
[Supplementary]

Reproducibility Report

Rigging the Lottery: Making All Tickets Winners

Anonymous Author(s)

Affiliation

Address

email

1 Architecture Specific Details—ResNet-50 on CIFAR100

Table 1: **ResNet-50 architecture used on CIFAR100**. Building blocks are shown in brackets, with the numbers of blocks stacked. Downsampling is performed by conv3_1, conv4_1, and conv5_1 with a stride of 2.

Layer Name	Output Size	ResNet-50
conv1	32×32	3×3 , 64, no stride
conv2_x	32×32	$\begin{bmatrix} 1 \times 1, 64 \\ 3 \times 3, 64 \\ 1 \times 1, 256 \end{bmatrix} \times 3$
conv3_x	16×16	$\begin{bmatrix} 1 \times 1, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 512 \end{bmatrix} \times 4$
conv4_x	8×8	$\begin{bmatrix} 1 \times 1, 256 \\ 3 \times 3, 256 \\ 1 \times 1, 1024 \end{bmatrix} \times 6$
conv5_x	4×4	$\begin{bmatrix} 1 \times 1, 512 \\ 3 \times 3, 512 \\ 1 \times 1, 2048 \end{bmatrix} \times 3$
	1×1	average pool, 100-d fc, softmax
FLOPs		2.59e9

2 We use a variant of the originally proposed ResNet architecture (He et al. [2016]). Particularly, we replace the initial
3 7×7 conv layer with a 3×3 conv layer. Here, “conv layer” refers to convolution followed by batchnorm (Ioffe
4 and Szegedy [2015]) and ReLU activation. This is intended to not excessively downsample the image—CIFAR-100
5 (Alex Krizhevsky [2009]) has images of dimensions 32×32 , compared to Imagenet’s (Russakovsky et al. [2015])
6 224×224 . Each block used (conv2_x, conv3_x, etc.) is a bottleneck block, and uses the conv-batchnorm-ReLU
7 ordering.

8 2 FLOP Counting Procedure

9 Following Evci et al. [2020], we base our counting procedure on the Micronet Challenge¹, which was conducted as
10 a part of NeurIPS 2019. Support for unstructured sparsity is assumed while computing the number of additions and
11 multiplication operations. The sum of these two gives us the theoretical FLOPs for a single forward pass through the
12 model.

13 Concretely, let the FLOPs required for a forward pass through a dense model be f_d and the corresponding for a sparse
14 model (or small-dense model) be f_s . Then, the FLOPs for training a dense model are $3f_d$ —since the backward pass
15 involves computing gradients with respect to each weight and activation. f_s can be computed for a model given its
16 sparsity distribution via the counting procedure. The FLOPs required to train a sparse model depend on the technique
17 used, as detailed below.

18 2.1 Inference FLOPs

19 **Small-Dense, RigL, SET, Static** These methods involve constant layer-wise sparsity throughout training, hence the
20 FLOP count can be determined during any step. The FLOP count for Random initialized models are $(1 - s)$ times the
21 Dense FLOPs.

22 **SNFS, Pruning** Both methods involve varying layer-wise sparsity during training, and hence non-constant FLOP
23 consumption. The final weights are used to determine inference FLOPs in this case.

24 2.2 Train FLOPs

25 **Small-Dense, Static** Dense gradients are not required by these models, and hence have a train FLOP count of $3f_s$.

26 **SET** Dense gradients are not required, and random growth can be implemented quite efficiently. Thus, the train FLOP
27 count is $3f_s$.

28 **RigL** Dense gradients are required only every ΔT steps, hence the corresponding train FLOP count is: $\frac{3\Delta T f_s + 2f_s + f_d}{\Delta T + 1}$.
29 We note that since ΔT is typically set between 100–1000, the preceding expression is quite close to $3f_s$.

30 **SNFS** Dense gradients are required at each training step, resulting in $2f_s + f_d$ FLOPs consumed at each step. Since
31 the sparse FLOP count varies as we train, the average FLOP count is: $2\mathbb{E}[f_{s,t}] + f_d$, where $f_{s,t}$ is the sparse inference
32 FLOPs at train step t .

33 **Pruning** Does not require dense gradients, but the sparsity increases smoothly from 0% to the target value as we train.
34 The FLOP consumption here is $3\mathbb{E}[f_{s,t}]$, where $f_{s,t}$ is the sparse inference FLOPs at train step t .

35 To determine $\mathbb{E}[f_{s,t}]$, we compute a running average of the FLOP consumption after every epoch. Notably, we find
36 that the inference cost of Pruning is often close to a Random initialized sparse network, while SNFS, regardless of
37 initialization, is compute-intensive.

38 3 Trial Space of Hyperparameter Tuning

39 Figure 1 shows the hyper-parameter study for tuning $(\alpha, \Delta T)$ as a contour plot. We observe that for multiple
40 initialization-density configurations, the reference choice $(\alpha = 0.3, \Delta T = 100)$, is quite close to the optimal hyper-
41 parameters. Furthermore, where they differ, the difference is within standard deviation bounds (Table 4 of the main
42 report).

43 4 Dynamic Structured Sparsity

44 Present hardware accelerators lack efficient implementations for unstructured sparsity. As a result, in practice, the
45 reduced FLOP requirement of sparse methods rarely translate to wall-clock improvements. In comparison, there are
46 efficient implementations available for structured (or block) sparsity which reach theoretical speedups (Gray et al.
47 [2017], Teja Vooturi et al. [2019]). Motivated by this, we try modifying *RigL* to explicitly work on structured sparsity.

¹<https://micronet-challenge.github.io/>

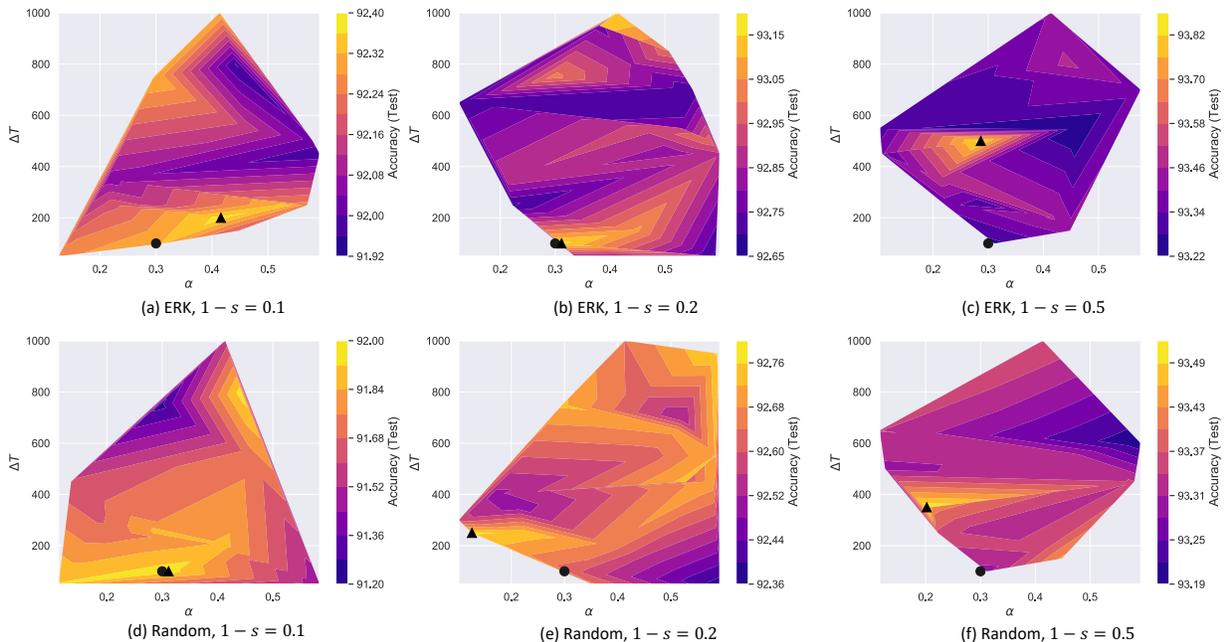


Figure 1: **Trial space of tuning** (α , ΔT), shown as a contour plot. Here, black circle corresponds to $(\alpha = 0.3, \Delta T = 100)$, while black triangle corresponds to the optimal hyper-parameter pair found. We plot the convex hull of the trial space, so in a few cases the reference point lies on the border of this space.

Table 2: **Modifying *RigL* for structured sparsity, compared on CIFAR-10 and CIFAR-100 datasets.** *RigL*-struct fails to match the accuracy of *RigL* and just matches Small-Dense in performance.

Method	CIFAR-10				CIFAR-100			
	$1-s = 0.1$		$1-s = 0.2$		$1-s = 0.1$		$1-s = 0.2$	
	Accuracy \uparrow (Test)	Wall Time \downarrow						
Small-Dense	89.0 ± 0.35	0.11x	91.0 ± 0.07	0.20x	70.8 ± 0.22	0.11x	72.6 ± 0.93	0.20x
Random Initialization								
RigL	91.7 ± 0.18	1.0x	92.9 ± 0.10	1.0x	71.8 ± 0.33	1.0x	73.5 ± 0.04	1.0x
RigL-Struct	87.0 ± 0.09	0.10x	90.4 ± 0.27	0.20x	69.1 ± 0.11	0.10x	71.9 ± 0.13	0.20x
ERK Initialization								
RigL	92.4 ± 0.06	1.0x	93.1 ± 0.09	1.0x	72.6 ± 0.37	1.0x	73.4 ± 0.15	1.0x
RigL-Struct	89.6 ± 0.16	0.17x	91.3 ± 0.18	0.35x	71.1 ± 0.15	0.23x	72.9 ± 0.08	0.38x

48 We promote channel sparsity for convolutional layers and keep fully connected layers dense. Mask update steps also
 49 operate at the channel level, based on *RigL*'s growth and pruning criterion. We name this method as *RigL*-struct. Such
 50 an approach is enticing, as we can remove masked-out channels, and obtain practical speedups on accelerators without
 51 needing support for unstructured sparsity.

52 Unfortunately, *RigL*-struct does not preserve the performance of originally proposed *RigL* (Table 2). In fact, it performs
 53 only as good as Small-Dense models, which negates the motivation behind such an experiment—Small-Dense models
 54 already achieve the intended speedups.

55 References

56 Geoffrey Hinton Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.

- 57 Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. Rigging the lottery: Making all tickets
58 winners. In *Proceedings of Machine Learning and Systems (ICML)*, July 2020.
- 59 Scott Gray, Alec Radford, and Diederik P Kingma. Gpu kernels for block-sparse weights. *arXiv preprint*
60 *arXiv:1711.09224*, 3, 2017.
- 61 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings*
62 *of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- 63 Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal
64 covariate shift. In *International Conference on Machine Learning (ICML)*, July 2015.
- 65 Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy,
66 Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Chal-
67 lenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- 68 Dharma Teja Vooturi, Girish Varma, and Kishore Kothapalli. Dynamic block sparse reparameterization of convolutional
69 neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*,
70 Oct 2019.