

Winning the Lottery Ahead of Time: Efficient Early Network Pruning

John Rachwan¹ Daniel Zügner¹ Bertrand Charpentier¹ Simon Geisler¹ Morgane Ayle¹
Stephan Günnemann¹

Abstract


Pruning, the task of sparsifying deep neural networks, received increasing attention recently. Although state-of-the-art pruning methods extract highly sparse models, they neglect two main challenges: (1) the process of finding these sparse models is often very expensive; (2) unstructured pruning does not provide benefits in terms of GPU memory, training time, or carbon emissions. We propose Early Compression via Gradient Flow Preservation (EarlyCroP), which efficiently extracts state-of-the-art sparse models before or early in training addressing challenge (1), and can be applied in a structured manner addressing challenge (2). This enables us to train sparse networks on commodity GPUs whose dense versions would be too large, thereby saving costs and reducing hardware requirements. We empirically show that EarlyCroP outperforms a rich set of baselines for many tasks (incl. classification, regression) and domains (incl. computer vision, natural language processing, and reinforcement learning). EarlyCroP leads to accuracy comparable to dense training while outperforming pruning baselines.

1. Introduction

State-of-the-art deep learning typically operates in the over-parametrized regime. However, a large body of literature has shown that a high number of carefully chosen parameters can be removed (i.e. pruned) while maintaining the network’s predictive performance (LeCun et al., 1990; Molchanov et al., 2017; Evci et al., 2020; Su et al., 2020; Lee et al., 2019; Wang et al., 2020).

It was first believed that sparse networks obtained from pruning pre-trained networks cannot be retrained from scratch. However, (Frankle & Carbin, 2019) presented

¹Technical University Munich, Germany. Correspondence to: John Rachwan <john.rachwan@tum.de>.



the earliest point at which we can prune without losing performance is challenging, and the authors present Linear Mode Connectivity (LMC), a computationally very expensive approach involving training multiple copies of the network. This suggests that to achieve a good trade-off between performance and efficiency of finding these subnetworks, pruning should follow the same spirit as before training methods but be applied *early* instead.

Besides the question of *when* to prune, an orthogonal dimension is structured vs. unstructured pruning. Unstructured pruning prunes individual weights (i.e., sets weight matrix elements to zero), while structured pruning removes entire neurons (i.e., rows/columns of weight matrices or convolution filters). Thus, structured pruning can reduce the training/inference time, memory footprint, and carbon emissions of the model; unstructured pruning has no significant impact on the above. On the other hand, structured pruning is generally much more challenging, and most previous works (including the LTH) perform unstructured pruning.

Unless we maintain the learning dynamics of a neural network, pruning will hinder the learning process. The learning dynamics of a feed forward neural network can be described through the Neural Tangent Kernel (NTK) which approximately remains constant after some epochs of training in networks (Goldblum et al., 2020). If we preserve the NTK while pruning, we expect the training process to not be affected. We develop a novel and principled pruning method that preserves the Gradient Flow (GF). By leveraging the close relation of the NTK to the GF, we show that we can prune a network while keeping the effect on the NTK minimal. Furthermore, we use the connection between GF and NTK to track when the learning dynamics become stable enough to perform *early* pruning.

We present Early Compression via Gradient Flow Preservation (EarlyCroP), a method for pruning a network early in training. EarlyCroP requires training the model only once, yet maintains the dense network’s performance at high levels of sparsity. Thus, in the unstructured setting, EarlyCroP is about 5 times less expensive than the LTH. In addition, our method can be applied *before* or *early* in training, and extended to *structured* pruning. Performing structured pruning before training provides a better accuracy/efficiency trade-off than most previous structured baselines, and enables us to train sparse networks *whose dense versions would not fit into the GPU*. Furthermore, EarlyCroP reduces carbon emissions by up to 70% without affecting dense performance and can thus help mitigate the environmental impact of deep learning and reduce training and inference costs.

Contributions. We approach neural network pruning with the explicit goal of unlocking real-world, practical improvements. Our key contributions are:

- **Why to prune?** We transfer a GF based pruning criterion to be applicable for structured pruning, which allows faster forward and backward passes using less GPU memory and computational cost, while surpassing baselines in test accuracy;
- **How to prune?** We leverage a connection between the NTK and GF by using a pruning criterion that aims to minimally affect the GF, and therefore the NTK and learning dynamics;
- **When to prune?** We further utilize the connection between GF and NTK to indicate the smooth transition to the *lazy kernel regime*, the phase during which we can prune the network with little effect on the training dynamics. Thus, this brings the cost saving of structured pruning during training as well. We also show that our method can be applied before training, reducing costs even further at only a small drop in accuracy.

These contributions unlock substantial real-world benefits for practitioners and researchers: we can train large sparse models on commodity GPUs whose dense counterparts would be too large. We evaluate our approach extensively over a diverse set of model architectures, datasets, and tasks.

2. Related work

Pruning Criterion. In order to prune network weights, they need to be ranked according to an importance score. This concept is not new, in fact, it was introduced in ‘Optimal Brain Damage’ (LeCun et al., 1990) and ‘Optimal Brain Surgeon’ (Hassibi et al., 1993). Yet, it only regained traction when (Han et al., 2015) showed successful deep compression by pruning based on weight magnitude. Most pruning research since then has followed this approach (Zhou et al., 2019; Evci et al., 2020; Mostafa & Wang, 2019; Bellec et al., 2018; Dettmers & Zettlemoyer, 2019; Mocanu et al., 2018; You et al., 2020; Chen et al., 2020). However, the biggest drawback of using weight magnitudes is that the network needs to be trained first to achieve a good accuracy. Therefore, more recent works have focused on scoring weights without the need for training using first order (Lee et al., 2019; Tanaka et al., 2020) and second order (Wang et al., 2020; Lubana & Dick, 2021) information. Note that the pruning process can be applied in one-shot or iteratively (de Jorge et al., 2021; Verdenius et al., 2020).

Pruning Time. Up until the introduction of the LTH (Frankle & Carbin, 2019), the consensus in the literature was that pruned models cannot be trained from scratch. Therefore, all sparse networks were extracted either from pre-trained networks (Han et al., 2015; LeCun et al., 1990; Hassibi et al., 1993; Wang et al., 2019; Li et al., 2016), or through-out training (Srinivas & Babu, 2016; Louizos et al., 2018; Evci et al., 2020; Mostafa & Wang, 2019; Bellec et al.,

2018; Dettmers & Zettlemoyer, 2019; Mocanu et al., 2018). However, the LTH showed that there exist sparse models within the original randomly initialized dense model that can achieve comparable performance to the dense model. That being said, the LTH’s pruning algorithm, Iterative Magnitude Pruning (IMP), requires multiple iterations of a train-prune cycle. Nevertheless, the LTH’s findings motivated works that strove to extract these sparse networks directly from the randomly initialized dense network (Su et al., 2020; Lee et al., 2019; Wang et al., 2020; de Jorge et al., 2020; Frankle et al., 2020b; Verdenius et al., 2020).

The first method that introduced pruning before training was SNIP (Lee et al., 2019), with the goal of preserving weights that have the highest effect on the loss. A subsequent work, GraSP (Wang et al., 2020), uses the Hessian-gradient product in its score and prunes the weights with the goal of increasing the Gradient Flow (GF). Finally, Lubana & Dick (2021) show that GraSP can lead to an increasing loss and instead propose to prune the weights that least affect the GF.

The performance of the LTH degrades with bigger networks and datasets (Frankle et al., 2020c). Subsequently, the LTH was updated to indicate that the best performing sparse models do not necessarily exist at initialization but rather that they appear early in training.

To the best of our knowledge, the only work that explores the extraction of sparse models early in training is Early Bird Tickets (You et al., 2020). They perform structured pruning early in training when the Hamming Distance between pruning masks at subsequent epochs becomes smaller than some threshold. However, they do not offer any theoretical justification for pruning early in training and they only show their results for a maximum pruning ratio of 70%, suggesting that the Hamming Distance is not an effective criterion to achieve high sparsities.

The Early Phase of DNN Training. Another line of work aims to analyze the early phase of neural network training. Gur-Ari et al. (2018) study the Hessian eigenspectrum and observe that during training, a few large eigenvalues emerge in which gradient descent happens, whereas the rest get close to zero. However, these observations depend on the architecture. Achille et al. (2019) found that the network goes through critical training periods during which perturbing the data can cause irreversible damage to the network’s final performance, after which the network becomes robust to these perturbations. However, the critical periods occur very late in the training process. Finally, Frankle et al. (2020a) propose the Linear Mode Connectivity (LMC) as a method for detecting when networks become stable to SGD noise. However, LMC is extremely expensive, requiring to train two copies of a network to completion at every epoch.

Structured Pruning. Pruning methods are divided into

two categories: (1) unstructured methods which generate a binary mask that is applied before every forward pass (Frankle & Carbin, 2019; Lee et al., 2019; Tanaka et al., 2020; Wang et al., 2020), and (2) structured methods that remove entire neurons or convolutional filters (Ding et al., 2019; Li et al., 2016; Louizos et al., 2018; Verdenius et al., 2020; You et al., 2019). Unstructured pruning is the more common variant for its simplicity and ease of implementation. However, since the entire dense network remains the same size, pruning does not provide improvements in GPU RAM, time, and carbon emissions. While these improvements can be obtained for unstructured pruning by using operations on sparse compressed matrices, they require significant changes to the network when dealing with advanced layers. Conversely, structured pruning reduces the size of weight matrices, thereby requiring less space, time and energy during training and inference. We highlight: (1) SNAP (Verdenius et al., 2020), which adapts the SNIP (Lee et al., 2019) score to the structured setting to prune before training, (2) Gate Decorators (You et al., 2019), which builds on top of (Liu et al., 2017) by adding a sensitivity-based criterion and pruning the network iteratively during training, and (3) EfficientConvNets (Li et al., 2016), which prunes a pre-trained network by scoring filters by their L_1 -norm. Other recent works include (He et al., 2020; 2019; Lym et al., 2019) but we omit them since GateDecorators outperforms them.

3. Background

Neural Tangent Kernel (NTK). The NTK is defined as $g_Y(\Theta_t)^T g_Y(\Theta_t)$ (Jacot et al., 2018) where $g_Y(\Theta_t)$ denotes the gradient of the model prediction Y w.r.t. the model parameters Θ_t at time t . The NTK is known to accurately describe the *dynamics of the network’s prediction during training*, under the assumption that the following Taylor expansion holds:

$$Y(\Theta_t) \approx Y(\Theta_0) + g_Y(\Theta_0)^T (\Theta_t - \Theta_0) \quad (1)$$

Under the NTK assumption, a neural network reduces to a linear model with the Neural Tangent Kernel (NTK). The NTK assumption is particularly accurate for wide neural networks. In practice, this assumption holds (i.e. the NTK remains approximately constant) after the model training dynamic has transitioned from the *rich active regime* to the *lazy kernel regime* (see Section 4.3).

Gradient Flow (GF). We define the GF as $g_L(\Theta_t)^T g_L(\Theta_t)$ (Lubana & Dick, 2021), where $g_L(\Theta_t)$ denotes the gradient of the model loss L w.r.t. the model parameters Θ_t . The GF is known to accurately describe the *dynamics of the network’s gradient norm during training*, under the

assumption that the following Taylor expansion holds:

$$g_L(\Theta_t)^T g_L(\Theta_t) = \|g_L(\Theta_t)\|_2^2 \quad (2)$$

$$\approx \|g_L(\Theta_0)\|_2^2 + 2(H_L(\Theta_0)g_L(\Theta_0))^T(\Theta_t - \Theta_0) \quad (3)$$

where $H_L(\Theta_t)$ denote the model’s Hessian at time t . In order to prune the weights which least affect the GF, [Lubana & Dick \(2021\)](#) propose to use the following weight importance score:

$$I(\Theta_t) = |\Theta_t^T H_L(\Theta_t) g_L(\Theta_t)|, \quad (4)$$

and remove $\rho\%$ of the parameters with the lowest scores. Preserving the GF stands in stark contrast to the importance score of GraSP ([Wang et al., 2020](#)) $I_{\text{GraSP}}(\Theta_t) = -\Theta_t^T H_L(\Theta_t) g_L(\Theta_t)$, that maximizes the GF. Note that while the importance score (4) was initially used before training, we propose to use this importance score to prune during training either in one-shot or iteratively.

4. Method

The core motivation of our work is to improve the applicability of sparse neural networks w.r.t. concrete real-world metrics such as carbon emissions, price, time or memory at both training and inference time. To this end, our method first transfers the pruning criterion (4) to structured pruning, thus allowing faster forward and backward passes (see Sec. 4.1). Second, we derive a relation between the NTK and the GF suggesting that preserving the GF also preserves the NTK (see Sec. 4.2). Hence, the pruning criterion (4) is a suitable importance weight score for pruning a neural network once the NTK assumption holds i.e. in the *lazy kernel regime*. Third, our method detects when we enter the *lazy kernel regime* to prune early in training without impacting the training dynamics (see Sec. 4.3), thereby extending the cost saving of our (structured) sparse neural networks to the training phase while achieving a high test accuracy.

4.1. Why to prune?

The main use case of unstructured pruning is to highlight the overparametrized nature of neural networks. In particular, while dense-like sparsity for deep learning ([Zhou et al., 2021](#)) is a promising research direction, it suffers from multiple downsides: **(a)** models are typically only sparsified in the forward pass and hence dense-like sparsity has limited potential in speeding up training. **(b)** Not all deep learning frameworks (e.g. PyTorch) support it. **(c)** Only the newest GPUs (starting Nvidia Ampere 2020) support dense-like sparsity.

In order to really benefit from pruning, we need to *prune full structures* (neurons and channels) instead. This reduction in dimensions/channels directly translates into lower computational cost on existing GPUs *without* further

implementation efforts or any usage of any specialized tensor operations. Thus, this leads to a sparse model that provides improvements in time, memory and carbon emissions. Combined with the fact that we can apply our pruning method before and early in training, we can drastically reduce not only model costs after training but during training as well (see Sec. 4.3).

In order to achieve structured pruning, we need to score entire nodes instead of individual weights, i.e. generate a score for a node’s activation function f_l . However, since f_l is simply a function and not a learnable parameter, we cannot use pruning score (4) directly. Instead, similarly to [Verdenius et al. \(2020\)](#), we define auxiliary gates of a layer l by $c_l = 1$ over each node’s input, which in turn will act as a learnable parameter whose gradient information represents the activation’s information. We can formally define this for a linear layer l with weight Θ_l and bias b_l , and an input x in the following way:

$$f_l(\Theta_l * x + b_l) = f_l(c_l(\Theta_l * x + b_l)) \quad (5)$$

$$I(f_l) = |H_L(c_l)g_L(c_l)| \quad (6)$$

After scoring the nodes using the auxiliary gates, the pruning process follows the original one by removing $\rho\%$ of nodes.

4.2. How to prune?

In this section, we draw an important connection between GF and NTK showing that pruning the weights with the lowest importance score (4) aims at preserving the training dynamics of both the network’s gradient norm and the network’s prediction during training. First, we observe that GF and NTK are connected by the following relation:

$$GF = g_L(\Theta_t)^T g_L(\Theta_t) \quad (7)$$

$$= g_L(Y)^T g_Y(\Theta_t)^T g_Y(\Theta_t) g_L(Y) \quad (8)$$

$$= g_L(Y)^T NTK g_L(Y) \quad (9)$$

Second, [Lubana & Dick \(2021\)](#) present evidence that preserving the GF also implicitly preserves the model loss $L(\Theta_t)$. In particular, preserving the GF also preserves the gradient of the loss w.r.t. the prediction $g_L(Y)$. Hence, the relation (7) and the preservation $g_L(Y)$ imply that the NTK is also preserved when the GF is preserved.

Furthermore, given that Taylor expansions (1) and (3) hold, the pruning criteria (4) which preserves the GF – which maintains the gradient-norm dynamics – is also preserving the NTK – which maintains the prediction dynamics. This remark is crucial since while the dynamics of the neural network’s predictions during training can be approximated well by (1) during the *lazy kernel regime*, this approximation might not be accurate during the *rich active regime*.

4.3. When to prune?

First, we show in an introductory experiment that the pruning time has an important impact on the final accuracy of the pruned model. Hence, we train multiple ResNet50 models on CIFAR100 and prune each to 98% from epoch 0 (i.e. before pruning) to epoch 80 (see Fig.2). We observe that (1) the longer we train the dense model before pruning, the higher the final accuracy of the sparse model, and most importantly (2) after a certain point in time, further training of the dense model before pruning does not bring significant improvement on the final accuracy. Indeed, we observe a 3% improvement in the final accuracy of the sparse model when pruning at epoch 1 of training instead of before training, an 11% improvement when pruning at epoch 26, and no great improvement when pruning after epoch 30.

We now introduce the pruning time detection score used by EarlyCroP which is motivated from both practical and theoretical perspectives. EarlyCroP aims to detect the best time for pruning in two steps: (1) at every epoch we compute the pruning time score

$$\Delta_0^t = \|\Theta(t) - \Theta(0)\|^2, \quad (10)$$

and, (2) if the difference of the scores at two subsequent epochs relative to the initial score Δ_0^1 is smaller than a defined threshold th ,

$$\frac{|\Delta_0^t - \Delta_0^{t-1}|}{|\Delta_0^1|} < th \quad (11)$$

we run the EarlyCroP pruning algorithm described in Algorithm 1. It can be clearly deduced that the smaller the pruning time score, the more negligible the second order term in eq. 3 will be, making the latter a good approximation. Additionally, by the triangle inequality, we can extract the following upper bound from 11

$$\frac{|\Delta_0^t - \Delta_0^{t-1}|}{|\Delta_0^1|} \leq \frac{\|\Theta(t) - \Theta(t-1)\|}{\|\Theta(1) - \Theta(0)\|} \quad (12)$$

which is expected to lie in $[0, 1]$ when weights change less significantly over time. Hence, the scale of the threshold is expected to be similar for different models and datasets. The complexity of computing the score is $\mathcal{O}(m)$ where m is the number of model parameters, thus incurring only minor computation overhead at every epoch to detect the pruning time.

From a theoretical perspective, the pruning time detection algorithm’s goal is to detect when the linearization of the prediction dynamics (1) assumed by the NTK holds during training. In the early phase of training called *rich active regime*, neural network parameters move to a significant distance from the initial weights. Thus the linearization of the prediction dynamics (1) usually does not hold in

early training epochs and the NTK quickly changes. This *rich active regime* is crucial to achieve high performance, in particular for deep models (Woodworth et al., 2020). In the second phase of training called *lazy kernel regime*, the parameters move by a small distance, thus making the linearization (1) a good approximation of the training dynamics of the predictions (Sun, 2019). Since our importance weight score (4) assumes the linearization (1), the best moment to prune is when the model transitions to the *lazy kernel regime* during which the NTK is approximately constant. Further, previous works (Sun, 2019; Ichi Amari, 2020; Ghorbani et al., 2020) showed that constancy of the NTK is a consequence of a constant weight norm during training. The transition to the lazy kernel training regime is gradual and can be detected when the relative change in the weight norm from initialization Δ_0^t becomes roughly constant i.e. when $\frac{|\Delta_0^t - \Delta_0^{t-1}|}{|\Delta_0^1|}$ becomes very close to 0.

In practice, we expect the pruning time criteria to be a reliable indicator of the final accuracy of the pruned model. Indeed, we observed that the detection score correlates well with the final test accuracy of the sparse model (see Fig.2). It can be clearly seen that the smaller the detection score at the moment of pruning, the higher the final test accuracy of the pruned model. In practice, we observed in Fig. 2 and in Fig. 8 in the appendix that pruning to higher sparsities benefits more from longer training. Therefore, we use $th = 1 - \rho$ which connects the time pruning threshold th to the target sparsity ρ . As desired, a higher target sparsity ρ leads to longer dense training (see Fig.8)




Figure 2. ResNet50 pruned to sparsity 98% at different points during training on CIFAR10 plotted against the final test accuracy of the pruned model after 200 epochs of training (purple) and the difference between the Relative Weight Change in two subsequent epochs (blue). \star denotes when the Relative Weight Change in two subsequent epochs is below the threshold $th = 1 - \rho$.

5. Empirical Evaluation

We now show the effectiveness of our EarlyCroP for structured pruning (EarlyCroP-S) and unstructured pruning

(EarlyCroP-U). For this we determine the point for early pruning as described in Section 4.3. We use CroP as our pruning criterion before training and CroPit if pruning is additionally performed iteratively. We use a cloud instance of GTX 1080 TIs for all experiments. Further details about the experimental setup can be found in the appendix. The code and further supplementary material is available online¹.

Image Classification. The datasets used for Image Classification are the common public benchmarks CIFAR10 (Krizhevsky & Hinton, 2009), CIFAR100 (Krizhevsky & Hinton, 2009), and Tiny-Imagenet (Deng et al., 2009). Regarding networks, we use ResNet18, VGG16, ResNeXt-101 32x16d and ResNeXt-101 32x48d. For unstructured pruning baselines, we use random pruning, SNIP (Lee et al., 2019), GraSP (Wang et al., 2020), and LTR (Frankle et al., 2020a). For structured pruning baselines, we use random pruning, EfficientConvNets (Li et al., 2016), GateDecorators (You et al., 2019), and SNAP (Verdenius et al., 2020). All models are trained for 80 epochs, except LTR which re-trains the network up to 10 times. We report train and test accuracy, weight and node sparsity, batch and total training time in seconds, GPU memory in GB, disk size in MB, carbon emitted from the extraction and training of the sparse model using CodeCarbon (Schmidt et al., 2021) in grams. Note that total training time includes the time to find and train the sparse model.

Regression. We evaluate a Fully Convolutional Residual Network (Laina et al., 2016) on the NYU Depth Estimation task (Nathan Silberman & Fergus, 2012). We compare EarlyCroP-S and EarlyCroP-U against all unstructured baselines since they are stronger than the structured baselines. All pruned models are trained for 10 epochs. We report the performance using the Root Mean Squared Error (RMSE).

Natural Language Processing (NLP). We evaluate the Pointer Sentinel Mixture Model (Merity et al., 2017) on the PTB language modeling dataset (Marcus et al., 1993). We compare EarlyCroP-S and EarlyCroP-U to all unstructured baselines since they are stronger than the structured baselines. We train the pruned models for 30 epochs and we report the achieved log perplexity.

Reinforcement Learning (RL). We use the FLARE framework (Akbik et al., 2019) to evaluate a simple 3-layer FCNN with layer size 256 using the A2C algorithm on the classic control game CartPole-v0 (Brockman et al., 2016). We run 20 agents with 640 games each. We compare our EarlyCroP-S and EarlyCroP-U against LTR and Random baselines. All pruned models are trained for 30 epochs. We report the performance of the pruned models using the average returned environment reward.

¹www.cs.cit.tum.de/daml/early-crop/

5.1. Image Classification

Accuracy. We present the accuracy over different sparsity levels for the model-dataset combinations ResNet18/ CIFAR10, ResNet18/Tiny-Imagenet, VGG16/CIFAR10, and VGG16/CIFAR100 in Figure 3 a-d, respectively. Our methods EarlyCroP-S and EarlyCroP-U consistently outperform all other methods but the LTR, where we perform on par. However, as will be discussed later, the LTR comes with a 3-5 times higher training time than the dense model while our methods reduce training time. There are two further exceptions when it comes to the best accuracy on ResNet18/Tiny-ImageNet. First, GateDecorators performs as well as EarlyCroP-S. Second, for lower sparsity rates, EarlyCroP-U is outperformed by some methods that prune before training.

Structured vs. Unstructured. EarlyCroP-S closes the accuracy gap of existing approaches between unstructured and structured pruning on CIFAR10 dataset. However, a gap remains for the larger and more complex datasets CIFAR100 and Tiny-Imagenet. Nevertheless, structured pruning can be used to reduce the training time and memory requirements. This also implies that with our EarlyCroP-S we can use a larger model while saving compute (see Sec. 5.2).

Training cost. In Tables 1 & 2 we complement the accuracy with the training time, batch time, GPU RAM, Disk space and CO₂ emissions for a sparsity of 95% on ResNet18/CIFAR10 and 98% on VGG16/CIFAR10 respectively. We see that our EarlyCroP-S is not only preserving the high accuracy but also comes with significant improvement in training time (33% and 36% resp.) and time per batch (39% and 61% resp.). It is as efficient as the other structured pruning methods or outperforms them. When only considering the CO₂ footprint, CroPit-S and SNAP outperform methods that prune later in training. In the appendix, we also give details about different model-dataset combinations. In summary, the stated observations also hold for the other evaluated model-dataset combinations.

Pruning early vs. before. Pruning early in training (i.e. when we enter the *lazy kernel regime*) outperforms pruning before training. Our EarlyCroP-S and EarlyCroP-U have a clear edge over the methods that prune before training. This is even true if we use the same pruning criterion (CroP) and additionally prune iteratively (CroPit). The only drawback of pruning early in training vs. before is that for the first epochs we either require more GPU RAM or need to reduce the batch size. For the model size on disk we do not find a significant difference among pruning methods.

5.2. Pruning a Large Model

The goal of this experiment is two-fold: we show that (1) our criterion can be used to prune large models that don't

Winning the Lottery Ahead of Time: Efficient Early Network Pruning

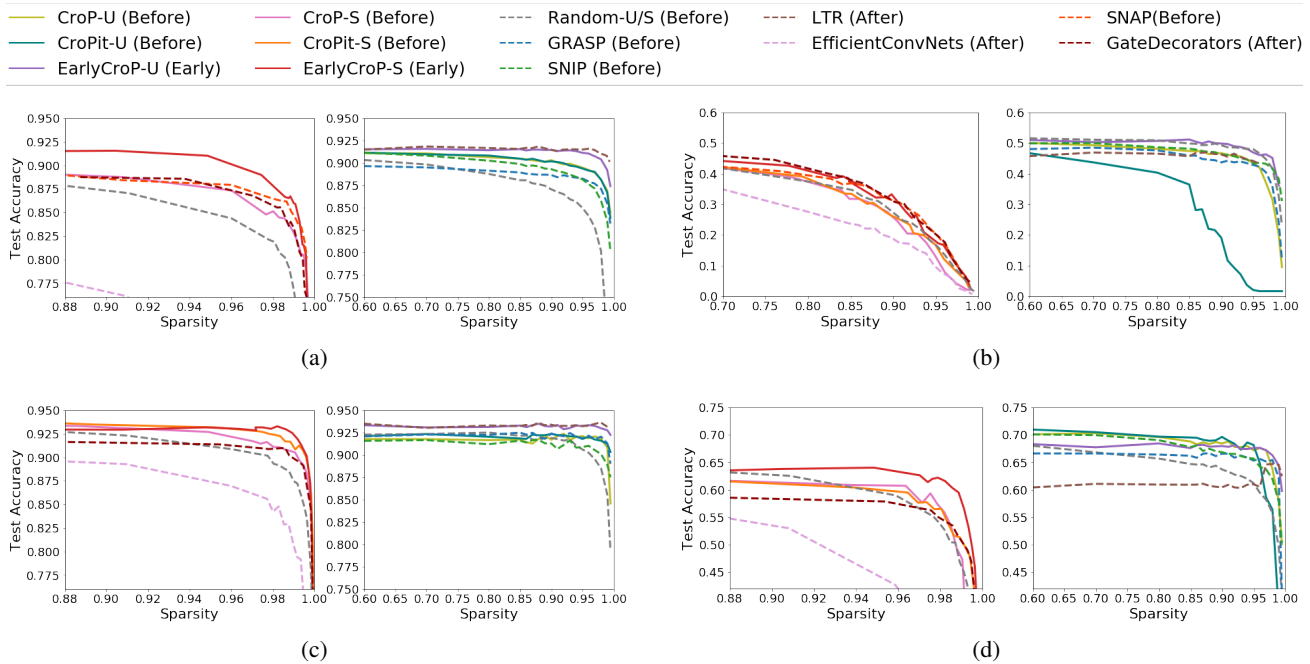


Figure 3. Structured (Left) and Unstructured (Right) test accuracy for ResNet18/CIFAR10 (a), ResNet18/Tiny-Imagenet (b), VGG16/CIFAR10 (c), and VGG16/CIFAR100 (d) with increasing weight sparsity.

Table 1. Comparison between different pruning criteria on ResNet18/CIFAR10 at 95% sparsity, averaged over three runs. Since for structured pruning we cannot precisely control the weight sparsity, we show results closest to 95% weight sparsity. \pm denotes standard deviation, and \uparrow/\downarrow indicate metrics where higher/lower is better. Bold/underline indicate best/second best results. GPU RAM and Disk correspond to those of the final pruned model.

	Method	Test accuracy \uparrow	Weight sparsity	Node sparsity	Training time (h) \downarrow	Batch time (ms) \downarrow	GPU RAM (GB) \downarrow	Disk (MB) \downarrow	Emissions (g) \downarrow
	Dense	91.5% \pm 0.12	-	-	0.78	109	2.38	398	83
Structured	Random-S	86.3% \pm 0.06	93.7%	75.0%	0.68	82	0.62	24.9	38
	SNAP	87.6% \pm 0.94	93.6%	72.6%	0.70	81	0.63	25.4	39
	CroP-S	87.5% \pm 0.36	93.6%	72.3%	0.67	91	0.63	25.4	43
	CroPit-S	<u>87.8%</u> \pm 0.33	95.0%	74.5%	<u>0.52</u>	<u>0.48</u>	0.59	19.6	35
	EarlyBird	84.3% \pm 0.32	95.3%	65.0%	0.48	72	0.58	19.1	55
	EarlyCroP-S	91.0% \pm 0.52	95.1%	65.8%	<u>0.52</u>	66	0.56	19.2	68
	GateDecorators	87.3% \pm 0.09	95.7%	73.7%	0.72	83	0.58	17.2	54
	EfficientConvNets	70.5% \pm 0.53	95.9%	79.7%	0.77	83	0.76	25.4	63
Unstructured	Random-U	84.9% \pm 0.24	95.0%	-	0.78	<u>102</u>	2.86	12.0	79
	SNIP	88.2% \pm 0.57	95.0%	-	0.79	105	2.86	12.0	80
	GRASP	88.4% \pm 0.13	95.0%	-	0.79	106	2.84	12.0	81
	CroP-U	87.9% \pm 0.16	95.0%	-	<u>0.75</u>	107	2.88	12.0	79
	CroPit-U	89.1% \pm 0.24	95.0%	-	0.80	113	2.88	12.0	87
	EarlyCroP-U	<u>91.1%</u> \pm 0.23	95.0%	-	0.74	97	2.86	12.0	83
	LTR	91.5% \pm 0.26	95.0%	-	1.94	111	2.51	12.0	202

Table 2. Comparison between different pruning criteria on VGG16/CIFAR10 at 98% sparsity. Since for structured pruning we cannot precisely control the weight sparsity, we show results closest to 98% weight sparsity. ↑/↓ indicate metrics where higher/ lower is better. Bold/ underline indicate best/ second best results. GPU RAM and Disk correspond to those of the final pruned model.

	Method	Test accuracy ↑	Weight sparsity	Node sparsity	Training time (h) ↓	Batch time (ms) ↓	GPU RAM (GB) ↓	Disk (MB) ↓	Emissions (g) ↓
-	Dense	90.2%	-	-	1.82	290	1.02	1720	246
Structured	Random-S	89.3%	98.0%	86.1%	0.67	82	0.23	33.6	43
	SNAP	89.8%	98.2%	89.0%	<u>0.68</u>	<u>89</u>	0.22	30	55
	CroP-S	91.1%	98.0%	88.0%	0.71	91	0.23	33.6	83
	CroPit-S	<u>92.4%</u>	98.0%	88.0%	0.81	112	0.23	30.4	100
	EarlyBird	85.9%	98%	89%	0.52	110	0.32	36.2	160
	EarlyCroP-S	93.0%	98.0%	89.0%	1.16	112	0.63	36.0	156
	GateDecorators	90.0%	98.0%	87.0%	1.07	111	0.23	37.8	143
	EfficientConvNets	84.2%	98.0%	86.0%	1.66	<u>89</u>	0.64	34.2	209
Unstructured	Random-U	88.5%	98.0%	-	2.03	159	1.22	35.0	247
	SNIP	90.1%	98.0%	-	<u>2.02</u>	157	1.22	35.0	248
	GRASP	92.0%	98.0%	-	2.03	157	1.23	35.0	249
	CroP-U	91.8%	98.0%	-	<u>2.02</u>	157	1.22	35.0	248
	CroPit-U	91.6%	98.0%	-	<u>2.02</u>	157	1.22	35.0	249
	EarlyCroP-U	<u>93.0%</u>	98.0%	-	2.01	157	1.22	35.0	250
	LTR	93.6%	98.0%	-	4.07	158	1.22	35.0	592

fit on commodity GPUs, and (2) the resulting sparse model matches the performance of the dense one, and outperforms a dense model of the same size. To this end, we introduce the ResNext101_32x48d as our large model, a network that has 829 Million parameters and requires 15.5 GB to be loaded into GPU memory; exceeding the memory of common GPUs such as the RTX 3080 Ti. Nevertheless, with our method we can still efficiently train such a large model. For this, we perform one initial pruning step before training using the CPU and then continue on a commodity GPU as usual. We also introduce the ResNext101_32x16d as our smaller dense model which has 193 Million parameters and requires 3.9 GB to be loaded into GPU memory. The results of the experiment are depicted in Table 3.

First, we observe that the large ResNext101_32x48d pruned to 98.5% weight sparsity matches the performance of its dense counterpart in test accuracy. Moreover, training the sparse subnetwork has a 14 times smaller carbon footprint, is 7 times faster to train, is 192 times smaller on disk, and takes 4.9 times less GPU memory than the large dense model. Interestingly, the pruned model also outperforms the ResNext101_32x16d model of the same size, while training 6.2 times faster and emitting 9.5 times less carbon. Finally, we show that when training for more epochs, the sparse model achieves an even bigger performance gap compared to both dense models while still taking less total training time. This experiment not only shows that CroP-S makes training large models on commodity machines possible, but can extract sparse models that are more efficient and more

accurate compared to dense models of the same size.

Table 3. Comparison of a pruned ResNext101_32x48d (RN48) model and a similar sized dense ResNext101_32x16d (RN16) model (CIFAR10). RN48-S are models pruned with CroP-S.

Model	Test acc.	Weight sparsity	Node sparsity	Epochs	Training time (h)	VRAM (GB)	Emissions (g)
RN48	92.4%	-	-	30	4.60	18.84	634
RN16	92.1%	-	-	30	4.02	3.89	445
RN48-S	92.5%	98.5%	89.9%	30	0.64	3.56	47
RN48-S	93.2%	98.5%	89.9%	80	2.60	3.56	194

5.3. Regression

Regarding Regression, we can see from Figure 5 that both variants of EarlyCroP preserve the dense model’s RMSE even at 99.9% weight sparsity. All Before training methods except GraSP have an instant increase to 0.20 RMSE with continuous decline at higher sparsities. Surprisingly, Random Pruning outperforms GraSP at all pruning ratios. This is due to GraSP pruning entire layers and limiting the network’s learning capabilities.

5.4. Natural Language Processing

NLP presents itself as the most challenging out of all evaluated tasks. Nevertheless, both versions of EarlyCroP outperform all other baselines until 89% sparsity (see Figure 6). Beyond that, the unstructured version is on par or slightly




Figure 4. Sparse model performance on the classic control game Cartpole-v0 in average reward (higher is better)

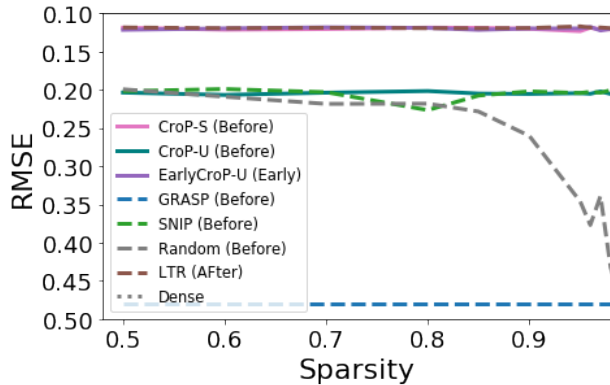


Figure 5. Sparse model performance on the NYU Depth Estimation task in RMSE (lower is better)

better than other unstructured baselines whereas the structured version continues to outperform all compared baselines. In this task, the importance of early pruning is accentuated by the large gap between the early and before versions of CroP. Interestingly, LTR performs very poorly compared to all other baselines on all reported pruning ratios. Indeed, certain layers in the PSMM network converge to small weight magnitudes during training compared to the rest of the network. This means that any pruning method relying solely on the weight magnitudes, and operating at a global scale in the network, would prune these layers entirely, leading to an untrainable network. Thus, this experiment highlights the importance of gradient-based information when evaluating the importance of model parameters. We show additional NLP results by evaluating BERT on multiple language tasks (see Appendix E.2).




Figure 6. Sparse model performance on the PTB language modelling task in log perplexity (lower is better)

5.5. Reinforcement Learning

We can observe from Figure 4 that EarlyCroP outperforms LTR in both the structured and unstructured setting. Note that the EarlyCroP-S once again outperforms its unstructured counterpart. However, if we allow the unstructured models to train for a longer time, they achieve similar performance to the structured version. This can be explained by the ease of training of structured models, which are still fully-connected models where all computed gradients contribute to the weight updates, whereas unstructured models compute gradients that are not used by the pruned weights, rendering the training slower.

6. Conclusion

We have demonstrated that, for vision, NLP, and RL tasks, EarlyCroP-U extracts winning tickets matching and often outperforming those found by the LTR by pruning early in training when the model enters *lazy kernel training*. Additionally, we showed that EarlyCroP-S outperforms other structured methods, providing the best trade-off between final test accuracy and efficiency in terms of time, space, and carbon emissions. Finally, we show that we can use CroP-S to train models that do not fit on commodity GPUs by extracting sparse models that preserve the initial model’s performance and outperform a similarly sized dense model for the same number of epochs. Thus, our methods bring tangible real-world benefits for researchers and practitioners. We hope that the results shown in this paper motivate more research on the study of structured pruning in the early phase of DNN training.

References

- Achille, A., Rovere, M., and Soatto, S. Critical learning periods in deep neural networks, 2019. 3
- Akbik, A., Bergmann, T., Blythe, D., Rasul, K., Schweter, S., and Vollgraf, R. Flair: An easy-to-use framework for state-of-the-art nlp. In *NAACL 2019, 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, 2019. 6
- Bellec, G., Kappel, D., Maass, W., and Legenstein, R. Deep rewiring: Training very sparse deep networks. In *International Conference on Learning Representations*, 2018. 2
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym, 2016. 6
- Buluç, A., Fineman, J. T., Frigo, M., Gilbert, J. R., and Leiserson, C. E. Parallel sparse matrix-vector and matrix-transpose-vector multiplication using compressed sparse blocks. In *Proceedings of the Twenty-First Annual Symposium on Parallelism in Algorithms and Architectures*, SPAA '09, 2009. 3
- Chen, T., Frankle, J., Chang, S., Liu, S., Zhang, Y., Wang, Z., and Carbin, M. The lottery ticket hypothesis for pre-trained bert networks. *Advances in Neural Information Processing Systems*, 2020. 2
- de Jorge, P., Sanyal, A., Behl, H. S., Torr, P. H., Rogez, G., and Dokania, P. K. Progressive skeletonization: Trimming more fat from a network at initialization. *arXiv preprint arXiv:2006.09081*, 2020. 3
- de Jorge, P., Sanyal, A., Behl, H., Torr, P., Rogez, G., and Dokania, P. K. Progressive skeletonization: Trimming more fat from a network at initialization. In *International Conference on Learning Representations*, 2021. 2
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009. 6, 2, 3
- Dettmers, T. and Zettlemoyer, L. Sparse networks from scratch: Faster training without losing performance. *arXiv preprint arXiv:1907.04840*, 2019. 2, 3
- Ding, X., Ding, G., Guo, Y., and Han, J. Centripetal sgd for pruning very deep convolutional networks with complicated structure. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- Evci, U., Gale, T., Menick, J., Castro, P. S., and Elsen, E. Rigging the lottery: Making all tickets winners. In *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research, 2020. 1, 2
- Frankle, J. and Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019. 1, 2, 3
- Frankle, J., Dziugaite, G. K., Roy, D., and Carbin, M. Linear mode connectivity and the lottery ticket hypothesis. In *Proceedings of the 37th International Conference on Machine Learning*, Proceedings of Machine Learning Research, 2020a. 1, 3, 6
- Frankle, J., Dziugaite, G. K., Roy, D. M., and Carbin, M. Pruning neural networks at initialization: Why are we missing the mark? *arXiv preprint arXiv:2009.08576*, 2020b. 3
- Frankle, J., Dziugaite, G. K., Roy, D. M., and Carbin, M. Stabilizing the lottery ticket hypothesis, 2020c. 3
- Ghorbani, B., Mei, S., Misiakiewicz, T., and Montanari, A. Linearized two-layers neural networks in high dimension, 2020. 5
- Goldblum, M., Geiping, J., Schwarzschild, A., Moeller, M., and Goldstein, T. Truth or backpropaganda? an empirical investigation of deep learning theory, 2020. 2
- Gur-Ari, G., Roberts, D. A., and Dyer, E. Gradient descent happens in a tiny subspace. *arXiv preprint arXiv:1812.04754*, 2018. 3
- Han, S., Pool, J., Tran, J., and Dally, W. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, 2015. 2
- Hassibi, B., Stork, D. G., and Wolff, G. J. Optimal brain surgeon and general network pruning. In *IEEE International Conference on Neural Networks*, 1993. 2
- He, Y., Liu, P., Wang, Z., Hu, Z., and Yang, Y. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 3
- He, Y., Dong, X., Kang, G., Fu, Y., Yan, C., and Yang, Y. Asymptotic soft filter pruning for deep convolutional neural networks. *IEEE Transactions on Cybernetics*, 50(8): 3594–3604, 2020. doi: 10.1109/TCYB.2019.2933477. 3

- ichi Amari, S. Any target function exists in a neighborhood of any sufficiently wide random network: A geometrical perspective, 2020. 5
- Jacot, A., Hongler, C., and Gabriel, F. Neural tangent kernel: Convergence and generalization in neural networks. In *NeurIPS*, 2018. 3
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015. 2
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. *Master's thesis, Department of Computer Science, University of Toronto*, 2009. 6, 2
- Laina, I., Rupprecht, C., Belagiannis, V., Tombari, F., and Navab, N. Deeper depth prediction with fully convolutional residual networks. In *3D Vision (3DV), 2016 Fourth International Conference on*. IEEE, 2016. 6
- LeCun, Y., Denker, J. S., and Solla, S. A. Optimal brain damage. In *Advances in Neural Information Processing Systems 2*. 1990. 1, 2
- Lee, N., Ajanthan, T., and Torr, P. SNIP: Single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations*, 2019. 1, 2, 3, 6
- Li, H., Kadav, A., Durdanovic, I., Samet, H., and Graf, H. P. Pruning filters for efficient convnets. 2016. 2, 3, 6
- Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., and Zhang, C. Learning efficient convolutional networks through network slimming, 2017. 3
- Louizos, C., Welling, M., and Kingma, D. P. Learning sparse neural networks through l_0 regularization. In *International Conference on Learning Representations*, 2018. 2, 3
- Lubana, E. S. and Dick, R. A gradient flow framework for analyzing network pruning. In *International Conference on Learning Representations*, 2021. 2, 3, 4
- Lym, S., Choukse, E., Zangeneh, S., Wen, W., Sanghavi, S., and Erez, M. Prunetrain: Fast neural network training by dynamic sparse model reconfiguration. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '19*, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362290. doi: 10.1145/3295500.3356156. URL <https://doi.org/10.1145/3295500.3356156>. 3
- Marcus, M. P., Marcinkiewicz, M. A., and Santorini, B. Building a large annotated corpus of english: The penn treebank. *Comput. Linguist.*, 1993. ISSN 0891-2017. 6
- Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. *ArXiv*, 2017. 6
- Mocanu, D. C., Mocanu, E., Stone, P., Nguyen, P. H., Gibescu, M., and Liotta, A. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature communications*, 2018. 2, 3
- Molchanov, P., Tyree, S., Karras, T., Aila, T., and Kautz, J. Pruning convolutional neural networks for resource efficient inference. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017. 1
- Mostafa, H. and Wang, X. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. In *Proceedings of the 36th International Conference on Machine Learning*, Proceedings of Machine Learning Research, 2019. 2
- Nathan Silberman, Derek Hoiem, P. K. and Fergus, R. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 6
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*. 2019. 2, 3
- Schmidt, V., Goyal, K., Joshi, A., Feld, B., Conell, L., Laskaris, N., Blank, D., Wilson, J., Friedler, S., and Lucioni, S. CodeCarbon: Estimate and Track Carbon Emissions from Machine Learning Computing. 2021. 6, 3
- Srinivas, S. and Babu, R. V. Generalized dropout. *arXiv preprint arXiv:1611.06791*, 2016. 2
- Su, J., Chen, Y., Cai, T., Wu, T., Gao, R., Wang, L., and Lee, J. D. Sanity-checking pruning methods: Random tickets can win the jackpot. *Advances in Neural Information Processing Systems*, 2020. 1, 3
- Sun, R. Optimization for deep learning: theory and algorithms, 2019. 5
- Tanaka, H., Kunin, D., Yamins, D. L., and Ganguli, S. Pruning neural networks without any data by iteratively conserving synaptic flow. *arXiv preprint arXiv:2006.05467*, 2020. 2, 3
- Verdenius, S., Stol, M., and Forré, P. Pruning via iterative ranking of sensitivity statistics. *arXiv preprint arXiv:2006.00896*, 2020. 2, 3, 4, 6

- Wang, C., Grosse, R. B., Fidler, S., and Zhang, G. Eigen-damage: Structured pruning in the kronecker-factored eigenbasis. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, Proceedings of Machine Learning Research, 2019. 2
- Wang, C., Zhang, G., and Grosse, R. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations*, 2020. 1, 2, 3, 4, 6
- Woodworth, B., Gunasekar, S., Lee, J. D., Moroshko, E., Savarese, P., Golan, I., Soudry, D., and Srebro, N. Kernel and rich regimes in overparametrized models. In *Proceedings of Thirty Third Conference on Learning Theory*, 2020. 5
- You, H., Li, C., Xu, P., Fu, Y., Wang, Y., Chen, X., Baraniuk, R. G., Wang, Z., and Lin, Y. Drawing early-bird tickets: Toward more efficient training of deep networks. In *International Conference on Learning Representations*, 2020. 2, 3
- You, Z., Yan, K., Ye, J., Ma, M., and Wang, P. Gate decorator: Global filter pruning method for accelerating deep convolutional neural networks, 2019. 3, 6
- Zhou, A., Ma, Y., Zhu, J., Liu, J., Zhang, Z., Yuan, K., Sun, W., and Li, H. Learning N:M Fine-grained Structured Sparse Neural Networks From Scratch. *9th International Conference on Learning Representations, {ICLR} 2021*, pp. 15, 2021. 4
- Zhou, H., Lan, J., Liu, R., and Yosinski, J. Deconstructing lottery tickets: Zeros, signs, and the supermask. In *Advances in Neural Information Processing Systems*, 2019. 2

Supplementary Material

A. Training and Inference Cost Computation

This section details the computations in Figure 1. We introduce the V100 16GB GPU (2.48\$/h) and the V100 32GB GPU (4.96\$/h). We use the total training time needed to train RN48 and RN48-S for 30 epochs each (see Table 3). Consequently, in order to train RN48, we need $4.96\$/h \times 4.6h = 22.816\%$. In order to train RN48-S, we need $2.48\$/h \times 0.64h = 1.59\%$

B. Algorithm

Algorithm 1 Early-CroP

```

1: Initialize the weights  $\Theta$ , weight mask  $Mask_0$  with all
   1s, pruning ratio  $\rho$ , maximum dense training time  $m$ ,
   score threshold  $th$ , and the number of pruning iterations
    $it$ ;
2:  $\Theta_0 = \Theta$ 
3:  $\Theta_{prev} = \Theta$ 
4: while  $t(epoch) < m$  do
5:   Do one SGD epoch;
6:    $diff = \frac{\|\Theta(t) - \Theta(0)\|^2}{\|\Theta(0)\|^2} - \frac{\|\Theta(t-1) - \Theta(0)\|^2}{\|\Theta(0)\|^2}$ 
7:   if  $diff \leq th$  then
8:     break;
9:   end if
10:   $\Theta_{prev} = \Theta$ 
11: end while
12: while  $i \leq it$  do
13:   $\rho_i = \rho_{final} - (\rho_{final} - 1/2) \times (1/2)^i$ 
14:   $I(\Theta) = |\Theta_t^T H_L(\Theta_t) g_L(\Theta_t)|$ 
15:   $Mask_i = CroP(I(\Theta), \rho_i)$ 
16:  Apply Mask on Network
17: end while
18: Apply Mask on Network
19: while  $t \leq t_{max}$  do
20:   Do one SGD epoch;
21: end while

```

C. Experimental Setup

C.1. Optimization

Image Classification. For all experiments, we use the ADAM(Kingma & Ba, 2015) optimizer and a learning rate of $2e-3$. The One Cycle Learning Rate scheduler is used to train all models except VGG16. The batch size used for CIFAR10 and CIFAR100 experiments is 256 while for Tiny-Imagenet it is 128. All sparse models are allowed to train the same amount of epochs (80) to converge which, except for LTR, includes the number of epochs required to extract the sparse model. In the case of LTR the final sparse

model is allowed to train for 80 epochs.

Regression. For all experiments, we use a batch size of 8, the ADAM optimizer with a learning rate of $1e-5$. All pruned models are trained for 10 epochs.

Natural Language Processing. For all experiments, we use a batch size of 128, and the ADAM optimizer with a learning rate of $1e-3$. All pruned models are trained for 30 epochs.

Reinforcement Learning. A description of the models used and number of runs used for each environment can be found in Table 4.

Table 4. A summary of the setup used in Reinforcement Learning

Name	Network	Algorithm	Agents	Games
CartPole-v0	MLP(128-128-128-out)	A2C	16	8000
Acrobot-v1	MLP(256-256-256-out)	A2C	16	8000
LunarLander-v2	MLP(256-256-256-out)	A2C	16	8000

C.2. Datasets Pre-Processing

CIFAR10 (Krizhevsky & Hinton, 2009). We augment the normalized CIFAR10 with Random Crop and Random Horizontal Flip. Images are additionally resized to 64×64 .

CIFAR100 (Krizhevsky & Hinton, 2009). We augment the normalized CIFAR100 with Random Crop, Random Horizontal Flip, and Random Rotation.

Tiny-Imagenet (Deng et al., 2009). We normalize the dataset and resize each image to 224×224 .

D. Evaluation Metrics

In this section we describe how specific metrics are calculated.

Time. We report time in two separate ways. First, we report the total time required (Training time), which is defined from the start of the experiment until the sparse model’s training is finished. Second, we report the time it takes to perform a full forward and backward pass (Batch time) on a given batch using the CUDA time measurement tool (Paszke et al., 2019).

GPU RAM. The RAM footprint of a process refers to how much memory it consumes on the GPU. This effectively includes the costs of loading the model and performing a

training step on it. We use the CUDA memory measurement tool to report this metric (Paszke et al., 2019).

Disk Storage. We estimate the storage needed to store a model on disk using the CSR sparse matrix format (Buluç et al., 2009). Similarly to (Verdenius et al., 2020), we used a ratio of 16:1 float precision on all vectors of the CSR format.

Energy Emissions. We estimate CO₂ emissions in g using CodeCarbon emissions tracker (Schmidt et al., 2021). These estimates consider all emissions from the start of experiments until the end of training.

over, for the same number of training epochs, EarlyCrop-S achieves 60.01% in 51 hours while the dense model achieves 58.78% in 62 hours.

E. Additional Results

E.1. Reinforcement Learning

See Figure 7 for experiments on the Acrobot-v1 and LunarLander-v2 environments.

E.2. Natural Language Processing

We evaluated BERT on multiple language tasks (see Table 8). At the same pruning sparsity, EarlyCroP-U outperforms LTR on 5 out of 8 tasks while training 10× faster.

E.3. Pruning Point Experiments

In Figure 8 we present more experiments on pruning models at different points in training. We can clearly observe a correlation between the desired pruning rate and the optimal time to prune. The higher the desired final sparsity, the longer the network should be trained before being pruned.

E.4. VGG16/CIFAR100

See Table 5 for comparisons between different pruning criteria at the same pruning level on VGG16/CIFAR100.

E.5. ResNet18/TinyImageNet

See Table 6 for comparisons between different pruning criteria at the same pruning level on ResNet18/Tiny-Imagenet.

E.6. VGG16/ImageNet

In Table 7 we present a comparison between a dense VGG16 and a sparse VGG16 pruned using EarlyCroP-S on the ImageNet-2012 (ILSVRC2012) (Deng et al., 2009) classification dataset. Given 62 hours of training on ImageNet-2012 (ILSVRC2012) on a single V100 GPU, EarlyCroP-S on VGG16 (50% weights pruned) achieves an accuracy of 61.43% while the dense model only achieves 58.78%. More-




Figure 7. Sparse model performance on classic control games (a) Acrobot-v1 and (b) LunarLander-2

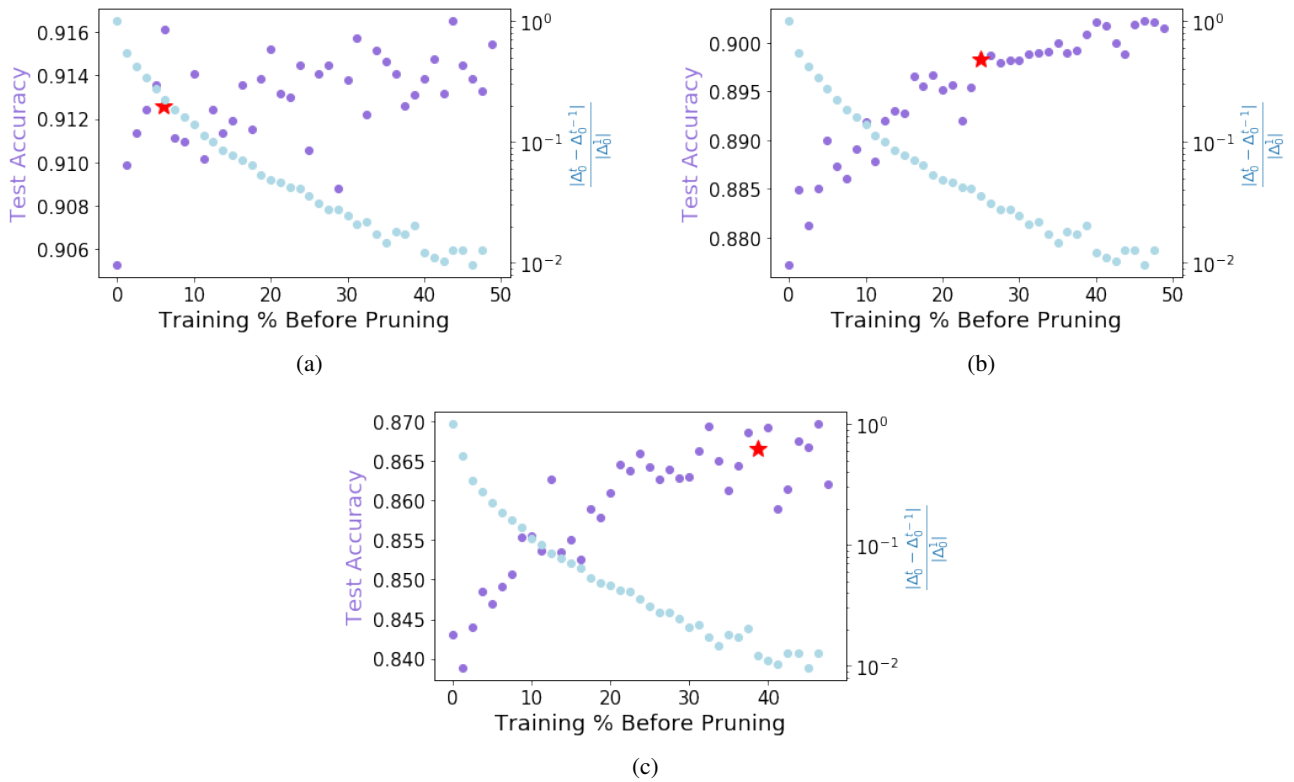


Figure 8. ResNet18 pruned to sparsity 80%(a), 98%(b), and 99.5% (c) at different points during training on CIFAR10 plotted against the final test accuracy of the pruned model after 80 epochs of training (purple) and the difference between the Relative Weight Change in two subsequent epochs (blue). All pruned models are trained for 80 epochs. ★ denotes when the Relative Weight Change in two subsequent epochs is below $th < 1 - \rho$

Winning the Lottery Ahead of Time: Efficient Early Network Pruning

Table 5. Comparison between different pruning criteria on VGG16/CIFAR100 at 98% sparsity. Since for structured pruning we cannot precisely control the weight sparsity, we show results closest to 98% weight sparsity. \uparrow/\downarrow indicate metrics where higher/ lower is better. Bold/ underline indicate best/ second best results. GPU RAM and Disk correspond to those of the final pruned model.

	Method	Test accuracy \uparrow	Weight sparsity	Node sparsity	Training time (h) \downarrow	Batch time (ms) \downarrow	GPU RAM (GB) \downarrow	Disk (MB) \downarrow	Emissions (g) \downarrow
-	Dense	62.1%	-	-	0.77	114	1.03	1745	88
Structured	Random-S	53.9%	98.0%	86.0%	0.59	53	0.23	35	29
	SNAP	49.3%	98.0%	89.0%	0.67	54	0.16	36	33
	CroP-S	<u>57.4%</u>	98.0%	89.0%	<u>0.61</u>	<u>46</u>	0.23	36	35
	CroPit-S	56.5%	98.1%	89.0%	0.62	44	0.23	33	30
	EarlyBird	60.7%	98.0%	89.0%	0.56	68	0.20	36	62
	EarlyCroP-S	62.2%	97.9%	88.0%	0.64	69	0.23	36	58
	GateDecorators	55.0%	97.9%	87.0%	<u>0.61</u>	78	0.23	36	68
	EfficientConvNets	29.5%	98.0%	86.0%	0.72	55	0.24	36	83
Unstructured	Random-U	55.8%	98.0%	-	0.74	118	1.23	35	99
	SNIP	61.9%	98.0%	-	0.79	109	1.24	35	90
	GRASP	63.4%	98.0%	-	0.79	113	1.24	35	91
	CroP-U	63.8%	98.0%	-	0.74	109	1.23	35	94
	CroPit-U	56.3%	98.0%	-	0.74	111	1.23	35	91
	EarlyCroP-U	65.1%	98.0%	-	0.74	109	1.23	35	91
	LTR	<u>64.7%</u>	98.0%	-	3.44	109	1.28	35	301

Table 6. Comparison between different pruning criteria on ResNet18/TinyImageNet at 90% sparsity. Since for structured pruning we cannot precisely control the weight sparsity, we show results closest to 90% weight sparsity. \uparrow/\downarrow indicate metrics where higher/ lower is better. Bold/ underline indicate best/ second best results. GPU RAM and Disk correspond to those of the final pruned model.

	Method	Test accuracy \uparrow	Weight sparsity	Node sparsity	Training time (h) \downarrow	Batch time (ms) \downarrow	GPU RAM (GB) \downarrow	Disk (MB) \downarrow	Emissions (g) \downarrow
-	Dense	51.3%	-	-	7.26	320	3.53	569	882
Structured	Random-S	37.3%	91.2%	80.0%	6.23	289	1.08	51	464
	SNAP	38.3%	90.4%	82.6%	6.06	268	0.84	55	514
	CroP-S	<u>39.1%</u>	90.1%	77.7%	6.72	237	1.11	54	615
	CroPit-S	<u>39.1%</u>	91.4%	79.3%	6.66	236	1.08	49	591
	EarlyCroP-S	39.2%	90.8%	84.1%	7.03	<u>202</u>	0.25	49	676
	GateDecorators	30.1%	89.2%	91.2%	<u>6.20</u>	193	0.87	61	930
	EfficientConvNets	27.7%	91.0%	79.8%	6.60	226	0.22	52	769
	Unstructured	Random-U	<u>49.3%</u>	90.0%	-	7.25	351	4.20	57
SNIP		46.2%	90.0%	-	7.27	314	4.18	57	854
GRASP		43.7%	90.0%	-	7.27	315	4.22	57	881
CroP-U		46.7%	90.0%	-	7.26	314	4.22	57	877
CroPit-U		19.1%	90.0%	-	7.26	313	4.22	57	890
EarlyCroP-U		49.8%	90.0%	-	7.26	314	4.23	57	880
LTR		46.3%	90.0%	-	44.7	603	3.68	57	5540

Table 7. EarlyCroP-S results on VGG16/ImageNet at 50% weight sparsity

Method	Top-1 Accuracy	Top-5 Accuracy	Train Time (hours)	Epochs	Batch Time (seconds)	GPU Memory (GB)
Dense	58.78%	82.55%	62	18	1.01	12.15
EarlyCroP-S	61.43%	87.01%	62	26	0.66	10.63
EarlyCroP-S	60.01%	83.38%	51	18	0.66	10.63

Table 8. Comparison between LTH and EarlyCroP-U on different learnable BERT tasks.

	MNLI	QQP	STS-B	WNLI	QNLI	RTE	SST-2	CoLA	Training time
Dense BERT	82.39	90.19	88.44	54.93	89.14	63.30	92.12	54.51	1x
Sparsity	70%	90%	50%	90%	70%	60%	60%	50%	
LTR (Rewind 0)%	82.45	89.20	88.12	54.93	88.05	63.06	91.74	52.05	10x
LTR (Rewind 50)%	82.94	89.54	88.41	53.32	88.72	62.45	92.66	52.00	10x
EarlyCroP-U	82.11	89.99	88.02	56.33	89.12	62.1	92.03	52.2	1x