

Connectivity Matters: Neural Network Pruning Through the Lens of Effective Sparsity

Artem Vysogorets

*Center for Data Science
New York University
New York, NY 10011, USA*

AMV458@NYU.EDU

Julia Kempe

*Center for Data Science
Courant Institute for Mathematical Sciences
New York University
New York, NY 10011, USA*

JK185@NYU.EDU

Editor: Pradeep Ravikumar

Abstract

Neural network pruning is a fruitful area of research with surging interest in high sparsity regimes. Benchmarking in this domain heavily relies on faithful representation of the sparsity of subnetworks, which has been traditionally computed as the fraction of removed connections (direct sparsity). This definition, however, fails to recognize unpruned parameters that detached from input or output layers of the underlying subnetworks, potentially underestimating actual effective sparsity: the fraction of inactivated connections. While this effect might be negligible for moderately pruned networks (up to $10\times$ – $100\times$ compression rates), we find that it plays an increasing role for sparser subnetworks, greatly distorting comparison between different pruning algorithms. For example, we show that effective compression of a randomly pruned LeNet-300-100 can be orders of magnitude larger than its direct counterpart, while no discrepancy is ever observed when using SynFlow for pruning (Tanaka et al., 2020). In this work, we adopt the lens of effective sparsity to reevaluate several recent pruning algorithms on common benchmark architectures (e.g., LeNet-300-100, VGG-19, ResNet-18) and discover that their absolute and relative performance changes dramatically in this new, and as we argue, more appropriate framework. To aim for effective, rather than direct, sparsity, we develop a low-cost extension to most pruning algorithms. Further, equipped with effective sparsity as a reference frame, we partially reconfirm that random pruning with appropriate sparsity allocation across layers performs as well or better than more sophisticated algorithms for pruning at initialization (Su et al., 2020). In response to this observation, using an analogy of pressure distribution in coupled cylinders from thermodynamics, we design novel layerwise sparsity quotas that outperform all existing baselines in the context of random pruning.

Keywords: Neural networks, pruning, sparsity, lottery tickets

1. Introduction

Recent successful advances of Deep Neural Networks are commonly attributed to their high architectural complexity and excessive size (*over-parameterization*) (Denton et al., 2014; Neyshabur et al., 2019; Arora et al., 2018). Modern state-of-the-art architectures exhibit

enormous parameter overhead, requiring prohibitive amounts of resources during both training and inference and leaving a significant environmental footprint (Shoeybi et al., 2019). In response to these challenges, much attention has turned to compression of neural networks and, in particular, parameter pruning. While initial approaches mostly focused on pruning models after training (LeCun et al., 1990; Hassibi et al., 1993), contemporary algorithms optimize the sparsity structure of a network while training its parameters (Mocanu et al., 2018; Evci et al., 2020) or even remove connections before any training whatsoever (Lee et al., 2019; Wang et al., 2020).

Compression rates considered in the pruning literature usually fall between $10\times$ and $100\times$ of the size of the original model. However, as contemporary model sizes grow into the billions of parameters, studying higher compression regimes becomes increasingly important. Recently, a new bold sparsity benchmark was set by Tanaka et al. (2020) with Iterative Synaptic Flow (SynFlow), a data-agnostic algorithm for pruning at initialization. Reportedly, it is capable of removing all but only a few hundreds of parameters (a $100,000\times$ compression for VGG-16) and still produce trainable subnetworks, while other pruning methods disconnect networks at much lower sparsity levels (Tanaka et al., 2020). Related work by de Jorge et al. (2021) proposes an iterative version of one-shot pruning algorithm, Single-shot Network Pruning (SNIP) (Lee et al., 2019), and evaluates it in a similar high sparsity regime, reaching more than $10,000\times$ compression ratio.

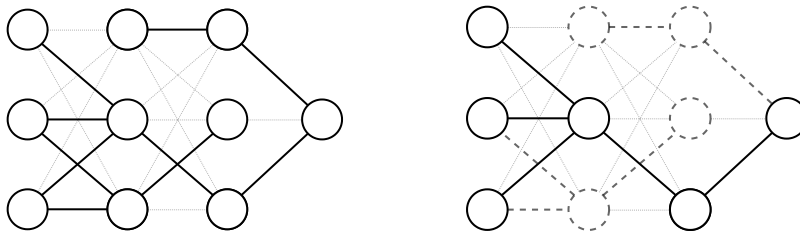


Figure 1: Pruning 11 edges from a fully-connected 21-edge network. Left: direct sparsity ($11/21$) does not account for disconnected edges (compression $21/10 = 2.1$). Right: effective sparsity ($16/21$) accounts for the 5 dashed connections incident to inactivated neurons (yielding twice as large effective compression $21/5 = 4.2$).

Effective sparsity. This increased focus on extreme sparsity leads us to consider *what sparsity is meant to represent* in neural networks and computational graphs at large. In the context of neural network pruning, sparsity to date is computed straightforwardly as the *fraction of removed connections (direct sparsity)*—and compression as the inverse fraction of unpruned connections (*direct compression*). We observe that this definition does not distinguish between connections that have actually been pruned, and those that have become *effectively* pruned because they have disconnected from the computational flow. Formally, an edge θ_i is considered inactive if, for any input x , the output of the neural network $f(\theta, x)$ does not depend on the value of θ_i . In this work, we propose to abandon direct sparsity in favor of *effective sparsity*—the *fraction of inactivated connections*, be it through direct pruning or through otherwise disconnecting from either input or output of a network (see Figure 1 for an illustration).

We advocate that effective sparsity (effective compression) be used universally in place of its direct counterpart since it more accurately depicts what one would reasonably consider the network’s sparsity state. Using the lens of effective compression for benchmarking allows for a fairer comparison between different unstructured pruning algorithms. Note that effective compression is lower bounded by direct compression, which means that some pruning algorithms will give improved sparsity-accuracy trade-offs in this new framework. In Section 3, we critically reexamine a plethora of recent pruning algorithms for a variety of architectures to find that, in this refined framework, conclusions drawn in previous works appear overstated or incorrect. Figure 2 gives a sneak-preview of this effect for three ab-initio pruning algorithms: SynFlow (Tanaka et al., 2020), SNIP (Lee et al., 2019) and plain random pruning for LeNet-300-100 on MNIST. While SynFlow appears superior to other methods when evaluated against direct compression, it loses its advantage in the effective framework. Such radical performance changes are partly explained by differing gaps between effective and direct compression inherent to different pruning algorithms (Figure 2). We can see that significant departure of direct from effective compression kicks in at relatively low rates below $100\times$, making our work relevant even in these moderate regimes. For example, using random pruning to compress LeNet-300-100 by $100\times$ (sparsity 99%) results in $\sim 1,000\times$ effective compression; yet, removing the same number of parameters with SynFlow yields an unchanged $100\times$ effective compression. What makes certain iterative algorithms like SynFlow less likely to amass disconnected edges? In Section 3, we show that they are fortuitously designed to achieve a close convergence of direct and effective sparsity, hinting that preserving connectivity is an important aspect in the strong performance of high-compression pruning algorithms (Tanaka et al., 2020; de Jorge et al., 2021). Moreover, the lens of effective compression gives access to more extreme compression regimes for some pruning algorithms, which appear to disconnect much earlier when not accounting for inactive connections. For these high effective compression ratios all three pruning methods from Figure 2 perform surprisingly similar, even though they use varying degrees of information on data and parameter values.

Layerwise Sparsity Quotas (LSQ) and Ideal Gas Quotas (IGQ). A recent thread of research by Frankle et al. (2021) and Su et al. (2020) shows that performance of trained subnetworks produced by algorithms for pruning at initialization is robust to randomly reshuffling unpruned edges within layers before training. This observation led to the conjecture that these algorithms essentially generate successful distributions of sparsity across layers, while the exact connectivity patterns are unimportant. In Section 4, we reexamine this conjecture through the lens of effective sparsity, confirm it for moderate compression regimes ($10\times$ – $100\times$) studied by Frankle et al. (2021) and Su et al. (2020), but find the truth to be more nuanced at higher compression rates. Nonetheless, this result highlights the importance of algorithms that carefully engineer *layerwise sparsity quotas (LSQ)* to obtain very simple and adequately performing random pruning algorithms. Furthermore, concurrently with our work, Liu et al. (2022b), find that random pruning with carefully allocated sparsity among layers can match the performance of their dense counterparts. Another important motivation to search for good LSQ is that global pruning algorithms frequently remove entire layers prematurely (Lee et al., 2020) (cf. layer-collapse in Tanaka et al. (2020)), even before any significant differences between direct and effective sparsity

emerge. Well-engineered LSQ could avoid this and enforce proper redistribution of compression across layers (see Gale et al. (2019); Mocanu et al. (2018); Evci et al. (2020) for existing baselines). In Section 4, we propose a novel LSQ coined *Ideal Gas Quotas (IGQ)* by drawing intuitive analogies from physics. Effortlessly computable for any network-sparsity combination, IGQ performs similarly or better than any other baseline in the context of random pruning at initialization and of magnitude pruning after training.

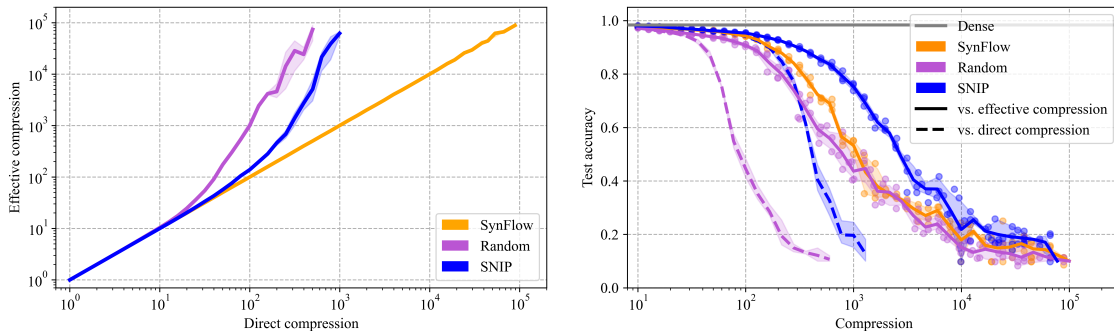


Figure 2: LeNet-300-100 trained on MNIST after pruning. Left: gaps between direct and effective compression. Right: SynFlow has a better sparsity-accuracy trade-off than SNIP when plotted against direct (dashed), but not against effective compression (solid curves). Dots represent individual experiments. Dashed and solid curves coincide for SynFlow.

Effective pruning. Pruning to any desired direct sparsity is straightforward: one simply needs to mask out the corresponding number of parameters from a network. Effective sparsity, unfortunately, is more unpredictable and difficult to control. In particular, several known pruning algorithms suffer from layer-collapse once reaching a certain sparsity level, leading to unstable effective sparsity just before the disconnection. As a result, most pruning methods are unable to deliver certain values of effective sparsity regardless of how many connections are pruned. When possible, however, one needs to carefully tune the number of pruned parameters so that effective sparsity lands near a desired value. In Section 5, we suggest a simple extension to algorithms for pruning at initialization or after training that helps bring effective sparsity close to any predefined achievable value while incurring costs that are at most logarithmic in model size.

Our contributions: summary. In this study, we (i) formulate and illustrate the importance of effective sparsity by reevaluating several recent pruning strategies; (ii) provide algorithms to prune according to and compute effective sparsity; (iii) reconfirm that networks pruned at initialization are robust to layerwise reshuffling of survived edges (Frankle et al., 2021) in the new sparsity framework, and (iv) design efficient layerwise sparsity quotas IGQ for random pruning that perform consistently well across all sparsity regimes.

2. Related work

Neural network compression encompasses a number of orthogonal approaches such as parameter regularization (Lebedev and Lempitsky, 2016; Louizos et al., 2018), variational

dropout (Molchanov et al., 2017), vector quantization and parameter sharing (Gong et al., 2014; Chen et al., 2015; Han et al., 2016), low-rank matrix decomposition (Denton et al., 2014; Jaderberg et al., 2014), and knowledge distillation (Buciluă et al., 2006; Hinton et al., 2015). Network pruning, however, is by far the most common technique for model compression, and can be partitioned into structured (at the level of entire neurons/units) and unstructured (at the level of individual connections). While the former offers resource efficiency unconditioned on use of specialized hardware (Liu et al., 2019) and constitutes a fruitful research area (Li et al., 2017; Liu et al., 2017), we focus on the more actively studied unstructured pruning, which is where differences between effective and direct sparsity emerge. In what follows we give a quick overview, naturally grouping pruning methods by the time they are applied relative to training (see Frankle and Carbin (2019) and Wang et al. (2020) for a similar taxonomy).

Pruning after training. These earliest pruning techniques were designed to remove the least “salient” learned connections without sacrificing predictive performance. Optimal Brain Damage (LeCun et al., 1990) and its sequel Optimal Brain Surgeon (Hassibi et al., 1993) use the Hessian of the loss to estimate sensitivity to removal of individual parameters. Han et al. (2015) popularized magnitude as a simple and effective pruning criterion. It proved to be especially successful when applied alternately with several finetuning cycles, which is commonly referred to as Iterative Magnitude Pruning (IMP), a modification of which was used by Frankle and Carbin (2019) to discover lottery tickets—sparse subnetworks that achieve the performance of their dense counterparts within a commensurate number of iterations. Later, Dong et al. (2017) showed that magnitude-based pruning minimizes ℓ_2 distortion of each layer’s output incurred by parameter removal. Recently, Lee et al. (2021) extend this idea and propose Layer-Adaptive Magnitude-Based Pruning (LAMP), which approximately minimizes the upper bound of the ℓ_2 distortion of the entire network. While equivalent to magnitude pruning within individual layers, LAMP automatically discovers excellent layerwise sparsity quotas (see Section 4) that yield better performance (as a function of *direct* compression) than existing alternatives in the context of IMP.

Pruning during training. Algorithms in this category learn sparsity structures together with parameter values, hoping that continued training will correct for damage incurred by pruning. To avoid inefficient prune-retrain cycles inherent to IMP, Narang et al. (2017) introduce gradual magnitude pruning over a single training round. Subsequently, Zhu and Gupta (2018) modify this algorithm by introducing a simpler pruning schedule and keeping layerwise sparsities uniform throughout training. Sparse Evolutionary Training (SET) (Mocanu et al., 2018) starts with an already sparse subnetwork and restructures it during training by pruning and randomly reviving connections. Unlike SET, Mostafa and Wang (2019) allow redistribution of sparsity across layers, while Dettmers and Zettlemoyer (2019) use gradient momentum as the criterion for parameter regrowth. Evci et al. (2020) rely on the instantaneous gradient to revive weights but follow SET to maintain the initial layerwise sparsity distribution during training. The In-time Over-parameterization (ITOP) framework provides insights into the underlying mechanisms of the above methods and leads to improved training protocols that boost their performance (Liu et al., 2021). A different body of works tackle the general optimization problem with an intractable ℓ_0 parameter sparsity constraint by designing and solving related continuous problems (Zhou et al., 2021;

Savarese et al., 2020; Kusupati et al., 2020). For example, Continuous Sparsification (CS) by Savarese et al. (2020) uses a sigmoid of learnable continuous variables as mask values and applies ℓ_1 regularization, effectively forcing them to either 0 or 1 during training.

Pruning before training. Pruning at initialization is especially alluring to deep learning practitioners as it promises lower costs of both optimization and inference. While this may seem too ambitious, the Lottery Ticket Hypothesis (LTH) postulates that randomly initialized dense networks do indeed contain highly trainable and equally well-performing sparse subnetworks (Frankle and Carbin, 2019). Inspired by the LTH, Lee et al. (2019) design SNIP, which uses connection sensitivity as a parameter saliency score. Wang et al. (2020) notice that SNIP creates bottlenecks or even removes entire layers and propose Gradient Signal Preservation (GraSP) as an alternative that aims to maximize gradient flow in a pruned network. de Jorge et al. (2021) improve SNIP by applying it iteratively, allowing for reassessment of saliency scores during pruning and helping networks stay connected at higher compression rates. A truly new compression benchmark was set by Tanaka et al. (2020); their algorithm, SynFlow, iteratively prunes subsets of parameters according to their ℓ_1 path norm and helps networks reach maximum compression without disconnecting. For example, SynFlow achieves non-random test accuracy on CIFAR-10 with a $100,000\times$ compressed VGG-16, while SNIP and GraSP fail already at $100\times$ and $1,000\times$, respectively. An extensive ablation study by Frankle et al. (2021) examines SNIP, GraSP and SynFlow within moderate compression rates (up to $100\times$) and reveals that performance of subnetworks produced by these methods is stable under random layerwise rearrangement of edges prior to training. Later, this result was independently confirmed by Su et al. (2020) for SNIP and GraSP only. This observation suggests that these algorithms perform as well as random pruning with corresponding layerwise quotas, putting the spotlight on designing competitive LSQ (Mocanu et al., 2018; Gale et al., 2019; Lee et al., 2021). Price and Tanner (2021) augment the functionality of sparse layers by precomputing a deterministic transformation of the input, thus maintaining information propagation and avoiding layer-collapse.

3. Effective sparsity

In this section, we present our comparisons of a variety of pruning algorithms under the lens of effective compression. To illustrate the striking difference between direct and effective sparsity and expose the often radical change in relative performance of pruning algorithms when switching from the former to the latter, we evaluate several recent methods (SNIP, GraSP, SynFlow, Magnitude & LAMP¹, CS², SNIP-iterative) and random pruning with uniform sparsity distribution across layers in both frameworks. Our experiments encompass modern architectures on commonly used computer vision benchmark datasets: LeNet-300-100 (Lecun et al., 1998) on MNIST, LeNet-5 (Lecun et al., 1998) on CIFAR-10, VGG-19 (Simonyan and Zisserman, 2015) on CIFAR-100, ResNet-18 (He et al., 2016) on TinyImageNet, and ResNet-50 and MobileNetV2 (Howard et al., 2017) on ImageNet. We place results of VGG-16 (Simonyan and Zisserman, 2015) on CIFAR-10 in Appendix B, as they closely resemble those of VGG-19. Further experimental details are listed in Appendix A.

-
1. as two versions of magnitude pruning after training and close siblings of lottery tickets (Frankle and Carbin, 2019).
 2. as a representative of methods that use learnable sparsity

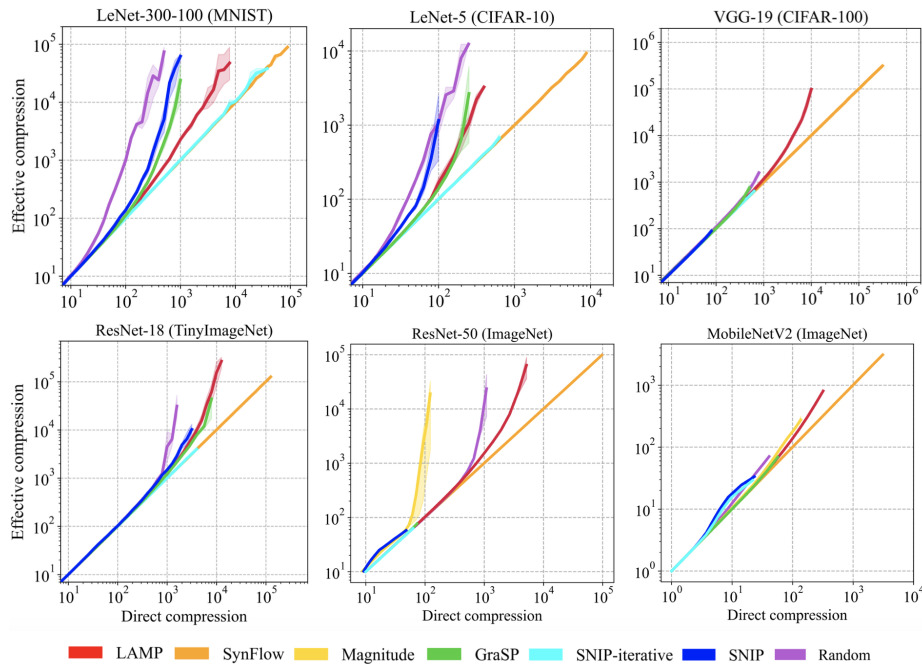


Figure 3: Effective versus direct compression across different pruning methods and architectures (curves and bands represent min/average/max across 3 seeds where subnetworks disconnect last among a total of 5 seeds).

Notation. Consider an L -layer neural network $f(\Theta; x)$ with weight tensors $\Theta = \{\Theta_\ell\}_{\ell=1}^L$ for $\ell \in [L]$. A subnetwork is specified by a set of binary masks that indicate unpruned parameters $M_\ell \in \{0, 1\}^{|\Theta_\ell|}$. With $\mathbf{M} = \{M_\ell\}_{\ell=1}^L$, it is given by $f(\Theta \odot \mathbf{M}; x)$ where \odot is Hadamard product. Note that biases and batchnorm parameters (Ioffe and Szegedy, 2015) are normally considered unprunable. Direct sparsity, the fraction of pruned weights, is given by $s(\mathbf{M}) = 1 - \sum_\ell \|M_\ell\|_0 / \sum_\ell |M_\ell|$ and direct compression rate is defined as $(1 - s(\mathbf{M}))^{-1}$.

Results. Figure 3 reveals that different algorithms tend to develop varying amounts of inactive connections. For example, effective compression of subnetworks pruned by LAMP consistently reaches $10\times$ of their direct compression across all architectures, at which point at least nine in ten unpruned connections are effectively inactivated. Other methods (e.g., SNIP on VGG-19) remove entire layers early on, before any substantial differences between effective and direct compression emerge. Similarly, magnitude pruning applied after training disconnects most models very quickly and hence is not shown (e.g., LeNet-5 and VGG-19). In contrast, the picture is different with ResNet-50 whose residual connections might have allowed higher compression rates for this method. For example, we observe that the network is, in fact, two orders of magnitude less dense when direct sparsity reads just above 0.99. SNIP-iterative and especially SynFlow demonstrate a truly unique property: subnetworks pruned by these two algorithms exhibit practically equal effective and direct compressions, and, in the case of SynFlow, disconnect only at very high compression rates. What makes them special? Both SynFlow and SNIP-iterative are multi-shot pruning algorithms that remove parameters over 100 and 300 iterations, respectively. SynFlow ranks connections

by their ℓ_1 path norm (sum of weighted paths passing through the edge, where the weight of a path is the product of magnitudes of weights of its edges). SNIP uses connection sensitivity scores from Lee et al. (2019) $|\frac{\partial \mathcal{L}}{\partial \theta_i} \theta_i|$ as a saliency measure where \mathcal{L} is the loss function. Both these pruning criteria assign the lowest possible score of zero to inactive connections, scheduling them for immediate removal in the subsequent pruning iteration. Thus, by virtue of their iterative design, these two methods produce subnetworks with little to no difference between effective and direct compression. They are fortuitously designed to prune inactivated edges, which might explain their performance in high compressions.

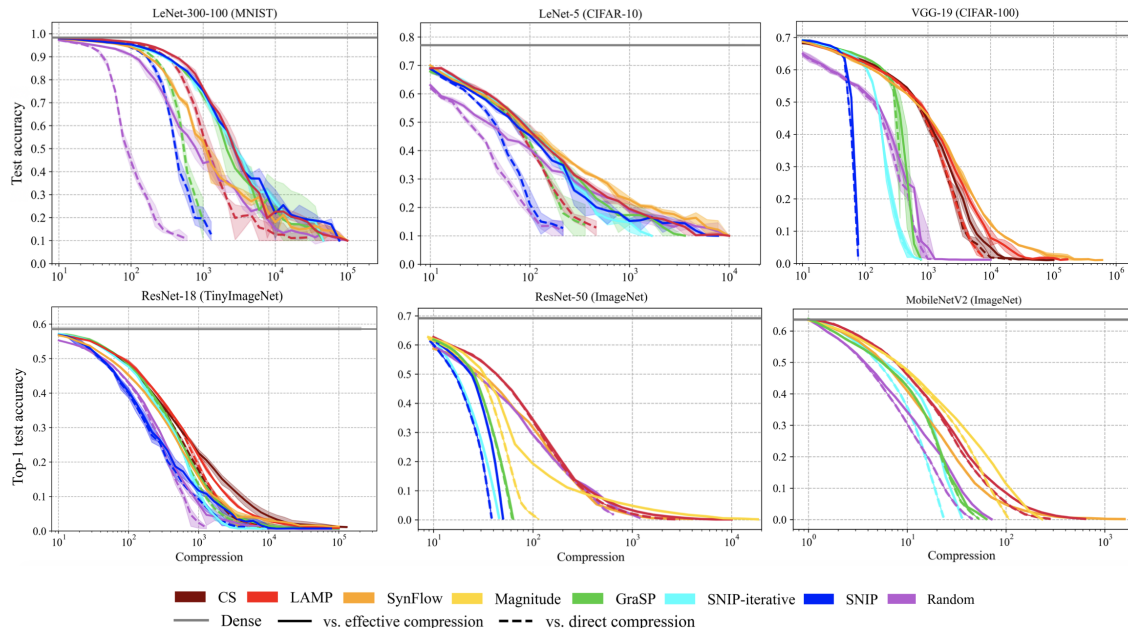


Figure 4: Test accuracy (min/average/max) of subnetworks trained from scratch after being pruned by different algorithms plotted against direct (dashed) and effective (solid) compression. Dashed and solid curves overlap for SynFlow and SNIP-iterative. Solid curves are fitted to scatter data (not shown for clarity of the presentation) as in Figure 2.

Tanaka et al. (2020) compare SynFlow to SNIP and GraSP using direct sparsity, claiming it vastly superior in high compression regimes. However, pruning methods that generate large amounts of inactivated connections are clearly at a significant disadvantage in the original direct framework. Figure 4 shows that the performance gap between SynFlow and other methods shrinks on all tested architectures under effective compression. The most dramatic changes are perhaps evident with LeNet-300-100 where SynFlow significantly dominates both SNIP and GraSP in direct comparison, but becomes strictly inferior when taken to the more telling effective compression. On the other hand, differences are not as pronounced on purely convolutional architectures such as VGG-19, and ResNet-18. Feature maps in convolutional layers are connected via groups of several parameters (kernels), making them more robust to inactivation compared to neurons in fully-connected layers.

Computing effective sparsity. In advocating the use of effective sparsity, we must make sure that it can be calculated efficiently. We propose an simple approach leveraging

SynFlow; note that a connection is inactive if and only if it is not part of any path from input to output. Assuming that unpruned weights are non-zero, this is equivalent to having zero ℓ_1 path norm. Tanaka et al. (2020) observe that path norms can be efficiently computed with one pass on the all-ones input as $|\frac{\partial \mathcal{R}}{\partial \theta_i}|$, where $\mathcal{R} = \mathbb{1}^\top f^*(|\Theta| \odot \mathbf{M}, \mathbb{1})$ and f^* is the linearized version of the original network f . For deep models, rescaling of weights may be required to avoid numerical instability (Tanaka et al., 2020).

4. Layerwise sparsity quotas (LSQ) and a novel allocation method (IGQ)

Inspired by Frankle et al. (2021) and Su et al. (2020), we wish to confirm that SNIP, GraSP, and SynFlow work no better than random pruning with corresponding layerwise sparsity allocation. While Frankle et al. (2021) and Su et al. (2020) only considered moderate compression rates up to $100\times$ and used direct sparsity as a reference frame, we reconfirm their conjecture in the effective framework and test it across the entire compression spectrum. We generate and train two sets of subnetworks: (i) pruned by either SNIP, GraSP, and SynFlow (*original*), and (ii) randomly pruned while preserving layerwise sparsity quotas provided by each of these three methods (*random*).

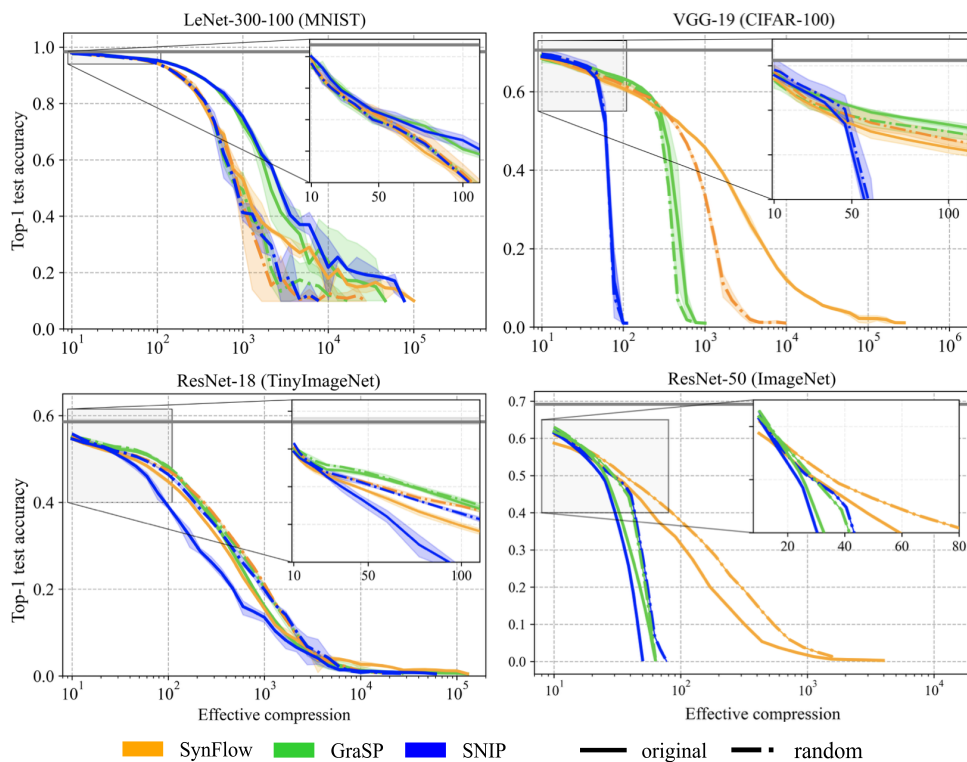


Figure 5: Original methods for pruning at initialization (solid) and random pruning with corresponding layerwise sparsity quotas (dashdot). Test accuracy of the unpruned network is shown in grey.

Our results in Figure 5 agree with observations made by Frankle et al. (2021) and Su et al. (2020): in the $10\times$ – $100\times$ compression range, all three random pruning algorithms perform

similarly (LeNet-300-100, VGG-19) or better (ResNet-18, ResNet-50) than their original counterparts. Effective sparsity allows us to faithfully examine higher compression, where the evidence is more equivocal. Similar patterns are still seen on ResNet-18; however, the original SNIP and GraSP beat random pruning with corresponding layerwise sparsities by a wide margin starting at $100\times$ compression on LeNet-300-100. Random pruning associated with SynFlow matches original SynFlow on the same network for longer, up to $1,000\times$ compression. On VGG-19, SynFlow bests the corresponding random pruning from about $500\times$ compression onward, while the original SNIP suffers from disconnection early on together with its random variant. Despite these nuances in the high compression regime, random pruning with specific layerwise sparsity quotas fares extremely well in the moderate sparsity regime (up to 99%) and is even competitive to full-fledged SynFlow (see Figure 8). Therefore, random pruning can be a cheap and competitive alternative to more sophisticated and resource-consuming algorithms. This phenomenon is also reconfirmed in a recent study, which states that randomly pruned networks with carefully crafted LSQ can match the performance of their dense counterparts while comparing favorably in terms of adversarial robustness, out-of-distribution detection, and uncertainty estimation (Liu et al., 2022b). In particular, they consider LSQ derived from SNIP and find it among the best performing sparsity distributions for random pruning. Alas, SNIP and other methods from Figure 5 require expensive computations just to retrieve the corresponding pruning ratios, which may still suffer from issues like layer-collapse. This motivates us to ask: can we engineer readily computable and consistently well-performing sparsity quotas?

To our knowledge, there are only a few *ab-initio* approaches in the literature to allocate sparsity in a principled fashion. *Uniform* is the simplest solution that keeps sparsity constant across all layers. Gale et al. (2019) give a modification (denoted *Uniform+* following Lee et al. (2021)) that retains all parameters in the first convolutional layer and caps sparsity of the last fully-connected layer at 80%. A more sophisticated approach, *Erdős-Rényi-Kernel (ERK)*, sets the density of a convolutional layer with kernel size $w \times h$, fan-in n_{in} and fan-out n_{out} proportional to $(w + h + n_{\text{in}} + n_{\text{out}})/(w \cdot h \cdot n_{\text{in}} \cdot n_{\text{out}})$. Although originally used as a sparsity distribution schema for methods with dynamic sparse structures (SET by Mocanu et al. (2018) and RigL by Evci et al. (2020)), we follow Lee et al. (2021) and use ERK as a baseline sparsity distribution for sparse-to-sparse training with a fixed subnetwork topology. The last two approaches are unable to support the entire range of sparsities: *Uniform+* can only achieve moderate *direct* compression because of the prunability constraints on its first and last layer, while both direct and effective sparsity levels achievable with ERK are often lower bounded. For example, the density of certain layers of VGG-16 set by ERK exceeds 1 when cutting less than 99% of parameters, unless excessive density is redistributed. Su et al. (2020) propose Smart-Ratios, which is an ad-hoc distribution method that requires the density of the i -th layer within an L -layer network to be proportional to $(L - l + 1)^2 + (L - l + 1)$. This method was developed exclusively for VGG-like networks and, like ERK and *Uniform+*, can be infeasible for certain sparsities.

To avoid problems that riddle *Uniform+*, ERK, and smart-ratios, we require that any layerwise sparsity quotas must be attainable for any level of network sparsity $s \in [0, 1]$. At the same time, neither layer should be removed in its entirety unless $s = 1$ to avoid layer-collapse inherent to SNIP and some other global pruning methods. These requirements lead

us to formulate a formal definition for layerwise sparsity quotas to guide principled future research into sparsity allocation.

Definition 1 (Layerwise Sparsity Quotas). A function $\mathcal{Q}: [0, 1] \rightarrow [0, 1]^L$ mapping a target sparsity s to layerwise sparsities $\{s_\ell\}_{\ell=1}^L$ is called *Layerwise Sparsity Quotas (LSQ)* if it satisfies the following properties: (i) *total sparsity*: for any $s \in [0, 1]$, $s \sum_\ell |\Theta_\ell| = \sum_\ell s_\ell |\Theta_\ell|$, and (ii) *layer integrity*: for all layers $\ell \in [L]$, $[\mathcal{Q}(s)]_\ell < 1$ if $s < 1$.

Ideal Gas Quotas (IGQ). Aiming to unfold the secret of well-performing layerwise compression quotas associated with such global pruning algorithms as SNIP, LAMP, and SynFlow, we note that they prune larger, parameter-heavy layers more aggressively than smaller layers (Figure 7), which has been already conjectured to be a desirable property (Su et al., 2020). To design a valid LSQ with this feature, we consult an intuitive (although lacking formal connection with neural network pruning) analogy from physics. In particular, we interpret compression of a multi-layer network as compression of stacked gas-filled weightless cylinders of unit volume and height equal to the size of the corresponding layer (Figure 6). As force is applied to the system, the Ideal Gas Law governs the compression rate of each cylinder, giving the final compression distribution which we interpret as the layerwise compression (sparsity) distribution within the given network. Using simple algebra, we arrive at compression quotas $\{F|\Theta_\ell| + 1\}_{\ell=1}^L$ (or sparsity quotas $\{1 - (F|\Theta_\ell| + 1)^{-1}\}_{\ell=1}^L$) parameterized by the force F that controls the overall sparsity of the network. Thus, we selected cylinder dimensions to encode our prior belief that larger layers can withstand higher pruning rates since “flatter” cylinders undergo lighter compression under the same external force (compression constraint). Given a target sparsity s , the needed value of F can be instantly found using binary search to any given precision. IGQ clearly satisfies all requirements of Definition 1 and applies layerwise higher compression to larger layers, as de-

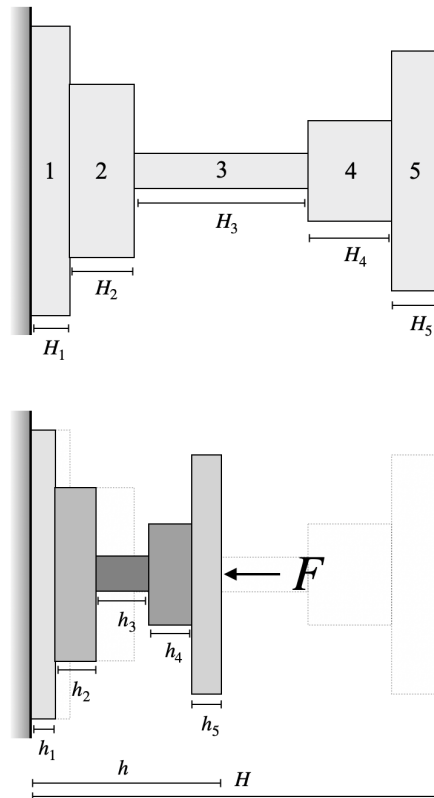


Figure 6: Schematic diagram of the cylinder system underlying IGQ compression of LeNet-5. Each of the five layers of the network is represented by a cylinder of unit volume and height H_ℓ proportional to the number of parameters $|\Theta_\ell|$ in that layer. As force F is applied to the outermost cylinder (5), the system transforms according to the Ideal Gas Law, yielding IGQ compression rates of H_ℓ/h_ℓ , while the overall network compression is $H/h = \sum_{\ell=1}^5 H_\ell / \sum_{\ell=1}^5 h_\ell$. Darker colors indicate higher compression. Note that cylinders are not drawn to scale.

sired. In principle, IGQ is applicable in a variety of contexts with use-cases in pruning before training (in conjunction with random pruning), during training (e.g., as default LSQ for RigL (Evci et al., 2020)), and after training (e.g., together with magnitude pruning). In this study, we adopt the first and the last scenarios to evaluate IGQ against baselines (Figures 8, 9).

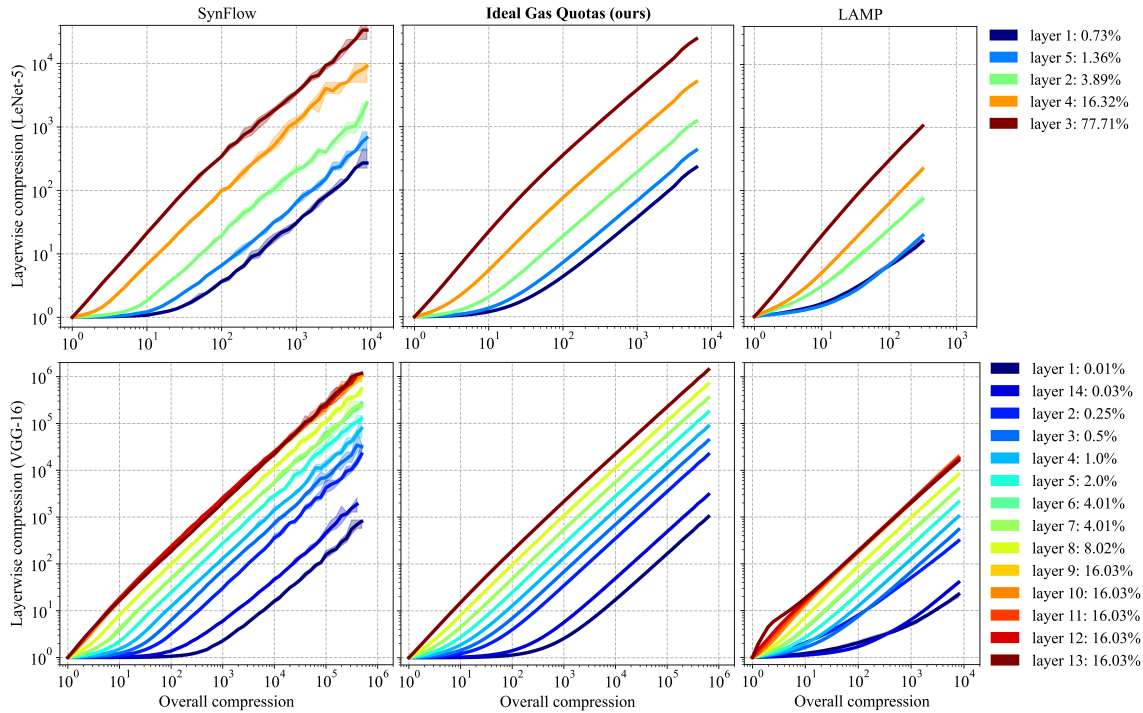


Figure 7: Layerwise direct compression quotas of LeNet-5 (top) and VGG-16 (bottom) associated with SynFlow (left), our IGQ (middle), and LAMP (right). Percentages indicate layer sizes relative to the total number of parameters; colors are assigned accordingly from blue (smaller layers) to red (larger layers). Curves of LAMP and SynFlow end when the underlying network disconnects.

Random pruning with IGQ. While Liu et al. (2022b) experiment with lower sparsities (up to 90%) and a slightly different set of LSQ, our results largely match their evidence. In particular, we also find that ERK consistently outperforms more naive baselines like Uniform and Uniform+. Although ERK sometimes exhibits similar (ResNet-18) or even better (VGG-19 compressed to 1,000× or higher) performance than IGQ, it yields invalid layerwise sparsity quotas when removing less than 98% and 99% of parameters from ResNet-18 and VGG-19, respectively, thus failing to satisfy Definition 1. Uniform+ produces invalid layerwise compressions from 40× onward for ResNet-50. In the moderate sparsity regime (up to 99%), subnetworks pruned by IGQ reach unparalleled performance after training, especially on ResNet-50. Across all architectures, random pruning with IGQ and SynFlow sparsity quotas are almost indistinguishable from each other, suggesting that IGQ successfully mimics the quotas produced by SynFlow, which require substantial effort to compute.

Therefore, judging by a tripartite criterion of test performance, compliance with Definition 1, and computational efficiency, IGQ beats all baselines.

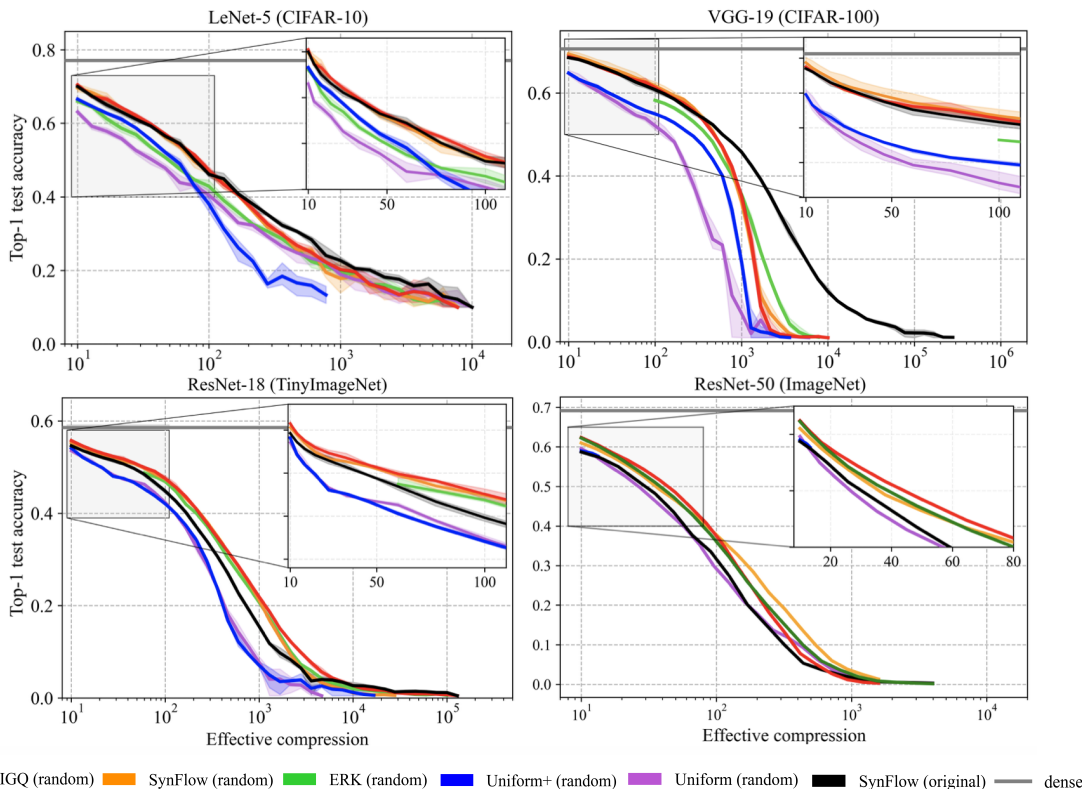


Figure 8: Test performance of trained subnetworks after random pruning with different layerwise sparsity distributions. Original SynFlow (black) is shown for reference.

Magnitude pruning with IGQ. In the second set of experiments, we pretrain fully-dense models and prune them by magnitude using global methods (Global Magnitude Pruning, LAMP) or layer-by-layer respecting sparsity allocation quotas (Uniform, Uniform+, ERK, and IGQ). Then, we revert the unpruned weights back to their original random values and fully retrain the resulting subnetworks to convergence. Results are displayed in Figure 9 in the framework of effective compression. Overall, our method for distributing sparsity in the context of magnitude pruning performs consistently well across all architectures and favorably compares to other baselines, especially in moderate compression regimes of $100\times$ or less. Even though Global magnitude pruning can marginally outperform IGQ, it is completely unreliable on VGG-19. ERK appears slightly better than IGQ on VGG-19, ResNet-18 and ResNet-50 at extreme sparsities, however, it performs much worse on LeNet-5 and has other general deficiencies as discussed earlier. Another close rival of IGQ is LAMP, which performs very similarly but is still unable to reach its performance on VGG-19, ResNet-18 and ResNet-50 in moderate compression regimes. Note, however, that all presented methods require practically equal compute and time; thus, the evidence in Figure 9 is not meant to advertise IGQ as a cheaper alternative to LAMP but rather to illustrate the effectiveness of IGQ.

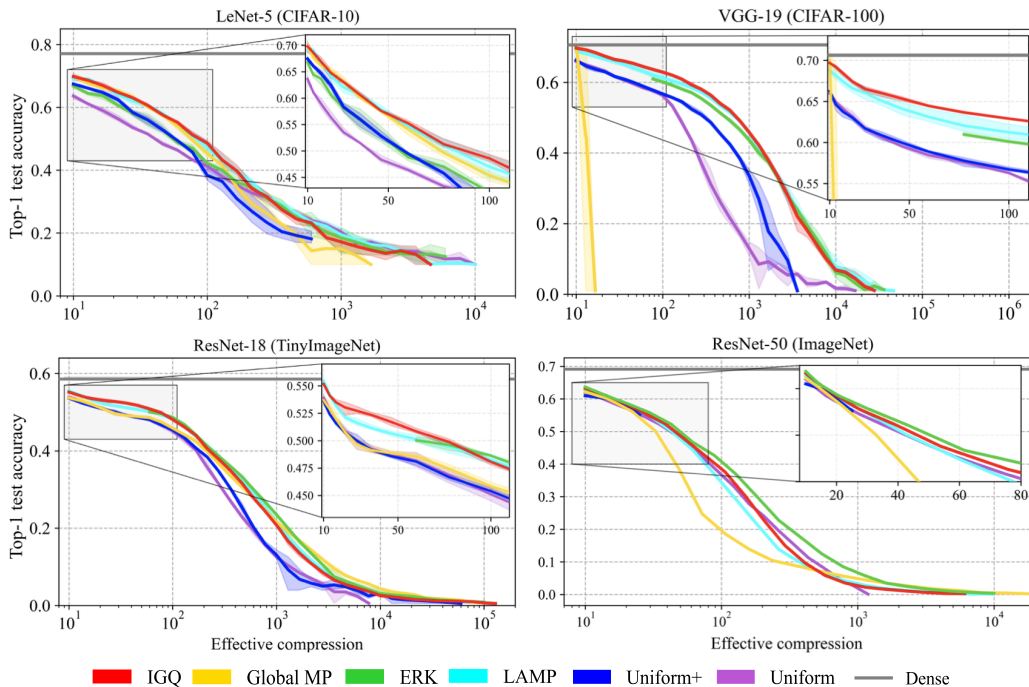


Figure 9: Test performance of retrained subnetworks after magnitude-based pruning. Uniform+ is not shown for LeNet-300-100 since it is designed for convolutional networks.

5. Effective pruning

Unlike pruning to a target direct sparsity, pruning to achieve a particular *effective* sparsity can be tricky. Here, we present an extension to algorithms for pruning at initialization or after training that achieves this goal efficiently, when possible (see Figure 10).

Effective ranking-based pruning. Algorithms like GraSP, SynFlow, and LAMP rank parameters by some notion of importance to guide pruning. When such a ranking $R: \Theta \rightarrow \mathbb{R}$ is available, we employ binary search for the appropriate cut-off threshold t in $\mathcal{O}(\log |\Theta|)$ time. This approach leverages the following monotonicity property: given two pruning thresholds $t_1, t_2 \in R$ and corresponding subnetworks S_1, S_2 , we have $t_1 \leq t_2$ if and only if $S_2 \subseteq S_1$, which implies $\text{EffectiveSparsity}(S_1) \leq \text{EffectiveSparsity}(S_2)$ (note that in general $\text{Sparsity}(S_1) \leq \text{Sparsity}(S_2)$ does not imply the last inequality above). Thus, binary search will branch in the correct direction.

Effective random pruning. In Section 4, we saw that random pruning with carefully crafted layerwise sparsity quotas $\mathcal{Q}: [0, 1] \rightarrow [0, 1]^L$ fares well (especially in the framework of effective sparsity) with more sophisticated pruning methods, proving to be a cheaper and simpler alternative. Effective pruning without parameter scores is more challenging because there is no obvious way to produce a neat chain of embedded subnetworks as above. For example, given two subnetworks S_1 and S_2 , $\text{Sparsity}(S_1) \leq \text{Sparsity}(S_2)$ does not imply $\text{EffectiveSparsity}(S_1) \leq \text{EffectiveSparsity}(S_2)$. Assigning random scores requires $\mathcal{O}(|\Theta|)$ time to ensure that any cut-off threshold yields LSQ according to \mathcal{Q} , which is not scalable.

To circumvent this issue, we design an improved algorithm that produces embedded subnetworks on each iteration, allowing binary search to work (see Algorithm 1). Starting from the extreme subnetworks S_1 (fully-dense, corresponding to masks $\mathbf{M}^{(1)}$) and S_2 (fully-sparse, corresponding to masks $\mathbf{M}^{(2)}$), we narrow the sparsity gap between them while preserving $S_2 \subseteq S_1$ so that $\text{EffectiveSparsity}(S_1) \leq \text{EffectiveSparsity}(S_2)$. For each layer, we keep track of unpruned connections U_ℓ of S_1 and pruned connections P_ℓ of S_2 , randomly sample parameters T_ℓ from $U_\ell \cap P_\ell$ according to \mathcal{Q} and form another network S by pruning out $\bigcup_\ell T_\ell$ from S_1 (or, equivalently, reviving in S_2). Depending on where effective sparsity of S lands relative to target s , we assign S to either S_1 or S_2 and branch. Since connections to be pruned from S_1 (or revived in S_2) are chosen randomly at each step, weights within the same layer have equal probability of being pruned. Once S_1 and S_2 are only 1 parameter away from each other, the algorithm returns S_1 , yielding a connected model. Note that this algorithm implicitly requires the LSQ function \mathcal{Q} to be layerwise monotone: if $s_1 \leq s_2$, then $[\mathcal{Q}(s_1)]_\ell \leq [\mathcal{Q}(s_2)]_\ell$ for each layer $\ell \in [L]$. This is a reasonable assumption and is satisfied in practice (see Figure 7).

6. Discussion

In our work, we argue that *effective sparsity* (*effective compression*) is the correct benchmarking measure for pruning algorithms since it discards effectively inactive connections and represents the true remaining connectivity pattern. Moreover, effective sparsity allows us to study extreme compression regimes for subnetworks that otherwise appear disconnected at much lower direct sparsities. We initiate the study of current pruning algorithms in this refined frame of reference and rectify previous benchmarks. To facilitate the use of effective sparsity in future research, we describe low-cost procedures to both compute and achieve desired effective sparsity when pruning. Lastly, with effective sparsity allowing us to zoom more fairly into higher compression regimes than previously possible, we examine random pruning with prescribed layerwise sparsities and propose our own readily computable quotas (IGQ) after establishing conditions reasonable LSQ should fulfill. We show that IGQ, while allowing for any level of sparsity, is more advantageous than all existing similar baselines (Uniform, ERK) and gives comparable performance to sparsity quotas derived from more sophisticated and computationally expensive algorithms like SynFlow.

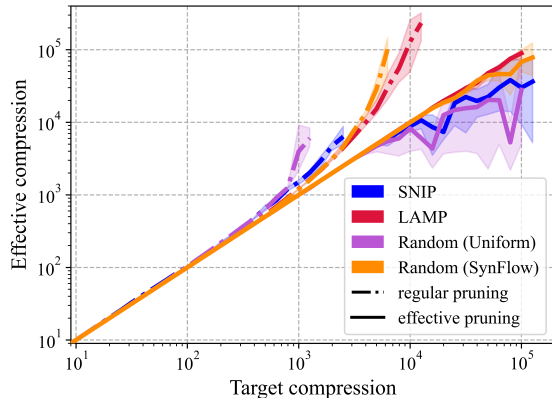


Figure 10: Effective compression produced by regular (dashdot) and our effective (solid) pruning on ResNet-18 according to ranking-based (left) and random (right) algorithms. Our procedures help pruning reach target effective sparsity, falling short only when the subnetwork is on the brink of disconnection.

Algorithm 1: Approximate Effective Random Pruning

Input: Desired effective sparsity s ; LSQ function $\mathcal{Q}: [0, 1] \rightarrow [0, 1]^L$.
 $i \leftarrow 0$; $j \leftarrow |\Theta|$; $\mathbf{M}^{(1)} \leftarrow \mathbf{1}$; $\mathbf{M}^{(2)} \leftarrow \mathbf{0}$; $P_\ell, U_\ell \leftarrow \Theta_\ell$ for all $\ell \in [L]$;
while $j - i > 1$ **do**
 $m \leftarrow \lfloor (i + j)/2 \rfloor$; $\{s_\ell\}_{\ell=1}^L \leftarrow \mathcal{Q}(m/|\Theta|)$;
 for $\ell \in [L]$ **do**
 CurrSparsity $_\ell \leftarrow (1 - |U_\ell|/|\Theta_\ell|)$;
 $T_\ell \leftarrow \text{RandomSelect}(\text{from} = U_\ell \cap P_\ell, \text{size} = |\Theta_\ell|(s_\ell - \text{CurrSparsity}_\ell))$;
 $M_\ell \leftarrow \text{CreateMask}(\text{pruned} = \Theta_\ell \setminus [U_\ell \setminus T_\ell], \text{unpruned} = U_\ell \setminus T_\ell)$;
 end
 $\mathbf{M} \leftarrow \{M_\ell\}_{\ell=1}^L$;
 if $\text{EffectiveSparsity}(\mathbf{M}) < s$ **then**
 $U_\ell \leftarrow U_\ell \setminus T_\ell$ for all $\ell \in [L]$; $\mathbf{M}^{(1)} \leftarrow \mathbf{M}$; $i \leftarrow m$;
 else
 $P_\ell \leftarrow P_\ell \setminus T_\ell$ for all $\ell \in [L]$; $\mathbf{M}^{(2)} \leftarrow \mathbf{M}$; $j \leftarrow m$;
 end
end
Return: Masks $\mathbf{M}^{(1)}$ s.t. $\text{EffectiveSparsity}(\mathbf{M}^{(1)}) \sim s$, $\|M_\ell^{(1)}\|_0 = |\Theta_\ell|(1 - [\mathcal{Q}(s)]_\ell)$.

Limitations and Broader Impacts. We hope that the lens of effective compression will spur more research in high compression regimes. One possible limitation is that it is harder to control effective compression exactly. In particular using different seeds might lead to slightly different effective compression rates. However, these perturbations are minor. Additionally, one might argue that for some architectures accuracy drops precipitously with higher compression thus making very sparse subnetworks less practical. We hope that opening the study of high compressions will allow to explore how to use sparse networks as building blocks, for instance using the power of ensembling (Liu et al., 2022a). Our framework allows a principled study of this regime. Finally, since effective compression strips away unnecessary computational units, it offers a potentially higher resource efficiency during both inference and training without compromising the flexibility of unstructured pruning or requiring specialized hardware (Liu et al., 2019).

Acknowledgments

Both authors were supported by the National Science Foundation under NSF Award 1922658. Neither of the authors has any competing interests to report.

Appendix A. Experimental details

Our experimental work encompasses seven different architecture-dataset combinations: LeNet-300-100 (Lecun et al., 1998) on MNIST (Creative Commons Attribution-Share Alike 3.0 license), LeNet-5 (Lecun et al., 1998) and VGG-16 (Simonyan and Zisserman, 2015) on CIFAR-10 (MIT license), VGG-19 (Simonyan and Zisserman, 2015) on CIFAR-100 (MIT

license), and ResNet-18 (He et al., 2016) on TinyImageNet (MIT license), ResNet-50 and MobileNetV2 (Howard et al., 2017) on ImageNet-2012 (Deng et al., 2009). Following Frankle et al. (2021), we do not reinitialize subnetworks after pruning (we revert back to the original initialization when pruning a pretrained model by LAMP). We use our own implementation of all pruning algorithms in TensorFlow except for GraSP, for which we use the original code in PyTorch published by Wang et al. (2020). All non-ImageNet runs were repeated 3 times for stability of results. Training was performed on an internal cluster equipped with NVIDIA RTX-8000, NVIDIA V-100, and AMD MI50 GPUs. Hyperparameters and training schedules used in our experiments are adopted from related works and are listed in Table 1. We apply standard augmentations to images during training. In particular, we normalize examples per-channel for all datasets and randomly apply: (i) shifts by at most 4 pixels in any direction and horizontal flips (CIFAR-10, CIFAR-100, and TinyImageNet), (ii) rotations by up to 4 degrees (MNIST), and (iii) random 224×224 crops and horizontal flips (ImageNet).

Model	Epochs	Drop epochs	Batch	LR	Decay	Source
LeNet-300-100	160	41/83/125	100	0.1	$5e-4$	Lee et al. (2019)
LeNet-5	307	76/153/230	128	0.1	$5e-4$	Lee et al. (2019)
VGG-16	160	80/120	128	0.1	$1e-4$	Frankle et al. (2021)
VGG-19	160	80/120	128	0.1	$5e-4$	Wang et al. (2020)
ResNet-18	200	100/150	256	0.2	$1e-4$	Frankle et al. (2021)
ResNet-50	90	30/60/80	512	0.4	$1e-4$	Frankle et al. (2021)
MobileNetV2	90	30/60/80	512	0.4	$1e-4$	Frankle et al. (2021)

Table 1: Summary of experimental work. All architectures include batch normalization layers followed by ReLU activations. Models are initialized using Kaiming normal scheme (fan-avg) and optimized by SGD (momentum 0.9) with a stepwise LR schedule ($10\times$ drop factor applied on specified *drop epochs*). The categorical cross-entropy loss function is used for all models.

Appendix B. Experiments with VGG-16

In Figure 11, we display the results of our experiments with VGG-16 on CIFAR-10. As we argued in Section 3, higher sparsities are required for purely convolutional architectures (such as VGG-16) to develop inactive connections since feature maps are harder to disconnect. At the same time, several algorithms (SNIP, SNIP-iterative, GraSP) suffer from layer-collapse at modest sparsities (99.9% or less) and, hence, fail to develop significant amounts of inactive parameters. For this reason, as evident from Figures 3, 4, and 11, VGG-16 arguably showcases the least differences between effective and direct compression among all tested architectures.

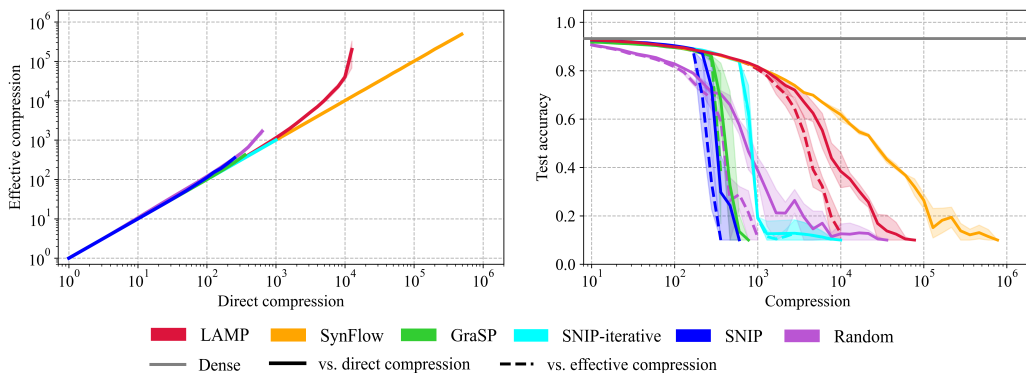


Figure 11: Left: effective versus direct compression of VGG-16 when pruned by different algorithms. Right: test accuracy (min/average/max) of VGG-16 trained from scratch after being pruned by different algorithms plotted against direct (dashed) and effective (solid) compression. Dashed and solid curves overlap for SynFlow and SNIP-iterative.

References

- Sanjeev Arora, Nadav Cohen, and Elad Hazan. On the optimization of deep networks: Implicit acceleration by overparameterization. In Andreas Krause and Jennifer Dy, editors, *35th International Conference on Machine Learning, ICML 2018*, 35th International Conference on Machine Learning, ICML 2018, pages 372–389. International Machine Learning Society (IMLS), jan 2018.
- Cristian Bucilua, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '06*, page 535–541, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595933395. doi: 10.1145/1150402.1150464. URL <https://doi.org/10.1145/1150402.1150464>.
- Wenlin Chen, James Wilson, Stephen Tyree, Kilian Weinberger, and Yixin Chen. Compressing neural networks with the hashing trick. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2285–2294, Lille, France, 07–09 Jul 2015. PMLR. URL <http://proceedings.mlr.press/v37/chenc15.html>.
- Pau de Jorge, Amartya Sanyal, Harkirat Behl, Philip Torr, Grégory Rogez, and Puneet K. Dokania. Progressive skeletonization: Trimming more fat from a network at initialization. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=9GsF0UyUPi>.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.

- Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/file/2afe4567e1bf64d32a5527244d104cea-Paper.pdf>.
- Tim Dettmers and Luke Zettlemoyer. Sparse networks from scratch: Faster training without losing performance. *CoRR*, abs/1907.04840, 2019. URL <http://arxiv.org/abs/1907.04840>.
- Xin Dong, Shangyu Chen, and Sinno Pan. Learning to prune deep neural networks via layer-wise optimal brain surgeon. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/c5dc3e08849bec07e33ca353de62ea04-Paper.pdf>.
- Utku Evci, Trevor Gale, Jacob Menick, Pablo Samuel Castro, and Erich Elsen. Rigging the lottery: Making all tickets winners. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 2943–2952. PMLR, 2020. URL <http://proceedings.mlr.press/v119/evci20a.html>.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=rJl-b3RcF7>.
- Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Pruning neural networks at initialization: Why are we missing the mark? In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Ig-VyQc-MLK>.
- Trevor Gale, Erich Elsen, and Sara Hooker. The state of sparsity in deep neural networks. *arXiv e-prints*, arXiv:1902.09574, 2019. URL <https://arxiv.org/abs/1902.09574>.
- Yunchao Gong, Liu Liu, Ming Yang, and Lubomir D. Bourdev. Compressing deep convolutional networks using vector quantization. *CoRR*, abs/1412.6115, 2014. URL <http://arxiv.org/abs/1412.6115>.
- Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL <https://proceedings.neurips.cc/paper/2015/file/ae0eb3eed39d2bcef4622b2499a05fe6-Paper.pdf>.
- Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. In Yoshua Bengio and

- Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1510.00149>.
- B. Hassibi, D.G. Stork, and G.J. Wolff. Optimal brain surgeon and general network pruning. In *IEEE International Conference on Neural Networks*, pages 293–299 vol.1, 1993. doi: 10.1109/ICNN.1993.298572.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi: 10.1109/CVPR.2016.90.
- Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015. URL <http://arxiv.org/abs/1503.02531>.
- Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR. URL <http://proceedings.mlr.press/v37/ioffe15.html>.
- Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. In *Proceedings of the British Machine Vision Conference*. BMVA Press, 2014. doi: <http://dx.doi.org/10.5244/C.28.88>.
- Aditya Kusupati, Vivek Ramanujan, Raghav Somani, Mitchell Wortsman, Prateek Jain, Sham Kakade, and Ali Farhadi. Soft threshold weight reparameterization for learnable sparsity. In *Proceedings of the International Conference on Machine Learning*, July 2020.
- Vadim Lebedev and Victor Lempitsky. Fast convnets using group-wise brain damage. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2554–2564, 2016. doi: 10.1109/CVPR.2016.280.
- Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- Yann LeCun, John S. Denker, and Sara A. Solla. Optimal brain damage. In *Advances in Neural Information Processing Systems*, pages 598–605. Morgan Kaufmann, 1990.
- Jaeho Lee, Sejun Park, Sangwoo Mo, Sungsoo Ahn, and Jinwoo Shin. Layer-adaptive sparsity for the magnitude-based pruning. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=H6ATjJ0TKdf>.

- Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Snip: single-shot network pruning based on connection sensitivity. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=B1VZqjAcYX>.
- Namhoon Lee, Thalaiyasingam Ajanthan, Stephen Gould, and Philip H. S. Torr. A signal propagation perspective for pruning neural networks at initialization. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=HJeTo2VFwH>.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=rJqFGTslg>.
- Shiwei Liu, Lu Yin, Decebal Constantin Mocanu, and Mykola Pechenizkiy. Do we actually need dense over-parameterization? in-time over-parameterization in sparse training. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 6989–7000. PMLR, 18–24 Jul 2021.
- Shiwei Liu, Tianlong Chen, Zahra Atashgahi, Xiaohan Chen, Ghada Sokar, Elena Mocanu, Mykola Pechenizkiy, Zhangyang Wang, and Decebal Constantin Mocanu. Deep ensembling with no overhead for either training or testing: The all-round blessings of dynamic sparsity. In *International Conference on Learning Representations*, 2022a. URL <https://openreview.net/forum?id=RLtqs6pzj1->.
- Shiwei Liu, Tianlong Chen, Xiaohan Chen, Li Shen, Decebal Constantin Mocanu, Zhangyang Wang, and Mykola Pechenizkiy. The unreasonable effectiveness of random pruning: Return of the most naive baseline for sparse training. In *International Conference on Learning Representations*, 2022b.
- 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. doi: 10.1109/ICCV.2017.298.
- Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. In *ICLR*, 2019.
- Christos Louizos, Max Welling, and Diederik P. Kingma. Learning sparse neural networks through l_0 regularization. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=H1Y8hhg0b>.
- Decebal Constantin Mocanu, Elena Mocanu, Peter Stone, Phuong H. Nguyen, Madeleine Gibescu, and Antonio Liotta. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature Communications*, 9(1):2383, 2018. doi: 10.1038/s41467-018-04316-3. URL <https://doi.org/10.1038/s41467-018-04316-3>.

- Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variational dropout sparsifies deep neural networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2498–2507. PMLR, 06–11 Aug 2017. URL <http://proceedings.mlr.press/v70/molchanov17a.html>.
- Hesham Mostafa and Xin Wang. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. In *Proceedings of the 36th International Conference on Machine Learning*, pages 4646–4655. PMLR, 2019.
- Sharan Narang, Greg Diamos, Shubho Sengupta, and Erich Elsen. Exploring sparsity in recurrent neural networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=BylSPv9gx>.
- Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. The role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=BygfghAcYX>.
- Ilan Price and Jared Tanner. Dense for the price of sparse: Improved performance of sparsely initialized networks via a subspace offset. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8620–8629. PMLR, 18–24 Jul 2021.
- Pedro Savarese, Hugo Silva, and Michael Maire. Winning the lottery with continuous sparsification. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11380–11390. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/83004190b1793d7aa15f8d0d49a13eba-Paper.pdf>.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *CoRR*, abs/1909.08053, 2019. URL <http://arxiv.org/abs/1909.08053>.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1409.1556>.
- Jingtong Su, Yihang Chen, Tianle Cai, Tianhao Wu, Ruiqi Gao, Liwei Wang, and Jason D Lee. Sanity-checking pruning methods: Random tickets can win the jackpot. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20390–20401. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/eae27d77ca20db309e056e3d2dcd7d69-Paper.pdf>.

Hidenori Tanaka, Daniel Kunin, Daniel L Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6377–6389. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/46a4378f835dc8040c8057beb6a2da52-Paper.pdf>.

Chaoqi Wang, Guodong Zhang, and Roger Grosse. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=SkgsACVKPH>.

Xiao Zhou, Weizhong Zhang, Hang Xu, and Tong Zhang. Effective sparsification of neural networks with global sparsity constraint. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3599–3608, June 2021.

Michael Zhu and Suyog Gupta. To prune, or not to prune: Exploring the efficacy of pruning for model compression. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=Sy1iIDkPM>.