
Implicit Filter Sparsification In Convolutional Neural Networks

Dushyant Mehta^{1,2} Kwang In Kim³ Christian Theobalt¹

Abstract

We show implicit *filter level* sparsity manifests in convolutional neural networks (CNNs) which employ Batch Normalization and ReLU activation, and are trained with adaptive gradient descent techniques and L2 regularization or weight decay. Through an extensive empirical study (Mehta et al., 2019) we hypothesize the mechanism behind the sparsification process, and find surprising links to certain filter sparsification heuristics proposed in literature. Emergence of, and the subsequent pruning of selective features is observed to be one of the contributing mechanisms, leading to feature sparsity at par or better than certain explicit sparsification / pruning approaches. In this workshop article we summarize our findings, and point out corollaries of selective-feature-penalization which could also be employed as heuristics for filter pruning.

1. Introduction

In this article we discuss the findings from (Mehta et al., 2019) regarding filter level sparsity which emerges in certain types of feedforward convolutional neural networks. Filter refers to the weights and the nonlinearity associated with a particular feature, acting together as a unit. We use filter and feature interchangeably throughout the document. We particularly focus on the implications of the findings on feature pruning for neural network speed up.

In networks which employ Batch Normalization and ReLU activation, after training, certain filters are observed to not activate for any input. Importantly, the sparsity emerges in the presence of regularizers such as L2 and weight decay (WD) which are in general understood to be non sparsity

inducing, and the sparsity vanishes when regularization is removed. We experimentally observe the following:

- The sparsity is much higher when using adaptive flavors of SGD vs. (m)SGD.
- Adaptive methods see higher sparsity with L2 regularization than with WD. No sparsity emerges in the absence of regularization.
- In addition to the regularizers, the extent of the emergent sparsity is also influenced by hyperparameters seemingly unrelated to regularization. The sparsity decreases with increasing mini-batch size, decreasing network size and increasing task difficulty.
- The primary hypothesis that we put forward is that selective features¹ see a disproportionately higher amount of regularization than non-selective ones. This consistently explains how parameters such as mini-batch size, network size, and task difficulty indirectly impact sparsity by affecting feature selectivity.
- A secondary hypothesis to explain the higher sparsity observed with adaptive methods is that Adam (and possibly other) adaptive approaches learn more selective features. Though there is evidence of highly selective features with Adam, this requires further study.
- Synthetic experiments show that the interaction of L2 regularizer with the update equation in adaptive methods causes stronger regularization than WD. This can explain the discrepancy in sparsity between L2 and WD.

Quantifying Feature Sparsity: Feature sparsity can be measured by per-feature activation and by per-feature scale. For sparsity by activation, the absolute activations for each feature are max pooled over the entire feature plane. If the value is less than 10^{-12} over the entire *training* corpus, the feature is inactive. For sparsity by scale, we consider the scale γ of the learned affine transform in the Batch Norm layer. We consider a feature inactive if $|\gamma|$ for the feature is less than 10^{-3} . Explicitly zeroing the features thus marked inactive does not affect the test error. The thresholds chosen are purposefully conservative, and comparable levels of sparsity are observed for a higher feature activation threshold of 10^{-4} , and a higher $|\gamma|$ threshold of 10^{-2} .

The implicit sparsification process can remove 70-80% of

¹Feature selectivity is the fraction of training exemplars for which a feature produces max activation less than some threshold.

¹Max Planck Institute For Informatics, Saarbrücken, Germany
²Saarland Informatics Campus, Germany ³Ulsan National Institute of Science and Technology, South Korea. Correspondence to: Dushyant Mehta <dmehta@mpi-inf.mpg.de>.

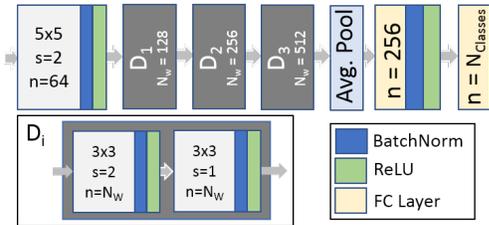


Figure 1. **BasicNet**: Structure of the basic convolution network studied in this paper. We refer to the convolution layers as C1-7.

the convolutional filters from VGG-16 on CIFAR-10/100, far exceeding that for (Li et al., 2017), and performs comparable to (Liu et al., 2017) for VGG-11 on ImageNet. Common hyperparameters such as mini-batch size can be used as knobs to control the extent of sparsity, with no tooling changes needed to the traditional NN training pipeline.

We present some of the experimental results, and discuss links of the hypothesis to pruning heuristics.

2. Related Work

Among the various explicit filter level sparsification heuristics and approaches (Li et al., 2017; Srinivas & Babu, 2016; Hu et al., 2016; Theis et al., 2017; Molchanov et al., 2017; Mozer & Smolensky, 1989; Liu et al., 2017; Ye et al., 2018), some (Ye et al., 2018; Liu et al., 2017) make use of the learned scale parameter γ in Batch Norm for enforcing sparsity on the filters. (Ye et al., 2018) argue that BatchNorm makes feature importance less susceptible to scaling reparameterization, and the learned scale parameters (γ) can be used as indicators of feature importance. We find that Adam with L2 regularization, owing to its implicit pruning of features based on feature selectivity, makes it an attractive alternative to explicit sparsification/pruning approaches. The link between ablation of selective features and explicit feature pruning is also established in prior work (Morcos et al., 2018; Zhou et al., 2018). We show that the emergence of selective features in Adam, and the increased propensity for pruning the said selective features when using L2 regularization may thus be helpful both for better generalization performance and network speedup.

3. Experimental Setup and Results

We use a 7-layer convolutional network with 2 fully connected layers as shown in Figure 1. We refer to this network as *BasicNet* in the rest of the document. We also compare VGG-11/16 (Simonyan & Zisserman, 2014) on various datasets. For the basic experiments on CIFAR-10/100, we use a variety of gradient descent approaches, a mini-batch size of 40, with a method specific base learning rate for 250 epochs which is scaled down by 10 for 75 more epochs. The base learning rates and other hyperparameters are as

Table 1. Convolutional filter sparsity in *BasicNet* trained on CIFAR10/100 for different combinations of regularization and gradient descent methods. Shown are the % of non-useful / inactive convolution filters, as measured by activation over training corpus (max act. $< 10^{-12}$) and by the learned BatchNorm scale ($|\gamma| < 10^{-03}$), averaged over 3 runs. The lowest test error per optimizer is highlighted, and sparsity (green) or lack of sparsity (red) for the best and near best configurations indicated via text color. L2: L2 regularization, WD: Weight decay (adjusted with the same scaling schedule as the learning rate schedule).

		CIFAR10			CIFAR100		
		% Sparsity		Test	% Sparsity		Test
L2		by Act	by γ	Error	by Act	by γ	Error
		SGD	1e-03	27	27	21.8	23
1e-04	0		0	11.8	0	0	37.4
1e-05	0		0	10.5	0	0	39.0
0	0		0	11.3	0	0	40.1
Adam	2e-03	88	86	14.7	82	81	42.7
	1e-04	71	70	10.5	47	47	36.6
	1e-05	48	48	10.7	5	5	40.6
	0	3	0	11.0	0	0	40.3
Adadelata	5e-04	82	82	13.6	61	61	39.1
	2e-04	40	40	11.3	3	3	35.4
	1e-04	1	1	10.2	1	1	35.9
Adagrad	2e-02	75	75	11.3	88	88	63.3
	1e-02	65	65	11.2	59	59	37.2
	5e-03	56	56	11.3	24	25	35.9
AMSGrad	1e-02	93	93	20.9	95	95	71.9
	1e-04	51	47	9.9	20	13	35.6
	1e-06	0	0	11.2	0	0	40.2
Adamax	1e-02	75	90	16.4	74	87	51.8
	1e-04	49	50	10.1	10	10	39.3
	1e-06	4	4	11.3	0	0	39.8
RMSProp	1e-02	95	95	26.9	97	97	78.6
	1e-04	72	72	10.4	48	48	36.3
	1e-06	29	29	10.9	0	0	40.6
WD		CIFAR10			CIFAR100		
		% Sparsity		Test	% Sparsity		Test
SGD		by Act	by γ	Error	by Act	by γ	Error
		1e-03	27	27	21.6	23	23
	2e-04	0	0	13.3	0	0	39.4
	1e-04	0	0	12.4	0	0	37.7
Adam	5e-04	81	81	18.1	59	59	43.3
	2e-04	60	60	13.4	16	16	37.3
	1e-04	40	40	11.2	3	3	36.2

follows: Adam (1e-3, $\beta_1=0.9$, $\beta_2=0.99$, $\epsilon=1e-8$), Adadelata (1.0, $\rho=0.9$, $\epsilon=1e-6$), SGD (0.1, mom.=0.9), Adagrad (1e-2), AMSGrad (1e-3), AdaMax (2e-3), RMSProp (1e-3).

Observing Sparsity Trends: We study the effect of varying the amount and type of regularization² on the extent

²Note that L2 regularization and weight decay are distinct. See (Loshchilov & Hutter, 2017) for a detailed discussion.

of sparsity and test error in Table 1. It shows significant convolutional filter sparsity emerges with adaptive gradient descent methods when combined with L2 regularization. The extent of sparsity is reduced when using Weight Decay instead, and absent entirely in the case of SGD with moderate levels of regularization. The extent of sparsity decreases with increasing mini-batch size, as shown in Tables 2, 3, 4, 5, 6 and 7.

Table 2. BasicNet sparsity variation on CIFAR10/100 trained with Adam and L2 regularization with different mini-batch sizes.

	Batch Size	CIFAR 10				CIFAR 100			
		Train Loss	Test Loss	Test Err	%Spar. by γ	Train Loss	Test Loss	Test Err	%Spar. by γ
L2: 1e-4	20	0.17	0.36	11.1	70	0.69	1.39	35.2	57
	40	0.06	0.43	10.5	70	0.10	1.98	36.6	46
	80	0.02	0.50	10.1	66	0.02	2.21	41.1	35
	160	0.01	0.55	10.6	61	0.01	2.32	44.3	29

Table 3. Convolutional filter sparsity for BasicNet trained on Tiny-ImageNet, with different mini-batch sizes.

Batch Size	Train Loss	Val Loss	Top 1 Val Err.	Top 5 Val Err.	% Spar. by γ
20	1.05	2.13	47.7	22.8	63
40	0.16	2.96	48.4	24.7	48
120	0.01	2.48	48.8	27.4	26

Table 4. Layerwise % Sparsity by γ for VGG-16 on CIFAR10 and 100, with different mini-batch sizes, compared to the handcrafted sparse structure of (Li et al., 2017)

	CIFAR 10				CIFAR 100		
	Adam, L2:1e-4 B: 40	B: 80	B: 160	Li et al.	Adam, L2:1e-4 B: 40	B: 80	B: 160
%Feat. Pruned	86	84	83	37	76	69	69
Test Err	7.2	7.0	6.5	6.6	29.2	28.1	27.8

Table 5. Sparsity by γ on VGG-16, trained on TinyImageNet, and on ImageNet. Also shown are the pre- and post-pruning top-1/top-5 single crop validation errors. Pruning using $|\gamma| < 10^{-3}$ criteria.

TinyImageNet	# Conv	Pre-pruning		Post-pruning	
	Feat. Pruned	top1	top5	top1	top5
L2: 1e-4, B: 20	3016 (71%)	45.1	21.4	45.1	21.4
L2: 1e-4, B: 40	2571 (61%)	46.7	24.4	46.7	24.4
ImageNet					
L2: 1e-4, B: 40	292	29.93	10.41	29.91	10.41

Comparison With Explicit Feature Sparsification Approaches: For VGG-16, we compare the network trained on CIFAR-10 with Adam using different mini-batch sizes against the handcrafted approach of Li et al. (Li et al., 2017). Similar to tuning the explicit sparsification hyperparameter in (Liu et al., 2017), the mini-batch size can be varied to find the sparsest representation with an acceptable level of test performance. We see from Table 4 that when trained with a batch size of 160, 83% of the features can be pruned away and leads to a better performance than the 37% of the

features pruned for (Li et al., 2017). For VGG-11 on ImageNet (Table 6), by simply varying the mini-batch size from 90 to 60, the number of convolutional features pruned goes from 71 to 140. This is in the same range as the number of features pruned by the explicit sparsification approach of (Li et al., 2017), and gives a comparable top-1 and top-5 validation error. For VGG-19 on CIFAR10 and CIFAR100 (Table 7), we see again that varying the mini-batch size controls the extent of sparsity. For the mini-batch sizes we considered, the extent of sparsity is much higher than that of (Liu et al., 2017), with consequently slightly worse performance. Tweaking the mini-batch size or other hyperparameters would further tradeoff sparsity for accuracy, and reach a comparable sparsity-accuracy as (Liu et al., 2017).

Table 6. Effect of different mini-batch sizes on sparsity (by γ) in VGG-11, trained on ImageNet. Same network structure employed as (Liu et al., 2017). * indicates finetuning after pruning

	# Conv	Pre-pruning		Post-pruning	
	Feat. Pruned	top1	top5	top1	top5
Adam, L2: 1e-4, B: 90	71	30.50	10.65	30.47	10.64
Adam, L2: 1e-4, B: 60	140	31.76	11.53	31.73	11.51
(Liu et al., 2017)	85	29.16		31.38*	-

Table 7. Sparsity by γ on VGG-19, trained on CIFAR10/100. Also shown are the post-pruning test error. Compared with explicit sparsification approach of (Liu et al., 2017)

	CIFAR 10			CIFAR 100		
	Adam, L2:1e-4 B: 64	B: 512	(Liu et al., 2017)	Adam, L2:1e-4 B: 64	B: 512	(Liu et al., 2017)
%Feat. Pruned	85	81	70	75	62	50
Test Err	7.1	6.9	6.3	29.9	28.8	26.7

4. Discussion

The penalization of selective features can be seen as a greedy local search heuristic. While the extent of implicit filter sparsity is significant, it obviously does not match up with some of the more recent approaches (He et al., 2018; Lin et al., 2017) which utilize more expensive model search and advanced heuristics such as filter redundancy.

One question that emerges from our findings is whether other filter sparsification heuristics can be implicitly incorporated into the network learning process, such that the training process jointly optimizes for the training objective and network capacity.

Another issue to consider is the selective-feature pruning criteria itself. On the opposite end of the spectrum from selective features lie non-selective features, which putatively have comparably low discriminative information as selective features and could also be pruned. These non-selective features are however not captured by greedy local search heuristics because pruning them can have a significant impact on the accuracy. The accuracy can presumably be recouped after fine-tuning, and this remains open as a possible line of future investigation.

References

- He, Y., Lin, J., Liu, Z., Wang, H., Li, L.-J., and Han, S. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 784–800, 2018.
- Hu, H., Peng, R., Tai, Y.-W., and Tang, C.-K. Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. *arXiv preprint arXiv:1607.03250*, 2016.
- Li, H., Kadav, A., Durdanovic, I., Samet, H., and Graf, H. P. Pruning filters for efficient convnets. In *ICLR*, 2017.
- Lin, J., Rao, Y., Lu, J., and Zhou, J. Runtime neural pruning. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 2181–2191. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/6813-runtime-neural-pruning.pdf>.
- Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., and Zhang, C. Learning efficient convolutional networks through network slimming. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pp. 2755–2763. IEEE, 2017.
- Loshchilov, I. and Hutter, F. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101*, 2017.
- Mehta, D., Kim, K.-I., and Theobalt, C. On implicit filter level sparsity in convolutional neural networks. In *Proc. IEEE CVPR*, 2019.
- Molchanov, P., Tyree, S., Karras, T., Aila, T., and Kautz, J. Pruning convolutional neural networks for resource efficient inference. 2017.
- Morcos, A. S., Barrett, D. G., Rabinowitz, N. C., and Botvinick, M. On the importance of single directions for generalization. 2018.
- Mozer, M. C. and Smolensky, P. Skeletonization: A technique for trimming the fat from a network via relevance assessment. In *Advances in neural information processing systems*, pp. 107–115, 1989.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Srinivas, S. and Babu, R. V. Data-free parameter pruning for deep neural networks. In *BMVC*, 2016.
- Theis, L., Korshunova, I., Tejani, A., and Huszár, F. Faster gaze prediction with dense networks and fisher pruning. In *ICLR*, 2017.
- Ye, J., Lu, X., Lin, Z., and Wang, J. Z. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. In *ICLR*, 2018.
- Zhou, B., Sun, Y., Bau, D., and Torralba, A. Revisiting the importance of individual units in cnns via ablation. *arXiv preprint arXiv:1806.02891*, 2018.