

# When Layers Play the Lottery, all Tickets Win at Initialization

Artur Jordao<sup>1</sup>, George de Araújo<sup>2</sup>, Helena de Almeida Maia<sup>2</sup> and Helio Pedrini<sup>2</sup>

<sup>1</sup>Escola Politécnica, Universidade de São Paulo

<sup>2</sup>Institute of Computing, University of Campinas

## Abstract

Pruning is a standard technique for reducing the computational cost of deep networks. Many advances in pruning leverage concepts from the Lottery Ticket Hypothesis (LTH). LTH reveals that inside a trained dense network exists sparse subnetworks (tickets) able to achieve similar accuracy (i.e., win the lottery – winning tickets). Pruning at initialization focuses on finding winning tickets without training a dense network. Studies on these concepts share the trend that subnetworks come from weight or filter pruning. In this work, we investigate LTH and pruning at initialization from the lens of layer pruning. First, we confirm the existence of winning tickets when the pruning process removes layers. Leveraged by this observation, we propose to discover these winning tickets at initialization, eliminating the requirement of heavy computational resources for training the initial (over-parameterized) dense network. Extensive experiments show that our winning tickets notably speed up the training phase and reduce up to 51% of carbon emission, an important step towards democratization and green Artificial Intelligence. Beyond computational benefits, our winning tickets exhibit robustness against adversarial and out-of-distribution examples. Finally, we show that our subnetworks easily win the lottery at initialization while tickets from filter removal (the standard structured LTH) hardly become winning tickets.

## 1. Introduction

The Lottery Ticket Hypothesis (LTH) conjectures that (pre-trained) dense networks contain sparse subnetworks capable of obtaining the same accuracy when trained from the original initialization of their dense counterpart [9]. Subnetworks (tickets) that satisfy this property are referred to as *winning tickets*. Many advances emerge from the LTH, for example, we can reduce the computational cost of learning a dense network by replacing it with a sparse subnetwork dur-




Figure 1. Lottery Ticket Hypothesis (LTH) views according to the structure (weights, neurons/filters or layers) the pruning is eliminating (transparent regions). Top-left. Original unstructured LTH: the pruning removes weights and yields unstructured tickets; thereby, the tickets only provide practical benefits on specialized frameworks for sparse computations. Top-right. Structured LTH: the pruning eliminates neurons/filters. In this setting, the tickets are structured and promote computational advantages to standard deep learning frameworks. Bottom-left: Ours structured LTH: the pruning eliminates entire layers, encouraging additional performance gains since it decreases the sequential processing (latency). Bottom-right. The highest gain (the higher, the better) obtained by a winning ticket regarding its dense counterpart. Our winning tickets successfully emerge at initialization, which means we can discover efficient subnetworks without training a dense network. In this direction, we can considerably speed up the learning phase by replacing a dense network with its sparse version before training begins. Our winning tickets also exhibit robustness against adversarial attacks.

ing [41, 4] or before the course of training [21, 40, 36]. In the context of adversarial attacks, we can increase the robustness of dense networks by considering their sparse versions [6, 25, 34].

The traditional LTH seeks winning tickets by removing (i.e., pruning) weights of a dense network and assigning the survival weights to their original initialization [9] – weight rewinding. Variants of this mecha-

nism propose relaxing the weight-rewinding constraint and enabling subnetworks to inherit weights from different training epochs [11, 32] and even from random (re) initialization [26]. Another variant in the traditional LTH is the replacement of weight pruning (unstructured) by filter pruning (structured) [26, 32, 41]. For example, Prasanna et al. [31] studied the LHN on structured pruning by removing self-attention modules (i.e., heads from multi-head attention layers [37]). One practical drawback of unstructured pruning is that existing deep learning frameworks (e.g., TensorFlow and Pytorch) do not support sparse tensor computations. Hence, in order to obtain performance gains, this family of pruning requires specialized frameworks or hardware (Nvidia A100 GPU) for optimizing sparse computations [12, 29, 43].

Regardless of the structure, we can categorize LTH according to the phase in which the pruning algorithm picks the subnetworks (tickets): after, during, or before training. The first class trains a dense network and, then, extracts a subnetwork according to an importance criterion, such steps compose the original LTH [9]. The second prunes a dense network during the course of training, obtaining a subnetwork when the learning phase is over [41, 4]. Finally, pruning before training constitutes a relatively recent form of pruning named *pruning at initialization* [21, 36, 40]. The idea behind this category is to find subnetworks prior to training, which means picking subnetworks using the randomly initialized parameters (without any update) to guide the pruning algorithm. Many efforts have been put into pruning at initialization as it provides all benefits of sparse subnetworks without spending computational resources to train a dense network (a process required by the other categories).

Importantly, all pruning categories in LTH share the same characteristic: the pruning eliminates small structures such as weights or filters. Outside the context of LTH, previous works have demonstrated notable benefits of removing entire layers instead of small structures. In particular, this type of structured pruning promotes additional performance gains since it reduces the sequential processing (latency) of a network [22, 44]. Hence, many studies have focused on removing layers instead of other structures [38, 8, 44, 42]. Unfortunately, none of these efforts have been done in the direction of LTH and pruning at initialization. To bridge this gap, we take a step towards understanding the behavior of LTH when the pruning process removes layers. In this setting, we first verify the existence of winning tickets. Then, we propose a systematic strategy for discovering winning tickets before training, which means finding sparse subnetworks (from layer re-

moval) at initialization that match the predictive ability of their dense equivalent.

**Contributions.** We list the following key contributions. First, we demonstrate the existence of winning tickets – sparse subnetworks that obtain the same accuracy as their dense equivalent – when the pruning process takes into account the removal of layers. From a practical perspective, such a contribution enables winning tickets to be more efficient in terms of memory consumption, inference time (latency) and carbon emission, as removing layers provides superior computational benefits than standard forms of pruning employed in LTH (see Figure 1). Second, we successfully find winning tickets from layer removal at initialization without any training. This contribution plays a role in low-cost and energy-efficient training, as we can replace the learning of a dense network with its sparse version. In contrast to previous LTH studies [11, 32], this contribution eliminates the requirement of heavy computational resources for learning the initial (over-parameterized) dense network. Finally, we show that this novel family of subnetworks shares some properties of the standard structured LTH (filter pruning), for example, they emerge early at training. On the other hand, we observe that subnetworks from filter pruning exhibit poor performance when applied to shallow (underparametrized) networks and before training, which is aligned with previous evidence [25]. In other words, tickets from filter pruning hardly become winning tickets. However, our tickets (layer removal) easily win the lottery, raising the question of whether the eliminated structure plays a role in the LTH and pruning at initialization.

According to our analysis, winning tickets from layer pruning are accurate and robust to many settings such as the pruning criteria and density, and weight rewinding. When found at initialization, these winning tickets speed up the training time by up to  $2\times$  and save notable FLOPs and memory consumption. On a specific pruning criterion, all tickets outperform the predictive ability of their dense counterpart – all tickets become winning tickets. Additionally, our winning tickets reduce up to 51% of carbon emission during the training phase, an important step towards democratization and sustainability of Artificial Intelligence (AI) – green AI. As suggested by previous works [25, 34, 4], we also evaluate our winning tickets on adversarial images and out-of-distribution examples using CIFAR-C [15] and CIFAR-10.2 [27] datasets, respectively. On the standard clean training (i.e., without any defense mechanism), tickets from layer pruning achieve higher robustness than their dense equivalent; thus, confirming their suitability for safety-critical applications.

## 2. Related Work

Frankle et al. [9] introduced the Lottery Ticket Hypothesis and observed that subnetworks from a dense network can obtain similar accuracy since trained from the same original initialization. Later, Frankle et al. [11] and Renda et al. [32] confirmed that rewinding the weights to other epochs of the training, instead of the original initialization, enables subnetworks to match the accuracy of their dense counterpart.

The work by Achille et al. [1] investigated the sensibility of networks during training. The authors observed that perturbations, such as blur and random label, affect less the training dynamic when applied in early epochs. By considering pruning as a form of perturbation (e.g., due to the removal of structures), such evidence reinforces the superior performance of subnetworks rewound to early training epochs, as argued by subsequent studies [41]. While weight rewinding lies at the heart of LTH, Liu et al. [26] demonstrated that the winning tickets exist even on a random initialization regime. Throughout this study, we show that our winning tickets (from layer pruning) emerge from different rewinding epochs, random initialization and before training, with the latter achieving the best gains in terms of accuracy and computational performance.

Liu et al. [24] suggested that the existence of winning tickets correlates with the transition from the initial and final parameters of the dense network. From a different perspective, Paul et al. [30] showed that data quantity and quality play a role in LTH. More concretely, the authors observed that training on easy or on a small fraction of randomly chosen data promotes a fine initialization to dense networks such that their subnetworks become winning tickets. In line with Liu et al. [24], You et al. [41] proposed drawing sparse subnetworks early at training by identifying when the magnitude of the weights becomes stable (i.e., suffer small changes) during the course of training. After locating a subnetwork, named early bird, the authors replaced the training of the dense with the subnetwork; thus leading to faster and lower-cost training. Built upon this idea, Chen et al. [4] proposed to find early birds to reduce the cost of adversarial training. However, the authors identify subnetworks by pruning weights (unstructured) while You et al. [41] focus on removing filters (structured pruning). Interestingly, the results by Chen et al. [4] corroborate the observation of previous works [6, 25], which state that pruning not only reduces the computational cost but also promotes robustness against adversarial images. Additionally, T et al. [34] confirmed that winning tickets generalize better than dense networks in limited-data regimes and out-of-distribution scenarios.

Beyond computational benefits, You et al. [41] and Chen et al. [4] showed that winning tickets drawn early at training obtain higher accuracy. From the LTH rewinding perspective, this means that resetting the weights of subnetworks to early epochs of the dense network allow subnetworks to achieve higher predictive ability, hence, they are more probable to become a winning ticket. Our results are aligned with these findings: subnetworks from layer pruning achieve superior accuracy when rewound to early epochs. In contrast to these works, we also propose to identify winning tickets at initialization (prior to training). Informally speaking, our subnetworks are the earliest as possible.

Closely related to LTH, several efforts have been put into pruning at initialization. In this strategy, the pruning algorithm removes unimportant structures before any training [21, 40, 36]. The study by Frankle et al. [10] pointed out the difficulties and inherent properties of many pruning at initialization strategies. The authors showed that this family of strategies is less effective than standard pruning (i.e., LTH) in terms of both accuracy and sparsity. Jorge et al. [5] corroborate such findings; further, they observed that pruning at initialization is no better than random pruning on high sparsity regimes. Liu et al. [25] suggested an opposite behavior: pruning at initialization works well on deep and wide networks, even randomly selecting the unimportant structures. Additionally, the authors showed that employing a layer-wise pruning density yields more accurate winning tickets. We observe that the evidence by Liu et al. [25] holds when pruning takes into account filters; however, we successfully find winning tickets at initialization on both shallow (i.e., ResNet32) and deep architectures (i.e., ResNet56), which suggest an architecture-agnostic form of discovering winning tickets prior to training. Though it looks like a counterintuitive phenomenon, there is a body of studies confirming the benefits of removing layers instead of filters [42, 44]. In summary, our work differs from previous studies on LTH and pruning at the initialization in terms of the structure we focus on removing – layers.

To the best of our knowledge, the idea behind pruning layers dates back to 2016, when studies [39, 16] demonstrated that residual-based architectures exhibit no degradation when we remove some of their layers. Later, Dong et al. [7] confirmed that self-attention architectures also share similar behavior. Since these pioneer works, many works have proposed either removing layers statically or dynamically. It is important to mention the difference between these categories of layer pruning. In the former, the pruning algorithm eliminates layers in the same way as weight and filter pruning does, yielding a permanently shallower net-

work [42, 44]. Differently, the rationale behind the dynamic layer pruning consists of skipping (i.e., deactivating) some layers according to the input sample the network receives [13]. Our work belongs to the first category of layer pruning. In particular, due to its dynamic nature, we believe it is impractical to study LTH and pruning at the initialization in the second group.

### 3. Preliminaries and Problem Statement

**Definitions.** Let  $\mathcal{F}$  and  $\mathcal{F}'$  be a dense and sparse (subnetwork) network, respectively. The latter is a version of  $\mathcal{F}$  without some structures (neurons, filters or layers), meaning  $\mathcal{F}$  after some pruning process. Assume  $\theta_i$  the parameters from  $\mathcal{F}$ , in which the subscript  $i$  indicates the weights at the training epoch  $i$ . We indicate the randomly initialized weights and the ones after the training stage as  $\theta_0$  and  $\theta_n$ , in this order.

**Pruning Algorithm.** A pruning algorithm identifies and removes unimportant structures composing a network. For this purpose, it measures the importance of each structure according to an importance criterion  $c$ . Let  $S$  be a set of (sorted) scores that indicates the importance of each structure of  $\mathcal{F}$ . Given  $S$ , the pruning algorithm removes the least important structures in order to satisfy a pruning density  $p$  (i.e., the percentage or number of structures removed.). In the literature on pruning, it is common to mask the eliminated structures with zeros. In contrast, we indeed remove the structures (layers) to achieve practical performance gains without needing specialized frameworks or hardware for sparse computations. We refer interested readers to supplementary material for additional details about this technical process.

**Lottery Ticket Hypothesis.** Following the previous definitions, the original Lottery Ticket Hypothesis (LTH) [9] states that inside  $\mathcal{F}$  exists sparse subnetworks  $\mathcal{F}'$  able to achieve similar (or, ideally, superior) accuracy since trained from the identical initialization  $\theta_0$ . In this definition, the subnetworks are named tickets and the ones that satisfy such property are named winning tickets. In other words, a winning ticket is a subnetwork with the same predictive ability (i.e., accuracy) as its dense equivalent. In practice, we can define the existence of winning tickets in terms of

$$Accuracy(\mathcal{F}') + \xi \geq Accuracy(\mathcal{F}), \xi \geq 0. \quad (1)$$

As we mentioned before, variants of LTH enable  $\mathcal{F}'$  subnetworks to inherit weights from different training epochs (i.e.,  $\theta_i$  with  $i > 0$ ) [11, 10]. This step composes the weight rewinding process. Algorithm 1 summarizes the steps of our LTH (layer pruning).

From the above description, we highlight two observations. (i) In Equation 1,  $\xi$  enables controlling (re-

---

#### Algorithm 1 LTH Removing Layers of Deep Networks

---

**Input:** Convolutional Network  $\mathcal{F}$  trained on  $n$  epochs  
**Input:** Weight Rewind  $\theta_i$   
**Input:** Pruning Criterion  $c$   
**Input:** Pruning Density  $p$   
**Output:** Subnetwork (Ticket)  $\mathcal{F}'$   
 $S \leftarrow c(\mathcal{F}, \theta_n) \triangleright$  Assigns importance for each layer  
 $I \leftarrow p$  unimportant layers based on  $S$   
 $\mathcal{F}' \leftarrow \mathcal{F} \setminus I \triangleright$  Removes the layers indexed by  $I$   
Set the weights of  $\mathcal{F}'$  as  $\theta_i \triangleright$  Weight rewinding  
Train  $\mathcal{F}'$  via standard training for  $n - i$  epochs

---

laxing) how much the accuracy of a ticket can differ from its dense network and still be considered a winning ticket, where common values are one percentage point or one standard deviation [3, 10]. (ii) Before extracting a potential winning ticket, we need to train  $\mathcal{F}$  to completion (i.e., training for  $n$  epochs to obtain  $\theta_n$ ) – see the first input in Algorithm 1.

**Pruning at Initialization.** This relatively recent category of pruning estimates which structures to remove before training. Technically speaking, it focuses on discovering subnetworks using the random initialization ( $\theta_0$ ) of  $\mathcal{F}$  to guide the pruning algorithm. From the lens of LTH, pruning at initialization aims at yielding winning tickets without training  $\mathcal{F}$ .

Algorithm 2 summarizes the steps of our pruning at initialization. In Algorithm 2, it is important to observe that it receives an untrained network, whereas in our standard LTH (Algorithm 1) the input is a trained network.

**Research Questions.** From the above definitions, our research questions are the following. (i) We ask if there are winning tickets when  $\mathcal{F}'$  derives from a *layer-pruning* process. In other words, we investigate if Equation 1 holds when the pruning removes layers. We confirm that the answer is positive; thus, raising our second question: (ii) Is it possible to discover such winning tickets (from layer pruning) at initialization? Formally, is it possible to extract an  $\mathcal{F}'$  that satisfies

---

#### Algorithm 2 Winning Tickets at Initialization

---

**Input:** Untrained Convolutional Network  $\mathcal{F}$   
**Input:** Pruning Criterion  $c \triangleright$  SNIP or GraSP  
**Input:** Pruning Density  $p$   
**Output:** Subnetwork (Ticket)  $\mathcal{F}'$   
 $S \leftarrow c(\mathcal{F}, \theta_0) \triangleright$  Assigns importance for each layer  
 $I \leftarrow p$  unimportant layers based on  $S$   
 $\mathcal{F}' \leftarrow \mathcal{F} \setminus I \triangleright$  Removes the layers indexed by  $I$   
Train  $\mathcal{F}'$  via standard training for  $n$  epochs

---

Equation 1 without training  $\mathcal{F}$ ? Answering this enables us to avoid the computationally expensive training of a dense network by replacing it directly with its sparse version before training begins. Overall, both questions focus on analyzing Equation 1 after performing Algorithm 1 and Algorithm 2.

## 4. Experiments

**Experimental Setup.** We conduct our experiments on CIFAR-10 [19] using the ResNet32 and ResNet56 networks [14]. Such settings are common choices for LTH and general pruning evaluation [2]. In our LTH, we employ the learning rate rewinding scheme suggested by Renda et al. [32], as it leads to better results than the one proposed in the original LTH [9]. In contrast to most LTH studies, we evaluate the LTH on multiple criteria instead of employing only the  $\ell_1$ -norm criterion (the iterative magnitude pruning — IMP — originally proposed by Frankle et al. [9]). The reason for this choice is to assess the existence of winning tickets on multiple pruning settings. Additionally, layers at different depths exhibit different magnitudes, hence, it is unfeasible to compare the  $\ell_1$ -norm from multiple layers (see supplementary material). Specifically, we employ the pruning criteria proposed by Lin et al. [23] (HRank), Tan and Motani [35] (expectABS), Jordao et al. [17] (PLS) and Luo and Wu [28] (KL). In addition, we use the commonly used random criterion.

We report (in percentage points – pp) the improvement and decline of the subnetworks (tickets) with respect to their dense counterpart using the symbols (+) and (-), in this order. Thereby, the symbol (+) stands for winning tickets since these networks match or outperform the accuracy of their dense equivalent. Since our pruning process considers removing layers, it cannot be evaluated on VGG-like architectures (plain and non-residual networks) due to technical and theoretical details. We refer interested readers to supplementary material for additional information.

Unless stated otherwise, the term subnetworks indicate networks yielded by the process of removing layers, which is the scope of our work. In this setting, the pruning density  $p$  stands for the number of layers removed, e.g.,  $p = 1$  indicates that the pruning removed 1 residual block. (In ResNet32-56 as well as its shallower and deeper variations, one residual block (layer) corresponds to two convolutional layers.)

**Existence of Winning Tickets and Weight Rewinding.** Our first experiment verifies if subnetworks become winning tickets when pruning removes layers. One of the most important facts to a subnetwork becoming a winning ticket is the epoch we rewind its weights (weight rewinding), which means the  $\theta_i$  that

the subnetwork inherits from its dense version.

Previous studies on LTH confirmed that inheriting weights from early epochs enables subnetworks successfully become winning tickets [41, 10]. From a practical perspective, resetting the weights to these epochs requires more training iterations to completion (see the last step in Algorithm 1). Here, we evaluate the effect of setting the weights of subnetworks from layer pruning to different epochs. In other words, we study the behavior of subnetworks when they inherit different  $\theta_i$ .

Table 1 summarizes the results. From this table, we highlight the following observations. First, rewinding subnetworks to the same (random) initialization of the dense network,  $\theta_0$ , enables all subnetworks to become winning tickets (except HRank). This means that, rewinding to  $\theta_0$ , all subnetworks satisfy Equation 1. Second, most subnetworks achieve the lowest accuracy rewinding after 50 epochs. This aligns with the findings of Achille et al. [1] and You et al. [41]: rewinding to late epochs, the subnetworks obtain lower accuracy, as they have a short period (epochs) to recover from damage in their structure. Overall, the weight rewinding of LTH when the pruning removes filters and layers shares similar behavior.

The previous discussion takes into account the original setup by Frankle et al. [9] and its variations [11, 32]. Liu et al. [26] suggested an alternative to recover the weights for some initialization composing the training trajectory ( $\theta_i$ ) of the dense network. The authors demonstrated that randomly re-initializing subnetworks allows them to become winning tickets. From the lens of layer pruning, we observe that subnetworks share the same trends. More specifically, on the rewinding by Liu et al. [26], our subnetworks outperform the dense network by up to 0.85 pp.

The previous results confirm the existence of winning tickets when the pruning process removes entire layers; thus answering our first research question.

**Picking Winning Tickets at Initialization.** So far, our experiments confirm the existence of winning tickets with the weight rewinding strategy, as in the

Table 1. Predictive ability of the tickets when rewinding their weights to  $i$ -th training epoch ( $\theta_i$ ) of the dense network (ResNet32). For each criterion, we highlight in bold and underline the rewind epoch that leads to the top-1 and top-2 best results, respectively.

	$\theta_0$	$\theta_{25}$	$\theta_{50}$	$\theta_{75}$
HRank	(-) 0.01	(+) <b>0.21</b>	(+) 0.10	(+) 0.12
expectABS	(+) <b>0.33</b>	(+) 0.13	(-) 0.48	(-) 0.62
PLS	(+) 0.06	(+) <b>0.31</b>	(+) 0.10	(+) 0.10
Random	(+) 0.38	(+) 0.15	(+) <u>0.26</u>	(-) 0.07
KL	(+) <b>0.74</b>	(+) 0.46	(-) 0.03	(-) 0.17

Table 2. Predictive ability of tickets from ResNet32 when our LTH discovers them at initialization.

Pruning Density	SNIP	GraSP
1	(+) 0.08	(+) 0.40
2	(+) 0.00	(+) 0.72
3	(-) 0.50	(+) 0.34
4	(-) 0.53	(+) 0.41
5	(-) 1.00	(+) 0.12

traditional LTH conjecture. Additionally, we observe that relaxing the weight rewinding and re-initializing the subnetworks enable them to become winning tickets. Both settings assume that winning tickets emerge from a well-trained dense network. In this experiment, we remove this constraint and propose to discover winning tickets without any training. In other words, we propose to find winning tickets before training begins. Such a setting eliminates the requirement of heavy computational resources for training the initial (over-parameterized) dense network, an important step towards democratization and green machine learning. As we discussed before, early works have proposed to pick winning tickets at initialization [21, 40, 36, 25]. However, these works focus on removing weights while we target removing layers, which is a more efficient and hardware-agnostic form of pruning.

One relevant question to find subnetworks from initialization is the choice of the criterion for estimating structure importance. It turns out that most criteria fail to measure importance without any training as the weights change drastically (hence the importance score). In summary, the criteria evaluated in Table 1 are unsuitable<sup>1</sup> to prune at initialization.

To face the above issue, we employ two criteria specifically designed for estimating importance at initialization. The first criterion, named SNIP, computes the importance by multiplying the weight by its gradient (at the initial – randomly – configuration); then, it takes the absolute value of the resulting operation [21]. The second criterion, termed GraSP, is a variant of SNIP, in which the importance considers the signal of the weight [40]. Importantly, SNIP and GraSP are data-driven criteria since both forward data (and labels) through the network to compute the gradient during the importance estimation.

Table 2 summarizes the results. From this table, we highlight the following observations. (i) SNIP is an inefficient criterion for discovering very sparse subnetworks (tickets) since its tickets are no longer win-

ning tickets after removing three layers. (ii) All tickets from GraSP win the lottery. Particularly, GraSP provides subnetworks with predictive ability superior to the dense network even on the highest pruning density we consider (5 layers – 66% of layers of ResNet32).

In order to confirm the superiority of Grasp over SNIP, we compare the subnetworks yielded by each criterion with the dense network through the lens of representation similarity among models. For this purpose, we employ the method by Kornblith et al. [18] (named CKA) that measures the representation similarity of two (architecturally equal or distinct) models. The idea is to verify if this similarity correlates with the results in Table 2. On the lowest pruning severity ( $p = 1$ ), the difference in CKA similarity between the dense network and the subnetworks from GraSP and SNIP is less than one percentage point, with the subnetwork from GRASP obtaining the higher similarity. However, on the highest pruning severity ( $p = 5$ ), the subnetwork from SNIP exhibited a difference of almost two percentage points. On the one hand, the representation similarity between the subnetworks yielded by SNIP and the dense network decreases as a function of the pruning severity. On the other hand, according to the CKA similarity, the internal representation of subnetworks from GraSP remains (partially) undamaged compared to the dense network. This observation reinforces the positive results of GraSP in Table 2 and indicates that we can effectively extract winning tickets at initialization using GraSP. Unless stated otherwise, we employ GraSP to discover winning tickets at initialization in the subsequent experiments.

### Comparison with standard (structured) LTH.

This experiment compares our winning tickets (layer pruning) with the standard structured LTH (filter pruning). During this evaluation, we extract winning tickets of layer and filter pruning at initialization. Here, a critical aspect is how to produce subnetworks with the same sparsity (in terms of filters) when the pruning removes different structures. It turns out that when it eliminates layers, the sparsity becomes inflexible. For example, by removing 1 and 2 layers, the number of remaining filters is 1200 and 1168 and we cannot obtain a filter sparsity among these values (i.e., 1190). For fairness of comparison, we adopt the following procedure. First, we remove layers from a dense network and calculate the number of filters in the obtained subnetwork (i.e., the number of kept filters). Then, to produce comparable subnetworks removing only filters, we run the pruning process forcing it to eliminate the closest number of filters of subnetworks from layer pruning. Thereby, subnetworks yielded from a dense network without layers and filters are as close as possible in

<sup>1</sup>In fact, we can always employ the random criterion; however, we are interested in a systemic process for selecting winning tickets at initialization.

Table 3. Comparison between tickets from the standard structured LTH (filter pruning) and our LTH (layer pruning). On both criteria for pruning at initialization, our subnetworks outperform their dense (ResNet32) counterpart. In other words, our tickets become winning tickets (they win the lottery). However, tickets from the standard structured LTH hardly become winning tickets.

	SNIP	GraSP
Filter (Standard LTH)	(-) 0.79	(-) 0.11
Layers (Our LTH)	(+) 0.08	(+) 0.40

terms of filter sparsity. We highlight that the opposite process – first removing filters and then layers – impairs a comparable sparsity.

Table 3 shows the results in the lowest sparsity regime (one layer removed –  $p = 1$ ). On both SNIP and GraSP criteria, our winning tickets at initialization yielded early winning tickets with better improvements. In particular, the subnetworks from the standard structured LTH become winning tickets only if we relax Equation 1 by increasing the value  $\xi$ .

Previous observations revealed the difficulty of extracting winning tickets at initialization and confirmed that it is no better than the standard *training and prune* paradigm [10, 5]; thus, corroborating the poor results of filter pruning in Table 3. Our winning tickets (layer pruning) pose a different perspective for such an issue: filters cannot be the most effective structure in structured LTH. Technically speaking, removing a layer consists of eliminating a group of filters (in particular, all filters) from a specific location, see supplementary material. Hence, it sounds like a counter-intuitive behavior that our winning tickets obtain accuracy superior to the standard LTH. Our results, however, are aligned with a body of studies that confirm the benefits of removing layers over other structures [42, 44, 13]. To reinforce these claims, we measure the representation similarity between the models of Table 3 and the dense network, similarly as we performed in previous experiments. On this metric, our subnetworks exhibited a higher similarity than subnetworks from filter pruning, regardless of the importance criterion (SNIP or GraSP). Such values suggest that our winning tickets kept an internal representation more similar to the original dense network than the ones found by the standard structured LTH, which supports their highest accuracies in Table 3.

Overall, we believe the findings above open a new direction for research in LTH: the influence of the structure taken into account during the pruning process.

**Computational Cost of Winning Tickets at Initialization.** The aforementioned discussion confirms

the existence of winning tickets at initialization when pruning removes layers (recall that our subnetworks come from layer pruning – see Algorithms 1 and 2). From a practical perspective, identifying winning tickets at initialization means that we can speed up the training stage and save computation in terms of different performance metrics, as we can replace the learning of a dense network with its sparse version. Such achievements become possible since we discover subnetworks that (i) are more efficient than the heavy and over-parametrized dense network; (ii) match the accuracy of their dense equivalent (i.e., the subnetworks are winning tickets) and (iii) spend no additional cost since we discover them before any training epoch.

Table 4 reports the gains of our winning tickets at initialization on standard pruning metrics such as floating point operations (FLOPs), memory consumption and training speed up. Following a modern trend [20, 33, 41], we also report the CO<sub>2</sub> emission during the training. According to Table 4, it is evident the advantages of discovering winning tickets at initialization. For example, our winning tickets speed up the training time of ResNet32 and ResNet56 by up to  $1.57\times$  and  $2.03\times$ , meaning that we can significantly reduce the costs involved during the training step such as energy consumption and CO<sub>2</sub> emission. Importantly, all these improvements come at no additional cost.

Previous attempts have proposed to find winning tickets early in the training phase [41, 4]. Our work differs from these works, as we discover winning tickets prior to any training (roughly speaking, our subnetworks are the earliest as possible), which is a more efficient strategy since the gains in Table 4 occur before the training starts.

#### Robustness of Winning Tickets at Initialization.

There is a growing body of work evaluating pruning methods in adversarial images since the technique emerges as a powerful and efficient defense mechanism against adversarial attacks and out-of-distribution examples [6, 25, 34, 4]. In this experiment, we demonstrate if there exists some sensibility of our winning tickets to these scenarios. For this purpose, we employ the CIFAR-C [15] and CIFAR-10.2 [27] datasets. It is important to mention that our training stage employs the standard clean training on CIFAR-10, which means that we do not employ any defense mechanism (i.e., adversarial training).

Table 5 (third column) shows the improvements in robustness of winning tickets at initialization (the ones in Table 4) over the deep network used to discover them. According to the results, our winning tickets effectively increase robustness to adversarial images. Specifically, our winning tickets outperformed their re-

Table 4. Performance gains when the pruning discovers winning tickets (layer pruning) at initialization. The values indicate the improvements (in percentage – the higher, the better) of tickets regarding the dense network.

	Pruning Density (p)	FLOPs	Memory Consumption	CO <sub>2</sub> Emission	Training Speed up
Winning Tickets at Initialization (ResNet32)	1	6.82	10.13	11.11	1.10
	3	20.47	25.29	25.92	1.35
	4	34.13	30.11	33.33	1.57
Winning Tickets at Initialization (ResNet56)	3	11.25	15.15	13.33	1.16
	4	15.00	19.42	20.00	1.24
	5	18.76	23.69	24.44	1.34
	12	45.00	45.82	51.11	2.03

spective dense network by up to 1.61 and 3.73 pp for ResNet32 and ResNet56, respectively. Intriguingly, for each architecture, one winning ticket exhibited low robustness, in which the highest decline is only 0.16 pp.

Regarding the out-of-distribution generalization (last column in Table 5), our winning ticket increase the generalization of ResNet56 in up to 1.75. On the ResNet32, however, we observe that the results become negative as a function of the pruning severity; thus raising the question of whether our tickets generalize well in out-of-distribution examples when extracted from shallow architectures. The analysis of such behavior is an interesting direction for future work.

For some settings (5 out of 14 configurations), our winning tickets underperform their dense equivalent. We did not observe any pattern in these low-performance tickets and left an in-depth investigation of this issue for future research. Overall, the achievements in robustness against adversarial images and out-of-distribution examples suggest that the benefits of our winning tickets are beyond reducing computational performance, which aligns with findings of previous works: removing structures from networks increases robustness and out-of-distribution genera-

Table 5. Robustness against adversarial and out-of-distribution examples of our winning tickets at initialization. On both CIFAR-C and CIFAR-10.2 datasets, we report the improvement (in pp – the higher, the better) of subnetworks regarding their dense counterpart.

	p	CIFAR-C	CIFAR-10.2
Winning Tickets at Initialization (ResNet32)	1	(+) 1.13	(+) 0.00
	3	(-) 0.05	(-) 1.15
	4	(+) 1.61	(-) 0.65
Winning Tickets at Initialization (ResNet56)	3	(+) 0.59	(+) 0.70
	4	(+) 3.73	(+) 1.75
	5	(-) 0.16	(+) 0.70
	12	(+) 1.98	(-) 1.95

tion [6, 4, 34].

## 5. Conclusions

In this work, we explore the concepts of the Lottery Ticket Hypothesis (LTH) and pruning at initialization from the lens of layer pruning. From this form of pruning, we bring both theoretical and practical benefits. First, we verify the behavior of LTH when the pruning yields subnetworks (tickets) by removing layers from a dense network, in which we confirm that the tickets successfully become winning tickets. Built upon this, we propose to systematically discover winning tickets at initialization, which means identifying sparse subnetworks (from layer pruning) without the need of training a dense network. According to extensive experiments, our winning tickets at initialization speed up the learning phase by up to 2×, reducing the carbon emission by up to 51%. Such achievements come at no additional price, as we extract winning tickets before training the dense network. Additionally, our winning tickets become more robust against adversarial images and generalize better in out-of-distribution examples than their dense counterpart. We hope these benefits attract more research on LTH from the lens of layer pruning.

## Acknowledgments

This research was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001 – and by the Brazilian National Council for Scientific and Technological Development (CNPq #304836/2022-2).

## References

- [1] Alessandro Achille, Matteo Rovere, and Stefano Soatto. Critical learning periods in deep networks. In *International Conference on Learning Representations (ICLR)*, 2019. 3, 5



- [2] Davis W. Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John V. Guttag. What is the state of neural network pruning? In *Conference on Machine Learning and Systems (MLSys)*, 2020. 5
- [3] Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, Zhangyang Wang, and Michael Carbin. The lottery ticket hypothesis for pre-trained BERT networks. In *Neural Information Processing Systems (NeurIPS)*, 2020. 4
- [4] Tianlong Chen, Zhenyu Zhang, Pengjun Wang, Santosh Balachandra, Haoyu Ma, Zehao Wang, and Zhangyang Wang. Sparsity winning twice: Better robust generalization from more efficient training. In *International Conference on Learning Representations (ICLR)*, 2022. 1, 2, 3, 7, 8
- [5] Pau de Jorge, Amartya Sanyal, Harkirat S. Behl, Philip H. S. Torr, Grégory Rogez, and Puneet K. Dokania. Progressive skeletonization: Trimming more fat from a network at initialization. In *International Conference on Learning Representations (ICLR)*, 2021. 3, 7
- [6] James Diffenderfer, Brian R. Bartoldson, Shreya Chaganti, Jize Zhang, and Bhavya Kailkhura. A winning hand: Compressing deep networks can improve out-of-distribution robustness. In *Neural Information Processing Systems (NeurIPS)*, 2021. 1, 3, 7, 8
- [7] Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. Attention is not all you need: pure attention loses rank doubly exponentially with depth. In *International Conference on Machine Learning (ICML)*, 2021. 3
- [8] Angela Fan, Edouard Grave, and Armand Joulin. Reducing transformer depth on demand with structured dropout. In *International Conference on Learning Representations (ICLR)*, 2020. 2
- [9] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations (ICLR)*, 2019. 1, 2, 3, 4, 5
- [10] 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 (ICLR)*, 2021. 3, 4, 5, 7
- [11] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning (ICML)*, 2020. 2, 3, 4, 5
- [12] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A. Horowitz, and William J. Dally. EIE: efficient inference engine on compressed deep neural network. In *International Symposium on Computer Architecture (ISCA)*, 2016. 2
- [13] Yizeng Han, Gao Huang, Shiji Song, Le Yang, Honghui Wang, and Yulin Wang. Dynamic neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 4, 7
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [15] Dan Hendrycks and Thomas G. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations (ICLR)*, 2019. 2, 7
- [16] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q. Weinberger. Deep networks with stochastic depth. In *European Conference on Computer Vision (ECCV)*, 2016. 3
- [17] Artur Jordao, Fernando Yamada, and William Robson Schwartz. Deep network compression based on partial least squares. *Neurocomputing*, 2020. 5
- [18] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey E. Hinton. Similarity of neural network representations revisited. In *International Conference on Machine Learning (ICML)*, 2019. 6
- [19] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). 2009. 5
- [20] Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. Quantifying the carbon emissions of machine learning. In *NeurIPS*. 2019. 7
- [21] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Snip: single-shot network pruning based on connection sensitivity. In *International Conference on Learning Representations (ICLR)*, 2019. 1, 2, 3, 6
- [22] Zhengang Li, Geng Yuan, Wei Niu, Pu Zhao, Yanyu Li, Yuxuan Cai, Xuan Shen, Zheng Zhan, Zhenglun Kong, Qing Jin, Zhiyu Chen, Sijia Liu, Kaiyuan Yang, Bin Ren, Yanzhi Wang, and Xue Lin. NPAS: A compiler-aware framework of unified network pruning and architecture search for beyond real-time mobile acceleration. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [23] Mingbao Lin, Rongrong Ji, Yan Wang, Yichen Zhang, Baochang Zhang, Yonghong Tian, and Ling Shao. Hrank: Filter pruning using high-rank feature map. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5
- [24] Ning Liu, Geng Yuan, Zhengping Che, Xuan Shen, Xiaolong Ma, Qing Jin, Jian Ren, Jian Tang, Sijia Liu, and Yanzhi Wang. Lottery ticket preserves weight correlation: Is it desirable or not? In *International Conference on Machine Learning (ICML)*, 2021. 3
- [25] 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 (ICLR)*, 2022. 1, 2, 3, 6, 7
- [26] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. In *International Conference on Learning Representations (ICLR)*, 2019. 2, 3, 5

- [27] Shangyun Lu, Bradley Nott, Aaron Olson, Alberto Todeschini, Hossein Vahabi, Yair Carmon, and Ludwig Schmidt. Harder or different? a closer look at distribution shift in dataset reproduction. In *International Conference on Machine Learning (ICML) Workshop on Uncertainty and Robustness in Deep Learning*, 2020. 2, 7
- [28] Jian-Hao Luo and Jianxin Wu. Neural network pruning with residual-connections and limited-data. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5
- [29] Wei Niu, Xiaolong Ma, Sheng Lin, Shihao Wang, Xuehai Qian, Xue Lin, Yanzhi Wang, and Bin Ren. Patdnn: Achieving real-time DNN execution on mobile devices with pattern-based weight pruning. In *Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2020. 2
- [30] Mansheej Paul, Brett W. Larsen, Surya Ganguli, Jonathan Frankle, and Gintare Karolina Dziugaite. Lottery tickets on a data diet: Finding initializations with sparse trainable networks. In *Neural Information Processing Systems (NeurIPS)*, 2022. 3
- [31] Sai Prasanna, Anna Rogers, and Anna Rumshisky. When BERT plays the lottery, all tickets are winning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020. 2
- [32] Alex Renda, Jonathan Frankle, and Michael Carbin. Comparing rewinding and fine-tuning in neural network pruning. In *International Conference on Learning Representations (ICLR)*, 2020. 2, 3, 5
- [33] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in NLP. In *ACL*, 2019. 7
- [34] Mukund Varma T, Xuxi Chen, Zhenyu Zhang, Tianlong Chen, Subhashini Venugopalan, and Zhangyang Wang. Sparse winning tickets are data-efficient image recognizers. In *Neural Information Processing Systems (NeurIPS)*, 2022. 1, 2, 3, 7, 8
- [35] Chong Min John Tan and Mehul Motani. Dropnet: Reducing neural network complexity via iterative pruning. In *International Conference on International Conference on Machine Learning (ICML)*. 2020. 5
- [36] Hidenori Tanaka, Daniel Kunin, Daniel L. Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. In *Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2, 3, 6
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Neural Information Processing Systems (NeurIPS)*, 2017. 2
- [38] Andreas Veit and Serge J. Belongie. Convolutional networks with adaptive inference graphs. *International Journal of Computer Vision (IJCV)*, 128:730–741, 2020. 2
- [39] Andreas Veit, Michael J. Wilber, and Serge J. Belongie. Residual networks behave like ensembles of relatively shallow networks. In *Neural Information Processing Systems (NeurIPS)*, 2016. 3
- [40] Chaoqi Wang, Guodong Zhang, and Roger B. Grosse. Picking winning tickets before training by preserving gradient flow. In *International Conference on Learning Representations (ICLR)*, 2020. 1, 2, 3, 6
- [41] Haoran You, Chaojian Li, Pengfei Xu, Yonggan Fu, Yue Wang, Xiaohan Chen, Yingyan Lin, Zhangyang Wang, and Richard G. Baraniuk. Drawing early-bird tickets: Toward more efficient training of deep networks. In *International Conference on Learning Representations (ICLR)*, 2020. 1, 2, 3, 5, 7
- [42] Ke Zhang and Guangzhe Liu. Layer pruning for obtaining shallower resnets. *IEEE Signal Processing Letters*, 2022. 2, 3, 4, 7
- [43] Aojun Zhou, Yukun Ma, Junnan Zhu, Jianbo Liu, Zhi-jie Zhang, Kun Yuan, Wenxiu Sun, and Hongsheng Li. Learning N: M fine-grained structured sparse neural networks from scratch. In *International Conference on Learning Representations (ICLR)*, 2021. 2
- [44] Yao Zhou, Gary G. Yen, and Zhang Yi. Evolutionary shallowing deep neural networks at block levels. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 2, 3, 4, 7