <b>0.1</b>         |
| Ours                 | <b>0</b>           | <b>0.8</b> | 1.5  |                   |                    |

**Table 2:** Acceleration of VGG-16 (left) and ResNet-50 (right, 2× speedup) on ImageNet. The values are increased single-view top-5 error on ImageNet.

| Method | CPU time (baseline: 1815 ms) |           |           | GPU time (baseline: 5.159 ms) |             |             |
|--------|------------------------------|-----------|-----------|-------------------------------|-------------|-------------|
|        | 2×                           | 4×        | 5×        | 2×                            | 4×          | 5×          |
| CP [9] | 826(2.2×)                    | 500(3.6×) | 449(4.0×) | 3.206(1.6×)                   | 2.202(2.3×) | 2.034(2.5×) |
| Ours   | 861(2.1×)                    | 469(3.9×) | 409(4.4×) | 3.225(1.6×)                   | 2.068(2.5×) | 1.991(2.6×) |

**Table 3:** Inference time of conv layers of CP and our method on VGG-16. Evaluation is carried out with batch size 10 and averaged by 50 runs on  $224 \times 224$  RGB images. CPU: Intel Xeon(R) E5-2620 v4 @ 2.10GHz, single thread; GPU: GeForce GTX 1080Ti, without cuDNN. Open-sourced models of CP are used for this evaluation.

## 4 Conclusion

We propose a new structured pruning method based on an incremental way of regularization, which helps CNNs to transfer their expressiveness to the rest parts during pruning by increasing the regularization factors of unimportant weight groups little by little. Our method is proved to be comparably effective on popular CNNs compared with state-of-the-art methods, especially in face of large pruning ratios and compact networks.

## References

- [1] S. Anwar and W. Sung. Compact deep convolutional neural networks with coarse pruning. *arXiv preprint*, arXiv:1610.09639, 2016.
- [2] K. Chellapilla, S. Puri, and P. Simard. High performance convolutional neural networks for document processing. In *International Workshop on Frontiers in Handwriting Recognition*, 2006.
- [3] W. Chen, J. T. Wilson, S. Tyree, K. Q. Weinberger, and Y. Chen. Compressing neural networks with the hashing trick. In *Proceedings of the International Conference on Machine Learning, ICML*, pages 1–10, Lille, France, 2015.
- [4] S. Chetlur, C. Woolley, and P. Vandermersch. cuDNN: Efficient primitives for deep learning. *arXiv preprint*, arXiv:1410.0759, 2014.
- [5] M. Courbariaux and Y. Bengio. BinaryNet: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv preprint*, arXiv:1602.02830, 2016.
- [6] J. Deng, W. Dong, R. Socher, L. J. Li, K. Li, and L. Feifei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the International Conference on Computer Vision and Pattern Recognition, CVPR'09*, pages 248–255, Miami, FL, 2009.
- [7] X. Ding, G. Ding, J. Han, and S. Tang. Auto-balanced filter pruning for efficient convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1–8, New Orleans, LA, 2018.
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, CVPR*, pages 770 – 778, Las Vegas, NV, 2016.
- [9] Y. He, X. Zhang, and J. Sun. Channel pruning for accelerating very deep neural networks. In *International Conference on Computer Vision (ICCV)*, volume 2, 2017.
- [10] G. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *Computer Science*, 14(7):38–39, 2015.
- [11] 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. 2017.
- [12] F. Iandola, M. Moskewicz, and K. Ashraf. SqueezeNet: Alexnet-level accuracy with 50x fewer parameters and <0.5MB model size. *arXiv preprint*, arXiv:1602.07360, 2016.
- [13] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrel. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint*, arXiv:1408.5093, 2014.
- [14] A. Krizhevsky and G. E. Hinton. Learning multiple layers of features from tiny images. 2009.
- [15] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Information Processing Systems, NIPS*, pages 1097–1105, Lake Tahoe, CA, 2012.
- [16] V. Lebedev, Y. Ganin, M. Rakhuba, I. Oseledets, and V. Lempitsky. Speeding-up convolutional neural networks using fine-tuned CP-decomposition. *arXiv preprint*, arXiv:1510.03009, 2016.
- [17] V. Lebedev and V. Lempitsky. Fast convnets using group-wise brain damage. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, CVPR*, pages 2554–2564, Las Vegas, NV, 2016.
- [18] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf. Pruning filters for efficient convnets. In *International Conference on Learning Representations, ICLR*, 2017.
- [19] Z. Lin, M. Courbariaux, R. Memisevic, and Y. Bengio. Neural networks with few multiplications. *arXiv preprint*, arXiv:1510.03009, 2016.
- [20] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz. Pruning convolutional neural networks for resource efficient inference. In *International Conference on Learning Representations, ICLR*, 2017.
- [21] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European Conference on Computer Vision, ECCV*, pages 525–542, Amsterdam, Netherland, 2016.
- [22] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *Computer Science*, 2014.
- [23] V. Sze, Y. H. Chen, T. J. Yang, and J. Emer. Efficient processing of deep neural networks: A tutorial and survey. *arXiv preprint*, arXiv:1703.09039, 2017.
- [24] H. Wang, Q. Zhang, Y. Wang, and H. Hu. Structured probabilistic pruning for convolutional neural network acceleration. In *British Machine Vision Conference*, 2018.
- [25] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li. Learning structured sparsity in deep neural networks. In *Advances in Neural Information Processing Systems*, pages 2074–2082, 2016.

- [26] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson. How transferable are features in deep neural networks? In *Advances in neural information processing systems*, pages 3320–3328, 2014.
- [27] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society*, 68(1):49–67, 2006.
- [28] X. Zhang, X. Zhou, M. Lin, and J. Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. 2017.
- [29] X. Zhang, J. Zou, K. He, and J. Sun. Accelerating very deep convolutional networks for classification and detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(10):1943–1955, 2016.