
Three Dimensional Convolutional Neural Network Pruning with Regularization-Based Method

Yuxin Zhang, Huan Wang, Yang Luo, Roland Hu*

Zhejiang University

{yuxinzhang, huanw, luoyang95, haoji_hu}@zju.edu.cn

Abstract

In recent years, three-dimensional convolutional neural network (3D CNN) is intensively applied in video analysis and receives good performance. However, 3D CNN leads to massive computation and storage consumption, which hinders its deployment on mobile and embedded devices. In this paper, we propose a three-dimensional regularization-based pruning method to assign different regularization parameters to different weight groups based on their importance to the network. Experiments show that the proposed method outperforms other popular methods in this area.

1 Introduction

In recent years, convolutional neural network (CNN) has developed rapidly and has achieved remarkable success in computer vision tasks such as identification, classification and segmentation. However, due to the lack of motion modeling, this image-based end-to-end feature can not directly apply to videos. In [1, 2], the authors use three-dimensional convolutional networks (3D CNN) to identify human actions in videos. Tran *et al.* proposed a 3D CNN for action recognition which contains 1.75 million parameters [3]. The development of 3D CNN also brings challenges because of its higher dimensions. This leads to massive computing and storage consumption, which hinders its deployment on mobile and embedded devices.

In order to reduce the computation cost, researchers propose methods to compress CNN models, including knowledge distillation [4], parameter quantization [5, 6], matrix decomposition [7] and parameter pruning [8]. However, all of the above methods are based on two-dimensional convolution. In this paper, we expand the idea of [9] to 3D CNN acceleration. The main idea is to add group regularization items to the objective function and prune weight groups gradually, where the regularization parameters for different weight groups are differently assigned according to some importance criteria.

2 The Proposed Method

For a three-dimensional convolutional neural network with L layers, the weights of the l th ($1 \leq l \leq L$) convolutional layer $\mathbf{W}^{(l)} \in \mathbb{R}^{N^l \times C^l \times W^l \times H^l \times D^l}$ is a sequences of 5-D tensors. Here N^l , C^l , W^l , H^l and D^l are the dimensions along the axes of filter, channel, spatial height, spatial width and spatial depth. The proposed objective function for structured sparsity regularization is defined by Eqn.(1).

$$E(\mathbf{W}) = L(\mathbf{W}) + \frac{\lambda}{2}R(\mathbf{W}) + \frac{\lambda_g}{2} \sum_{l=1}^L R_g(\mathbf{W}^{(l)}) \quad (1)$$

*Contact author




Figure 1: The im2col implementation of 3D CNN is to expand tensors into matrices, so that convolutional operations are transformed to matrix multiplication. The weights at the blue squares are to be pruned. (a) Pruning a filter. (b) Pruning all the weights at the same position. (c) Pruning a channel.

Here $L(\mathbf{W})$ is the loss on data; $R(\mathbf{W})$ is the non-structured regularization (L_2 norm in this paper). R_g is the structured sparsity regularization on each layer. In [10, 11], the authors used the same λ_g for all groups and adopted Group LASSO for R_g . Recently Wang et al. [9] use the squared L_1 norm for R_g and vary the regularization parameters λ_g for different groups. We build on top of that approach but extend it from two dimensions to three dimensions.

The structure learned is determined by the way of splitting groups of $\mathbf{W}_g^{(l)}$. There are normally filter-wise, channel-wise, shape-wise, and depth-wise structured sparsity with different ways of grouping [10]. Pruning of different weight groups for 3D CNN is shown in Fig.1.

In [9], Wang *et al.* theoretically proved that by increasing the regularization parameter λ_g , the magnitude of weights tends to be minimized. The more λ_g increases, the more magnitude of weights are compressed to zero. Therefore, we can assign different λ_g for the weight groups based on their importance to the network. Here, we use the L_1 norm as a criterion of importance.

Our goal is to prune RN_g weight groups in the network, where R is the pruning ratio to each layer and N_g is total number of weight groups in the layer. In other words, we need to prune RN_g weight groups which ranks lower in the network. We sort the weight groups in ascending order of the L_1 norms. In order to remove the oscillation of ranks during one training iteration, we averaged the rank through training iterations to obtain the average rank \bar{r}_{avg} in N training iterations: $\bar{r}_{avg} = \frac{1}{N} \sum_{n=1}^N r_n$.

The final average rank \bar{r} is obtained by sorting \bar{r}_{avg} of different weight groups in ascending order, making its range from 0 to $N_g - 1$. The update of λ_g is determined by the following formula: $\lambda_g^{(new)} = \lambda_g^{(old)} + \Delta\lambda_g$.

Here $\Delta\lambda_g$ is the function of average rank \bar{r} , we follow the formula proposed by Wang [9] as follows:

$$\Delta\lambda_g(\bar{r}) = \begin{cases} -\frac{A}{RN_g}\bar{r} + A & \text{if } \bar{r} \leq RN_g \\ -\frac{A}{N_g(1-R)-1}(\bar{r} - RN_g) & \text{if } \bar{r} > RN_g \end{cases} \quad (2)$$

Here A is a hyperparameter which controls the speed of convergence. According to Eqn.(2), we can see that $\Delta\lambda_g$ is zero when $\bar{r} = RN_g$ because we need to increase the regularization parameters of the weight groups whose ranks are below RN_g to further decrease their L_1 norms; and for those with greater L_1 norms and rank above RN_g , we need to decrease their regularization parameters to further increase their L_1 norms. Thus, we can ensure that exactly RN_g weight groups are pruned at the final stage of the algorithm. When we obtain $\lambda_g^{(new)}$, the weights can be updated through back-propagation deduced from Eqn.(1). Further details can be found in [9].

3 Experiments

Our experiments are carried out by Caffe[12]. We set the weight decay factor λ to be the same as the baseline and set hyper-parameter A to half of λ . We only compress the weights in convolutional

Table 1: The increased error when accelerating C3D on UCF101(baseline: 79.94%).

Method	Increased err. (%)	
	2×	4×
TP (our impl.)	11.50	21.19
FP (our impl.)	4.92	10.96
Ours	3.56	7.02

Table 2: The increased error when accelerating 3D-ResNet18 on UCF101(baseline: 72.50%).

Method	Increased err. (%)	
	2×	4×
TP (our impl.)	5.72	14.24
FP (our impl.)	1.60	6.92
Ours	0.91	3.50

layers and leave the fully connected layers unchanged because we focus on network acceleration. The pruning ratios of the convolutional layers are set to the same for convenience. The methods used for comparison are Taylor Pruning (TP) [13] and Filter Pruning (FP) [14]. For all experiments, the ratio of speedup is calculated by GFLOPS reduction.

3.1 C3D on UCF101

We apply the proposed method to C3D [3], which is composed of 8 convolution layers, 5 max-pooling layers, and 2 fully connected layers. We download the open Caffe model as our pre-trained model, whose accuracy on UCF101 dataset is 79.94%. UCF101 contains 101 types of actions and a total of 13320 videos with a resolution of 320×240 . All videos are decoded into image files with 25 fps rate. Frames are resized into 128×171 and randomly cropped to 112×112 . Then frames are split into non-overlapped 16-frame clips which are then used as input to the networks. The results are shown in Table 1. With different speedup ratios, our approach is always better than TP and FP.

3.2 3D-ResNet18 on UCF101

We further demonstrate our method on 3D-ResNet18 [3], which has 17 convolution layers and 1 fully-connected layer. The network is initially trained on the Sport-1M database. We download the model and then fine-tune it by UCF101 for 30000 iterations, obtaining an accuracy of 72.50%. The video preprocessing method is the same as above. The training settings are similar to that of C3D.

Experimental results are shown in Table 2. Our approach only suffers 0.91% increased error while achieving 2× acceleration, obtaining better results than TP and FP. Fig.2 shows the loss during the pruning process for different methods. As the number of iterations increases, the losses of TP and FP change dramatically, while the loss of our method remains at a lower level consistently. This is probably because the proposed method imposes gradual regularization, making the network changes little-by-little in the parameter space, while both the TP and FP direct prune less important weights once for all.




Figure 2: The training losses on 3D-ResNet18 for TP, FP and the proposed method.

4 Conclusion

In this paper, we implement the regularization based method for 3D CNN acceleration. By assigning different regularization parameters to different weight groups according to the importance estimation, we gradually prune weight groups in the network. The proposed method achieves better performance than other two popular methods in this area.

References

- [1] S. Ji, W. Xu, M. Yang, and K. Yu, “3d convolutional neural networks for human action recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 221–231, 2012.
- [2] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and F. F. Li, “Large-scale video classification with convolutional neural networks,” in *Proceedings of the International Conference on Computer Vision and Pattern Recognition, CVPR’14*, 2014, pp. 1725–1732.
- [3] D. Tran, L. D. Bourdev, R. Fergus, L. Torresani, and M. Paluri, “C3D: generic features for video analysis,” *ArXiv preprint: 1412.0767*, 2014.
- [4] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *ArXiv preprint: 1503.02531*, 2015.
- [5] M. Courbariaux, Y. Bengio, and J. P. David, “Binaryconnect: training deep neural networks with binary weights during propagations,” in *Proceedings of the International Conference on Neural Information Processing Systems, NIPS’15*, 2015, pp. 3123–3131.
- [6] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, “Xnor-net: Imagenet classification using binary convolutional neural networks,” in *Proceedings of the European Conference on Computer Vision, ECCV’16*, 2016, pp. 525–542.
- [7] 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*, vol. 38, no. 10, pp. 1943–1955, 2015.
- [8] S. Han, H. Mao, and W. J. Dally, “Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding,” *ArXiv preprint: 1510.00149*, 2015.
- [9] H. Wang, Q. Zhang, Y. Wang, and R. Hu, “Structured Deep Neural Network Pruning by Varying Regularization Parameters,” *ArXiv preprint: 1804.09461*, 2018.
- [10] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, “Learning structured sparsity in deep neural networks,” *ArXiv preprint: 1608.03665*, 2016.
- [11] V. Lebedev and V. Lempitsky, “Fast convnets using group-wise brain damage,” in *Proceedings of the International Conference on Computer Vision and Pattern Recognition, CVPR’16*, 2016, pp. 2554–2564.
- [12] 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*, vol. arXiv:1408.5093, 2014.
- [13] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz, “Pruning convolutional neural networks for resource efficient transfer learning,” *ArXiv preprint: 1611.06440*, 2016.
- [14] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, “Pruning filters for efficient convnets,” *ArXiv preprint: 1608.08710*, 2016.