

FIND A WINNING SIGN: SIGN IS ALL WE NEED TO WIN THE LOTTERY

Junghun Oh¹, Sungyong Baik^{3,4}, and Kyoung Mu Lee^{1,2}

¹Dept. of ECE&ASRI, ²IPAI, Seoul National University

³Dept. of Artificial Intelligence, ⁴Dept. of Data Science, Hanyang University

{dh6dh, kyoungmu}@snu.ac.kr, dsybaik@hanyang.ac.kr

ABSTRACT

The Lottery Ticket Hypothesis (LTH) posits the existence of a sparse subnetwork (a.k.a. winning ticket) that can generalize comparably to its over-parameterized counterpart when trained from scratch. The common approach to finding a winning ticket is to preserve the original strong generalization through Iterative Pruning (IP) and transfer information useful for achieving the learned generalization by applying the resulting sparse mask to an untrained network. However, existing IP methods still struggle to generalize their observations beyond ad-hoc initialization and small-scale architectures or datasets, or they bypass these challenges by applying their mask to trained weights instead of initialized ones. In this paper, we demonstrate that the parameter sign configuration plays a crucial role in conveying useful information for generalization to any randomly initialized network. Through linear mode connectivity analysis, we observe that a sparse network trained by an existing IP method can retain its basin of attraction if its parameter signs and normalization layer parameters are preserved. To take a step closer to finding a winning ticket, we alleviate the reliance on normalization layer parameters by preventing high error barriers along the linear path between the sparse network trained by our method and its counterpart with initialized normalization layer parameters. Interestingly, across various architectures and datasets, we observe that any randomly initialized network can be optimized to exhibit low error barriers along the linear path to the sparse network trained by our method by inheriting its sparsity and parameter sign information, potentially achieving performance comparable to the original. The code is available at https://github.com/JungHunOh/AWS_ICLR2025.git

1 INTRODUCTION

In the field of deep learning, over-parameterization is viewed as a key to enhancing network capacity and improving generalization (Neyshabur et al., 2019; Belkin et al., 2019). It is well known that after training an over-parameterized dense network, many redundant parameters arise that can be removed without affecting performance, leading to the emergence of network pruning techniques (Liu et al., 2017; Lin et al., 2020). However, pruning an initialized dense network before training often leads to a sparse network that is difficult to optimize and fails to match the original generalization (Li et al., 2017; Evci et al., 2022). This phenomenon led to a question, first posed by Frankle & Carbin (2019) as the lottery ticket hypothesis (LTH): Is there a sparse subnetwork (i.e., a winning lottery ticket) that can achieve generalization comparable to its dense counterpart when trained from initialization? This challenging research question has garnered significant attention and inspired many follow-up studies.

Iterative magnitude pruning (Frankle & Carbin, 2019) (IMP) is one of the representative methods to identify a winning ticket through iterating three phases: training, pruning, and rewinding. Many researchers have sought to understand how IMP finds a winning ticket. Among several insightful findings, perspectives from the loss landscape have provided valuable insights. Frankle et al. (2020a); Evci et al. (2022); Paul et al. (2023) have shown that IMP can find a winning ticket only when the network obtained after the training phase maintains its basin of attraction¹ after the pruning and

¹Following previous works (Evci et al., 2022; Paul et al., 2023), we define a basin of attraction as a set of points connected by low-loss paths, where gradient descent from any point converges to the same minimum.




Figure 1: **Illustration of our motivation and method.** ψ and ϕ denote network parameters of normalization layers and parameters excluding those of normalization layers, respectively. The ‘LMC region’ refers to a region of solutions that are linearly mode-connected to the LRR or AWS solution.

rewinding phases, thereby preserving strong generalization potential (i.e., generalization ability after training) throughout subsequent iterations. Since this condition is difficult to satisfy in relatively large-scale settings, variants of IMP bypass the challenges by either rewinding to warm-up trained parameters (*weight rewinding*) (Frankle et al., 2020a) or skipping the parameter rewinding phase (*learning rate rewinding*) (Renda et al., 2020). Although they found a subnetwork that performs comparably to a dense network after training, it is not at initialization, and thus not a winning ticket.

In this paper, we empirically show that an effective signed mask, a sparse mask with parameter sign information, is a key to satisfying the challenging condition for finding a winning ticket. Specifically, we leverage learning rate rewinding (LRR) for its ability to find effective parameter sign and sparsity configuration (Gadhikar & Burkholz, 2024). Then, with a slight modification to LRR, we demonstrate that if parameter signs are preserved, the subnetwork obtained through our LRR variant remains within its basin of attraction even after randomly initializing its parameters. This implies that the generalization potential of the subnetwork can be transferred to any randomly initialized network via the signed mask, possibly allowing it to generalize comparably to the dense network after training.

We observe that the original LRR fails to achieve this. As illustrated on the left side of Figure 1, we observe that the LRR subnetwork leaves its basin after randomly initializing its parameters while preserving their signs, indicated by the red ball and the high error barrier between the yellow and red balls, similar to the findings in Frankle et al. (2020b). This failure stems from the significant influence of initializing normalization layer parameters. Interestingly, when we exclude normalization parameters from initialization and only randomize the other parameters while preserving their signs, the resulting network retains the original basin. This, in turn, leads to convergence to a solution with a low error barrier along the linear path to the LRR solution (i.e., the solution obtained by training the LRR subnetwork), as illustrated by the green ball moving towards the blue ball on the orange line. These results suggest that the signed mask and normalization parameters of the LRR subnetwork enable any randomly initialized network to inherit its generalization potential.

To take a step closer to finding a winning ticket, we eliminate the need for trained normalization parameters by addressing the adverse effects of initializing them. To this end, we propose AWS, a variation of LRR to find **A Winning Sign**, which prevents high error barriers along the linear path between the AWS subnetwork and its counterpart with initialized normalization parameters. During training, AWS randomly interpolates between the current and initialized normalization parameters linearly, using the interpolated values for each forward pass. As illustrated on the right side of Figure 1, we argue that the AWS subnetwork stays within its basin even after randomizing parameters while preserving their signs, indicated by the low error barrier between the yellow and red balls. This helps the resulting network converge to a solution linearly mode-connected to the AWS solution, with performance close to the dense network (gray dotted line). Experiments across architectures and datasets show that the signed mask from the AWS subnetwork enables any randomly initialized network to generalize comparably to the dense counterpart after training.

We summarize the contributions of our work as follows:

- We observe that any randomly initialized network can inherit the generalization potential of the LRR subnetwork through its signed mask and the normalization layer parameters.
- We propose AWS that alleviates the dependence on trained normalization parameters by preventing high error barriers between the AWS subnetwork and its counterpart with initialized normalization parameters.
- In contrast to existing methods that are limited to finding a winning ticket with an ad-hoc initialization, we show that any randomly initialized network can generalize comparably to a dense network after training by applying the AWS-driven signed mask to it.

2 RELATED WORKS

Lottery Ticket Hypothesis (LTH). Frankle & Carbin (2019) propose LTH which states that within a dense network, there exists a sparse subnetwork that, when trained from initialization, can achieve performance comparable to the dense counterpart. To find such a winning lottery ticket, the authors proposed *iterative magnitude pruning* (IMP) and demonstrated that IMP successfully finds a winning ticket in a relatively small-scale setting. Follow-up works have delved into a broad range of topics related to LTH, such as theoretical support for the existence of the winning ticket (Malach et al., 2020; Orseau et al., 2020; Burkholz, 2022; da Cunha et al., 2022), efficient alternatives to IMP (You et al., 2020), searching for a winning ticket without weight training (Chen et al., 2022; Sreenivasan et al., 2022; Koster et al., 2022), and empirical analyses on winning ticket (Zhou et al., 2019; Frankle et al., 2020a; Ma et al., 2021; Sakamoto & Sato, 2022; Evci et al., 2022; Paul et al., 2023).

Insights into a Winning Ticket. One of the most important topics is investigating what makes a sparse network win the lottery. Frankle et al. (2020a) introduced the notion of mode connectivity (Freeman & Bruna, 2017; Nguyen, 2019; Draxler et al., 2018; Garipov et al., 2018; Lubana et al., 2023) into LTH to investigate the conditions under which IMP finds a winning ticket. They consider a network stable against SGD noise if, under different SGD randomness, it converges to a region of solutions that exhibit low error barriers along the linear path connecting them. Based on this definition, they demonstrated that IMP succeeds only when the rewound network is stable against SGD noise. Evci et al. (2022) found that a winning ticket can be found only when it resides in the same basin as the pruning solution used to obtain the sparse pruning mask. Paul et al. (2023) demonstrated that a sparse mask obtained in an IMP iteration guides the subsequent pruning solution to be linearly mode-connected to the previous IMP solution, leading the consecutive pruning solutions to be piece-wise linearly mode-connected. In summary, these findings suggest that IMP can identify a winning ticket only when the network obtained from the training phase remains within its basin of attraction after the pruning and rewinding phases, thereby preserving the generalization of the original dense network throughout all iterations. Variants of IMP, such as *weight rewinding* (Frankle et al., 2020a) and *learning rate rewinding* (Renda et al., 2020), bypass this challenging condition by applying the found sparse mask to trained parameters rather than initialized ones, failing to find a winning ticket.

Significance of Parameter Signs in LTH. Recently, several works reported the importance of parameter signs from the perspective of representation capacity (Wang et al., 2023a;b). Zhou et al. (2019) are the first to discover the role of parameter signs in the context of LTH. They empirically showed that in the parameter rewinding stage of IMP, rewinding parameter signs has a greater impact on the performance than the magnitudes. By contrast, Frankle et al. (2020b) showed that transferring the signed mask obtained by IMP to the original initialization performs worse than transferring them along with the respective magnitudes. Gadhikar & Burkholz (2024) demonstrated that IMP can fail to find a winning ticket because it loses crucial sign information during parameter rewinding and struggles to learn an effective sign configuration again due to reduced network capacity. They claimed that learning rate rewinding (Renda et al., 2020) (LRR), on the other hand, identifies more performant sparse networks by finding and maintaining the effective sign configuration during training.

We observe that the ineffectiveness of transferring parameter signs to an initialized network, as noted in Frankle et al. (2020b), is due to the adverse effect of initializing the normalization layer parameters. To address this issue, we propose a slight variant of LRR and demonstrate that the challenging conditions for finding a winning ticket suggested by Frankle et al. (2020a); Evci et al. (2022); Paul et al. (2023) can be satisfied using the effective signed mask acquired by our LRR variant.

3 METHOD

3.1 NOTATIONS AND BACKGROUND

Lottery Ticket Hypothesis. Let $\theta \in \mathbb{R}^d$ denote the parameters of a neural network. We use θ to also represent the neural network parameterized by θ . Consider a binary mask $m \in \{0, 1\}^d$, having the same shape as θ . Note that the mask values for parameters not targeted for pruning, such as biases, are fixed to be 1. Then, the mask m defines a sparse subnetwork of the original dense network as $\theta \odot m$. The lottery ticket hypothesis (LTH) (Frankle & Carbin, 2019) posits the existence of a mask with non-trivial sparsity (i.e., $\sum_i m_i \ll d$) that allows $\theta_{\text{init}} \odot m$, where θ_{init}

represents initialized parameters, to perform comparably to the dense network after training. Such a sparse network at initialization is referred to as a *winning ticket*. *Iterative magnitude pruning* (IMP) (Frankle & Carbin, 2019) was suggested as a way to find the winning ticket through iterative `training` \rightarrow `pruning` \rightarrow `rewinding` procedures. Let $\mathcal{A}(\theta, u)$ represent an SGD learning algorithm with an initialized learning rate scheduler that updates θ until convergence using SGD randomness $u \sim U$ (e.g. randomness from a data loader or data augmentations). We omit u from $\mathcal{A}(\theta, u)$ if unnecessary. At `training` phase of the t -th iteration, IMP trains the masked initial parameters, obtaining $\theta_t^{\text{IMP}} \odot m_{t-1}^{\text{IMP}} = \mathcal{A}(\theta_0^{\text{IMP}} \odot m_{t-1}^{\text{IMP}})$, where m_{t-1}^{IMP} denotes the mask obtained from the $(t-1)$ -th iteration. At `pruning` phase, IMP produces the t -th mask by removing a portion of the non-zero weights (commonly 20%) in $\theta_t^{\text{IMP}} \odot m_{t-1}^{\text{IMP}}$ with the smallest magnitudes: $m_t^{\text{IMP}} = \text{prune}(\theta_t^{\text{IMP}} \odot m_{t-1}^{\text{IMP}})$. At `rewinding` phase, θ_t^{IMP} is rewound to the initial parameters, θ_0^{IMP} , and the entire process is repeated until $t = T$. Finally, IMP produces $\theta_0^{\text{IMP}} \odot m_T^{\text{IMP}}$, referred to as the *IMP subnetwork*, and after training the IMP subnetwork, IMP obtains $\mathcal{A}(\theta_0^{\text{IMP}} \odot m_T^{\text{IMP}})$, referred to as the *IMP solution*.

Variants of IMP. Several variants of IMP have been proposed to address the failure of IMP in generalizing to more challenging settings, especially focusing on the `rewinding` phase. Based on the analysis of stability against SGD noise, Frankle et al. (2020a) proposed *weight rewinding* (WR) that rewinds θ_t^{IMP} to the warm-up trained parameters rather than to θ_0^{IMP} . *Learning rate rewinding* (LRR) (Renda et al., 2020) is another variant of IMP that skips parameter rewinding and instead rewinds only the learning rate schedule. At the `training` phase of the t -th iteration, LRR trains $\theta_{t-1}^{\text{LRR}} \odot m_{t-1}^{\text{LRR}}$, obtaining $\theta_t^{\text{LRR}} \odot m_{t-1}^{\text{LRR}} = \mathcal{A}(\theta_{t-1}^{\text{LRR}} \odot m_{t-1}^{\text{LRR}})$, where m_{t-1}^{LRR} represents the mask obtained from the $(t-1)$ -th iteration. After obtaining $m_t^{\text{LRR}} = \text{prune}(\theta_t^{\text{LRR}} \odot m_{t-1}^{\text{LRR}})$ at the `pruning` phase, LRR skips the `rewinding` phase and continues the subsequent iterations until $t = T$. Finally, LRR produces $\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}}$, referred to as the *LRR subnetwork*, and after training the LRR subnetwork, LRR obtains $\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$, referred to as the *LRR solution*.

Linear Mode Connectivity and stability against SGD Noise. Linear mode connectivity and stability against SGD noise have been adopted as useful tools for analyzing winning tickets from the loss landscape perspective (Frankle et al., 2020a; Evci et al., 2022; Paul et al., 2023). Let $\mathcal{E}(\theta)$ denote the test error of θ . We define the error barrier between two parameters, θ and θ' , when interpolating them by a factor of α as the difference between the error of the interpolated network and the mean error:

$$\mathcal{E}_\alpha(\theta, \theta') = \mathcal{E}(\alpha\theta + (1-\alpha)\theta') - (\mathcal{E}(\theta) + \mathcal{E}(\theta'))/2. \quad (1)$$

Then, we define $\mathcal{E}(\theta, \theta') = \sup_\alpha \mathcal{E}_\alpha(\theta, \theta')$. If $\mathcal{E}(\theta, \theta')$ is smaller than a sufficiently small value ϵ , we consider θ and θ' to be linearly mode-connected (LMC). ϵ is often determined empirically such as the standard deviation of errors of a dense network across different training seeds (Paul et al., 2023). Based on the definition of LMC, we define the stability against SGD noise as the condition where a pair of networks are LMC when trained from the same initial parameters but with different SGD randomness (Frankle et al., 2020a). Formally, θ is considered stable against SGD noise if $\mathcal{A}(\theta, u)$ and $\mathcal{A}(\theta, u')$ are LMC with $u, u' \sim U$.

As demonstrated by previous works (Frankle et al., 2020a; Evci et al., 2022; Paul et al., 2023), IMP can identify a winning ticket only when at the t -th iteration, $\theta_0^{\text{IMP}} \odot m_{t-1}^{\text{IMP}}$ still resides in the basin of attraction of the $(t-1)$ -th solution after the `training` phase, $\theta_{t-1}^{\text{IMP}} \odot m_{t-2}^{\text{IMP}}$, even after updating m_{t-2}^{IMP} to m_{t-1}^{IMP} and rewinding $\theta_{t-1}^{\text{IMP}}$ to θ_0^{IMP} . Paul et al. (2023) claim that if this condition is satisfied at every iteration, the consecutive solutions after the training phase are piece-wise LMC, allowing the final IMP solution to generalize comparably to the dense network. To examine this condition, previous works investigate whether a rewound network is stable against SGD noise and converges to a solution with linear mode connectivity to the network before the rewinding and pruning phases. WR and LRR bypass the challenging condition by rewinding to warm-up trained parameters or skipping the `rewinding` phase. In contrast, this paper argues that any randomly initialized network can inherit strong generalization potential through a mask with parameter sign information, progressing toward the goal of LTH.

3.2 MOTIVATION

Gadhikar & Burkholz (2024) demonstrated that learning rate rewinding (LRR) successfully maintains the performance of dense networks by learning and maintaining an effective parameter sign

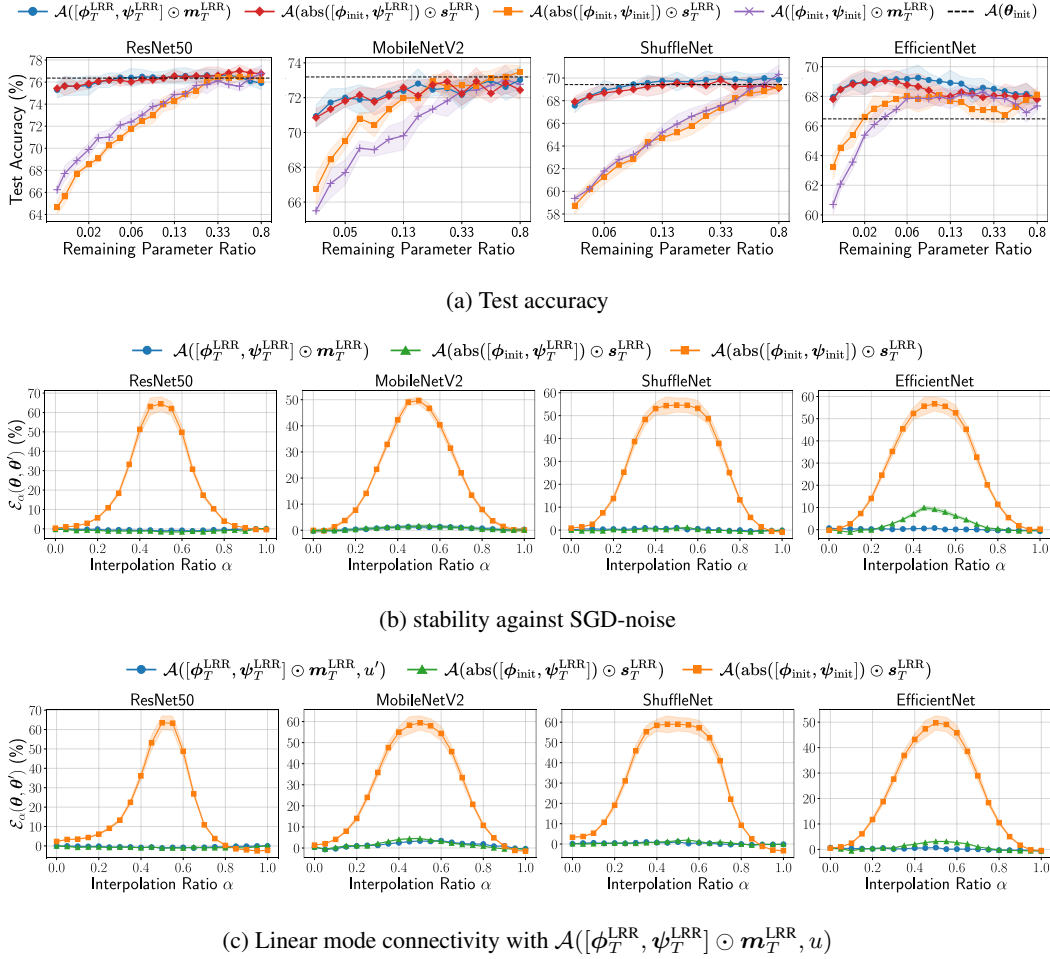


Figure 2: **Motivational experiments on CIFAR-100.** We investigate the effect of parameter initialization in the LRR subnetwork while preserving their signs with respect to (a) **test accuracy**, (b) **SGD-noise stability**, and (c) **linear mode connectivity with the LRR subnetwork**. In (b) and (c), we use a pruned network with a remaining parameter ratio of approximately 0.09. We show the mean (each point) and standard deviation (shaded area) across 3 trials.

configuration while pruning unimportant parameters, which iterative magnitude pruning (IMP) and weight rewinding (WR) fail. Several other works (Zhou et al., 2019; Frankle et al., 2020b; Chen et al., 2022; Sreenivasan et al., 2022; Koster et al., 2022; Wang et al., 2023b;a) also found the importance of parameter signs in the context of LTH or representation learning. In this work, motivated by these studies, we hypothesize that **the parameter sign information obtained through LRR can transfer the generalization potential of the LRR subnetwork to any randomly initialized network**. Let $\text{sign}_0(\cdot)$ denote a function that outputs the sign of each input element if it is non-zero, and 0 otherwise. Then, $s_T^{\text{LRR}} = \text{sign}_0(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$ represents the signed mask of the LRR subnetwork. More specifically, our hypothesis states that applying s_T^{LRR} to a randomly initialized network, θ_{init} , will enable the resulting network to match the performance of the LRR solution after training: $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}}) \approx \mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$ where $\text{abs}(\cdot)$ denotes the modulus function.

A similar idea was studied by Frankle et al. (2020b). They showed that replacing the signs of initialized parameters with those of the IMP subnetwork does not improve performance compared to using magnitude information in conjunction. We point out that this failure is attributed to ignoring the impact of parameters that may rely more on magnitudes than their signs. In the case of weight parameters, such as the weights in a convolutional layer, the sign configuration has a critical role in determining the functional mechanism of the layer as discussed in Wang et al. (2023b;a); Gadhikar & Burkholz (2024). By contrast, for parameters in a normalization layer, such as the batch normalization (Ioffe & Szegedy, 2015) or layer normalization (Ba et al., 2016), the magnitude may be much more important

Algorithm 1 AWS: a slight modification to LRR to find winning sign. The modification is highlighted in red.

Require: Initialize ϕ_0, ψ_0 , and $\mathbf{m}_0 \leftarrow (1, \dots, 1) \in \mathbb{R}^d$

```

1: for  $t = 1$  to  $T$  do
2:   while not converge do
3:      $(\psi_{t-1}, \psi_{\text{init}})_\alpha = \alpha \cdot \psi_{t-1} + (1 - \alpha) \cdot \psi_{\text{init}}$  where  $\alpha \sim U(0, 1)$  ▷ Interpolating  $\psi_{t-1}$  and  $\psi_{\text{init}}$ 
4:     Forward pass using  $[\phi_{t-1}, (\psi_{t-1}, \psi_{\text{init}})_\alpha] \odot \mathbf{m}_{t-1}$  ▷ Forward pass with the interpolated parameters
5:     Update  $\phi_{t-1}$  and  $\psi_{t-1}$  via gradient descent
6:   end while
7:    $\phi_t \leftarrow \phi_{t-1}$  and  $\psi_t \leftarrow \psi_{t-1}$ 
8:    $\mathbf{m}_t \leftarrow \text{prune}([\phi_t, \psi_t] \odot \mathbf{m}_{t-1})$  ▷ Update the sparse mask
9:   Rewind learning rate scheduler
10: end for
11: return  $\text{sign}_0([\phi_T, \psi_T] \odot \mathbf{m}_T)$  ▷ Obtain the signed mask

```

than the sign since the scaling parameter, initialized to 1, is nearly always positive after training, and the bias parameter loses the replaced signs since it is initialized to 0. We also observe that transferring the signs of parameters of normalization layers is not beneficial as analyzed in Appendix A.

Let ϕ and ψ denote the parameters excluding those of normalization layers and those of normalization layers, respectively. Then, we represent a randomly initialized network and the LRR subnetwork as $\theta_{\text{init}} = [\phi_{\text{init}}, \psi_{\text{init}}]$ and $\theta_T^{\text{LRR}} = [\phi_T^{\text{LRR}}, \psi_T^{\text{LRR}}]$, respectively. To test our conjecture, we compare two cases: in the LRR subnetwork, randomly initializing the magnitudes of both ϕ_T^{LRR} and ψ_T^{LRR} versus only ϕ_T^{LRR} while maintaining the signed mask (i.e., $\mathcal{A}(\text{abs}([\phi_{\text{init}}, \psi_{\text{init}}]) \odot s_T^{\text{LRR}})$ vs. $\mathcal{A}(\text{abs}([\phi_{\text{init}}, \psi_T^{\text{LRR}}]) \odot s_T^{\text{LRR}})$). Figure 2 shows the results on CIFAR-100 with various architectures. For details on the experiments, please refer to Section 4.1. Figure 2a shows that similar to the results in Frankle et al. (2020b), preserving the signed mask while initializing the magnitudes of all parameters randomly (indicated by the orange plots) results in performance similar to the case where sign information is not used (indicated by the purple plots), lagging far behind the performance of the LRR solution (indicated by the blue plots). This result indicates that sign information is not beneficial when all parameters are randomly initialized. On the other hand, interestingly, when the parameters of normalization layers are kept intact (indicated by the red plots), the performance is comparable to the LRR solution after training, indicating that using the sign information of the LRR subnetwork is beneficial when excluding the influence of initializing the normalization layer parameters. To examine whether the resulting networks reside in the basin of attraction of the LRR subnetwork, we analyze their stability against SGD noise and linear mode connectivity with the LRR subnetwork. Figure 2b demonstrates that while randomly initializing the magnitudes of all parameters significantly ruins the stability of the LRR subnetwork (indicated by the orange plots), it is effectively preserved when ignoring the influence of initializing the normalization layer parameters (indicated by the green plots). As shown in Figure 2c, we also observe that the preserved stability, in turn, leads the resulting network, $\text{abs}([\phi_{\text{init}}, \psi_T^{\text{LRR}}]) \odot s_T^{\text{LRR}}$, to converge to a solution with a low error barrier along the linear path connecting it to the LRR solution (indicated by the green plots), suggesting they are in the same basin of attraction. This contrasts with the high error barrier observed when initializing all parameter magnitudes (indicated by the orange plots).

The left side of Figure 1 illustrates the observations from our motivational experiments:

- In the LRR subnetwork, randomly initializing all parameters while preserving their signs causes the resulting network to lose SGD noise stability of the LRR subnetwork, potentially leading to a suboptimal solution, as indicated by the red ball.
- On the other hand, when the normalization layer parameters are kept intact, the resulting network, $\text{abs}([\phi_{\text{init}}, \psi_T^{\text{LRR}}]) \odot s_T^{\text{LRR}}$, exhibits significant stability against SGD noise and converges to a solution with a low error barrier along the linear path connecting it to the LRR solution, resulting in performance comparable to that of dense network, as indicated by the green ball.

Our observations suggest that when the parameter signs are preserved, the LRR subnetwork stays within its basin of attraction even after randomly initializing the parameters, excluding those of normalization layers. In other words, a randomly initialized network can inherit the generalization potential of the LRR subnetwork through the signed mask and normalization layer parameters of the LRR subnetwork.

3.3 AWS: FINDING A WINNING SIGN

Our observations provide valuable insights into the role of the signed mask in transferring strong generalization potential to any randomly initialized network. However, the need for the trained normalization parameters still limits LRR in achieving the goal of LTH. In this subsection, we introduce a method that addresses the adverse impact of initializing the normalization layer parameters to further progress toward the goal of LTH. Our goal is to maintain the basin of attraction in which the LRR subnetwork resides when the normalization layer parameters are initialized. Two networks residing in the same basin may indicate that no high error barrier exists along the linear path connecting them (Evci et al., 2022). Thus, we propose a simple variation of LRR to find a winning sign, referred to as AWS, that prevents any high error barriers along the linear path connecting the LRR subnetwork and its counterpart with initialized normalization parameters. Specifically, AWS randomly and linearly interpolates the parameters of normalization layers with their initialization and uses the interpolated parameters instead of the original parameters during training. Formally, at every network forward pass during the t -th iteration, AWS obtains

$$(\psi_t^{\text{AWS}}, \psi_{\text{init}})_\alpha = \alpha \cdot \psi_t^{\text{AWS}} + (1 - \alpha) \cdot \psi_{\text{init}}, \quad (2)$$

where ψ_t^{AWS} denotes the parameters of normalization layers during the t -th iteration of AWS and $\alpha \sim U(0, 1)$. Then, AWS uses $(\psi_t^{\text{AWS}}, \psi_{\text{init}})_\alpha$ instead of ψ_t^{AWS} for network forwarding. We present the pseudo-code of AWS in Algorithm 1, omitting the superscript ‘AWS’ for simplicity. After all iterations, we transfer the resulting signed mask, s_T^{AWS} , to a randomly initialized network, obtaining $\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}}$, and train it using normal training set-up until convergence.

The right side of Figure 1 illustrates the effectiveness of AWS. In contrast to LRR, AWS can allow a randomly initialized network to lie in the basin of attraction of the AWS subnetwork through the learned signed mask, s_T^{AWS} , possibly leading it to converge to a solution that is linearly mode-connected to the AWS solution (indicated by the red and blue ball). Thus, we argue that **s_T^{AWS} can transfer the generalization potential of the AWS subnetwork to any randomly initialized network, possibly resulting in performance comparable to a dense network after training.**

4 EXPERIMENTS

4.1 IMPLEMENTATION DETAILS

Datasets and models. Following the previous works (Ma et al., 2021; Gadhikar & Burkholz, 2024), we conduct experiments on CIFAR-100 (Krizhevsky & Hinton, 2009), Tiny-ImageNet (Le & Yang, 2015), and ImageNet (Russakovsky et al., 2015). For both CIFAR-100 and Tiny-ImageNet, we adopt ResNet-50 (He et al., 2016), MobileNetV2 (Sandler et al., 2018), ShuffleNet (Zhang et al., 2018), and EfficientNet (Tan & Le, 2019) to validate our method across various architectures. For ImageNet experiments, we use ResNet50, MobileNetV2, and MLP-Mixer (Tolstikhin et al., 2021). We adopt MLP-Mixer, which includes layer normalization, to demonstrate the generalization of our method to different types of normalization layers rather than the batch normalization layer.

Implementation. In both learning rate rewinding (LRR) and the proposed AWS method, we observe that many training epochs and learning rate scheduling are unnecessary during the training phase, as the network converges quickly due to the absence of parameter rewinding, and learning rate scheduling has little impact on the performance. Thus, during the training phase for both LRR and AWS, we train a network for 10 epochs for CIFAR-100 and Tiny-ImageNet experiments, and 5 epochs for ImageNet experiments without learning rate scheduling. To ensure a network can converge at the early iterations, we perform warm-up training before the first iteration: 10 epochs for CIFAR-100 and Tiny-ImageNet, and 20 epochs for ImageNet. After the T -th iteration, we conduct final training of 100 epochs for CIFAR-100 and Tiny-ImageNet, 120 epochs for ResNet-50 and MobileNetV2 on ImageNet, and 200 epochs for MLP-Mixer.

Optimization. For the CIFAR-100 and Tiny-ImageNet experiments, we use the SGD optimizer with a momentum of 0.9, a weight decay of $5e-4$, and an initial learning rate of 0.1, except for MobileNetV2, which uses a learning rate of 0.05. The learning rate decays by a factor of 0.1 at the 50th and 75th epochs over 100 epochs. For ImageNet experiments, we use the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and an initial learning rate of 0.001 for both ResNet-50 and MLP-Mixer, and the SGD optimizer with a momentum of 0.9 and an initial learning rate of 0.05 for MobileNetV2. The weight decay is set to $5e-5$ for MobileNetV2 and MLP-Mixer, and $1e-5$ for ResNet-50. We use cosine annealing for learning rate scheduling. The batch size is set to 128 for CIFAR-100 and Tiny-ImageNet, 256 for ResNet-50 and MobileNetV2 on ImageNet, and 512 for MLP-Mixer.

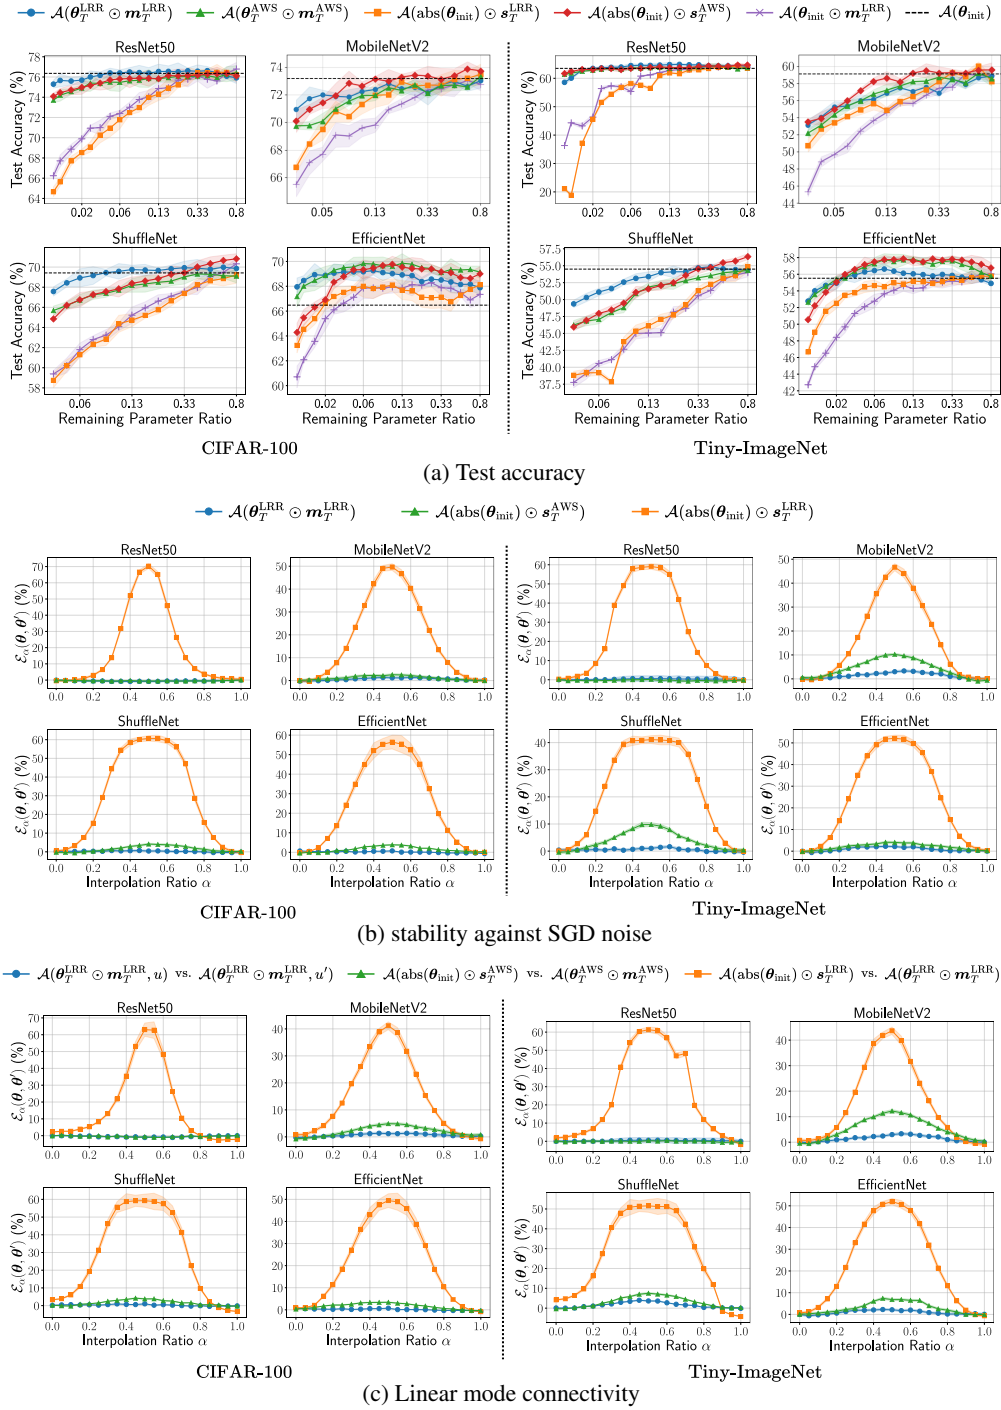


Figure 3: **Main results on CIFAR-100 and Tiny-ImageNet. (a):** Test accuracy of the LRR solution (blue), the AWS solution (green), a randomly initialized network trained with the LRR-driven signed mask (orange), and a randomly initialized network trained with the AWS-driven signed mask (red). **(b) and (c):** Analysis of SGD noise stability and linear mode connectivity, respectively. A randomly initialized network trained with the AWS-driven signed mask exhibits high SGD noise stability and low error barriers along the linear path to the AWS solution (green), contrasting to the case of LRR (orange). In (b) and (c), we use a pruned network with a remaining parameter ratio of approximately 0.09. We report the mean (each point) and standard deviation (shaded area) across 3 trials.

4.2 RESULTS ON CIFAR-100 AND TINY-IMAGENET

Test Performance. In Figure 3a, we report the test performance on CIFAR-100 and Tiny-ImageNet. We compare the performance of six networks after training: (1) initialized dense network, $\mathcal{A}(\theta_{\text{init}})$, (2) LRR subnetwork, $\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$, (3) AWS subnetwork, $\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$, (4) randomly initialized network masked with a signed mask of a LRR subnetwork, $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$, (5) randomly initialized network with a signed mask of a AWS subnetwork, $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$, and (6) randomly initialized network masked with a mask of a LRR subnetwork, $\mathcal{A}(\theta_{\text{init}} \odot m_T^{\text{LRR}})$. Note that $\mathcal{A}(\cdot)$ indicates a normal training algorithm without interpolating normalization layer parameters. First, we note that in most cases, the AWS solution (indicated by the green plots) achieves performance comparable to or better than the LRR solution (indicated by the blue plots). This addresses the potential concern that randomly interpolating normalization parameters in our method could adversely affect the performance of the AWS solution. Then, we transfer the signed mask from LRR, s_T^{LRR} , and AWS, s_T^{AWS} , to a randomly initialized network θ_{init} . In the case of LRR, the performance of a randomly initialized network masked with s_T^{LRR} after training (indicated by the orange plots) is similar to the case when the sign information is not used (indicated by the purple plots), lagging far behind the performance of the LRR solution. On the other hand, we observe that in most cases, the performance of a randomly initialized network masked with s_T^{AWS} after training (indicated by the red plots) is comparable to that of the AWS solution. Finally, We observe that the signed mask obtained through AWS allows a randomly initialized network to perform comparably to a dense network (indicated by the dotted line) after training at non-trivial sparsity as long as the AWS solution performs comparably to the dense network.

stability against SGD-Noise and Linear Mode Connectivity. We further compare the effectiveness of s_T^{AWS} and s_T^{LRR} by examining whether they can transfer the strong generalization potential of the AWS or LRR subnetwork to a randomly initialized network. To this end, we first examine the SGD noise stability of $\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}}$ and $\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}}$. The results in Figure 3b demonstrate that $\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}}$, shown by the orange plots, exhibit significantly high error barriers between a pair of networks trained with different SGD randomness. On the other hand, $\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}}$, represented by the green plots, exhibits comparably low error barriers between a pair of networks trained with different SGD randomness, relative to the reference stability of $\theta_T^{\text{LRR}} \odot s_T^{\text{LRR}}$ shown by the blue plots. Moreover, we examine the linear mode connectivity between $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$ and $\mathcal{A}(\text{abs}(\theta_T^{\text{AWS}}) \odot s_T^{\text{AWS}})$, as well as between $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$ and $\mathcal{A}(\text{abs}(\theta_T^{\text{LRR}}) \odot s_T^{\text{LRR}})$ to confirm whether they lie in the same basin. The results in Figure 3c demonstrate that transferring s_T^{LRR} to a randomly initialized network causes the network to converge to a solution with significantly high error barriers between the LRR solution, as indicated by the orange plots. By contrast, a random initialized network masked with s_T^{AWS} exhibits relatively low error barriers between the AWS solution, as indicated by the green plots, compared to the reference shown by the blue plots.

Discussion. Our motivational experiments in Figure 2 show that the signed mask of the LRR subnetwork alone cannot transfer the generalization potential of the LRR subnetwork to a randomly initialized network; it is also necessary to transfer the normalization layer parameters of the LRR subnetwork. For the goal of LTH, we propose AWS that eliminates the need for the trained normalization parameters. The test performance in Figure 3a and the analysis of the SGD noise stability and linear mode connectivity in Figure 3b and Figure 3c show that a randomly initialized network masked with the signed mask from the AWS subnetwork lies in the basin of attraction of the AWS subnetwork. This demonstrates that s_T^{AWS} itself can transfer the generalization potential of the AWS subnetwork.

4.3 GENERALIZATION TO IMAGENET AND LAYER NORMALIZATION

To validate the effectiveness of AWS on a more challenging task, we evaluate AWS on ImageNet dataset. The results in Table 1 show that across various model and sparsity (i.e., 1- ‘Remaining Params’), a randomly initialized network masked with the AWS-driven signed mask achieves performance comparable to the AWS solution (i.e., $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$ vs. $\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$), whereas masking with the LRR-driven signed mask fails to reach the performance of the LRR solution (i.e., $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$ vs. $\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$). Notably, we observe a similar trend in the results on MLP-Mixer, which uses only layer normalization for normalization. These results show that the signed mask of the AWS subnetwork can transfer the generalization potential of the AWS subnetwork to a randomly initialized network on a more challenging dataset and with layer normalization, further demonstrating the effectiveness of AWS.

Table 1: **Experimental results on ImageNet.** θ_{init} indicates randomly initialized parameters and ‘Remaining Params.’ refers to the remaining parameter ratio. We present more results in Table 2.

Model	Method	Use AWS subnetwork	Trained from θ_{init}	Top-1 Acc. (%)	Remaining Params.
ResNet50	Dense Network	-	-	75.26 ± 0.27	1
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	75.11 ± 0.27	0.27
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	74.48 ± 0.16	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	73.20 ± 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	74.65 ± 0.20	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	74.58 ± 0.25	0.14
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	73.63 ± 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	71.94 ± 0.22	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	73.65 ± 0.12	
MobileNetV2	Dense Network	-	-	69.42 ± 0.41	1
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	69.25 ± 0.33	0.41
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	68.75 ± 0.22	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	67.99 ± 0.11	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	68.66 ± 0.24	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	67.76 ± 0.17	0.21
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	66.17 ± 0.20	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	64.03 ± 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	66.02 ± 0.29	
MLP-Mixer	Dense Network	-	-	59.57 ± 0.12	1
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	59.18 ± 0.26	0.21
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	59.23 ± 0.18	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	57.54 ± 0.30	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	59.17 ± 0.19	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	58.61 ± 0.23	0.11
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	58.77 ± 0.11	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	56.91 ± 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	58.62 ± 0.35	

4.4 COMPARISON TO RELATED WORK

There are existing works that also aim to search for an effective signed mask (Koster et al., 2022; Sreenivasan et al., 2022). Among them, we compare our AWS to GM (Sreenivasan et al., 2022), as GM is evaluated on a more diverse and large-scale architecture than the method of Koster et al. (2022). Figure 4 demonstrates that while GM performance falls short of that of the dense network across all levels of sparsity, a randomly initialized network masked with s_T^{AWS} can perform comparably to both the AWS subnetwork and the dense network until a sparsity of about 0.9. We also highlight that our work provides valuable insights into the role of parameter signs and the influence of normalization layers concerning finding a winning ticket. Moreover, it is noteworthy that the signed mask obtained by AWS can be applied to any random initialization. In contrast, GM trains a signed mask tailored for a specific initialization, which likely limits its applicability to other initialization.

5 CONCLUSION

Finding a winning ticket is still an open problem in the field of the lottery ticket hypothesis (LTH). In this work, we show that an effective signed mask, a sparse mask with sign information, is crucial for conveying information that enables an initialized network to achieve strong generalization. We observe that the signed mask and normalization parameters from the subnetwork trained by learning rate rewinding (LRR) can transfer the generalization potential of the LRR subnetwork to any randomly initialized network. To progress towards the goal of LTH, we propose AWS, a slight variation of LRR, that mitigates the reliance on the trained normalization layer parameters by encouraging low error barriers between the AWS subnetwork and its counterpart with initialized normalization parameters. In contrast to the existing methods limited to finding a winning ticket with an ad-hoc initialization, we demonstrate that the signed mask from the AWS subnetwork can allow any randomly initialized network to reside within the basin of the AWS subnetwork, possibly leading the resulting network to generalize as well as the dense network. For future work, we will investigate the effectiveness of the signed mask acquired through AWS in the transfer learning scenario.




Figure 4: **Comparison to GM on CIFAR-100 with ResNet-32.** The results of GM are approximated from (Sreenivasan et al., 2022).

ACKNOWLEDGMENTS

This work was supported in part by the IITP grants [No.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University), No. 2021-0-02068, and No.2023-0-00156], and the NOTIE grant (No. RS-2024-00432410) by the Korean government.

REFERENCES

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016. URL <https://arxiv.org/abs/1607.06450>.
- Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019. doi: 10.1073/pnas.1903070116. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1903070116>.
- Rebekka Burkholz. Most activation functions can win the lottery without excessive depth. In *NeurIPS*, 2022.
- Xiaohan Chen, Jason Zhang, and Zhangyang Wang. Peek-a-boo: What (more) is disguised in a randomly weighted neural network, and how to find it efficiently. In *ICLR*, 2022.
- Arthur da Cunha, Emanuele Natale, and Laurent Viennot. Proving the lottery ticket hypothesis for convolutional neural networks. In *ICLR*, 2022.
- Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred A. Hamprecht. Essentially no barriers in neural network energy landscape. In *ICML*, 2018.
- Utku Evci, Yani A. Ioannou, Cem Keskin, and Yann Dauphin. Gradient flow in sparse neural networks and how lottery tickets win. In *AAAI*, 2022.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: finding sparse, trainable neural networks. In *ICLR*, 2019.
- Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *ICML*, 2020a.
- Jonathan Frankle, David J. Schwab, and Ari S. Morcos. The early phase of neural network training. In *ICLR*, 2020b.
- C. Daniel Freeman and Joan Bruna. Topology and geometry of half-rectified network optimization. In *ICLR*, 2017.
- Advait Gadhikar and Rebekka Burkholz. Masks, signs, and learning rate rewinding. In *ICLR*, 2024.
- Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry Vetrov, and Andrew Gordon Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. In *NIPS*, 2018.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- Sergey Ioffe and Christian Szegedy. Batch normalization: accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- Nils Koster, Oliver Grothe, and Achim Rettinger. Signing the supermask: Keep, hide, invert. In *ICLR*, 2022.
- Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. University of Toronto, 2009.
- Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *ICLR*, 2017.

- Mingbao Lin, Rongrong Ji, Yan Wang, Yichen Zhang, Baochang Zhang, Yonghong Tian, and Ling Shao. Hrank: filter pruning using high-rank feature map. In *CVPR*, 2020.
- Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *ICCV*, 2017.
- Ekdeep Singh Lubana, Eric J. Bigelow, Robert P. Dick, David Krueger, and Hidenori Tanaka. Mechanistic mode connectivity. In *ICML*, 2023.
- Xiaolong Ma, Geng Yuan, Xuan Shen, Tianlong Chen, Xuxi Chen, Xiaohan Chen, Ning Liu, Minghai Qin, Sijia Liu, Zhangyang Wang, and Yanzhi Wang. Sanity checks for lottery tickets: Does your winning ticket really win the jackpot? In *NeurIPS*, 2021.
- Eran Malach, Gilad Yehudai, Shai Shalev-Shwartz, and Ohad Shamir. Proving the lottery ticket hypothesis: Pruning is all you need. In *ICML*, 2020.
- Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. The role of over-parametrization in generalization of neural networks. In *ICLR*, 2019.
- Quynh Nguyen. On connected sublevel sets in deep learning. In *ICML*, 2019.
- Laurent Orseau, Marcus Hutter, and Omar Rivasplata. Logarithmic pruning is all you need. In *NeurIPS*, 2020.
- Mansheej Paul, Feng Chen, Brett W. Larsen, Jonathan Frankle, Surya Ganguli, and Gintare Karolina Dziugaite. Unmasking the lottery ticket hypothesis: What’s encoded in a winning ticket’s mask? In *ICLR*, 2023.
- Alex Renda, Jonathan Frankle, and Michael Carbin. Comparing rewinding and fine-tuning in neural network pruning. In *ICLR*, 2020.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *IJCV*, 2015.
- Keitaro Sakamoto and Issei Sato. Analyzing lottery ticket hypothesis from pac-bayesian theory perspective. In *NeurIPS*, 2022.
- Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: inverted residuals and linear bottlenecks. In *CVPR*, 2018.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- Kartik Sreenivasan, Jy yong Sohn, Liu Yang, Matthew Grinde, Alliot Nagle, Hongyi Wang, Eric Xing, Kangwook Lee, and Dimitris Papailiopoulos. Rare gems: Finding lottery tickets at initialization. In *NeurIPS*, 2022.
- Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019.
- Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision. In *NeurIPS*, 2021.
- Qingyang Wang, Michael A. Powell, Ali Geisa, Eric Bridgeford, Carey E. Priebe, and Joshua T. Vogelstein. Why do networks have inhibitory/negative connections? In *ICCV*, 2023a.
- Qingyang Wang, Michael A. Powell, Ali Geisa, Eric W. Bridgeford, and Joshua T. Vogelstein. Polarity is all you need to learn and transfer faster. In *ICML*, 2023b.
- Haoran You, Chaojian Li, Pengfei Xu, Yonggan Fu, Yue Wang, Xiaohan Chen, Richard G. Baraniuk, Zhangyang Wang, and Yingyan Lin. Drawing early-bird tickets: Towards more efficient training of deep networks. In *ICLR*, 2020.

Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, 2018.

Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. In *NeurIPS*, 2019.

A IMPORTANCE OF PARAMETER SIGNS IN NORMALIZATION LAYERS.

In Section 3.2, we observe that preserving the signs of parameters in the LRR subnetwork while initializing their magnitudes, excluding those in normalization layers, allows the resulting network to stay within the basin of attraction of the LRR subnetwork, but it fails when the normalization layer parameters are initialized together. We claim that this occurs since the parameters in normalization layers rely more on the magnitude of parameters rather than their signs. The scaling factor in a normalization layer is nearly always positive, thus its sign information may be useless. In the case of the bias factor, it loses its sign information after initialized to 0, but it does not necessarily mean its sign information is not beneficial. To further validate our claim, we also compare the case where the bias factor in normalization layers is initialized to a constant, thus their sign information does not disappear after initialization. In Figure 5, we observe that maintaining the signs of bias factors (indicated by the purple plots) is not beneficial compared to the original initialization case (indicated by the orange plots), demonstrating our claim.

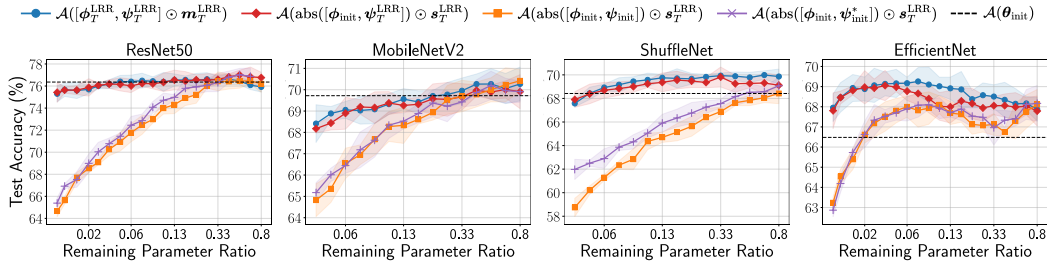


Figure 5: **Effect of transferring the sign of normalization layer parameters.** ψ_{init}^* denotes the initialized normalization layer parameters whose scaling and bias factors are set to 1 and 0.1. We conduct the experiments on CIFAR-100.

B IS AWS ALWAYS BETTER THAN LRR?

In the main manuscript, we show that after training, the post-training performance of a randomly initialized network masked with s_T^{LRR} (i.e., $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$) lags far behind that of the LRR solution. However, we found that for several network architectures, the performance of $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$ is similar to that of the LRR solution. In Figure 6a, we present the test performance on VGG11-bn (Simonyan & Zisserman, 2015). We observe that $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$ (indicated by the orange plots) achieves performance similar to that of the LRR solution (indicated by the blue plots). Thus, the performance difference between $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$ and $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$ (indicated by the red plots) is trivial. We also investigate the SGD noise stability of $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$ and $\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$ and their linear mode connectivity to the corresponding LRR or AWS solution in Figure 6b and Figure 6c, respectively. We observe that the signed masks from both LRR and AWS enable a randomly initialized network to remain stable against SGD noise and converge to a solution with linear mode connectivity to the corresponding LRR or AWS solution. Thus, in some cases, LRR can yield an effective signed mask that transfers the generalization potential of the LRR subnetwork to a randomly initialized network. However, we argue that AWS is more effective and generalizable to a wider range of more complex architectures as demonstrated in Figure 3.




Figure 6: **Results of LRR on CIFAR-100 with VGG11-bn.**

Table 2: **Experimental results on ImageNet.** θ_{init} indicates randomly initialized parameters and ‘Remaining Params.’ refers to the remaining parameter ratio.

Model	Method	Use AWS subnetwork	Trained from θ_{init}	Top-1 Acc. (%)	Remaining Params.
ResNet50	Dense Network	-	-	75.26 \pm 0.27	1
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	75.11 \pm 0.27	0.27
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	74.48 \pm 0.16	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	73.20 \pm 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	74.65 \pm 0.20	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	74.92 \pm 0.10	0.21
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	74.06 \pm 0.18	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	72.96 \pm 0.20	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	74.28 \pm 0.19	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	74.62 \pm 0.12	0.17
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	73.76 \pm 0.26	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	72.46 \pm 0.18	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	73.99 \pm 0.23	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	74.58 \pm 0.25	0.14
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	73.63 \pm 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	71.94 \pm 0.22	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	73.65 \pm 0.12	
MobileNetV2	Dense Network	-	-	69.42 \pm 0.41	1
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	69.25 \pm 0.33	0.41
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	68.75 \pm 0.22	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	67.99 \pm 0.11	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	68.66 \pm 0.24	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	68.88 \pm 0.18	0.33
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	68.13 \pm 0.32	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	67.01 \pm 0.19	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	68.04 \pm 0.28	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	68.36 \pm 0.20	0.27
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	67.39 \pm 0.24	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	64.92 \pm 0.25	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	67.11 \pm 0.16	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	67.76 \pm 0.17	0.21
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	66.17 \pm 0.20	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	64.03 \pm 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	66.02 \pm 0.29	
MLP-Mixer	Dense Network	-	-	59.57 \pm 0.12	1
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	59.18 \pm 0.26	0.21
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	59.23 \pm 0.18	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	57.54 \pm 0.30	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	59.17 \pm 0.19	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	58.95 \pm 0.21	0.17
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	59.01 \pm 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	57.21 \pm 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	58.93 \pm 0.19	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	58.83 \pm 0.23	0.14
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	58.79 \pm 0.24	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	57.09 \pm 0.19	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	58.73 \pm 0.25	
	$\mathcal{A}(\theta_T^{\text{LRR}} \odot m_T^{\text{LRR}})$	\times	\times	58.61 \pm 0.23	0.11
	$\mathcal{A}(\theta_T^{\text{AWS}} \odot m_T^{\text{AWS}})$	\checkmark	\times	58.77 \pm 0.11	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{LRR}})$	\times	\checkmark	56.91 \pm 0.21	
	$\mathcal{A}(\text{abs}(\theta_{\text{init}}) \odot s_T^{\text{AWS}})$	\checkmark	\checkmark	58.62 \pm 0.35	