



HAL
open science

EXTRACTING EFFECTIVE SUBNETWORKS WITH GUMBEL-SOFTMAX

Robin Dupont, Mohammed Amine Alaoui, Hichem Sahbi, Alice Lebois

► **To cite this version:**

Robin Dupont, Mohammed Amine Alaoui, Hichem Sahbi, Alice Lebois. EXTRACTING EFFECTIVE SUBNETWORKS WITH GUMBEL-SOFTMAX. IEEE International Conference on Image Processing (ICIP), Oct 2022, Bordeaux, France. pp.931-935, 10.1109/ICIP46576.2022.9897718 . hal-03838823

HAL Id: hal-03838823

<https://hal.science/hal-03838823>

Submitted on 3 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

EXTRACTING EFFECTIVE SUBNETWORKS WITH GUMBEL-SOFTMAX

Robin Dupont^{*†}

Mohammed Amine Alaoui[†]

Hichem Sahbi^{*}

Alice Lebois[†]

^{*} Sorbonne Université, LIP6, Paris

[†] Netatmo, Boulogne-Billancourt

ABSTRACT

Large and performant neural networks are often overparameterized and can be drastically reduced in size and complexity thanks to pruning. Pruning is a group of methods, which seeks to remove redundant or unnecessary weights or groups of weights in a network. These techniques allow the creation of lightweight networks, which are particularly critical in embedded or mobile applications.

In this paper, we devise an alternative pruning method that allows extracting effective subnetworks from larger untrained ones. Our method is stochastic and extracts subnetworks by exploring different topologies which are sampled using Gumbel Softmax. The latter is also used to train probability distributions which measure the relevance of weights in the sampled topologies. The resulting subnetworks are further enhanced using a highly efficient rescaling mechanism that reduces training time and improves performance. Extensive experiments conducted on CIFAR show the outperformance of our subnetwork extraction method against the related work.

Index Terms— Lightweight networks, pruning, efficient computation, topology selection

1. INTRODUCTION

Deep neural networks are nowadays becoming mainstream in solving many image processing tasks including visual category recognition. The success of these models has been reached at the expense of an increase in their inference time, memory consumption and energy footprint. With the era of intelligent embedded systems (provided with limited energy and computational resources), a current trend is to make these models *lightweight and frugal* while maintaining their high accuracy. Existing solutions in lightweight network design are targeted toward creating small and efficient architectures from scratch [1, 2, 3, 4] while others derive highly compact yet effective neural networks from larger ones. These methods predominantly include knowledge distillation [5, 6, 7, 8, 9, 10] and pruning [11, 12, 13].

Pruning methods, either structured or unstructured, are particularly successful, and seek to remove connections with the least perceptible impact on classification accuracy. Structured pruning consists in *jointly* removing groups of weights,

entire channels or subnetworks [14, 15], whereas unstructured pruning aims at removing weights *individually* [13, 16]. Unstructured pruning has witnessed a recent surge in interest in the wake of the Lottery Ticket Hypothesis [17]; an empirical study in [17] shows that large pretrained networks encompass subnetworks, called *Lottery Tickets*, whose training with initial weights taken from the large networks yields comparably accurate classifiers. Another study [18] pushes that finding further and concludes that only the topology of these subnetworks is actually important in order to reach comparable performances. In general, extracting an efficient subnetwork is still an open problem and is computationally demanding as this amounts to full training of large networks (till convergence) prior to their pruning. Existing alternatives approach this problem using early pruning [19, 20, 21], but still require to train the weights. In contrast to these works, our proposed solution in this paper identifies effective subnetworks by training only their topology and without any weights tuning.

A theoretical analysis in [22, 23, 24] has established the sufficient conditions about the existence of efficient and effective subnetworks in over-parameterized large networks, nonetheless, no constructive proof has been provided in order to identify these subnetworks. In this context, Zhou et al. [25] proposed the first attempt to extract efficient subnetworks using stochastic mask training. A probability of selecting each weight is defined (as the sigmoid of a mask) and trained using the Straight Through Estimator (STE) [26]. During training, weights are frozen and only the masks are allowed to vary. However, the major drawback of this method resides in the vanishing gradient of the sigmoid which makes mask training numerically challenging. Ramanujan et al. [27] proposed another alternative, based on binarized saliency indicators learned with STE, which selects the most prominent weights in the resulting subnetworks. Nevertheless, since this method enforces the pruning rate *a priori*, finding the pruning rate giving the highest performances has to be made through a cumbersome and time-consuming binary search or grid-search.

Considering the limitation of the aforementioned related work, we introduce in this paper a new stochastic subnetwork selection method based on Gumbel Softmax. The latter allows sampling subnetworks whose weights are the most

relevant for classification. The proposed contribution also relies on a new mask parametrization, dubbed as Arbitrarily Shifted Log Parametrization (ASLP), that allows a better conditioning of the gradient and thereby mitigates numerical instability during mask optimization. Besides, when combining ASLP with a learned weight rescaling mechanism, training is accelerated and the accuracy of the resulting sub-networks improves as shown later in experiments.

2. PROPOSED METHOD

Let f_θ be a deep neural network whose weights defined as $\theta = \{\mathbf{w}_1, \dots, \mathbf{w}_L\}$, with L being its depth, $\mathbf{w}_\ell \in \mathbb{R}^{d_\ell \times d_{\ell-1}}$ its ℓ^{th} layer weights, and d_ℓ the dimension of ℓ . The output of a given layer ℓ is defined as

$$\mathbf{z}_\ell = g_\ell(\mathbf{w}_\ell \otimes \mathbf{z}_{\ell-1}), \quad (1)$$

being g_ℓ an activation function and \otimes the usual matrix product. Without a loss of generality, we omit the bias in the definition of (1).

2.1. Stochastic Weight Sampling

Given a network f_θ , weight pruning consists in removing connections in the graph of f_θ . A node in this graph refers to a neural unit while an edge corresponds to a cross-layer connection. Pruning is usually obtained by freezing and zeroing-out a subset of weights in θ , and this is achieved by multiplying \mathbf{w}_ℓ by a binary mask $\mathbf{m}_\ell \in \{0, 1\}^{\dim(\mathbf{w}_\ell)}$. The binary entries of \mathbf{m}_ℓ are set depending on whether the underlying layer connections are kept or removed, so Equation (1) becomes

$$\mathbf{z}_\ell = g_\ell((\mathbf{m}_\ell \odot \mathbf{w}_\ell) \otimes \mathbf{z}_{\ell-1}). \quad (2)$$

Here \odot stands for the element-wise matrix product. In this definition, the masks $\{\mathbf{m}_\ell\}_\ell$ are stochastic and sampled from a Bernoulli distribution.

Straight Through Estimator. Zhou et al. [25] consider a Bernoulli parametrization of $\{\mathbf{m}_\ell\}_\ell$ in order to sample masks in Equation (2). However, due to sampling which is not a differentiable operation, optimizing directly $\{\mathbf{m}_\ell\}_\ell$ is not possible. Existing solutions, including [25], rely on the Straight Trough Estimator (STE), already described in [26]. The definition of $\{\mathbf{m}_\ell\}_\ell$ is instead based on another *latent* parametrization $\{\hat{\mathbf{m}}_\ell\}_\ell$, detailed subsequently, and obtained by applying a sigmoid function $\sigma(\cdot)$ to $\hat{\mathbf{m}}_\ell$. This allows optimizing $\hat{\mathbf{m}}_\ell$ using gradient descent while considering the following surrogate of Equation (2)

$$\mathbf{z}_\ell = g_\ell((\sigma(\hat{\mathbf{m}}_\ell) \odot \mathbf{w}_\ell) \otimes \mathbf{z}_{\ell-1}). \quad (3)$$

Authors in [25] use the STE in order to back-propagate the gradient and to update the parameters of the Bernoulli distribution $\hat{\mathbf{m}}_\ell$ with gradient descent.

Gumbel-Softmax. In what follows, we consider an alternative STE based on Gumbel Softmax (GS) [28]. The proposed method, dubbed as Straight Through Gumbel Softmax (STGS), is based (i) on a variant of GS, and also (ii) on the argmax operator which allows sampling from a categorical distribution, as the limit of GS (i.e., when its softmax temperature approaches zero). Let z be a categorical random variable, associated with n class probability distribution $\mathcal{P} = [\pi_1, \dots, \pi_n]$. The Gumbel Softmax estimator (i) takes a vector of log-probabilities $\log(\mathcal{P}) = [\log(\pi_1), \dots, \log(\pi_n)]$ as an input, (ii) disrupts the latter with a random additive noise sampled from the Gumbel distribution, and (iii) takes the argmax, yielding a categorical variable. More formally, following [28], the value q of our categorical variable z is obtained as

$$q = \underset{k}{\operatorname{argmax}} [\log(\pi_k) + g_k], \quad (4)$$

with g_k being i.i.d sampled from the Gumbel distribution. In what follows, and unless stated otherwise, we omit ℓ from \mathbf{w}_ℓ and we write it for short as \mathbf{w} . Let w_{ij} be the weight associated to the i -th and j -th neurons respectively belonging to layers $\ell - 1$ and ℓ ; we define a two-class categorical distribution \mathcal{P}_{ij} on $\{0, 1\}$ as $\mathcal{P}_{ij}(z = 1) = \pi_1^{ij}$, and $\mathcal{P}_{ij}(z = 0) = \pi_2^{ij}$ with $\pi_1^{ij} = p_{ij}$, $\pi_2^{ij} = 1 - p_{ij}$ and p_{ij} being the probability to keep the underlying connection. In other words, keeping the weight w_{ij} (or not) in the sampled topology is a Bernoulli trial with a probability p_{ij} . Considering Equation (4), a binary mask m_{ij} is defined as $1_{\{q_{ij}=1\}}$, $1_{\{\cdot\}}$ being the indicator function and $q_{ij} = \underset{k \in \{1, 2\}}{\operatorname{argmax}} [\log(\pi_k^{ij}) + g_k^{ij}]$. Thanks to STGS, it becomes possible to learn p_{ij} for each weight through stochastic gradient descent (SGD). However, optimizing p_{ij} (with SGD) raises a major issue as p_{ij} may not be appropriately bounded and thereby $\log(p_{ij})$ and $\log(1 - p_{ij})$ would also be undefined. On another hand, solving constrained SGD, besides being computationally expensive and challenging, may result into worse local minimum. In order to overcome all these issues, one may consider an alternative reparametrization $p_{ij} = \sigma(\hat{m}_{ij})$, with \hat{m}_{ij} being a latent mask variable and σ the sigmoid function which bounds p_{ij} in $[0, 1]$. However, this workaround suffers (in practice) from numerical instability in gradient estimation (due to the log and the sigmoid) and is also computationally demanding.

Arbitrarily Shifted Log Parametrization. Another alternative is to consider $\hat{m}_{ij} = \log(p_{ij})$ and $\log(1 - p_{ij}) = \log(1 - \exp(\hat{m}_{ij}))$ and learn the underlying mask. However, this reparametrization is also flawed in the same way as the aforementioned sigmoid reparametrization. In what follows, we propose an equivalent formulation which turns out to be highly effective and numerically more stable. Considering

$$\begin{bmatrix} \hat{m}_{ij} \\ 0 \end{bmatrix} = \log(\mathcal{P}_{ij}(\cdot)) + c = \begin{bmatrix} \log(p_{ij}) + c \\ \log(1 - p_{ij}) + c \end{bmatrix}, \quad (5)$$

in the above definition, instead of using $\log(\mathcal{P}_{ij}(\cdot))$, we consider $\log(\mathcal{P}_{ij}(\cdot)) + c$ as an input of the argmax in Equation (4). The constant $c \in \mathbb{R}$ ensures that if $\hat{m}_{ij} > 0$, then $\log(p_{ij}) \in] - \infty, 0] \Leftrightarrow p_{ij} \in [0, 1]$. This is enforced by setting the second coefficient of \mathcal{P}_{ij} to 0, rather than computing it explicitly. The formulation of Equation (5) is theoretically equivalent to the aforementioned sigmoid reparametrization. Indeed, solving the system of Equation (5) w.r.t. \hat{m}_{ij} yields $p_{ij} = \sigma(\hat{m}_{ij})$. Differently put, the formulation in Equation (5) considers a reparametrization $\hat{m}_{ij} = \log(p_{ij}) + c$ which is strictly equivalent to the sigmoid one while being computationally more efficient and also stable. Note that adding any arbitrary constant c to the log-probability makes the outcome of Gumbel-Softmax sampling and argmax invariant.

2.2. Weight Rescaling

Subnetwork selection may disrupt the dynamic of the forward pass [29, 27], and thereby requires adapting weights accordingly. Dynamic weight rescale (DWR) [25], and scaled Kaiming distribution [27] are two known mechanisms that adapt the weights of the selected subnetworks. However, some of these heuristics, besides being handcrafted, rely on the strong assumption that rescaling should be proportional to the pruning rate. In what follows, we consider a new weight adaptation mechanism, referred to as Smart Rescale (SR). Instead of handcrafting this rescaling factor proportionally to the pruning rate (as achieved for instance in [25]), SR is learned layerwise and provides an effective (and also efficient) way to adapt the dynamic of the forward pass without retraining the entire weights of the selected subnetwork. Indeed, this rescaling ends up reducing the amount of epochs needed to reach convergence and also improving accuracy (at some extent) as shown later in experiments.

With SR, the ℓ -th layer network output becomes

$$\mathbf{z}_\ell = g_\ell(s_\ell \times (\mathbf{m}_\ell \odot \mathbf{w}_\ell) \otimes \mathbf{z}_{\ell-1}), \quad (6)$$

where s_ℓ refers to the rescaling factor of the ℓ -th layer (see also algorithm 1). Smart Rescale increases the flexibility of subnetwork selection and adaptation compared to DWR (which is bound to the pruning rate). Moreover, scaling factors obtained with SR vary smoothly — and this makes training more stable with stochastic gradient descent (SGD) — compared to the ones obtained with DWR which are again set to the observed *pruning rates*, and changes of the latter are more abrupt due to stochastic mask sampling.

3. EXPERIMENTS

In this section, we show the performance of our method on the standard CIFAR10 and CIFAR100 datasets. They consist of 60k colored images of 32×32 pixels each. Training,

Algorithm 1 Forward pass for our method

Require: A network f_θ , with weights $\{\mathbf{w}_\ell\}_\ell$, ASLP $\{\hat{m}_\ell\}_\ell$, and input training data $\{(\mathbf{x}_k, \mathbf{y}_k)\}_k$

- 1: $q_{i,j} \leftarrow \operatorname{argmax} \begin{bmatrix} \hat{m}_{i,j} + g_{i,j} \\ 0 + g'_{i,j} \end{bmatrix} \triangleright$ Sampling of a topology
 - 2: $m_{i,j} \leftarrow 1_{\{q_{i,j}=1\}} \triangleright$ Giving the masks $m_{i,j}$ their values
 - 3: **Return** $\mathcal{L}(f_\theta(\{\mathbf{x}_k\}_k; \{s_\ell(\mathbf{m}_\ell \odot \mathbf{w}_\ell)\}_\ell), \{\mathbf{y}_k\}_k)$ \triangleright Computing the loss with masked weights and SR
-

validation and test sets include 45k, 5k and 10k images respectively.

In order to demonstrate the effectiveness of our method, we chose the widely used SGD optimizer with a momentum of 0.9 and a learning rate of 50. Faster convergence is obtained with higher learning rates, however, the latter also lead to worse observed accuracy. During training, the maximum number of epochs is set to 1000 and early stopping is triggered if the accuracy on the validation set stops improving during 100 epochs. In all these experiments, neither weight decay nor ℓ_2 regularization are applied. See implementation details and our code on the ASLP GitHub [?].

3.1. Performance and comparison

The accuracy of our method is evaluated on subnetworks whose topology corresponds to connections with (trained) probabilities larger than 0.5; in other words, if *the binary event of keeping a connection is more likely than its removal*. This setting is referred to as *thresholding*. As a matter of comparison, we also consider the setting in [25] which consists in sampling ten different subnetworks and evaluating an average accuracy over these subnetworks. This setting is referred to as *averaging*. In these experiments, we use the same networks as [25, 27] (originally introduced by Frankle and Carbin [17]) namely Conv2, Conv4 and Conv6 which are variants of VGG16.

Tab. 1 shows a comparison of our method against [25, 27]. These results show means of five independent runs; each run corresponds either to “thresholding” or “averaging”. These performances show a consistent gain (in accuracy) of our subnetwork selection. We also observe that “thresholding” is already effective compared to “averaging”; indeed, our method reaches a high accuracy despite learning a single subnetwork topology, and this makes it also highly efficient for training compared to the related work [25, 27].

Furthermore, our method and [25] do not impose a pruning rate. The optimal pruning rate is found during optimization and is around 51%, whereas [27] enforces a 50% pruning rate ($k = 50\%$). Thus, the networks capacities are comparable.

		Cifar 10								Cifar 100
		w/o data augmentation				with data augmentation (w.d.a)				w.d.a
		\emptyset	WR	SC	WR+SC	\emptyset	WR	SC	WR+SC	WR+SC
Conv2	[25] (averaging)	64.4	65.0	66.3	66.0	-	-	-	-	-
	[27] ¹ ($k = 50\%$)	-	-	-	-	-	-	71.5	71.7	40.9
	Our ASLP (averaging)	68.2	66.9	68.3	66.5	76.0	76.6	76.8	77.3	-
	Our ASLP (thresholding)	68.7	67.8	68.4	67.1	75.9	76.4	77.5	77.5	43.3
Conv4	[25] (averaging)	65.4	71.1	66.2	72.5	-	-	-	-	-
	[27] ¹ ($k = 50\%$)	-	-	-	-	-	-	81.6	80.5	51.1
	Our ASLP (averaging)	70.6	71.8	69.5	71.8	83.4	84.4	83.7	84.1	-
	Our ASLP (thresholding)	71.5	72.8	70.2	72.7	83.7	85.0	84.5	84.8	51.7
Conv6	[25] (averaging)	63.5	76.3	65.4	76.5	-	-	-	-	-
	[27] ¹ ($k = 50\%$)	-	-	-	-	-	-	85.4	85.1	53.8
	Our ASLP (averaging)	72.9	76.1	71.9	75.6	85.3	86.2	85.3	86.2	-
	Our ASLP (thresholding)	73.7	77.0	72.6	76.6	86.0	86.9	86.3	86.9	52.8

Table 1. Comparison of our method against [25] and [27] on Conv2, Conv4 and Conv6. These results are averaged through five independent runs. “WR” (Weight Rescale) refers to “Dynamic Weight Rescale” or “Smart Rescale” depending on which methods is used (respectively [25] or our proposed ASLP). Again, “SC” refers to the “Signed Constant” distribution. The latest results on CIFAR 100 were recently obtained with data augmentation and WR+SC.

3.2. Ablation study

In this section, we discuss the impact of all the components of the method when taken individually and combined, namely the use of weight rescaling (WR): either DWR or our proposed SR. We also consider another criterion: signed constant (SC) which consists in replacing weights in a given layer by the products of their signs and the standard deviation of their original weight distribution. We show all these results with and without data augmentation, which is composed of the combination of zero-padding, random crops and random horizontal flips. Note that pixel intensities are normalized from their original values in $[0, 255]$ to $[0, 1]$.

From the results in table 1, we observe a clear gain of our method alone w.r.t. [25] and the use of SR increases further its accuracy (excepting Conv2 w/o data augmentation). The gain in performances increases significantly with Conv6 and reaches up to 4 points even when no data augmentation is used. Note that the use of data augmentation attenuates, at some extent, the effect of SR on larger networks (conv4 and Conv6). Nonetheless, as discussed in Sec. 3.3, the positive impact of SR resides also in training efficiency. In contrast to SR, signed constant improves accuracy by a small margin when combined with data augmentation.

3.3. Computational efficiency

DWR requires rectifying weights layerwise using the inverse of the observed (computed) pruning rates. These layerwise evaluations introduce a significant overhead at each training epoch. In contrast, SR consists in simple products involving one scalar per layer. When training Conv4, we found

(on average) that enabling DWR increases epoch runtime by 0.2s while our SR by 0.13s only, so SR speeds up training overhead by 35% compared to DWR. When data augmentation and signed constant are used, SR allows a significant gain in the number of training epochs. Indeed, enabling SR on Conv4 saves (on average) 19.7% training epochs (8.2%, 14.0% on Conv2 and Conv6 respectively) before converging to its highest accuracy. Finally, our “thresholding” setting not only improves accuracy but makes subnetwork selection (training) and also inference more efficient compared to the related work [25, 27], as this selection is again achieved once and thereby only one subnetwork is applied during inference.

4. CONCLUSION

In this paper, we introduce a novel method that extracts effective subnetworks from larger networks without training its weights. The proposed method optimizes a probability distribution which measures the relevance of weights, and only those with the highest relevance define the topology of the selected subnetworks. An efficient and effective weight rescaling mechanism is also introduced and allows rectifying the parameters of the selected subnetworks which improves performances and reduces the number epochs needed to reach convergence. Experiments conducted on the standard CIFAR10 and CIFAR100 datasets show the effectiveness of our subnetwork selection method w.r.t. the related work. Future work includes the study of the scalability of the proposed method on more complex datasets and other larger networks.

Acknowledgement. This work has been achieved within a partnership between Sorbonne University and Netatmo.

Code. Our code is available at:

<https://github.com/N0ciple/ASLP>

¹Performances for [27] are reported with the optimizer described in Sec. 3. It is possible to improve performances by tuning the learning rate scheduler but this is out of the scope of this paper.

5. REFERENCES

- [1] G. Huang, S. Liu, L. van der Maaten, and K. Q. Weinberger, "Condensenet: An efficient densenet using learned group convolutions," in *CVPR*, 2018.
- [2] M. Sandler, A. G. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *CVPR*, 2018.
- [3] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *CoRR*, vol. abs/1704.04861, 2017.
- [4] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *ICML*. 2019, vol. 97 of *Proceedings of Machine Learning Research*, PMLR.
- [5] G. E. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *CoRR*, vol. abs/1503.02531, 2015.
- [6] S. Zagoruyko and N. Komodakis, "Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer," in *ICLR*, 2017.
- [7] A. Romero, N. Ballas, S. Ebrahimi Kahou, A. Chassang, C. Gatta, and Y. Bengio, "Fitnets: Hints for thin deep nets," in *ICLR*, 2015.
- [8] S.-I. Mirzadeh, M. Farajtabar, A. Li, N. Levine, A. Matsukawa, and H. Ghasemzadeh, "Improved knowledge distillation via teacher assistant," in *AAAI*, 2020.
- [9] Y. Zhang, T. Xiang, T. M. Hospedales, and H. Lu, "Deep mutual learning," in *CVPR*, 2018.
- [10] S. Ahn, S. X. Hu, A. C. Damianou, N. D. Lawrence, and Z. Dai, "Variational information distillation for knowledge transfer," in *CVPR*, 2019.
- [11] Y. LeCun, J. S. Denker, and S. A. Solla, "Optimal brain damage," in *NIPS*, 1989.
- [12] B. Hassibi and D. G. Stork, "Second order derivatives for network pruning: Optimal brain surgeon," in *NIPS*, 1992.
- [13] S. Han, J. Pool, J. Tran, and W. J. Dally, "Learning both weights and connections for efficient neural network," in *NIPS*, 2015.
- [14] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, "Pruning filters for efficient convnets," in *ICLR*, 2017.
- [15] Z. Liu, J. Li, Z. Shen, G. Huang, S. Yan, and C. Zhang, "Learning efficient convolutional networks through network slimming," in *ICCV*. 2017, IEEE Computer Society.
- [16] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding," in *ICLR*, 2016.
- [17] J. Frankle and M. Carbin, "The lottery ticket hypothesis: Finding sparse, trainable neural networks," in *ICLR*, 2019.
- [18] Z. Liu, M. Sun, T. Zhou, G. Huang, and T. Darrell, "Rethinking the value of network pruning," in *ICLR*, 2019.
- [19] N. Lee, T. Ajanthan, and P. H. S. Torr, "Snip: single-shot network pruning based on connection sensitivity," in *ICLR*, 2019.
- [20] C. Wang, G. Zhang, and R. B. Grosse, "Picking winning tickets before training by preserving gradient flow," in *ICLR*, 2020.
- [21] H. Tanaka, D. Kunin, D. L. Yamins, and S. Ganguli, "Pruning neural networks without any data by iteratively conserving synaptic flow," in *NeurIPS*, 2020.
- [22] E. Malach, G. Yehudai, S. Shalev-Shwartz, and O. Shamir, "Proving the lottery ticket hypothesis: Pruning is all you need," in *ICML*, 2020.
- [23] A. Pensia, S. Rajput, A. Nagle, H. Vishwakarma, and D. S. Papailiopoulos, "Optimal lottery tickets via subset sum: Logarithmic over-parameterization is sufficient," in *NeurIPS 2020*, 2020.
- [24] L. Orseau, M. Hutter, and O. Rivasplata, "Logarithmic pruning is all you need," in *NeurIPS 2020*, 2020.
- [25] H. Zhou, J. Lan, R. Liu, and J. Yosinski, "Deconstructing lottery tickets: Zeros, signs, and the supermask," in *NeurIPS*, 2019.
- [26] Y. Bengio, N. Léonard, and A. C. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," *CoRR*, vol. abs/1308.3432, 2013.
- [27] V. Ramanujan, M. Wortsman, A. Kembhavi, A. Farhadi, and M. Rastegari, "What's hidden in a randomly weighted neural network?," in *CVPR*. 2020, Computer Vision Foundation / IEEE.
- [28] E. Jang, S. Gu, and B. Poole, "Categorical reparameterization with gumbel-softmax," in *ICLR*, 2017.
- [29] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *ICCV*, 2015.