

---

# FLOPs as a Direct Optimization Objective for Learning Sparse Neural Networks

---

Raphael Tang, Ashutosh Adhikari, Jimmy Lin

David R. Cheriton School of Computer Science

University of Waterloo

{r33tang, ashutosh.adhikari, jimmylin}@uwaterloo.ca

## Abstract

There exists a plethora of techniques for inducing structured sparsity in parametric models during the optimization process, with the final goal of resource-efficient inference. However, to the best of our knowledge, none target a specific number of floating-point operations (FLOPs) as part of a single end-to-end optimization objective, despite reporting FLOPs as part of the results. Furthermore, a one-size-fits-all approach ignores realistic system constraints, which differ significantly between, say, a GPU and a mobile phone—FLOPs on the former incur less latency than on the latter; thus, it is important for practitioners to be able to specify a target number of FLOPs during model compression. In this work, we extend a state-of-the-art technique to directly incorporate FLOPs as part of the optimization objective and show that, given a desired FLOPs requirement, different neural networks can be successfully trained for image classification.

## 1 Introduction

Neural networks are a class of parametric models that achieve the state of the art across a broad range of tasks, but their heavy computational requirements hinder practical deployment on resource-constrained devices, such as mobile phones, Internet-of-things (IoT) devices, and offline embedded systems. Many recent works focus on alleviating these computational burdens, mainly falling under two non-mutually exclusive categories: manually designing resource-efficient models, and automatically compressing popular architectures. In the latter, increasingly sophisticated techniques have emerged [4, 5, 6], which have achieved respectable accuracy–efficiency operating points, some even Pareto-better than that of the original network; for example, network slimming [4] reaches an error rate of 6.20% on CIFAR-10 using VGGNet [10] with a 51% FLOPs reduction—an error decrease of 0.14% over the original.

However, to the best of our knowledge, none of the methods impose a FLOPs constraint as part of a single end-to-end optimization objective. MorphNets [1] apply an  $L_1$  norm, shrinkage-based relaxation of a FLOPs objective, but for the purpose of searching and training multiple models to find good network architectures; in this work, we learn a sparse neural network in a single training run. Other papers directly target device-specific metrics, such as energy usage [16], but the pruning procedure does not explicitly include the metrics of interest as part of the optimization objective, instead using them as heuristics. Falling short of continuously deploying a model candidate and measuring actual inference time, as in time-consuming neural architectural search [12], we believe that the number of FLOPs is reasonable to use as a proxy measure for actual latency and energy usage; across variants of the same architecture, Tang et al. suggest that the number of FLOPs is a stronger predictor of energy usage and latency than the number of parameters [13].

Indeed, there are compelling reasons to optimize for the number of FLOPs as part of the training objective: First, it would permit FLOPs-guided compression in a more principled manner. Second,

practitioners can directly specify a desired target of FLOPs, which is important in deployment. Thus, our main contribution is to present a novel extension of the prior state of the art [7] to incorporate the number of FLOPs as part of the optimization objective, furthermore allowing practitioners to set and meet a desired compression target.

## 2 FLOPs Objective

Formally, we define the FLOPs objective  $L_{flops} : f \times \mathbb{R}^m \mapsto \mathbb{N}_0$  as follows:

$$L_{flops}(h, \boldsymbol{\theta}) := g(h(\cdot; \mathbb{I}(\boldsymbol{\theta}_1 \neq 0), \mathbb{I}(\boldsymbol{\theta}_2 \neq 0), \dots, \mathbb{I}(\boldsymbol{\theta}_m \neq 0))) \quad |\boldsymbol{\theta}| = m \quad (1)$$

where  $L_{flops}$  is the FLOPs associated with hypothesis  $h(\cdot; \boldsymbol{\theta}) := p(\cdot | \boldsymbol{\theta})$ ,  $g(\cdot)$  is a function with the explicit dependencies, and  $\mathbb{I}$  is the indicator function. We assume  $L_{flops}$  to depend only on whether parameters are non-zero, such as the number of neurons in a neural network. For a dataset  $\mathcal{D}$ , our empirical risk thus becomes

$$\mathcal{R}(h; \boldsymbol{\theta}) = -\log p(\mathcal{D} | \boldsymbol{\theta}) + \lambda_f \max(0, L_{flops}(h, \boldsymbol{\theta}) - T) \quad \mathcal{D} = ((x_1, y_1), \dots, (x_n, y_n)) \quad (2)$$

Hyperparameters  $\lambda_f \in \mathbb{R}_0^+$  and  $T \in \mathbb{N}_0$  control the strength of the FLOPs objective and the target, respectively. The second term is a black-box function, whose combinatorial nature prevents gradient-based optimization; thus, using the same procedure in prior art [7], we relax the objective to a surrogate of the evidence lower bound with a fully-factorized spike-and-slab posterior as the variational distribution, where the addition of the clipped FLOPs objective can be interpreted as a sparsity-inducing prior  $p(\boldsymbol{\theta}) \propto \exp(-\lambda_f \max(0, L_{flops}(h, \boldsymbol{\theta}) - T))$ . Let  $\mathbf{z} \sim p(\mathbf{z} | \boldsymbol{\pi})$  be Bernoulli random variables parameterized by  $\boldsymbol{\pi}$ :

$$\mathcal{L}(h; \boldsymbol{\theta}) = \mathbb{E}_{p(\mathbf{z} | \boldsymbol{\pi})} [-\log p(\mathcal{D} | \boldsymbol{\theta} \odot \mathbf{z}) + \lambda_f \max(0, L_{flops}(h, \boldsymbol{\theta} \odot \mathbf{z}) - T)] \quad (3)$$

where  $\odot$  denotes the Hadamard product. To allow for efficient reparameterization and exact zeros, Louizos et al. [7] propose to use a hard concrete distribution as the approximation, which is a stretched and clipped version of the binary Concrete distribution [8]: if  $\hat{z} \sim \text{BinaryConcrete}(\alpha, \beta)$ , then  $\tilde{z} := \max(0, \min(1, (\zeta - \gamma)\hat{z} + \gamma))$  is said to be a hard concrete r.v., given  $\zeta > 1$  and  $\gamma < 0$ . Define  $\boldsymbol{\phi} := (\boldsymbol{\alpha}, \beta)$ , and let  $\psi(\boldsymbol{\phi}) = \text{Sigmoid}(\log \boldsymbol{\alpha} - \beta \log \frac{-\gamma}{\zeta})$  and  $\mathbf{z} \sim \text{Bernoulli}(\psi(\boldsymbol{\phi}))$ . Then, the approximation becomes

$$\mathcal{L}(h; \boldsymbol{\theta}) \approx \mathbb{E}_{p(\tilde{\mathbf{z}} | \boldsymbol{\phi})} [-\log p(\mathcal{D} | \boldsymbol{\theta} \odot \tilde{\mathbf{z}})] + \lambda_f \mathbb{E}_{p(\mathbf{z} | \psi(\boldsymbol{\phi}))} [\max(0, L_{flops}(h, \boldsymbol{\theta} \odot \mathbf{z}) - T)] \quad (4)$$

$\psi(\cdot)$  is the probability of a gate being non-zero under the hard concrete distribution. It is more efficient in the second expectation to sample from the equivalent Bernoulli parameterization compared to hard concrete, which is more computationally expensive to sample multiple times. The first term now allows for efficient optimization via the reparameterization trick [3]; for the second, we apply the score function estimator (REINFORCE) [15], since the FLOPs objective is, in general, non-differentiable and thus precludes the reparameterization trick. High variance is a non-issue because the number of FLOPs is fast to compute, hence letting many samples to be drawn. At inference time, the deterministic estimator is  $\hat{\boldsymbol{\theta}} := \boldsymbol{\theta} \odot \max(0, \min(1, \text{Sigmoid}(\log \boldsymbol{\alpha})(\zeta - \gamma) + \gamma))$  for the final parameters  $\hat{\boldsymbol{\theta}}$ .

### 2.1 Computing number of FLOPs under group sparsity

In practice, computational savings are achieved only if the model is sparse across “regular” groups of parameters, e.g., each filter in a convolutional layer. Thus, each computational group uses one hard concrete r.v. [7]—in fully-connected layers, one per input neuron; in 2D convolution layers, one per output filter. Under convention in the literature where one addition and one multiplication each count as a FLOP, the FLOPs for a 2D convolution layer  $h_{conv}(\cdot; \boldsymbol{\theta})$  given a random draw  $\mathbf{z}$  is then defined as  $L_{flops}(h_{conv}, \mathbf{z}) = (K_w K_h C_{in} + 1)(I_w - K_w + P_w + 1)(I_h - K_h + P_h + 1) \|\mathbf{z}\|_0$  for kernel width and height  $(K_w, K_h)$ , input width and height  $(I_w, I_h)$ , padding width and height  $(P_w, P_h)$ , and number of input channels  $C_{in}$ . The number of FLOPs for a fully-connected layer  $h_{fc}(\cdot; \boldsymbol{\theta})$  is  $L_{flops}(h_{fc}, \mathbf{z}) = (I_n + 1) \|\mathbf{z}\|_0$ , where  $I_n$  is the number of input neurons. Note that these are conventional definitions in neural network compression papers—the objective can easily use instead a number of FLOPs incurred by other device-specific algorithms. Thus, at each training step, we compute the FLOPs objective by sampling from the Bernoulli r.v.’s and using the aforementioned definitions, e.g.,  $L_{flops}(h_{conv}, \cdot)$  for convolution layers. Then, we apply the score function estimator to the FLOPs objective as a black-box estimator.

### 3 Experimental Results

We report results on MNIST, CIFAR-10, and CIFAR-100, training multiple models on each dataset corresponding to different FLOPs targets. We follow the same initialization and hyperparameters as Louizos et al. [7], using Adam [2] with temporal averaging for optimization, a weight decay of  $5 \times 10^{-4}$ , and an initial  $\alpha$  that corresponds to the original dropout rate of that layer. We similarly choose  $\beta = 2/3$ ,  $\gamma = -0.1$ , and  $\zeta = 1.1$ . For brevity, we direct the interested reader to their repository<sup>1</sup> for specifics. In all of our experiments, we replace the original  $L_0$  penalty with our FLOPs objective, and we train all models to 200 epochs; at epoch 190, we prune the network by weights associated with zeroed gates and replace the r.v.’s with their deterministic estimators, then finetune for 10 more epochs. For the score function estimator, we draw 1000 samples at each optimization step—this procedure is fast and has no visible effect on training time.

Table 1: Comparison of LeNet-5-Caffe results on MNIST

Model	Architecture	Err.	FLOPs
GL [14]	3-12-192-500	1.0%	205K
GD [11]	7-13-208-16	1.1%	254K
SBP [9]	3-18-284-283	0.9%	217K
BC-GNJ [6]	8-13-88-13	1.0%	290K
BC-GHS [6]	5-10-76-16	1.0%	158K
$L_0$ [7]	20-25-45-462	0.9%	1.3M
$L_0$ -sep [7]	9-18-65-25	1.0%	403K
$L_{flops}, T = 400K$	3-13-208-500	0.9%	218K
$L_{flops}, T = 200K$	3-8-128-499	1.0%	153K
$L_{flops}, T = 100K$	2-7-112-478	1.1%	111K

We choose  $\lambda_f = 10^{-6}$  in all of the experiments for LeNet-5-Caffe, the Caffe variant of LeNet-5.<sup>1</sup> We observe that our methods (Table 1, bottom three rows) achieve accuracy comparable to those from previous approaches while using fewer FLOPs, with the added benefit of providing a tunable “knob” for adjusting the FLOPs. Note that the convolution layers are the most aggressively compressed, since they are responsible for most of the FLOPs in this model.

Table 2: Comparison of WideResNet-28-10 results on CIFAR-10 and CIFAR-100

Method	CIFAR-10			CIFAR-100		
	Err.	$\mathbb{E}[\text{FLOPs}]$	FLOPs	Err.	$\mathbb{E}[\text{FLOPs}]$	FLOPs
Orig.	4.00%	5.9B	5.9B	21.18%	5.9B	5.9B
Orig. w/dropout	3.89%	5.9B	5.9B	18.85%	5.9B	5.9B
$L_0$	3.83%	5.3B	5.9B	<b>18.75%</b>	5.3B	5.9B
$L_0$ -small	3.93%	5.2B	5.9B	19.04%	5.2B	5.9B
$L_{flops}, T = 4B$	<b>3.82%</b>	3.9B	4.6B	18.93%	3.9B	4.6B
$L_{flops}, T = 2.5B$	3.91%	2.4B	2.4B	19.48%	2.4B	2.4B

Orig. in Table 2 denotes the original WRN-28-10 model [17], and  $L_0$ -\* refers to the  $L_0$ -regularized models [7]; likewise, we augment CIFAR-10 and CIFAR-100 with standard random cropping and horizontal flipping. For each of our results (last two rows), we report the median error rate of five different runs, executing a total of 20 runs across two models for each of the two datasets; we use  $\lambda_f = 3 \times 10^{-9}$  in all of these experiments. We also report both the expected FLOPs and actual FLOPs, the former denoting the number of FLOPs, on average, at training time under stochastic gates and the latter denoting the number of FLOPs at inference time. We restrict the FLOPs calculations to the penalized non-residual convolution layers only. For CIFAR-10, our approaches result in Pareto-better models with decreases in both error rate and the actual number of inference-time FLOPs. For CIFAR-100, we do not achieve a Pareto-better model, since our approach trades accuracy for improved efficiency. The acceptability of the tradeoff depends on the end application.

<sup>1</sup>[https://github.com/AMLab-Amsterdam/L0\\_regularization](https://github.com/AMLab-Amsterdam/L0_regularization)

## References

- [1] Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Hao Wu, Tien-Ju Yang, and Edward Choi. MorphNet: Fast & simple resource-constrained structure learning of deep networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [2] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [3] Diederik P. Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv:1312.6114*, 2013.
- [4] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *International Conference on Learning Representations*, 2017.
- [5] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2755–2763, 2017.
- [6] Christos Louizos, Karen Ullrich, and Max Welling. Bayesian compression for deep learning. In *Advances in Neural Information Processing Systems*, pages 3288–3298, 2017.
- [7] Christos Louizos, Max Welling, and Diederik P. Kingma. Learning sparse neural networks through  $l_0$  regularization. In *International Conference on Learning Representations*, 2018.
- [8] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *International Conference on Learning Representations*, 2017.
- [9] 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*, pages 6775–6784, 2017.
- [10] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- [11] Suraj Srinivas and R Venkatesh Babu. Generalized dropout. *arXiv:1611.06791*, 2016.
- [12] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, and Quoc V Le. MnasNet: Platform-aware neural architecture search for mobile. *arXiv:1807.11626*, 2018.
- [13] Raphael Tang, Weijie Wang, Zhucheng Tu, and Jimmy Lin. An experimental analysis of the power consumption of convolutional neural networks for keyword spotting. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5479–5483, 2018.
- [14] 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*, pages 2074–2082, 2016.
- [15] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [16] Tien-Ju Yang, Yu-Hsin Chen, and Vivienne Sze. Designing energy-efficient convolutional neural networks using energy-aware pruning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6071–6079, 2017.
- [17] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv:1605.07146*, 2016.