

L_0 -ARM: Network Sparsification via Stochastic Binary Optimization

Yang Li and Shihao Ji ✉

Georgia State University, USA
yli93@student.gsu.edu, sji@gsu.edu

Abstract. We consider network sparsification as an L_0 -norm regularized binary optimization problem, where each unit of a neural network (e.g., weight, neuron, or channel, etc.) is attached with a stochastic binary gate, whose parameters are jointly optimized with original network parameters. The Augment-Reinforce-Merge (ARM) [28], a recently proposed unbiased gradient estimator, is investigated for this binary optimization problem. Compared to the hard concrete gradient estimator from Louizos et al. [20], ARM demonstrates superior performance of pruning network architectures while retaining almost the same accuracies of baseline methods. Similar to the hard concrete estimator, ARM also enables conditional computation during model training but with improved effectiveness due to the *exact* binary stochasticity. Thanks to the flexibility of ARM, many smooth or non-smooth parametric functions, such as scaled sigmoid or hard sigmoid, can be used to parameterize this binary optimization problem and the unbiasedness of the ARM estimator is retained, while the hard concrete estimator has to rely on the hard sigmoid function to achieve conditional computation and thus accelerated training. Extensive experiments on multiple public datasets demonstrate state-of-the-art pruning rates with almost the same accuracies of baseline methods. The resulting algorithm L_0 -ARM sparsifies the Wide-ResNet models on CIFAR-10 and CIFAR-100 while the hard concrete estimator cannot. The code is public available at <https://github.com/leo-yangli/10-arm>.

Keywords: Network Sparsification · L_0 -norm Regularization · Binary Optimization.

1 Introduction

Deep Neural Networks (DNNs) have achieved great success in a broad range of applications in image recognition [3], natural language processing [4], and games [24]. Latest DNN architectures, such as ResNet [9], DenseNet [10] and Wide-ResNet [29], incorporate hundreds of millions of parameters to achieve state-of-the-art predictive performance. However, the expanding number of parameters not only increases the risk of overfitting, but also leads to high computational costs. Many practical real-time applications of DNNs, such as for smart phones, drones and the IoT (Internet of Things) devices, call for compute and memory

efficient models as these devices typically have very limited computation and memory capacities.

Fortunately, it has been shown that DNNs can be pruned or sparsified significantly with minor accuracy losses [8, 7], and sometimes sparsified networks can even achieve higher accuracies due to the regularization effects of the network sparsification algorithms [23, 20]. Driven by the widely spread applications of DNNs in real-time systems, there has been an increasing interest in pruning or sparsifying networks recently [8, 7], [26, 18, 19, 22, 23, 20]. Earlier methods such as the magnitude-based approaches [8, 7] prune networks by removing the weights of small magnitudes, and it has been shown that this approach although simple is very effective at sparsifying network architectures with minor accuracy losses. Recently, the L_0 -norm based regularization method [20] is getting attraction as this approach *explicitly* penalizes number of non-zero parameters and can drive redundant or insignificant parameters to be exact zero. However, the gradient of the L_0 regularized objective function is intractable. Louizos et al. [20] propose to use the hard concrete distribution as a close surrogate to the Bernoulli distribution, and this leads to a differentiable objective function while still being able to zeroing out redundant or insignificant weights during training. Due to the hard concrete substitution, however, the resulting hard concrete estimator is biased with respect to the original objective function.

In this paper, we propose L_0 -ARM for network sparsification. L_0 -ARM is built on top of the L_0 regularization framework of Louizos et al. [20]. However, instead of using a biased hard concrete gradient estimator, we investigate the Augment-Reinforce-Merge (ARM) [28], a recently proposed unbiased gradient estimator for stochastic binary optimization. Because of the unbiasedness and flexibility of the ARM estimator, L_0 -ARM exhibits a significantly faster rate at pruning network architectures and reducing FLOPs than the hard concrete estimator. Extensive experiments on multiple public datasets demonstrate the superior performance of L_0 -ARM at sparsifying networks with fully connected layers and convolutional layers. It achieves state-of-the-art prune rates while retaining similar accuracies compared to baseline methods. Additionally, it sparsifies the Wide-ResNet models on CIFAR-10 and CIFAR-100 while the original hard concrete estimator cannot.

The remainder of the paper is organized as follows. In Sec. 2 we describe the L_0 regularized empirical risk minimization for network sparsification and formulate it as a stochastic binary optimization problem. A new unbiased estimator to this problem L_0 -ARM is presented in Sec. 3, followed by related work in Sec. 4. Example results on multiple public datasets are presented in Sec. 5, with comparisons to baseline methods and the state-of-the-art sparsification algorithms. Conclusions and future work are discussed in Sec. 6.

2 Formulation

Given a training set $D = \{(\mathbf{x}_i, y_i), i = 1, 2, \dots, N\}$, where \mathbf{x}_i denotes the input and y_i denotes the target, a neural network is a function $h(\mathbf{x}; \boldsymbol{\theta})$ parametrized by $\boldsymbol{\theta}$ that fits to the training data D with the goal of achieving good generalization

to unseen test data. To optimize θ , typically a regularized empirical risk is minimized, which contains two terms – a data loss over training data and a regularization loss over model parameters. Empirically, the regularization term can be weight decay or Lasso, i.e., the L_2 or L_1 norm of model parameters.

Since the L_2 or L_1 norm only imposes shrinkage for large values of θ , the resulting model parameters θ are often manifested by smaller magnitudes but none of them are exact zero. Intuitively, a more appealing alternative is the L_0 regularization since the L_0 -norm measures *explicitly* the number of non-zero elements, and minimizing of it over model parameters will drive the redundant or insignificant weights to be exact zero. With the L_0 regularization, the empirical risk objective can be written as

$$\mathcal{R}(\theta) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}(h(\mathbf{x}_i; \theta), y_i) + \lambda \|\theta\|_0 \quad (1)$$

where $\mathcal{L}(\cdot)$ denotes the data loss over training data D , such as the cross-entropy loss for classification or the mean squared error (MSE) for regression, and $\|\theta\|_0$ denotes the L_0 -norm over model parameters, i.e., the number of non-zero weights, and λ is a regularization hyper-parameter that balances between data loss and model complexity.

To represent a sparsified model, we attach a binary random variable z to each element of model parameters θ . Therefore, we can re-parameterize the model parameters θ as an element-wise product of non-zero parameters $\tilde{\theta}$ and binary random variables \mathbf{z} :

$$\theta = \tilde{\theta} \odot \mathbf{z}, \quad (2)$$

where $\mathbf{z} \in \{0, 1\}^{|\theta|}$, and \odot denotes the element-wise product. As a result, Eq. 1 can be rewritten as:

$$\mathcal{R}(\tilde{\theta}, \mathbf{z}) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}(h(\mathbf{x}_i; \tilde{\theta} \odot \mathbf{z}), y_i) + \lambda \sum_{j=1}^{|\tilde{\theta}|} \mathbf{1}_{[z_j \neq 0]}, \quad (3)$$

where $\mathbf{1}_{[c]}$ is an indicator function that is 1 if the condition c is satisfied, and 0 otherwise. Note that both the first term and the second term of Eq. 3 are not differentiable w.r.t. \mathbf{z} . Therefore, further approximations need to be considered.

According to stochastic variational optimization [2], given any function $\mathcal{F}(\mathbf{z})$ and any distribution $q(\mathbf{z})$, the following inequality holds

$$\min_{\mathbf{z}} \mathcal{F}(\mathbf{z}) \leq \mathbb{E}_{\mathbf{z} \sim q(\mathbf{z})} [\mathcal{F}(\mathbf{z})], \quad (4)$$

i.e., the minimum of a function is upper bounded by the expectation of the function. With this result, we can derive an upper bound of Eq. 3 as follows.

Since $z_j, \forall j \in \{1, \dots, |\theta|\}$ is a binary random variable, we assume z_j is subject to a Bernoulli distribution with parameter $\pi_j \in [0, 1]$, i.e. $z_j \sim \text{Ber}(z; \pi_j)$. Thus,

we can upper bound $\min_{\mathbf{z}} \mathcal{R}(\tilde{\boldsymbol{\theta}}, \mathbf{z})$ by the expectation

$$\begin{aligned} \hat{\mathcal{R}}(\tilde{\boldsymbol{\theta}}, \boldsymbol{\pi}) &= \mathbb{E}_{\mathbf{z} \sim \text{Ber}(\mathbf{z}; \boldsymbol{\pi})} \mathcal{R}(\tilde{\boldsymbol{\theta}}, \mathbf{z}) \\ &= \mathbb{E}_{\mathbf{z} \sim \text{Ber}(\mathbf{z}; \boldsymbol{\pi})} \left[\frac{1}{N} \sum_{i=1}^N \mathcal{L} \left(h(\mathbf{x}_i; \tilde{\boldsymbol{\theta}} \odot \mathbf{z}), y_i \right) \right] + \lambda \sum_{j=1}^{|\tilde{\boldsymbol{\theta}}|} \pi_j. \end{aligned} \quad (5)$$

As we can see, now the second term is differentiable w.r.t. the new model parameters $\boldsymbol{\pi}$, while the first term is still problematic since the expectation over a large number of binary random variables \mathbf{z} is intractable and so its gradient. Since \mathbf{z} are binary random variables following a Bernoulli distribution with parameters $\boldsymbol{\pi}$, we now formulate the original L_0 regularized empirical risk (1) to a stochastic binary optimization problem (5).

Existing gradient estimators for this kind of discrete latent variable models include REINFORCE [27], Gumble-Softmax [11, 21], REBAR [25], RELAX [6] and the Hard Concrete estimator [20]. However, these estimators either are biased or suffer from high variance or computationally expensive due to auxiliary modeling. Recently, the Augment-Reinforce-Merge (ARM) [28] gradient estimator is proposed for the optimization of binary latent variable models, which is unbiased and exhibits low variance. Extending this gradient estimator to network sparsification, we find that ARM demonstrates superior performance of pruning network architectures while retaining almost the same accuracies of baseline models. More importantly, similar to the hard concrete estimator, ARM also enables conditional computation [1] that not only sparsifies model architectures for inference but also accelerates model training.

3 L_0 -ARM: Stochastic Binary Optimization

To minimize Eq. 5, we propose L_0 -ARM, a stochastic binary optimization algorithm based on the Augment-Reinforce-Merge (ARM) gradient estimator [28]. We first introduce the main theorem of ARM. Refer readers to [28] for the proof and other details.

Theorem 1. (ARM) [28]. For a vector of V binary random variables $\mathbf{z} = (z_1, \dots, z_V)$, the gradient of

$$\mathcal{E}(\boldsymbol{\phi}) = \mathbb{E}_{\mathbf{z} \sim \prod_{v=1}^V \text{Ber}(z_v; g(\phi_v))} [f(\mathbf{z})] \quad (6)$$

w.r.t. $\boldsymbol{\phi} = (\phi_1, \dots, \phi_V)$, the logits of the Bernoulli distribution parameters, can be expressed as

$$\nabla_{\boldsymbol{\phi}} \mathcal{E}(\boldsymbol{\phi}) = \mathbb{E}_{\mathbf{u} \sim \prod_{v=1}^V \text{Uniform}(u_v; 0, 1)} \left[\left(f(\mathbf{1}_{[\mathbf{u} > g(-\boldsymbol{\phi})]}) - f(\mathbf{1}_{[\mathbf{u} < g(\boldsymbol{\phi})]}) \right) (\mathbf{u} - 1/2) \right], \quad (7)$$

where $\mathbf{1}_{[\mathbf{u} > g(-\boldsymbol{\phi})]} := (\mathbf{1}_{[u_1 > g(-\phi_1)]}, \dots, \mathbf{1}_{[u_V > g(-\phi_V)]})^T$ and $g(\phi) = \sigma(\phi) = 1/(1 + \exp(-\phi))$ is the sigmoid function.

Parameterizing $\pi_j \in [0, 1]$ as $g(\phi_j)$, Eq. 5 can be rewritten as

$$\begin{aligned} \hat{\mathcal{R}}(\tilde{\boldsymbol{\theta}}, \boldsymbol{\phi}) &= \mathbb{E}_{\mathbf{z} \sim \text{Ber}(\mathbf{z}; g(\boldsymbol{\phi}))} [f(\mathbf{z})] + \lambda \sum_{j=1}^{|\tilde{\boldsymbol{\theta}}|} g(\phi_j) \\ &= \mathbb{E}_{\mathbf{u} \sim \text{Uniform}(\mathbf{u}; 0, 1)} [f(\mathbf{1}_{[\mathbf{u} < g(\boldsymbol{\phi})]})] + \lambda \sum_{j=1}^{|\tilde{\boldsymbol{\theta}}|} g(\phi_j), \end{aligned} \quad (8)$$

where $f(\mathbf{z}) = \frac{1}{N} \sum_{i=1}^N \mathcal{L}(h(\mathbf{x}_i; \tilde{\boldsymbol{\theta}} \odot \mathbf{z}), y_i)$. Now according to Theorem 1, we can evaluate the gradient of Eq. 8 w.r.t. $\boldsymbol{\phi}$ by

$$\begin{aligned} \nabla_{\boldsymbol{\phi}}^{ARM} \hat{\mathcal{R}}(\tilde{\boldsymbol{\theta}}, \boldsymbol{\phi}) &= \mathbb{E}_{\mathbf{u} \sim \text{Uniform}(\mathbf{u}; 0, 1)} \left[(f(\mathbf{1}_{[\mathbf{u} > g(-\boldsymbol{\phi})]}) - f(\mathbf{1}_{[\mathbf{u} < g(\boldsymbol{\phi})]})) (\mathbf{u} - 1/2) \right] \\ &\quad + \lambda \sum_{j=1}^{|\tilde{\boldsymbol{\theta}}|} \nabla_{\phi_j} g(\phi_j), \end{aligned} \quad (9)$$

which is an unbiased and low variance estimator as demonstrated in [28].

Note from Eq. 9 that we need to evaluate $f(\cdot)$ twice to compute the gradient, the second of which is the same operation required by the data loss of Eq. 8. Therefore, one extra forward pass $f(\mathbf{1}_{[\mathbf{u} > g(-\boldsymbol{\phi})]})$ is required by the L_0 -ARM gradient estimator. This additional forward pass might be computationally expensive, especially for networks with millions of parameters. To reduce the computational complexity of Eq. 9, we further consider another gradient estimator – Augment-Reinforce (AR) [28]:

$$\begin{aligned} \nabla_{\boldsymbol{\phi}}^{AR} \hat{\mathcal{R}}(\tilde{\boldsymbol{\theta}}, \boldsymbol{\phi}) &= \mathbb{E}_{\mathbf{u} \sim \text{Uniform}(\mathbf{u}; 0, 1)} \left[f(\mathbf{1}_{[\mathbf{u} < g(\boldsymbol{\phi})]}) (1 - 2\mathbf{u}) \right] \\ &\quad + \lambda \sum_{j=1}^{|\tilde{\boldsymbol{\theta}}|} \nabla_{\phi_j} g(\phi_j), \end{aligned} \quad (10)$$

which requires only one forward pass $f(\mathbf{1}_{[\mathbf{u} < g(\boldsymbol{\phi})]})$ that is the same operation as in Eq. 8. This L_0 -AR gradient estimator is still unbiased but with higher variance. Now with L_0 -AR, we can trade off the variance of the estimator with the computational complexity. We will evaluate the impact of this trade-off in our experiments.

3.1 Choice of $g(\boldsymbol{\phi})$

Theorem 1 of ARM defines $g(\boldsymbol{\phi}) = \sigma(\boldsymbol{\phi})$, where $\sigma(\cdot)$ is the sigmoid function. For the purpose of network sparsification, we find that this parametric function isn't very effective due to its slow transition between values 0 and 1. Thanks to the flexibility of ARM, we have a lot of freedom to design this parametric function $g(\boldsymbol{\phi})$. Apparently, it's straightforward to generalize Theorem 1 for any parametric

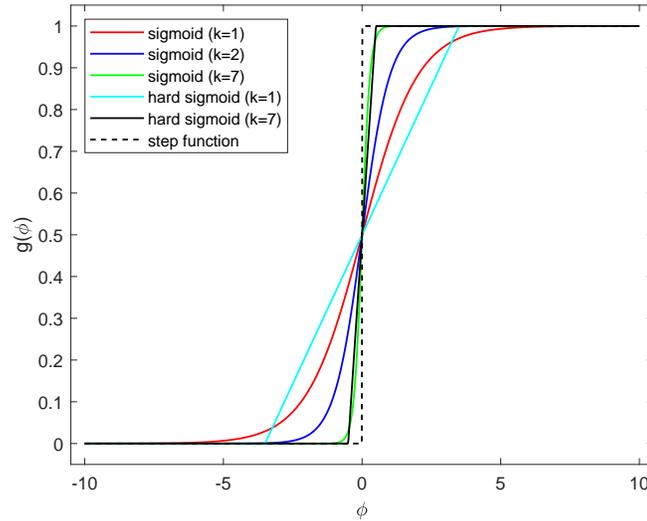


Fig. 1. The plots of $g(\phi)$ with different k for sigmoid and hard sigmoid functions. A large k tends to be more effective at sparsifying networks. Best viewed in color.

functions (smooth or non-smooth) as long as $g : \mathcal{R} \rightarrow [0, 1]$ and $g(-\phi) = 1 - g(\phi)$ ¹. Example parametric functions that work well in our experiments are the scaled sigmoid function

$$g_{\sigma_k}(\phi) = \sigma(k\phi) = \frac{1}{1 + \exp(-k\phi)}, \quad (11)$$

and the centered-scaled hard sigmoid

$$g_{\bar{\sigma}_k}(\phi) = \min(1, \max(0, \frac{k}{7}\phi + 0.5)), \quad (12)$$

where 7 is introduced such that $g_{\bar{\sigma}_1}(\phi) \approx g_{\sigma_1}(\phi) = \sigma(\phi)$. See Fig. 1 for some example plots of $g_{\sigma_k}(\phi)$ and $g_{\bar{\sigma}_k}(\phi)$ with different k . Empirically, we find that $k = 7$ works well for all of our experiments.

One important difference between the hard concrete estimator from Louizos et al. [20] and L_0 -ARM is that the hard concrete estimator has to rely on the hard sigmoid gate to zero out some parameters during training (a.k.a. conditional computation [1]), while L_0 -ARM achieves conditional computation naturally by sampling from the Bernoulli distribution, parameterized by $g(\phi)$, where $g(\phi)$ can be any parametric function (smooth or non-smooth) as shown in Fig. 1. We validate this in our experiments.

¹ The second condition is not necessary. But for simplicity, we will impose this condition to select parametric function $g(\phi)$ that is antithetic. Designing $g(\phi)$ without this constraint could be a potential area that is worthy of further investigation.

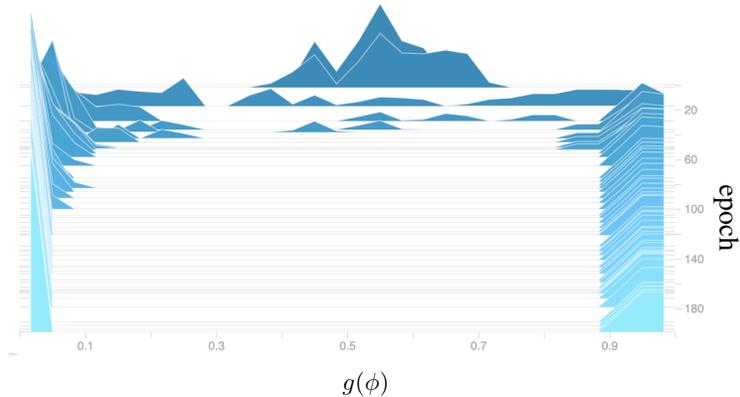


Fig. 2. Evolution of the histogram of $g(\phi)$ over training epochs. All $g(\phi)$ are initialized by random samples from a normal distribution $\mathcal{N}(0.5, 0.01)$, which are split into two spikes during training.

3.2 Sparsifying Network Architectures for Inference

After training, we get model parameters $\tilde{\theta}$ and ϕ . At test time, we can use the expectation of $\mathbf{z} \sim \text{Ber}(\mathbf{z}; g(\phi))$ as the mask $\hat{\mathbf{z}}$ for the final model parameters $\hat{\theta}$:

$$\hat{\mathbf{z}} = \mathbb{E}[\mathbf{z}] = g(\phi), \quad \hat{\theta} = \tilde{\theta} \odot \hat{\mathbf{z}}. \tag{13}$$

However, this will not yield a sparsified network for inference since none of the element of $\hat{\mathbf{z}} = g(\phi)$ is exact zero (unless the hard sigmoid gate $g_{\sigma_k}(\phi)$ is used). A simple approximation is to set the elements of $\hat{\mathbf{z}}$ to zero if the corresponding values in $g(\phi)$ are less than a threshold τ , i.e.,

$$\bar{z}_j = \begin{cases} 0, & g(\phi_j) \leq \tau \\ g(\phi_j), & \text{otherwise} \end{cases} \quad j = 1, 2, \dots, |\mathbf{z}| \tag{14}$$

We find that this approximation is very effective in all of our experiments as the histogram of $g(\phi)$ is widely split into two spikes around values of 0 and 1 after training because of the sharp transition of the scaled sigmoid (or hard sigmoid) function. See Fig. 2 for a typical plot of the histograms of $g(\phi)$ evolving during training process. We notice that our algorithm isn't very sensitive to τ , tuning which incurs negligible impacts to prune rates and model accuracies. Therefore, for all of our experiments we set $\tau = 0.5$ by default. Apparently, better designed τ is possible by considering the histogram of $g(\phi)$. However, we find this isn't very necessary for all of our experiments in the paper. Therefore, we will consider this histogram-dependent τ as our future improvement.

3.3 Imposing Shrinkage on Model Parameters θ

The L_0 regularized objective function (8) leads to sparse estimate of model parameters without imposing any shrinkage on the magnitude of θ . In some cases it might still be desirable to regularize the magnitude of model parameters with other norms, such as L_1 or L_2 (weight decay), to improve the robustness of model. This can be achieved conveniently by computing the expected L_1 or L_2 norm of θ under the same Bernoulli distribution: $\mathbf{z} \sim \text{Ber}(\mathbf{z}; g(\phi))$ as follows:

$$\mathbb{E}_{\mathbf{z} \sim \text{Ber}(\mathbf{z}; g(\phi))} [|\theta|_1] = \sum_{j=1}^{|\theta|} \mathbb{E}_{z_j \sim \text{Ber}(z_j; g(\phi_j))} [z_j |\tilde{\theta}_j|] = \sum_{j=1}^{|\theta|} g(\phi_j) |\tilde{\theta}_j|, \quad (15)$$

$$\mathbb{E}_{\mathbf{z} \sim \text{Ber}(\mathbf{z}; g(\phi))} [|\theta|_2^2] = \sum_{j=1}^{|\theta|} \mathbb{E}_{z_j \sim \text{Ber}(z_j; g(\phi_j))} [z_j^2 \tilde{\theta}_j^2] = \sum_{j=1}^{|\theta|} g(\phi_j) \tilde{\theta}_j^2, \quad (16)$$

which can be incorporated to Eq. 8 as additional regularization terms.

3.4 Group Sparsity Under L_0 and L_2 Norms

The formulation so far promotes a weight-level sparsity for network architectures. This sparsification strategy can compress model and reduce memory footprint of a network. However, it will usually not lead to effective speedups because weight-sparsified networks require sparse matrix multiplication and irregular memory access, which make it extremely challenging to effectively utilize the parallel computing resources of GPUs and CPUs. For the purpose of computational efficiency, it's usually preferable to perform group sparsity instead of weight-level sparsity. Similar to [26, 23, 20], we can achieve this by sharing a stochastic binary gate z among all the weights in a group. For example, a group can be all fan-out weights of a neuron in fully connected layers or all weights of a convolution filter. With this, the group regularized L_0 and L_2 norms can be conveniently expressed as

$$\mathbb{E}_{\mathbf{z} \sim \text{Ber}(\mathbf{z}; g(\phi))} [|\theta|_0] = \sum_{g=1}^{|G|} |g| g(\phi_g) \quad (17)$$

$$\mathbb{E}_{\mathbf{z} \sim \text{Ber}(\mathbf{z}; g(\phi))} [|\theta|_2^2] = \sum_{g=1}^{|G|} \left(g(\phi_g) \sum_{j=1}^{|g|} \tilde{\theta}_j^2 \right) \quad (18)$$

where $|G|$ denotes the number of groups and $|g|$ denotes the number of weights of group g . For the reason of computational efficiency, we perform this group sparsity in all of our experiments.

4 Related Work

It is well-known that DNNs are extremely compute and memory intensive. Recently, there has been an increasing interest to network sparsification [8, 7,

26, 18, 19, 22, 23, 20] as the applications of DNNs to practical real-time systems, such as the IoT devices, call for compute and memory efficient networks. One of the earliest sparsification methods is to prune the redundant weights based on the magnitudes [17], which is proved to be effective in modern CNN [8]. Although weight sparsification is able to compress networks, it can barely improve computational efficiency due to unstructured sparsity [26]. Therefore, magnitude-based group sparsity is proposed [26, 18], which can compress networks while reducing computation cost significantly. These magnitude-based methods usually proceed in three stages: pre-train a full network, prune the redundant weights or filters, and fine-tune the pruned model. As a comparison, our method L_0 -ARM trains a sparsified network from scratch without pre-training and fine-tuning, and therefore is more preferable.

Another category of sparsification methods is based on Bayesian statistics and information theory [22, 23, 19]. For example, inspired by variational dropout [13], Molchanov et al. propose a method that unbinds the dropout rate, and also leads to sparsified networks [22].

Recently, Louizos et al. [20] propose to sparsify networks with L_0 -norm. Since the L_0 regularization explicitly penalizes number of non-zero parameters, this method is conceptually very appealing. However, the non-differentiability of L_0 norm prevents an effective gradient-based optimization. Therefore, Louizos et al. [20] propose a hard concrete gradient estimator for this optimization problem. Our work is built on top of their L_0 formulation. However, instead of using a hard concrete estimator, we investigate the Augment-Reinforce-Merge (ARM) [28], a recently proposed unbiased estimator, to this binary optimization problem.

5 Experimental Results

We evaluate the performance of L_0 -ARM and L_0 -AR on multiple public datasets and multiple network architectures. Specifically, we evaluate MLP 500-300 [16] and LeNet 5-Caffe ² on the MNIST dataset [15], and Wide Residual Networks [29] on the CIFAR-10 and CIFAR-100 datasets [14]. For baselines, we refer to the following state-of-the-art sparsification algorithms: Sparse Variational Dropout (Sparse VD) [22], Bayesian Compression with group normal-Jeffreys (BC-GNJ) and group horseshoe (BC-GHS) [19], and L_0 -norm regularization with hard concrete estimator (L_0 -HC) [20]. For a fair comparison, we closely follow the experimental setups of L_0 -HC ³.

5.1 Implementation Details

We incorporate L_0 -ARM and L_0 -AR into the architectures of MLP, LeNet-5 and Wide ResNet. As we described in Sec. 3.4, instead of sparsifying weights, we apply group sparsity on neurons in fully-connected layers or on convolution filters

² <https://github.com/BVLC/caffe/tree/master/examples/mnist>

³ https://github.com/AMLab-Amsterdam/L0_regularization

in convolutional layers. Once a neuron or filter is pruned, all related weights are removed from the networks.

The Multi-Layer Perceptron (MLP) [16] has two hidden layers of size 300 and 100, respectively. We initialize $g(\phi) = \pi$ by random samples from a normal distribution $\mathcal{N}(0.8, 0.01)$ for the input layer and $\mathcal{N}(0.5, 0.01)$ for the hidden layers, which activate around 80% of neurons in input layer and around 50% of neurons in hidden layers. LeNet-5-Caffe consists of two convolutional layers of 20 and 50 filters interspersed with max pooling layers, followed by two fully-connected layers with 500 and 10 neurons. We initialize $g(\phi) = \pi$ for all neurons and filters by random samples from a normal distribution $\mathcal{N}(0.5, 0.01)$. Wide-ResNets (WRNs) [29] have shown state-of-the-art performance on many image classification benchmarks. Following [20], we only apply L_0 regularization on the first convolutional layer of each residual block, which allows us to incorporate L_0 regularization without further modifying residual block architecture. The architectural details of WRN are listed in Table 1. For initialization, we activate around 70% of convolutional filters.

Table 1. Architectural details of WRN incorporated with L_0 -ARM. The number in parenthesis is the size of activation map of each layer. For brevity, only the modified layers are included.

Group name	Layers
conv1	[Original Conv (16)]
conv2	[L_0 ARM (160); Original Conv (160)] \times 4
conv3	[L_0 ARM (320); Original Conv (320)] \times 4
conv4	[L_0 ARM (640); Original Conv (640)] \times 4

For MLP and LeNet-5, we train with a mini-batch of 100 data samples and use Adam [12] as optimizer with initial learning rate of 0.001, which is halved every 100 epochs. For Wide-ResNet, we train with a mini-batch of 128 data samples and use Nesterov Momentum as optimizer with initial learning rate of 0.1, which is decayed by 0.2 at epoch 60 and 120. Each of these experiments run for 200 epochs in total. For a fair comparison, these experimental setups closely follow what were described in L_0 -HC [20] and their open-source implementation ³.

5.2 MNIST Experiments

We run both MLP and LeNet-5 on the MNIST dataset. By tuning the regularization strength λ , we can control the trade off between sparsity and accuracy. We can use one λ for all layers or a separate λ for each layer to fine-tune the sparsity preference. In our experiments, we set $\lambda = 0.1/N$ or $\lambda = (0.1, 0.3, 0.4)/N$ for MLP, and set $\lambda = 0.1/N$ or $\lambda = (10, 0.5, 0.1, 10)/N$ for LeNet-5, where N denotes to the number of training datapoints.

We use three metrics to evaluate the performance of an algorithm: prediction accuracy, prune rate, and expected number of floating point operations (FLOPs). Prune rate is defined as the ratio of number of pruned weights to number of all weights. Prune rate manifests the memory saving of a sparsified network, while expected FLOPs demonstrates the training / inference cost of a sparsification algorithm.

Table 2. Performance comparison on MNIST. Each experiment was run five times and the median (in terms of accuracy) is reported. All the baseline results are taken from the corresponding papers.

Network	Method	Pruned Architecture	Prune rate (%)	Accuracy (%)
MLP 784-300-100	Sparse VD	219-214-100	74.72	98.2
	BC-GNJ	278-98-13	89.24	98.2
	BC-GHS	311-86-14	89.45	98.2
	L_0 -HC ($\lambda = 0.1/N$)	219-214-100	73.98	98.6
	L_0 -HC (λ sep.)	266-88-33	89.99	98.2
	L_0 -AR ($\lambda = 0.1/N$)	453-150-68	70.39	98.3
	L_0 -ARM ($\lambda = 0.1/N$)	143-153-78	87.00	98.3
	L_0 -ARM (λ sep.)	464-114-65	77.10	98.2
	L_0 -ARM (λ sep.)	159-74-73	92.96	98.1
LeNet-5-Caffe 20-50-800-500	Sparse VD	14-19-242-131	90.7	99.0
	GL	3-12-192-500	76.3	99.0
	GD	7-13-208-16	98.62	99.0
	SBP	3-18-284-283	80.34	99.0
	BC-GNJ	8-13-88-13	99.05	99.0
	BC-GHS	5-10-76-16	99.36	99.0
	L_0 -HC ($\lambda = 0.1/N$)	20-25-45-462	91.1	99.1
	L_0 -HC (λ sep.)	9-18-65-25	98.6	99.0
	L_0 -AR ($\lambda = 0.1/N$)	18-28-46-249	93.73	98.8
	L_0 -ARM ($\lambda = 0.1/N$)	20-16-32-257	95.52	99.1
	L_0 -AR (λ sep.)	5-12-131-22	98.90	98.4
	L_0 -ARM (λ sep.)	6-10-39-11	99.49	98.7

We compare L_0 -ARM and L_0 -AR to five state-of-the-art sparsification algorithms on MNIST, with the results shown in Table 2. For the comparison between L_0 -HC and L_0 -AR(M) when $\lambda = 0.1/N$, we use the exact same hyper-parameters for both algorithms (the fairest comparison). In this case, L_0 -ARM achieve the same accuracy (99.1%) on LeNet-5 with even sparser pruned architectures (95.52% vs. 91.1%). When separated λ s are considered (λ sep.), since L_0 -HC doesn't disclose the specific λ s for the last two fully-connected layers, we tune them by ourselves and find that $\lambda = (10, 0.5, 0.1, 10)/N$ yields the best performance. In this case, L_0 -ARM achieves the highest prune rate (99.49% vs. 98.6%) with very similar accuracies (98.7% vs. 99.1%) on LeNet-5. Similar patterns are also observed on MLP. Regarding L_0 -AR, although its performance is not as good

as L_0 -ARM, it’s still very competitive to all the other methods. The advantage of L_0 -AR over L_0 -ARM is its lower computational complexity during training. As we discussed in Sec. 3, L_0 -ARM needs one extra forward pass to estimate the gradient w.r.t. ϕ ; for large DNN architectures, this extra cost can be significant.

To evaluate the training cost and network sparsity of different algorithms, we compare the prune rates of L_0 -HC and L_0 -AR(M) on LeNet-5 as a function of epoch in Fig. 3 (a, b). Similarly, we compare the expected FLOPs of different algorithms as a function of epoch in Fig. 3 (c, d). As we can see from (a, b), L_0 -ARM yields much sparser network architectures over the whole training epochs, followed by L_0 -AR and L_0 -HC. The FLOPs vs. Epoch plots in (c, d) are more complicated. Because L_0 -HC and L_0 -AR only need one forward pass to compute gradient, they have the same expected FLOPs for training and inference. L_0 -ARM needs two forward passes for training. Therefore, L_0 -ARM is computationally more expensive during training (red curves), but it leads to sparser / more efficient architectures for inference (green curves), which pays off its extra cost in training.

5.3 CIFAR Experiments

We further evaluate the performance of L_0 -ARM and L_0 -AR with Wide-ResNet [29] on CIFAR-10 and CIFAR-100. Following [20], we only apply L_0 regularization on the first convolutional layer of each residual block, which allows us to incorporate L_0 regularization without further modifying residual block architecture.

Table 3 shows the performance comparison between L_0 -AR(M) and three baseline methods. We find that L_0 -HC cannot sparsify the Wide-ResNet architecture (prune rate 0%)⁴, while L_0 -ARM and L_0 -AR prune around 50% of the parameters of the impacted subnet. As we activate 70% convolution filters in initialization, the around 50% prune rate is not due to initialization. We also inspect the histograms of $g(\phi)$: As expected, they are all split into two spikes around the values of 0 and 1, similar to the histograms shown in Fig. 2. In terms of accuracies, both L_0 -ARM and L_0 -AR achieve very similar accuracies as the baseline methods.

To evaluate the training and inference costs of different algorithms, we compare the expected FLOPs of L_0 -HC and L_0 -AR(M) on CIFAR-10 and CIFAR-100 as a function of iteration in Fig. 4. Similar to Fig. 3, L_0 -ARM is more computationally expensive for training, but leads to sparser / more efficient architectures for inference, which pays off its extra cost in training. It’s worth to emphasize that for these experiments L_0 -AR has the lowest training FLOPs and inference FLOPs (since only one forward pass is needed for training and inference), while achieving very similar accuracies as the baseline methods (Table 3).

Finally, we compare the test accuracies of different algorithms as a function of epoch on CIFAR-10, with the results shown in Fig. 5. We apply the exact same hyper-parameters of L_0 -HC to L_0 -AR(M). As L_0 -AR(M) prunes around

⁴ This was also reported recently in the appendix of [5], and can be easily reproduced by using the open-source implementation of L_0 -HC³.

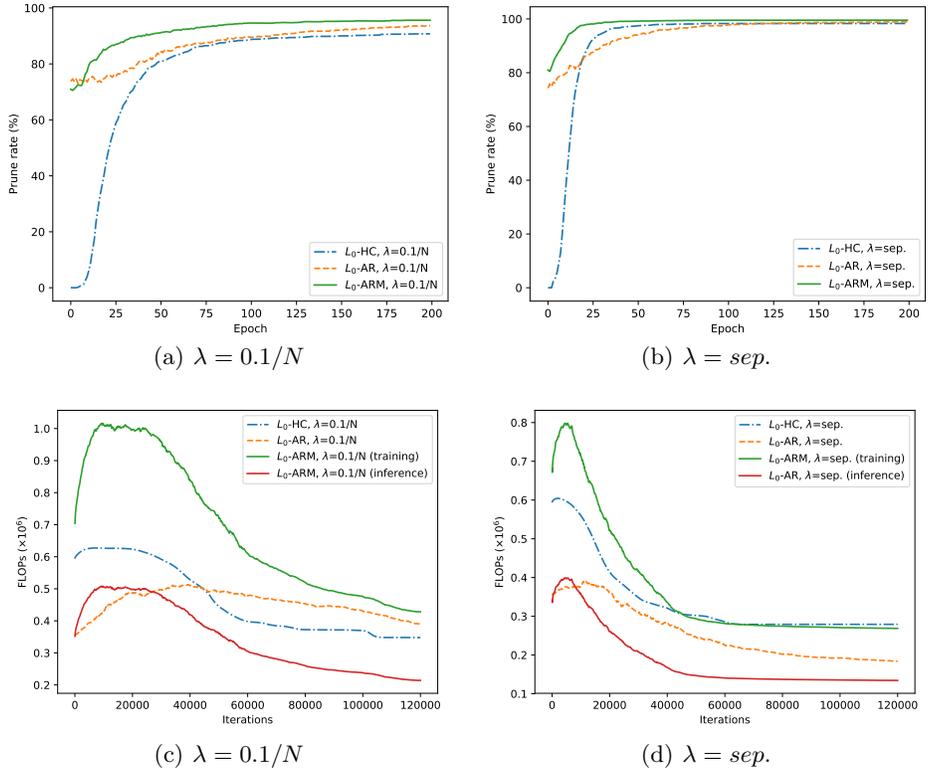


Fig. 3. (a, b) Comparison of prune rate of sparsified network as a function of epoch for different algorithms. (c, d) Comparison of expected FLOPs as a function of epoch for different algorithms during training and inference. The results are on LeNet-5 with L_0 -HC and L_0 -AR(M). Because L_0 -HC and L_0 -AR only need one forward pass to compute gradient, they have the same expected FLOPs for training and inference. L_0 -ARM needs two forward passes for training. Therefore, L_0 -ARM is computationally more expensive during training (red curves), but it leads to sparser / more efficient architectures for inference (blue curves), which pays off its extra cost in training.

50% parameters during training (while L_0 -HC prunes 0%), the test accuracies of the former are lower than the latter before convergence, but all the algorithms yield very similar accuracies after convergence, demonstrating the effectiveness of L_0 -AR(M).

6 Conclusion

We propose L_0 -ARM, an unbiased and low-variance gradient estimator, to sparsify network architectures. Compared to L_0 -HC [20] and other state-of-the-art

Table 3. Performance comparison of WRN on CIFAR-10 and CIFAR-100. Each experiment was run five times and the median (in terms of accuracy) is reported. All the baseline results are taken from the corresponding papers. Only the architectures of pruned layers are shown.

Network	Method	Pruned Architecture	Prune rate (%)	Accuracy (%)	
WRN-28-10 CIFAR-10	Original WRN [29]	full model	0	96.00	
	Original WRN-dropout [29]	full model	0	96.11	
	L_0 -HC ($\lambda = 0.001/N$) [20]	full model	0	96.17	
	L_0 -HC ($\lambda = 0.002/N$) [20]	full model	0	96.07	
		83-77-83-88-			
	L_0 AR ($\lambda = 0.001/N$)	169-167-153-165- 324-323-314-329	49.49	95.58	
		74-86-83-83-			
	L_0 ARM ($\lambda = 0.001/N$)	164-145-167-153- 333-333-310-330	49.46	95.68	
		82-75-82-87-			
	L_0 AR ($\lambda = 0.002/N$)	164-169-156-161- 317-317-317-324	49.95	95.60	
		75-72-78-78-			
	L_0 ARM ($\lambda = 0.002/N$)	157-165-131-162- 336-325-331-343	49.63	95.70	
WRN-28-10 CIFAR-100	Original WRN [29]	full model	0	78.82	
	Original WRN-dropout [29]	full model	0	81.15	
	L_0 -HC ($\lambda = 0.001/N$) [20]	full model	0	81.25	
	L_0 -HC ($\lambda = 0.002/N$) [20]	full model	0	80.96	
			78-78-79-85-		
		L_0 -AR ($\lambda = 0.001/N$)	168-168-162-164- 308-326-319-330	49.37	80.50
			75-83-80-58-		
		L_0 -ARM ($\lambda = 0.001/N$)	172-156-160-165- 324-311-313-318	50.51	80.74
		75-76-72-80-			
	L_0 -AR ($\lambda = 0.002/N$)	158-158-137-168- 318-295-327-324	50.93	80.09	
		81-74-77-73-			
	L_0 -ARM ($\lambda = 0.002/N$)	149-157-156-152- 299-332-305-325	50.78	80.56	

sparsification algorithms, L_0 -ARM demonstrates superior performance of sparsifying network architectures while retaining almost the same accuracies of the baseline methods. Extensive experiments on multiple public datasets and multiple network architectures validate the effectiveness of L_0 -ARM. Overall, L_0 -ARM yields the sparsest architectures and the lowest inference FLOPs for all the networks considered with very similar accuracies as the baseline methods.

As for future extensions, we plan to design better (possibly non-antithetic) parametric function $g(\phi)$ to improve the sparsity of solutions. We also plan to

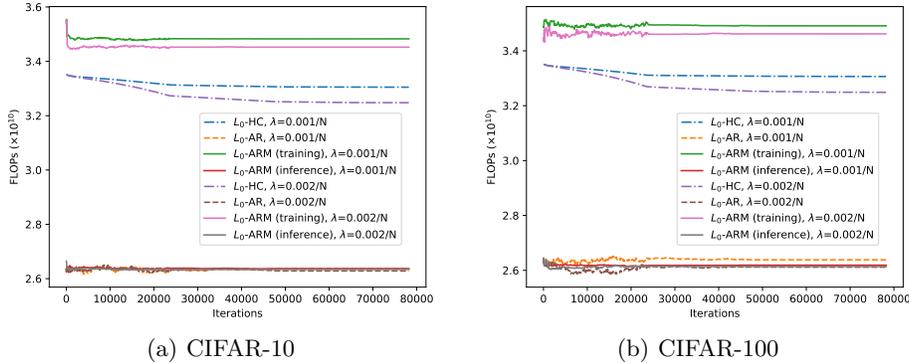


Fig. 4. Comparison of expected FLOPs as a function of iteration during training and inference. Similar to Fig. 3, L_0 -ARM is more computationally expensive for training, but leads to sparser / more efficient architectures for inference. For these experiments, L_0 -AR has the lowest training FLOPs and inference FLOPs, while achieving very similar accuracies as the baseline methods (Table 3).

investigate more efficient algorithm to evaluate L_0 -ARM gradient (9) by utilizing the antithetic structure of two forward passes.

References

1. Bengio, Y., Leonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv preprint arXiv:1308.3432 (2013)
2. Bird, T., Kunze, J., Barber, D.: Stochastic variational optimization. arXiv preprint arXiv:1809.04855 (2018)
3. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2009)
4. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
5. Gale, T., Elsen, E., Hooker, S.: The state of sparsity in deep neural networks. arXiv preprint arXiv:1902.09574 (2019)
6. Grathwohl, W., Choi, D., Wu, Y., Roeder, G., Duvenaud, D.: Backpropagation through the void: Optimizing control variates for black-box gradient estimation. In: International Conference on Learning Representations (ICLR) (2018)
7. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In: International Conference on Learning Representations (ICLR) (2016)
8. Han, S., Pool, J., Tran, J., Dally, W.: Learning both weights and connections for efficient neural network. In: Advances in neural information processing systems. pp. 1135–1143 (2015)

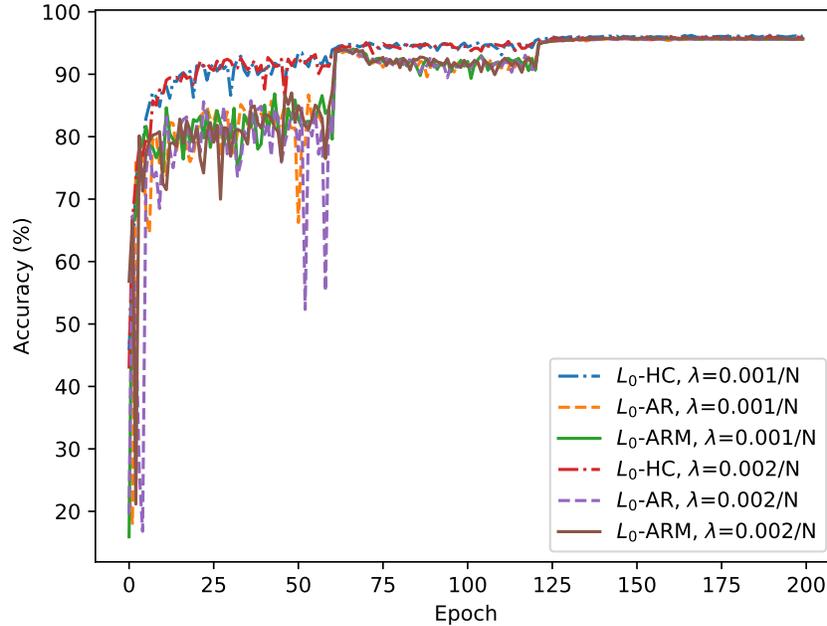


Fig. 5. Comparison of test accuracy as a function of epoch for different algorithms on CIFAR-10. We apply the exact same hyper-parameters of L_0 -HC to L_0 -AR(M), which yield similar accuracies for converged models even though the latter prunes around 50% parameters while the former prunes 0%.

9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
10. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
11. Jang, E., Gu, S., Poole, B.: Categorical reparameterization with gumbel-softmax. In: International Conference on Learning Representations (ICLR) (2017)
12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: International Conference on Learning Representations (ICLR) (2015)
13. Kingma, D.P., Salimans, T., Welling, M.: Variational dropout and the local reparameterization trick. In: Advances in Neural Information Processing Systems. pp. 2575–2583 (2015)
14. Krizhevsky, A.: Learning multiple layers of features from tiny images. Tech. rep. (2009)
15. Lecun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. In: Proceedings of the IEEE. pp. 2278–2324 (1998)
16. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., et al.: Gradient-based learning applied to document recognition. Proceedings of the IEEE **86**(11), 2278–2324 (1998)
17. LeCun, Y., Denker, J.S., Solla, S.A.: Optimal brain damage. In: Advances in neural information processing systems. pp. 598–605 (1990)

18. Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. arXiv preprint arXiv:1608.08710 (2016)
19. Louizos, C., Ullrich, K., Welling, M.: Bayesian compression for deep learning. In: Advances in Neural Information Processing Systems. pp. 3288–3298 (2017)
20. Louizos, C., Welling, M., Kingma, D.P.: Learning sparse neural networks through l_0 regularization. In: International Conference on Learning Representations (ICLR) (2018)
21. Maddison, C.J., Mnih, A., Teh, Y.W.: The concrete distribution: A continuous relaxation of discrete random variables. In: International Conference on Learning Representations (ICLR) (2017)
22. Molchanov, D., Ashukha, A., Vetrov, D.: Variational dropout sparsifies deep neural networks. In: Proceedings of the 34th International Conference on Machine Learning-Volume 70. pp. 2498–2507. JMLR. org (2017)
23. Neklyudov, K., Molchanov, D., Ashukha, A., Vetrov, D.: Structured bayesian pruning via log-normal multiplicative noise. In: Advances in Neural Information Processing Systems (NIPS) (2017)
24. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D.: Mastering the game of go with deep neural networks and tree search. *Nature* **529**, 484–503 (2016)
25. Tucker, G., Mnih, A., Maddison, C.J., Lawson, J., Sohl-Dickstein, J.: Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models. In: Advances in Neural Information Processing Systems (NIPS) (2017)
26. Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. In: Advances in Neural Information Processing Systems (NIPS) (2016)
27. Williams, R.J.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* **8**(3-4), 229–256 (May 1992)
28. Yin, M., Zhou, M.: Arm: Augment-REINFORCE-merge gradient for stochastic binary networks. In: International Conference on Learning Representations (ICLR) (2019)
29. Zagoruyko, S., Komodakis, N.: Wide residual networks. In: The British Machine Vision Conference (BMVC) (2016)