

GASL: GUIDED ATTENTION FOR SPARSITY LEARNING IN DEEP NEURAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

The main goal of network pruning is imposing sparsity on the neural network by increasing the number of parameters with zero value in order to reduce the architecture size and the computational speedup. In most of the previous research works, sparsity is imposed stochastically without considering any prior knowledge of the weights distribution or other internal network characteristics. Enforcing too much sparsity may induce accuracy drop due to the fact that a lot of important elements might have been eliminated. In this paper, we propose *Guided Attention for Sparsity Learning* (GASL) to achieve **(1)** model compression by having less number of elements and speedup; **(2)** prevent the accuracy drop by supervising the sparsity operation via a guided attention mechanism and **(3)** introduce a generic mechanism that can be adapted for any type of architecture; Our work is aimed at providing a framework based on an interpretable attention mechanisms for imposing structured and non-structured sparsity in deep neural networks. For *Cifar-100* experiments, we achieved the state-of-the-art sparsity level and $2.91 \times$ speedup with competitive accuracy compared to the best method. For *MNIST* and *LeNet* architecture we also achieved the highest sparsity and speedup level.

1 INTRODUCTION

Recent advances in deep neural networks came with ideas to train deep architectures that have led to near-human accuracy for image recognition, object categorization and a wide variety of other applications LeCun et al. (2015); Maturana & Scherer (2015); Schmidhuber (2015); Mnih et al. (2013); Hinton et al. (2012). One possible issue is that an over-parameterized network may make the architecture overcomplicated for the task at hand and it might be prone to over-fitting as well. In addition to the model complexity, a huge amount of computational power is required to train such deep models due to having billions of weights. Moreover, even if a huge model is trained, it cannot be effectively employed for model evaluation on low-power devices mainly due to having exhaustive matrix multiplications Courbariaux et al. (2015).

So far, a wide variety of approaches have been proposed for creating more compact models. Traditional methods include model compression Ba & Caruana (2014); Hinton et al. (2015), network pruning Han et al. (2015b), sparsity-inducing regularizer Collins & Kohli (2014), and low-rank approximation Jaderberg et al. (2014); Denton et al. (2014); Ioannou et al. (2015); Tai et al. (2015). The aforementioned methods usually induce random connection pruning which yields to few or no improvement in the computational cost. On the other hand, structured pruning methods proposed to compress the architecture with significant computational efficiency Wen et al. (2016); Neklyudov et al. (2017).

One of the critical subjects of interest in sparsity learning is to maintain the accuracy level. In this paper, we discuss the *intuitive reasons behind the accuracy drop and propose a method to prevent it*. The important step is to determine how the sparsity and accuracy are connected together in order to be able to propose a mechanism for controlling the sparsity to prevent severe accuracy drop. In order to connect the sparsity to accuracy, intuitively, the accuracy drop is caused by imposing too much sparsity on the network in a way that the remaining elements cannot transfer enough information for optimal feature extraction for the desired task. Another intuitive reasoning is to argue that the sparsity is not supervised with any attention towards the model performance during optimization.

For effective network pruning and feature selection, different approaches such as employing the group lasso for sparse structure learning Yuan & Lin (2006), structure scale constraining Liu et al. (2015), and structured regularizing deep architectures known as Structured Sparsity Learning (SSL) Wen et al. (2016) have previously been proposed. For most of the previous research efforts, there is lack of addressing the direct effect of the proposed method on the combination of the sparsity and accuracy drop. One may claim that successful sparsity imposing with negligible accuracy drop might be due to the initial over-parameterizing the network. Moreover, there is *no control mechanism to supervise the sparsity operation connected to the model performance* which limits the available methods to intensive hyper-parameter tuning and multiple stages of training.

Our contribution. We designed and employed a supervised attention mechanism for sparsity learning which: **(1)** performs model compression for having less number of parameters **(2)** prevents the accuracy drop by sparsity supervision by paying an attention towards the network using variance regularization and **(3)** is a generic mechanism that is not restricted by the sparsity penalty or any other limiting assumption regarding the network architecture. To the best of our knowledge, *this is the first research effort which proposes a supervised attention mechanism for sparsity learning.*

Paper Organization. At first, we provide a review of the related research efforts (Section 2). Then, we introduce the attention mechanism which is aimed at forcing some sections of the network to be active (Section 3). Later in Section 4, we propose an algorithm only for the attention supervision. We complement our proposed method in Section 5, by providing experimental results for which we target the sparsity level, accuracy drop and robustness of the model to hyper-parameter tuning. As will be observed, the proposed mechanism prevents the severe accuracy drop in higher levels of sparsity. We will empirically show the robustness to exhaustive hyper-parameter tuning in addition to performance superiority of the proposed method in higher sparsity levels.

2 RELATED WORKS

Network weights pruning. Network compression for parameter reduction has been of great interest for a long time and a large number of research efforts are dedicated to it. In Han et al. (2015b;a); Ullrich et al. (2017); Molchanov et al. (2017), network pruning has been performed with a significant reduction in parameters, although they suffer from computational inefficiency due to the mere weight pruning rather than the structure pruning.

Network structure pruning. In Louizos et al. (2017a); Wen et al. (2016); Neklyudov et al. (2017), pruning the unimportant parts of the structure¹ rather than simple weight pruning has been proposed and significant computational speedup has been achieved. However, for the aforementioned methods, the architecture must be fully trained at first and the potential training speedup regarding the sparsity enforcement cannot be attained. A solution for training speedup has been proposed by ℓ_0 -regularization technique by using online sparsification Louizos et al. (2017b). Training speedup is of great importance but adding a regularization for solely speeding up the training (because of the concurrent network pruning) is not efficient due to adding a computational cost for imposing ℓ_0 -regularization itself. Instead, we will use an adaptive gradient clipping Pascanu et al. (2013) for training speedup.

Attention. In this paper, the goal is to impose the sparsity in an accurate and interpretable way using the attention mechanism. So far, attention-based deep architectures has been proposed for image Fu et al. (2017); Jia et al. (2015); Mnih et al. (2014); Xu et al. (2015) and speech domains Bahdanau et al. (2016); Chorowski et al. (2015); Toro et al. (2005), as well as machine translation Bahdanau et al. (2014); Luong et al. (2015); Vaswani et al. (2017). Recently, the supervision of the attention mechanism became a subject of interest as well Liu et al. (2016; 2017); Chen et al. (2016); Mi et al. (2016) for which they proposed to supervise the attention using some external guidance. We propose the use of guided attention for enforcing the sparsity to map the sparsity distribution of the targeted elements² to the desired target distribution.

¹This can be neurons in fully-connected layers or channels/filters in convolutional layers.

²Elements on which the sparsity enforcement is desired: Weights, channels, neurons and etc.

3 PROPOSED ATTENTION MECHANISM

The main objective of the attention mechanism is to control and supervise the sparsity operation. For this aim, it is necessary to propose a method which is neither dependent on the architecture of the model nor to any layer type while maintaining the model accuracy and enforcing the compression objective. Considering the aforementioned goals, we propose the *variance loss* as an auxiliary cost term to force the distribution of the weights³ to be skewed. A skewed distribution with a high variance and a concentration on zero (to satisfy the sparsity objective) is desired. Our proposed scheme supervises the sparsity operation to keep a portion of the targeted elements (such as weights) to be dominant (with respect to their magnitude) as opposed to the other majority of the weights to simultaneously impose and control sparsity. Intuitively, this mechanism may force a portion of weights to survive for sufficient information transmission through the architecture.

Assume enforcing the sparsity is desired on a parametric model; let's have the training samples with $\{x_i, y_i\}$ as pairs. We propose the following objective function which is aimed to create a sparse structure in addition to the *variance regularization*:

$$\begin{aligned} \Omega(\theta) &= \frac{1}{N} \left(\sum_{i=1}^N \Gamma(F(\mathbf{x}_i; \theta), \mathbf{y}_i) \right) + R(\theta) + \lambda_s \cdot \Theta(G(\theta)) + \lambda_v \cdot \Psi^{-1}(H(\theta)), \\ \theta_{opt} &= \underset{\theta}{\operatorname{argmin}} \{ \Omega(\theta) \}, \end{aligned} \quad (1)$$

in which $\Gamma(\cdot)$ corresponds to the cross-entropy loss function and θ can be any combination of the target parameters. Model function is defined as $F(\cdot)$, $R(\cdot)$ is some regularization function, $G(\cdot)$ and $H(\cdot)$ are some arbitrary functions⁴ on parameters (such as grouping parameters), N is the number of samples, λ parameters are the weighting coefficients for the associated losses and $\Theta(\cdot)$ and $\Psi(\cdot)$ are the sparsity and variance functions⁵, respectively. The variance function is the utilized regularization term for any set of θ parameters⁶. The inverse operation on top of the $\Psi(\cdot)$ in Eq. 1 is necessary due to the fact that the higher variance is the desired objective. The power of the variance as a regularizer has been investigated in Namkoong & Duchi (2017). In this work, we expand the scope of variance regularization to the sparsity supervision as well.

3.1 MODEL COMPLEXITY

Adding a new term to the loss function can increase the model complexity due to adding a new hyper-parameter (the coefficient of the variance loss). For preventing this issue, we propose to have a dependent parameter as the variance loss coefficient. If the new hyperparameter is defined in terms of a variable dependent upon another hyperparameter, then it does not increase the model complexity. Considering the suggested approach, a dependency function must be defined over the hyperparameters definition. The dependency is defined as $\lambda_v = f(\lambda_s) = \alpha \times \lambda_s$ in which α is a scalar multiplier.

3.2 STRUCTURED ATTENTION FOR GROUP LASSO REGULARIZATION

Group Sparsity. Group sparsity has widely been utilized mostly for its feature selection ability by deactivating neurons⁷ via imposing sparsity on the whole cluster of weights in a group Yuan & Lin (2006); Meier et al. (2008). Regarding the Eq. 1, the group sparsity objective function can be defined by following expression:

³Or any elements on which we are enforcing sparsity.

⁴Diffrentiable in general.

⁵The $\Psi(\theta)$ is simply taking the variance on the set of θ parameters if $H(\cdot)$ is the identity function.

⁶Such as groups, weights and etc.

⁷Or channels in convolutional layer.

$$\ominus(G(\mathbf{w})) = \sum_{l=1}^{N_l} \frac{1}{\sqrt{|G(W^l)|}} \left(\sum_{j=1}^M \sqrt{\sum_{i=1}^{|w^{(j)}|} (w_i^{(j)})^2} \right)_l, \quad (2)$$

in which $w^{(j)}$ is the j_{th} group of weights in w and $|w^{(j)}|$ is the number of weights in the associated group in the case of having M groups. The l indicates the layer index, $|G(W^l)|$ is a normalizer factor which is in essence the number of groups for the l_{th} layer and $(\cdot)_l$ demonstrates the elements (weights) belonging to the the l_{th} layer.

Structured attention. We propose a *Structured Attention (SA)* regularization, which adds the attention mechanism on group sparsity learning (the sparsity imposition operation is similar to SSL Wen et al. (2016)). The attention is on the predefined groups. Under our general framework, it can be expressed by the following substitutions in Eq. 1:

$$\Psi(H(\mathbf{w})) = \sum_{l=1}^{N_l} \frac{1}{\sqrt{|G(W^l)|}} \left(\frac{1}{M} \sum_{j=1}^M \left(\sqrt{\sum_{i=1}^{|w^{(j)}|} (w_i^{(j)})^2} - \frac{1}{M} \sum_{k=1}^M \sqrt{\sum_{i=1}^{|w^{(k)}|} (w_i^{(k)})^2} \right)^2 \right)_l, \quad (3)$$

which is simply the variance of the group values for each layer, normalized by a factor and aggregated for all the layers.

Generalizability. It is worth noting that our proposed mechanism is not limited to the suggested structured method. It can operate on any $\ominus(\cdot)$ function as sparsity objective because the definition of the attention is independent of the type of the sparsity. As an example, one can utilize an *unstructured attention* which is simply putting the attention function $\Psi(\cdot)$ on all the network weights without considering any special groups of weights or prior objectives such as pruning unimportant channels or filters.

4 GASL: GUIDED ATTENTION IN SPARSITY LEARNING

The attention mechanism observes the areas of structure⁸ on which the sparsity is supposed to be enforced. we propose the *Guided Attention in Sparsity Learning (GASL)* mechanism, which aims at the attention supervision toward mapping the distribution of the elements' values to a certain target distribution. The target distribution characteristics must be aligned with the attention objective function with the goal of increasing the variance of the elements for sparsity imposition.

4.1 INCREASING VARIANCE BY ADDITIVE RANDOM SAMPLES

Assume we have the vector $V(\boldsymbol{\theta}) = [\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_{|\boldsymbol{\theta}|}]^T$ that is the values of the elements in the group $[\boldsymbol{\theta}] = \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_{|\boldsymbol{\theta}|}\}$ and for which we want to maximize the variance. In Paisley et al. (2012), variational Bayesian inference has been employed for the gradient computation of variational lower bound. Inspired by Wang et al. (2013), in which random vectors are used for stochastic gradient optimization, we utilize the additive random vectors for variance regularization. The random vector is defined as $V^r(\boldsymbol{\theta}) = [V_1^r, V_2^r, \dots, V_{|\boldsymbol{\theta}|}^r]^T$. The formulation is as below:

$$\hat{V}(\boldsymbol{\theta}) = V(\boldsymbol{\theta}) + M \cdot (V^r(\boldsymbol{\theta}) - E(V^r(\boldsymbol{\theta}))), \quad (4)$$

where M is a $|\boldsymbol{\theta}| \times |\boldsymbol{\theta}|$ matrix. The resulted vector $\hat{V}(\boldsymbol{\theta})$ does not make any changes in the mean of the parameters distribution since it has the same mean as the initial $V(\boldsymbol{\theta})$ vector.

Proposition 1. Assume the variance maximization of the $\hat{V}(\boldsymbol{\theta})$ is desired. The choice of matrix M , does not enforce any upper bound on the variance vector $\hat{V}(\boldsymbol{\theta})$.

⁸Groups, weights or elements.

Proof. For that aim, the task breaks to the subtask of finding the optimal M for which the trace of the $\hat{V}(\theta)$ is maximized. For the mini-batch optimization problem, we prove that the proposed model is robust to the selection of M . The maximizer M can be obtained when the trace of the variance matrix is maximized and it can be demonstrated as follows:

$$\begin{aligned}
 M_* &= \operatorname{argmax}_M \left\{ \operatorname{Tr} \left(\operatorname{Var}[\hat{V}(\theta)] \right) \right\} \\
 &= \operatorname{argmax}_M \left\{ \operatorname{Tr} \left(\operatorname{Cov}[\hat{V}(\theta), \hat{V}(\theta)] \right) \right\} \\
 &= \operatorname{argmax}_M \left\{ \operatorname{Tr} \left(\operatorname{Var}[V(\theta)] + \Upsilon + \Upsilon^T + M \cdot \operatorname{Var}[V^r(\theta)] \cdot M^T \right) \right\} \\
 &\quad , \Upsilon = M \cdot \operatorname{Cov}[V^r(\theta), V(\theta)].
 \end{aligned} \tag{5}$$

As can be observed from Eq. 5, as long as M is a positive definite matrix, the additive random can add to the value of the matrix trace without any upper bound. The detailed mathematical proof is available in the Appendix. \blacktriangle

Considering the mathematical proof, one can infer that the mere utilization of the variance loss term in Eq. 1, as the attention mechanism and without the additive random vector, can supervise the sparsity operation. However, we will empirically show that the additive random vectors can improve the accuracy due to the supervision of the attention. *The supervision of the variance is important regarding the fact that the high variance of the parameters may decrease the algorithm speed and performance for sparsity enforcement. This is due to the large number of operations that is necessary for gradient descent to find the trade-off between the sparsity and variance regularization.* From now one, without the loss of generality, we assume M to be *identity matrix* in our experiments.

4.2 THE CHOICE OF RANDOM VECTOR V^r

In practice, the choice of V^r should be highly correlated with V . Furthermore, Eq. 5 shows that without being correlated, the terms associated with $\operatorname{Cov}[V^r(\theta), V(\theta)]$ may go to zero which affects the high variance objective, negatively. The algorithm for random vector selection is declared in Algorithm. 1. *The distribution pdf(.) should be specified regarding the desired output distribution.* We chose *log-normal distribution* due to its special characteristics which create a concentration around zero and a skewed tail Limpert et al. (2001). If the variance of the random vector $V^r(\theta)$ is less than the main vector $V(\theta)$, no additive operation will be performed. In case the $[\theta]$ parameters variance is high-enough compared to the $V(\theta)$ vector, then there is no need for additive random samples. This preventing mechanism has been added due to the practical speedup.

Algorithm 1: GASL algorithm.

Data: Extract mini-batch;

Parameters group: Form $V(\theta) = [\theta_1, \theta_2, \dots, \theta_{|\theta|}]^T$;

Random sampling: Draw $\beta_k \sim \operatorname{pdf}(\beta)$ for k in $\{1, 2, \dots, \theta\}$;

Random vector creation: Form $V^r(\theta) = [\beta_1, \beta_2, \dots, \beta_{|\theta|}] \cdot V(\theta)$;

if $\operatorname{Var}[V^r(\theta)] > \operatorname{Var}[V(\theta)]$ **then**

Replacement vector: Calculate $\hat{V}(\theta) = V(\theta) + M \cdot (V^r(\theta) - E(V^r(\theta)))$;

Replacement operation: Replace $\hat{V}(\theta)$ with $V(\theta)$;

Return: $\hat{V}(\theta)$;

else

Return: $V(\theta)$;

Computation: Update gradient;

4.3 COMBINATION OF GASL AND SA

GASL algorithm can operate on the top of the structured attention for attention supervision. The schematic is depicted in Fig. 1. Furthermore, a visualized example of the output channels from the second convolutional layer in the MNIST experiments has also demonstrated in Fig. 1. The structured attention is dedicated to the output channels of a convolutional layer.

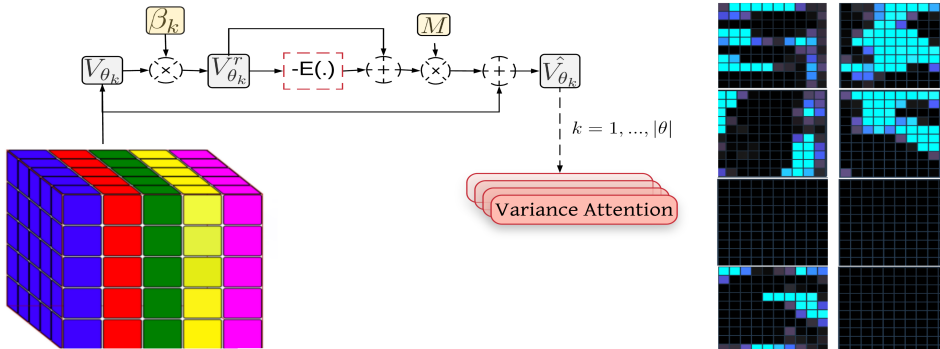


Figure 1: The combination of GASL and structured attention. The cube demonstrates the output feature map of a convolutional layer. The weights associated with each channel, form a group. For visualization, the right column is the activation visualization of the attention-based sparsity enforcement on output channels and the left one is the results of imposing sparsity without attention. As can be observed, some of the channels are turned off and the remaining ones are intensified.

5 EXPERIMENTAL RESULTS

We use three databases for the evaluation of our proposed method: MNIST LeCun et al. (1998), CIFAR-10 and CIFAR-100 Krizhevsky & Hinton (2009). For increasing the convergence speed without degrading the overall performance, we used gradient clipping Pascanu et al. (2013). A common approach is to clip individual gradients to some fixed predefined range $[-\zeta, \zeta]$. As the learning rate becomes smaller continuously, the effective gradient⁹ will approach zero and training convergence may become extremely slow. For tackling this issue, we used the method proposed in Kim et al. (2016) for gradient clipping which defined the range dynamically as $[-\zeta/\gamma, \zeta/\gamma]$ for which γ is the current learning rate. We chose $\zeta = 0.1$ in our experiments. Hyper-parameters are selected by cross-validation. For all our experiments, *the output channels of convolutional layers and neurons in fully connected layers are considered as groups*.

5.1 MNIST DATASET

For experiments on MNIST dataset, we use ℓ_2 -regularization with the default hyperparameters. Two network architectures have been employed: LeNet-5-Caffe¹⁰ and a multilayer perceptron (MLP). For the MLP network, the group sparsity is enforced on each neuron’s outputs for feature selection; Same can be applied on the neurons’ inputs as well. The results are shown in Table. 1. The percentage of sparsity is reported layer-wise. One important observation is the superiority of the SA method to the SSL in terms of accuracy, while the sparsity objective function for both is identical and the only difference is the addition of structural attention (Eq. 3). As a comparison to Louizos et al. (2017a), we achieved closed sparsity level with better accuracy for the MLP network. For the LeNet network, we obtained the highest sparsity level with competitive accuracy compared to the best method proposed in Molchanov et al. (2017).

⁹Which is $gradient \times learning_rate$

¹⁰<https://github.com/BVLC/caffe/blob/master/examples/mnist>

Table 1: Experiments on LeNet-5-Caffe architecture with 20-50-800-500 number of output filters and hidden layers and MLP with the architecture of 784-500-300 as the number of hidden units for each layer. The sparsity level is reported layer-wise. The sparsity and error are both reported as %.

Method	LeNet		MLP	
	Error	Sparsity / speedup	Error	Sparsity / speedup
baseline	0.8	0-0-0-0 / 1.00×	1.54	0-0-0 / 1.00×
ℓ_1 -regularization	2.44	15-31-37-53 / 1.07×	3.26	21-23-15 / 1.01×
Network Pruning Han et al. (2015b)	1.21	61-58-80-67 / 1.17×	1.71	20-32-69 / 1.03×
Bayesian Compression Louizos et al. (2017a)	0.9	60-74-89-97 / 2.31×	1.8	71-81-94 / 1.18×
Structured BP Neklyudov et al. (2017)	0.86	85-62-64-43 / 2.03×	1.55	68-77-89 / 1.28 ×
Structured Sparsity Learning Wen et al. (2016)	1.00	71-58-61-34 / 1.83×	1.49	52-61-74 / 1.12×
Sparse Variational Dropout Molchanov et al. (2017)	0.75	66-36-59-75 / 1.43×	1.57	31-56-57 / 1.05×
ℓ_0 -regularization Louizos et al. (2017b)	1.02	8-62-96-17 / 1.31×	1.41	32-34-37 / 1.07×
Structured Attention	1.05	78-62-72-50 / 1.91×	1.56	22-31-62 / 1.04×
Structured Attention + GASL	0.92	76-88-86-95 / 2.41 ×	1.53	64-80-95 / 1.23×

5.2 EXPERIMENTS ON CIFAR-10 AND CIFAR-100

For experiments in this section, we used VGG-16 Simonyan & Zisserman (2014) as the baseline architecture. Random cropping, horizontal flipping, and per-image standardization have been performed for data augmentation in the training phase and in the evaluation stage, only center cropping has been used Krizhevsky et al. (2012). Batch-normalization has also been utilized after each convolutional layer and before the activation Ioffe & Szegedy (2015). The initial learning rate of 0.1 has been chosen and the learning rate is dropped by a factor of 10 when the error plateaus for 5 consecutive epochs. As can be observed from Table. 2, the combination of the GASL algorithm and SA dominates regarding the achieved sparsity level and demonstrates competitive results in terms of accuracy for Cifar-100. We terminate the training after 300 epochs or if the averaged error is not improving for 20 consecutive epochs, whichever comes earlier. For Cifar-10 Krizhevsky et al. (2010), we obtained the second best results for both accuracy and sparsity level.

Table 2: Experiments on Cifar-10 and Cifar-100 using VGG-16 network. The sparsity and error are both reported as %.

Method	Cifar-10		Cifar-100	
	Error	Sparsity / speedup	Error	Sparsity / speedup
baseline	8.75	0 / 1.00×	27.41	0 / 1.00×
ℓ_1 -regularization	11.43	22 / 1.36×	31.75	17 / 1.35×
Network Pruning Han et al. (2015b)	10.76	32 / 1.54×	28.46	22 / 1.46×
Bayesian Compression Louizos et al. (2017a)	8.42	82 / 3.16×	25.72	41 / 2.02×
Structured BP Neklyudov et al. (2017)	8.62	46 / 2.86×	25.47	29 / 2.45×
Structured Sparsity Learning Wen et al. (2016)	9.12	74 / 3.31 ×	26.42	46 / 2.43×
Sparse Variational Dropout Molchanov et al. (2017)	7.79	61 / 2.14×	24.91	42 / 2.11×
ℓ_0 -regularization Louizos et al. (2017b)	8.83	52 / 2.41×	26.73	38 / 1.92×
Structured Attention	8.62	48 / 2.80×	25.76	32 / 2.37×
Structured Attention + GASL	8.32	76 / 3.06×	25.41	48 / 2.91 ×

The advantage of the proposed method for higher sparsity levels. For Cifar-100 experiments, we continued the process of enforcing sparsity for achieving the desired level of compression¹¹. We chose three discrete level of sparsity and for any of which, the accuracy drop for different methods is reported. Table. 3 demonstrates the comparison of different methods with regard to their accuracy drops at different levels of sparsity. For some levels of sparsity, it was observed that some methods performed better than the baseline. We deliberately selected higher levels of sparsity for having some performance drop as opposed to the baseline for all the implemented methods. As can be observed, our method shows its performance superiority in accuracy for the higher levels of sparsity. *In another*

¹¹The increasing of the λ_s might have been necessary.

word, the proposed method outperforms in preventing the accuracy drop in the situation of having high sparsity level.

Table 3: The accuracy drop, compared to the baseline, for different levels of sparsity using Cifar-100. Our method here is SA + GASL. The accuracy drop is measured compared to the baseline results with no sparsity.

Sparsity	Method					
	ℓ_1	Neklyudov et al.	Wen et al.	Molchanov et al.	Louizos et al.	Ours
60%	1.51	0.67	0.74	0.59	0.61	0.71
70%	3.46	1.79	1.89	1.81	1.74	1.68
80%	4.21	2.68	2.64	2.73	2.81	2.23

Robustness to the hyperparameter tuning. Regarding the discussion in Section 3.1, it is worth to investigate the effect of λ_v on the accuracy drop. *In another word, we investigate the relative importance of tuning the variance loss coefficient.* The accuracy drop is reported for Cifar-100 experiments using different α values and sparsity levels. The results depicted in Table 4, empirically shows the robustness of the proposed method to the selection of α , as the dependent factor, for which in the dynamic range of $[0.1, 10]$, the accuracy drop is not changing drastically. *This clearly demonstrates the robustness of the proposed method to the selection of the new hyperparameter associated with the attention mechanism as it is only a dependent factor to the sparsity penalty coefficient.*

Table 4: Experiments on Cifar-100 for investigating the robustness of the proposed method to the hyperparameter selection.

Sparsity	$\alpha = \lambda_v/\lambda_s$				
	0.01	0.1	1.0	10	100
60%	2.1	0.86	0.71	0.67	1.61
70%	3.46	1.91	1.68	1.79	3.21
80%	4.21	2.51	2.13	2.09	4.41

6 CONCLUSION

In this paper, we proposed a guided attention mechanism for controlled sparsity enforcement by keeping a portion of the targeted elements to be alive. The GASL algorithm has been utilized on top of the structured attention for attention supervision to prune unimportant channels and neurons of the convolutional and fully-connected layers. We demonstrated the superiority of the method for preventing the accuracy drop in high levels of sparsity. Moreover, it has been shown that regardless of adding a new term to the loss function objective, the model complexity remains the same and the proposed approach is relatively robust to exhaustive hyper-parameter selection. Without the loss of generality, the method can be adapted to any layer type and different sparsity objectives such as weight pruning for unstructured sparsity or channel, neuron or filter cancellation for structured sparsity.

REFERENCES

- Jimmy Ba and Rich Caruana. Do deep nets really need to be deep? In *Advances in neural information processing systems*, pp. 2654–2662, 2014.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Dzmitry Bahdanau, Jan Chorowski, Dmitriy Serdyuk, Philemon Brakel, and Yoshua Bengio. End-to-end attention-based large vocabulary speech recognition. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pp. 4945–4949. IEEE, 2016.

- Wenhu Chen, Evgeny Matusov, Shahram Khadivi, and Jan-Thorsten Peter. Guided alignment training for topic-aware neural machine translation. *arXiv preprint arXiv:1607.01628*, 2016.
- Jan K Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio. Attention-based models for speech recognition. In *Advances in neural information processing systems*, pp. 577–585, 2015.
- Maxwell D Collins and Pushmeet Kohli. Memory bounded deep convolutional networks. *arXiv preprint arXiv:1412.1442*, 2014.
- Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pp. 3123–3131, 2015.
- Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *Advances in neural information processing systems*, pp. 1269–1277, 2014.
- Jianlong Fu, Heliang Zheng, and Tao Mei. Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition. In *Conf. on Computer Vision and Pattern Recognition*, 2017.
- Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015a.
- Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, pp. 1135–1143, 2015b.
- Geoffrey Hinton, Li Deng, Dong Yu, George E Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Tara N Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Yani Ioannou, Duncan Robertson, Jamie Shotton, Roberto Cipolla, and Antonio Criminisi. Training cnns with low-rank filters for efficient image classification. *arXiv preprint arXiv:1511.06744*, 2015.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866*, 2014.
- Xu Jia, Efstratios Gavves, Basura Fernando, and Tinne Tuytelaars. Guiding the long-short term memory model for image caption generation. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pp. 2407–2415. IEEE, 2015.
- Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1646–1654, 2016.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.
- Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). URL <http://www.cs.toronto.edu/kriz/cifar.html>, 2010.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.
- Eckhard Limpert, Werner A Stahel, and Markus Abbt. Log-normal distributions across the sciences: Keys and clues: On the charms of statistics, and how mechanical models resembling gambling machines offer a link to a handy way to characterize log-normal distributions, which can provide deeper insight into variability and probability—normal or log-normal: That is the question. *AIBS Bulletin*, 51(5):341–352, 2001.
- Baoyuan Liu, Min Wang, Hassan Foroosh, Marshall Tappen, and Marianna Pensky. Sparse convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 806–814, 2015.
- Chenxi Liu, Junhua Mao, Fei Sha, and Alan L Yuille. Attention correctness in neural image captioning. In *AAAI*, pp. 4176–4182, 2017.
- Lemao Liu, Masao Utiyama, Andrew Finch, and Eiichiro Sumita. Neural machine translation with supervised attention. *arXiv preprint arXiv:1609.04186*, 2016.
- Christos Louizos, Karen Ullrich, and Max Welling. Bayesian compression for deep learning. In *Advances in Neural Information Processing Systems*, pp. 3290–3300, 2017a.
- Christos Louizos, Max Welling, and Diederik P Kingma. Learning sparse neural networks through l_0 regularization. *arXiv preprint arXiv:1712.01312*, 2017b.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015.
- Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pp. 922–928. IEEE, 2015.
- Lukas Meier, Sara Van De Geer, and Peter Bühlmann. The group lasso for logistic regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1):53–71, 2008.
- Haitao Mi, Zhiguo Wang, and Abe Ittycheriah. Supervised attentions for neural machine translation. *arXiv preprint arXiv:1608.00112*, 2016.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Volodymyr Mnih, Nicolas Heess, Alex Graves, et al. Recurrent models of visual attention. In *Advances in neural information processing systems*, pp. 2204–2212, 2014.
- Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variational dropout sparsifies deep neural networks. *arXiv preprint arXiv:1701.05369*, 2017.
- Hongseok Namkoong and John C Duchi. Variance-based regularization with convex objectives. In *Advances in Neural Information Processing Systems*, pp. 2975–2984, 2017.
- Kirill Neklyudov, Dmitry Molchanov, Arsenii Ashukha, and Dmitry P Vetrov. Structured bayesian pruning via log-normal multiplicative noise. In *Advances in Neural Information Processing Systems*, pp. 6778–6787, 2017.
- John Paisley, David Blei, and Michael Jordan. Variational bayesian inference with stochastic search. *arXiv preprint arXiv:1206.6430*, 2012.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*, pp. 1310–1318, 2013.
- Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.

- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Cheng Tai, Tong Xiao, Yi Zhang, Xiaogang Wang, et al. Convolutional neural networks with low-rank regularization. *arXiv preprint arXiv:1511.06067*, 2015.
- Juan M Toro, Scott Sinnett, and Salvador Soto-Faraco. Speech segmentation by statistical learning depends on attention. *Cognition*, 97(2):B25–B34, 2005.
- Karen Ullrich, Edward Meeds, and Max Welling. Soft weight-sharing for neural network compression. *arXiv preprint arXiv:1702.04008*, 2017.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 6000–6010, 2017.
- Chong Wang, Xi Chen, Alexander J Smola, and Eric P Xing. Variance reduction for stochastic gradient optimization. In *Advances in Neural Information Processing Systems*, pp. 181–189, 2013.
- Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 2074–2082, 2016.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*, pp. 2048–2057, 2015.
- Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.