Strong Lottery Ticket Hypothesis with ε–perturbation

author names withheld

Under Review for OPT 2022

Abstract

The strong Lottery Ticket Hypothesis (LTH) [18, 25] claims that there exists a subnetwork in a sufficiently large, randomly initialized neural network that approximates some target neural network without the need of training. This work extends the theoretical guarantee of the strong LTH literature to a scenario more similar to the original LTH, by generalizing the weight change achieved in the pre-training step to some perturbation around the initialization. In particular, we focus on the following open questions: *By allowing an ε-scale perturbation on the random initial weights, can we reduce the over-parameterization requirement for the candidate network in the strong LTH? Furthermore, does the weight change by SGD coincide with a good set of such perturbation?*

We answer the first question by first extending the theoretical result on subset sum [14] to a scenario that allows perturbation on the candidates, and forms the conjecture of the perturbed strong LTH by applying our generalized theoretical result to a neural network formulation. To answer the second question, we show via experiments that, when a larger perturbation is allowed, the required over-parameterization of the strong LTH decreases, and the final accuracy after pruning increases.

1. Introduction

Pruning techniques for over-parameterized neural networks have drawn growing attention in recent years [1, 10–13, 19–21, 26]. Amongst them, the *Lottery Ticket Hypothesis* (LTH) [8, 9] claims the existence of a small (sparse) subnetwork within a large (dense) neural network such that, when trained in isolation, achieves comparable or even better performance than the original dense network. Such subnetworks can be identified by pre-training the dense network and pruning it based on the magnitude of the learned weights [8]. Currently, to the best of our knowledge, the LTH lacks of any rigorous theoretical guarantees that justify superior performance of the subnetwork, especially under the pretraining-based pruning; yet, it has been proven to be effective in practice.

The *Strong Lottery Ticket Hypothesis* [18, 25] leverages a different pruning scheme: given a target dense neural network, and a randomly initialized, *sufficiently over-parameterized* candidate network, there exists a subnetwork in the latter that approximates the former arbitrarily well without the need of training. While we require a significant over-parameterization in the randomly initialized network, the strong LTH enjoys extensive theoretical guarantees [15–17]. Yet, the same theory hardly applies to the LTH, as strong LTH assumes that the candidate weights pruned are fixed at initialization. *That LTH pruning is based on weights modified by pre-training motivates us to analyze the approximation behavior that emerges beyond the randomness in the candidate weights.*

Further study on the pre-training process of LTH shows that the lottery ticket emerges in the early stage of training [23]. This implies that converging to a small training loss is not necessarily the intent of pre-training in the LTH procedure; *in other words, achieving small loss does not necessarily*

© .

explain how and why pre-training helps pruning in LTH. Instead, one could hypothesize that the pre-training –based on loss minimization– could *guide* the weight perturbation to a direction that facilitate the pruning. Based on this hypothesis, our work extends the theoretical guarantee of strong LTH to a scenario more similar to the original LTH, by generalizing the weight change in the pre-training step to some perturbation around the initialization. Our central question is as follows:

> *"By allowing an ε-scale perturbation on the random initial weights, can we reduce the over-parameterization requirement for the candidate network in the strong LTH? Furthermore, does the weight change by SGD coincide with a good set of such perturbation?"*

To be more specific, let $f_{\boldsymbol{\theta}}$ be a neural network with a set of parameters $\boldsymbol{\theta} \in \mathbb{R}^d$. Formally, an ε-perturbation is a mapping $\mathcal{P} : \mathbb{R}^d \to \mathbb{R}^d$ such that the maximum entrywise perturbation is bounded in absolute value by ε, i.e., $\|\mathcal{P}(\boldsymbol{\theta}) - \boldsymbol{\theta}\|_\infty \leq \varepsilon$. We point out that this definition of perturbation generalizes to two existing scenarios: when $\varepsilon = 0$ –i.e., we allow no weight perturbation– the question above reduces to the original strong LTH. When $\varepsilon = \infty$ –i.e., we allow arbitrarily large weight perturbation– the required over-parameterization for the candidate network is at most the size of the target network. In this case, one would use gradient based optimizers such as SGD to find such weight perturbation, but often without the need of pruning. Yet, both cases cover only one aspect in pruning and perturbation.

In this paper, we study the inter-dependence of the two aspects above by treating ε as a variable. In particular, we show that a larger perturbation scale ε, which corresponds to a larger amount of training, would alleviate the over-parameterization requirement, while keeping the accuracy of the pruned neural network the same. Our contributions can be summarized as below:

- We consider a generalized version of the subset sum problem where each candidate in the summation is allowed perturbation bounded by a fixed scale ε. We extend the analysis of the subset sum [14] to our generalized version, and show that when a larger perturbation is allowed, the required size of the candidate set can be reduced. We empirically validate our theoretical result on the perturbed subset sum problem.
- Applying the theoretical result above to neural networks, we conjecture that, when an ε-scale perturbation is allowed, the strong LTH on randomly initialized neural network requires less over-parameterization to achieve a specific approximation error.
- On neural networks, we empirically show that $i)$ the perturbation that alleviates the overparameterization requirement of the strong LTH can be obtained by projected SGD on the initialized weights; and, $ii)$under fixed overparameterization, neural networks with a larger freedom over the level of perturbation achieves a higher accuracy after pruning. This result establish the connection between the amount of pre-training and the accuracy of the pruned network.

2. Related Works

Lottery Ticket Hypothesis. Several works attempt to explain the LTH theoretically. [7] empirically study the behavior of gradient flow in the pruned network. [24] assumes that the optimal mask is given, and proves that the pruned network achieves faster convergence and better generalization when trained from initialization. [22] provides a theoretical guarantee of the loss in a pruning-after-training fashion. However, these works differ from ours in the following: $i)$ they usually consider neuron pruning on small neural network archictures (e.g., [22] focuses on a two-layer neural network with smooth activations), while our work considers weight pruning of a deep

ReLU neural network; $ii)$ they consider minimizing the loss on a specific dataset, and require an over-parameterization that scales quadratically with the number of samples, while we approximate a target network with a fixed architecture in terms of the function norm, and require an over-parameterization that scales with the width of the target network.

Strong Lottery Ticket Hypothesis. The strong LTH originates from the empirical observation that, by fixing the weights at initialization and learning the mask the weights, one can identify subnetworks that achieve comparable accuracy to the dense one with learned weights [25]. [18] improved this idea by proposing the `edge-popup` algorithm to efficiently learn the mask. [15] first proved such hypothesis under the assumption that the dense network's size scales polynomially with the target network's width and depth. Leveraging the advantage of weight decomposition and theoretical results on the subset sum problem [14], [16] and [17] improved the over-parameterization requirement to a logarithm factor times the size of the target network. Later work explores different variations of the strong LTH: [6] and [2] show that the strong LTH holds in the case of convolutional neural networks. [3] extends the strong LTH proof to a universial family of functions. Finally, [4] further reduces the over-parameterization requirement by employing the iterative randomization.

3. Notations and Setup

Notation. For a vector $\mathbf{a}$, we use $\|\mathbf{a}\|_2$ to denote its ℓ_2 (Euclidean) norm, and $\|\mathbf{a}\|_\infty$ to denote its ℓ_∞ norm. For a matrix $\mathbf{A}$, we use $\|\mathbf{A}\|_{\max} = \max_{ij} |A_{ij}|$ to denote its max norm. Moreover, $\texttt{Unif}\,(I)$ denotes the uniform distribution on the interval I, and $\texttt{Geom}\,(\cdot)$ and $\texttt{Bin}\,(\cdot,\cdot)$ denotes the geometric and binomial distributions, respectively. We use $\sigma\,(a) = \max\{0, a\}$ to denote the ReLU activation. We use $\min\{\cdot\}$ refers to the entrywise minimum and $\text{abs}(\cdot)$ refers to the entrywise absolute value.
Setup. Similar to [17], our focus is to approximate an L-layer, ReLU activated target multi-layer perceptron (MLP) $f(\mathbf{x})$ by pruning a $2L$-layer, ReLU activated candidate MLP $g_{\boldsymbol{\theta}}(\mathbf{x})$. For a input vector $\mathbf{x} \in \mathbb{R}^{d_0}$, we assume $f(\mathbf{x})$ and $g_{\boldsymbol{\theta}}(\mathbf{x})$ are represented by:

$$f(\mathbf{x}) = \mathbf{W}^L \sigma\left(\mathbf{W}^{L-1} \ldots \sigma\left(\mathbf{W}^1 \mathbf{x}\right)\right) g(\mathbf{x}) = \mathbf{U}^{2L} \sigma\left(\mathbf{U}^{2L-1} \ldots \sigma\left(\mathbf{U}^1 \mathbf{x}\right)\right)$$

We consider the pruning of $g_{\boldsymbol{\theta}}(\mathbf{x})$ with masks for the weights $\mathcal{M} = \{\mathbf{M}^\ell\}_{\ell=1}^{2L}$, denoted as $g_{\mathcal{M},\boldsymbol{\theta}}(\mathbf{x})$:

$$g_{\mathcal{M},\boldsymbol{\theta}}(\mathbf{x}) = \left(\mathbf{M}^{2L} \odot \mathbf{U}^{2L}\right) \sigma\left(\left(\mathbf{M}^{2L-1} \odot \mathbf{U}^{2L-1}\right) \ldots \sigma\left(\left(\mathbf{M}^1 \odot \mathbf{U}^1\right) \mathbf{x}\right)\right)$$

For a set of weights $\boldsymbol{\theta}$, we consider its ε-perturbation $\mathcal{P}(\boldsymbol{\theta}) = \{\mathcal{P}\left(\mathbf{U}^\ell\right)\}_{\ell=1}^{2L}$, with $\mathcal{P}$ applied to each $\mathbf{U}^\ell$ such that $\|\mathcal{P}(\mathbf{U}^\ell) - \mathbf{U}^\ell\|_{\max} = \max_{ij} |\mathcal{P}(\mathbf{U}^\ell)_{ij} - \mathbf{U}^\ell_{ij}| \leq \varepsilon$. Our focus in this paper is on the approximation error defined as:

$$\mathcal{L}(f, g) = \min_{\mathcal{M},\mathcal{P}} \sup_{\mathbf{x}\in\mathcal{B}} \left\|f(\mathbf{x}) - g_{\mathcal{M},\mathcal{P}(\theta)}\left(\mathbf{x}\right)\right\|_2. \tag{1}$$

4. Subset Sum with ε-Perturbation

For each layer in the target network, [17] constructed a two-layer subnetwork with block structure, such that each block approximates a single entry in the weight matrix of the target network. In particular, they obtain a logarithmic-scale over-parameterization by formulating the approximation as a subset sum problem [5, 14] Given a candidate set of values $\{x_i\}_{i=1}^n$ of size n and a target value z, the solution to the subset sum problem finds the best approximation of z using the sum of a subset

of $\{x_i\}_{i=1}^n$. From an optimization perspective, the optimal approximation error $\eta^\star$ is the solution to the following problem $\eta^\star = \min_{\boldsymbol{\delta}\in\{0,1\}^n} |\sum_{i=1}^n \delta_i x_i - z|$ where $\delta_i \in \{0,1\}$ is the indicator variable on whether x_i is selected in the sum to approximate z.

From the perspective of strong LTH, we can treat z as the weight entry in the target network that we wish to approximate, and $\{x_i\}_{i=1}^n$ as the weights in the candidate network we will prune. Here, $\delta_i = 1$ means that the i-th weight is kept, while $\delta_i = 0$ means that the i-th weight is pruned. [14] shows that, with high probability over the randomness of $\mathbf{x}_i \sim \texttt{Unif}\left([-1,1]\right)$, a candidate set with size of the order $n = \Omega\left(\log \eta^{-1}\right)$ is enough to guarantee that $\eta^* \leq \eta$ for all $z \in [-1/2, 1/2]$.

As an extension to the strong LTH, our setup incorporates an ε-perturbation on the weights of the random neural network. This calls for the attention of extending the subset sum problem to a version with ε-perturbation. In particular, we consider the following joint minimization problem:

$$\eta^\star = \min_{\boldsymbol{\delta}\in\{0,1\}^n, \mathbf{y}\in[-\varepsilon,\varepsilon]^n} \left| \sum_{i=1}^n \delta_i \left(x_i + y_i\right) - z \right|. \tag{2}$$

We denote the values that lead to the optimal approximation error as $\boldsymbol{\delta}^\star$ and $\mathbf{y}^\star$, respectively. In words, the above problem aims to select values from the set $\{x_i\}_{i=1}^n$ such that, after potential entrywise perturbation by some tunable $y_i \in [-\varepsilon, \varepsilon]$, the summation of the selected and perturbed $\sum_{i=1}^n \delta_i \left(x_i + y_i\right)$ will approximate z.

A central technical difficulty in our work is to extend the result of [14] to incorporate ε-perturbation. Intuitively, as the perturbation scale ε becomes larger, each candidate is susceptible to a larger change in order to better approximate the objective z. Intuitively, this implies that we should require a smaller size for the candidate size. We formally show this in the theorem below.

Theorem 1 *Given a candidate set $\{x_i\}_{i=1}^n$ with $x_i \sim \texttt{Unif}([-1,1])$ for all $i \in [n]$. Consider the ε-perturbed random subset sum problem in Equation 2. If $n \geq K_1 + K_2$ with*

$$K_1 = \Omega\left(\frac{\log \eta^{-1}}{\log\left(5+\varepsilon\right) - \log 4}\right); \quad K_2 = \Omega\left(1 + \frac{\log \eta^{-1}}{(1+2\varepsilon)}\right).$$

Then with probability at least $1 - e^{-\frac{K_1}{8}} - e^{-\frac{K_2}{7}}$, we have that $\eta^* \leq \eta$ holds for all $z \in [-1/2, 1/2]$.

The proof of Theorem (1) is provided in Appendix (A). We sketch the proof below:

1. We start by recursively defining an indicator function $f_{k,\eta}(z)$ corresponding to the event that z has an η-approximation by the first k candidates. Moreover, as in [14], we define p_k to be the fraction of z on the interval $[-1/2, 1/2]$ such that $f_{k,\eta}(z) = 1$. Differently, we show that the expectation of $p_{k+1} - p_k$ is lower bounded by $(1/2+\varepsilon)(1-p_k)p_k$. This demonstrates the expected growth p_{k+1} enjoys from p_k. Next, we apply several techniques to remove the expectation.

2. We first show the lower bound on K_1 such that $p_{K_1} \geq 1/4$ with a high probability by partitioning the interval $[0, 1/4]$ into sub-intervals that represents geometric growth. We then show that the number of steps to achieve such growth can be represented as a binomial random variable. Lastly, we apply Hoeffding's inequality to bound this random variable.

3. Next, starting from $p_k \geq 1/4$, we lower-bound the summation of $Z_{k+1} = \frac{p_{k+1}-p_k}{p_k(1-p_k)}$ using Azuma's inequality. To relate the summation of Z_{k+1} to the growth of p_k, we define a function $\psi(p)$ such that $\psi(p_{k+1}) - \psi(p_k) \geq Z_{k+1}$. In this way, we arrive at a lower bound on $\psi(p_{K_1+K_2}) - \psi(1/4)$. Enforcing a lower bound on $p_{K_1+K_2}$ gives a lower bound on K_2.

Remark. Notice that the lower bound on the candidate set n depends on K_1 and K_2, where K_1 scales inversely with $\log(5+\varepsilon)$ and K_2 scales inversely with $1+2\varepsilon$. This implies that n decreases monotonically as ε increases. We utilize this result to analyze the approximation error in Equation 1.

5. Strong Lottery Ticket Hypothesis with ε Perturbation

Lemma 1 *Let $g:\mathbb{R}\to\mathbb{R}$ be a randomly initialized network of the form $g(x)=\boldsymbol{v}^\top\sigma(\boldsymbol{u}x)$, where $\boldsymbol{v},\boldsymbol{u}\in\mathbb{R}^n$, $n=K_1+K_2$,*

$$K_1 \geq C_1\left(\frac{\log(2(\eta^{-1}))}{\log(5+\varepsilon)-\log 4}\right), K_2 \geq \left(1+\frac{\log(2(\eta^{-1})))}{1+6\varepsilon}\right),$$

and $v_i's$ are drawn i.i.d. from $\texttt{Unif}[-1,1]$, $u_i=1$. Then, with probability at least $1-\exp\left(-\frac{K_2-1}{32}\right)-\exp\left(-\frac{K_1}{7}\right)$, we have:

$$\forall w\in[-1/2,1/2],\ \exists\, \boldsymbol{s}\in\{0,1\}^n, \boldsymbol{y}\in[-\varepsilon,+\varepsilon]^n \sup_{x:|x|\leq 1}\left|wx-(\boldsymbol{v}\odot\boldsymbol{s})^\top((\boldsymbol{u}+\boldsymbol{y})x)\right|<\eta.$$

Proof Note that $\boldsymbol{v}^\top\boldsymbol{u}=\boldsymbol{u}$ and $u_i\sim\texttt{Unif}[-1,1]$. ■

The above lemma shows the number of over-parameterization required for approximating a linear target function using a over-parameterized linear network. Note that the lower bound of the required over-parameterization decreases monotonically as ε increases. Based on the Lemma above, we posit the ε-perturbed strong LTH as a conjecture.

Conjecture 1 *Let f be the target network defined in Section 3. Consider a randomly initialized neural network $g(\boldsymbol{x})$ with $2L$ layers, weights $\{U_l\}_{l=1}^{2L}$ where every entry in weight weight is drawn from $\texttt{Unif}[-1,1]$, and $\mathbf{U}_{2i},\mathbf{U}_{2i-1}$ has dimension:*

$$\texttt{dim}(\mathbf{U}_{2i})=d_i\times d_{i-1}\left(K_1+K_2\right);\quad \texttt{dim}(\mathbf{U}_{2i-1})=d_{i-1}\left(K_1+K_2\right)\times d_{i-1},$$

where

$$K_1\geq C_1\frac{\log\left(d_{i-1}d_iL\eta^{-1}\right)}{\log(5+\varepsilon)-\log(4)},\ K_2\geq C_2\left(1+\frac{\log(d_{i-1}d_iL\eta^{-1})}{1+6\varepsilon}\right)$$

Then, with probability at least $1-\exp\left(-\frac{K_2-1}{32}\right)-\exp\left(-\frac{K_1}{7}\right)$, we have $\mathcal{L}(f,g)\leq\eta$.

6. Experiments

6.1. Approximating Neural Nets with SubsetSum and ε Perturbation

We study the effect of weight perturbation on the required overparametrization by approximating a two-layer, 500 hidden node target. Each weight was approximated using a subset sum of n randomly initialized candidates where each candidate was allowed to perturbed by at most ε. In particular, for some given ε, η, we say n satisfies the over-parametrization requirement if η^* in Equation 2 satisfies $\eta^*\leq\eta$. We set randomly generate 10 sets of x_i's, and e record the minimum n such that 8 of such sets make n satisfying the over-parametrization requirement. We vary η from $1e-2$ to $1e-4$ and choose ε such that ε/η varies between 0 and 10. We are interested how such n changes as η and ε/η change. From Figure 1(*a*)subfigure, we can observe that for fixed η, n decreases as ε/η increases. More specifically, as ε increases, more changes in ε are required to make a decrease in the minimum over-parameterization requirement n, which coincides with Theorem 1.

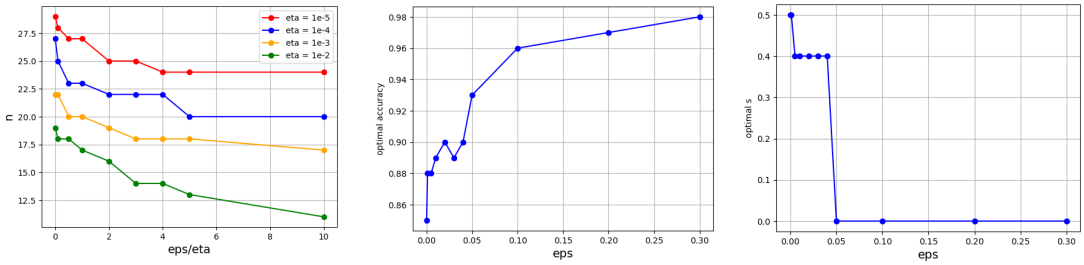

(a) (b) (c)

Figure 1: (a). Change of the required size of the candidate set (n) v.s. relative perturbation scale (ε/η); (b). change of accuracy on the pruned network v.s. the perturbation scale; (c). change of the optimal sparsity in pruning v.s. the perturbation scale.

6.2. Perturbation Using SGD

We hypothesize that *weight perturbation using SGD coincide with the desirable perturbation* $\boldsymbol{\delta}^*$ *in equation (2)*. We propose a two-stage algorithm to validate this hypothesis. With a given perturbation scale ε, we first train an over-parameterized neural network using projected SGD to convergence. Note that by applying the projected SGD we guarantee that each value of the neural network weight stays in an ε-neighborhood of initialization. Then, we run `edge-popup` for a range of pruning (sparsity) levels. We then consider the best accuracy amongst all pruning levels to be the optimal approximation scale. Details referred to Algorithm 1.

We train a four-layer fully connected networks on MNIST and each layer has 500 hidden nodes. We use Algorithm 1 to train the network: we use a learning rate of 0.03 for projected SGD and train the network for 100 epochs; for pruning we use `edge-popup` with a learning rate of 0.1 and train the network for 50 epochs with cosine annealing. All weights in the network are initialized from $\texttt{Unif}\left([-1/2, 1/2]\right)$, and ε ranges from 0 to 0.4. The results are shown in Table 1 and Figure 1(c)subfigure,1(b)subfigure. Note that as the perturbation scale ϵ increase, the optimal approximation error decreases. Also, as ϵ increase, the pruning level that achieves the best approximation error decreases.

s \ ε	0	0.001	0.005	0.01	0.02	0.03	0.04	0.05	0.1	0.2	0.4
0	0.12	0.14	0.25	0.42	0.68	0.84	**0.90**	**0.93**	**0.96**	**0.97**	**0.98**
0.1	0.49	0.48	0.65	0.70	0.78	0.82	0.87	0.87	0.94	0.97	0.98
0.2	0.75	0.76	0.77	0.79	0.84	0.86	0.88	0.87	0.93	0.96	0.97
0.3	0.83	0.82	0.82	0.82	0.88	0.88	0.86	0.90	0.92	0.94	0.93
0.4	0.82	0.86	**0.88**	**0.89**	**0.90**	**0.89**	0.90	0.90	0.88	0.91	0.86
0.5	**0.85**	**0.88**	0.86	0.89	0.87	0.88	0.89	0.89	0.90	0.89	0.76
0.6	0.83	0.87	0.87	0.83	0.86	0.88	0.87	0.88	0.87	0.85	0.54
0.7	0.81	0.85	0.84	0.83	0.86	0.82	0.81	0.81	0.79	0.74	0.29
0.8	0.73	0.71	0.71	0.75	0.77	0.75	0.73	0.68	0.77	0.55	0.17

Table 1: Test accuracy for different sparsity level s and perturbation ε. In each column, the highest accuracy is marked bold.

REFERENCES

[1] Davis Blalock, Jose Javier Gonzalez Ortiz, Jonathan Frankle, and John Guttag. What is the state of neural network pruning?, 2020. URL https://arxiv.org/abs/2003.03033.

[2] Rebekka Burkholz. Convolutional and residual networks provably contain lottery tickets. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 2414–2433. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/burkholz22a.html.

[3] Rebekka Burkholz, Nilanjana Laha, Rajarshi Mukherjee, and Alkis Gotovos. On the existence of universal lottery tickets. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=SYB4WrJql1n.

[4] Daiki Chijiwa, Shin'ya Yamaguchi, Yasutoshi Ida, Kenji Umakoshi, and Tomohiro INOUE. Pruning randomly initialized neural networks with iterative randomization. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL https://openreview.net/forum?id=QCPY2eMXYs.

[5] Arthur da Cunha, Francesco d'Amore, Frédéric Giroire, Hicham Lesfari, Emanuele Natale, and Laurent Viennot. Revisiting the random subset sum problem, 2022. URL https://arxiv.org/abs/2204.13929.

[6] Arthur da Cunha, Emanuele Natale, and Laurent Viennot. Proving the lottery ticket hypothesis for convolutional neural networks. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=Vjki79-619-.

[7] Utku Evci, Yani Ioannou, Cem Keskin, and Yann Dauphin. Gradient flow in sparse neural networks and how lottery tickets win. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(6):6577–6586, Jun. 2022. doi: 10.1609/aaai.v36i6.20611. URL https://ojs.aaai.org/index.php/AAAI/article/view/20611.

[8] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=rJl-b3RcF7.

[9] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M. Roy, and Michael Carbin. Stabilizing the lottery ticket hypothesis, 2019. URL https://arxiv.org/abs/1903.01611.

[10] Song Han, Jeff Pool, John Tran, and William J. Dally. Learning both weights and connections for efficient neural networks, 2015. URL https://arxiv.org/abs/1506.02626.

[11] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks, 2017. URL https://arxiv.org/abs/1707.06168.

[12] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Snip: Single-shot network pruning based on connection sensitivity, 2018. URL https://arxiv.org/abs/1810.02340.

[13] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets, 2016. URL https://arxiv.org/abs/1608.08710.

[14] George S. Lueker. Exponentially small bounds on the expected optimum of the partition and subset sum problems. *Random Structures & Algorithms*, 12(1):51–62, 1998. doi: https://doi.org/10.1002/(SICI)1098-2418(199801)12:1⟨51::AID-RSA3⟩3.0.CO;2-S. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291098-2418%28199801%2912%3A1%3C51%3A%3AAID-RSA3%3E3.0.CO%3B2-S.

[15] Eran Malach, Gilad Yehudai, Shai Shalev-Schwartz, and Ohad Shamir. Proving the lottery ticket hypothesis: Pruning is all you need. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6682–6691. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/malach20a.html.

[16] Laurent Orseau, Marcus Hutter, and Omar Rivasplata. Logarithmic pruning is all you need. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 2925–2934. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/file/1e9491470749d5b0e361ce4f0b24d037-Paper.pdf.

[17] Ankit Pensia, Shashank Rajput, Alliot Nagle, Harit Vishwakarma, and Dimitris Papailiopoulos. *Optimal Lottery Tickets via SUBSETSUM: Logarithmic over-Parameterization is Sufficient*. Curran Associates Inc., Red Hook, NY, USA, 2020. ISBN 9781713829546.

[18] Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, and Mohammad Rastegari. What's hidden in a randomly weighted neural network?, 2019. URL https://arxiv.org/abs/1911.13299.

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

[20] Yulong Wang, Xiaolu Zhang, Lingxi Xie, Jun Zhou, Hang Su, Bo Zhang, and Xiaolin Hu. Pruning from scratch, 2019. URL https://arxiv.org/abs/1909.12579.

[21] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks, 2016. URL https://arxiv.org/abs/1608.03665.

[22] Cameron R. Wolfe, Qihan Wang, Junhyung Lyle Kim, and Anastasios Kyrillidis. How much pre-training is enough to discover a good subnetwork?, 2021. URL https://arxiv.org/abs/2108.00259.

[23] 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, 2019. URL https://arxiv.org/abs/1909.11957.

[24] Shuai Zhang, Meng Wang, Sijia Liu, Pin-Yu Chen, and Jinjun Xiong. Why lottery ticket wins? a theoretical perspective of sample complexity on sparse neural networks. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 2707–2720. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper/2021/file/15f99f2165aa8c86c9dface16fefd281-Paper.pdf.

[25] Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/2019/file/1113d7a76ffceca1bb350bfe145467c6-Paper.pdf.

[26] Michael Zhu and Suyog Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression, 2017. URL https://arxiv.org/abs/1710.01878.

Appendix A. Proof of Theorem 1

The subset sum problem considers finding $\boldsymbol{\delta} \in \{0,1\}^n$ that minimizes $\ell(z, \boldsymbol{\delta}) = |z - \sum_{i=1}^n \delta_i x_i|$ for a given z and given x_i's. Previous work finds that, with $n = \Omega(\log \frac{1}{\eta})$, it holds with high probability that there exists $\boldsymbol{\delta} \in \{0,1\}^n$ such that $\ell(z, \boldsymbol{\delta}) \leq \eta$. Alternatively, this problem can be started as finding the smallest n such that $\eta^* \leq \eta$ with

$$\eta^* = \min_{\boldsymbol{\delta} \in \{0,1\}^n} \ell(z, \boldsymbol{\delta}).$$

In our case, we would like to give the freedom of each x_i to be perturbed for a small degree ε. In particular, we extend the definition of ℓ to

$$\ell(z, \boldsymbol{\delta}, \mathbf{y}) = \left| z - \sum_{i=1}^n \delta_i (x_i + y_i) \right|$$

and seeks condition of n such that $\eta^* \leq \eta$ with

$$\eta^* = \min_{\boldsymbol{\delta} \in \{0,1\}^n, \mathbf{y} \in [-\varepsilon, \varepsilon]^n} \ell(z, \boldsymbol{\delta}, \mathbf{y}). \tag{3}$$

If this condition is met for a fixed z, we say that such z has an η approximation.

Assumption 1 *Let the candidate values $x_i \sim \texttt{Unif}([-1,1])$ for all $i \in [n]$, and the target value $z \in [-\frac{1}{2}, \frac{1}{2}]$. Let $0 \leq \varepsilon \leq \eta \leq 1$ be given.*

Under this assumption, we attempts to prove the following theorem

Theorem 2 *Let the number of candidates satisfy $n = K_1 + K_2$ with*

$$K_1 = O\left(\frac{\log \eta^{-1}}{\log(5+\varepsilon) - \log 4}\right) \qquad K_2 = O\left(1 + \frac{\log \eta^{-1}}{(1+6\varepsilon)}\right).$$

Then with probability at least $1 - e^{-\frac{K_1}{7}} - e^{-\frac{K_2 - 1}{8}}$ we have that all $z \in [-\frac{1}{2}, \frac{1}{2}]$ has a 2η-approximation.

We define the indicator function for the existence of $\hat{\eta}$-approximation within the first k candidate.

$$f_{k,\hat{\eta}}(z) = \mathbb{I}\left\{ \exists \delta \in \{0,1\}^k, y \in [-\varepsilon, \varepsilon]^k \text{ s.t. } \left| \sum_{i=1}^k \delta_i (x_i + y_i) - z \right| \leq \hat{\eta} \right\}, z \in [-\frac{1}{2}, \frac{1}{2}]$$
$$f_{k,\hat{\eta}}(z) = 0, \text{otherwise}.$$

This indicator function has the following recurrence

$$f_{0,\eta} = \mathbb{I}\{|z| \leq \eta\}; \quad f_{k+1,\eta} = f_{k,\eta}(z) + (1 - f_{k,\eta}(z))\, f_{k,\eta+\varepsilon}(z - x_{k+1}).$$

Define the following random variable (depending on $\{x_k\}_{i=1}^k$)

$$p_k = \int_{-\frac{1}{2}}^{\frac{1}{2}} f_{k,\eta}(z)\, dz.$$

A.1. Martingale Analysis of p_k

We can compute its expectation with respect to x_{k+1} (and consider $\{x_i\}_{i=1}^k$ fixed), as

$$\begin{aligned}
\mathbb{E}\left[p_{k+1}\right] &= \mathbb{E}\left[\int_{-\frac{1}{2}}^{\frac{1}{2}} f_{k+1,\eta}(z)dz\right] \\
&= \frac{1}{2}\left(\int_{-1}^{1}\int_{-\frac{1}{2}}^{\frac{1}{2}} f_{k,\eta}(z)dzdx + \int_{-\frac{1}{2}}^{\frac{1}{2}}\left(1-f_{k,\eta}(z)\right)\int_{-1}^{1} f_{k,\eta+\varepsilon}(z-x)dxdz\right) \\
&\geq p_k(z) + \frac{1}{2}\int_{-\frac{1}{2}}^{\frac{1}{2}}(1-f_{k,\eta}(z))\int_{-1}^{1}\max_{y\in[-\varepsilon,\varepsilon]} f_{k,\eta}(z-x-y)\mathbb{I}\left\{z-x\in\left[-\frac{1}{2},\frac{1}{2}\right]\right\}dxdz
\end{aligned}$$

Notice that, for all $z-x\in\left[-\frac{1}{2},\frac{1}{2}\right]$, we have that $2\varepsilon(z-x)\in[-\varepsilon,\varepsilon]$. Therefore

$$\max_{y\in[-\varepsilon,\varepsilon]} f_{k,\eta}(z-x-y) \geq f_{k,\eta}(z-x+2\varepsilon(z-x)) = f_{k,\eta}((1+2\varepsilon)(z-x)).$$

Thus, by changing variable $u=(1+2\varepsilon)(z-x)$, we have

$$\begin{aligned}
\mathbb{E}\left[p_{k+1}\right] &\geq p_k + \frac{1}{2}\int_{-\frac{1}{2}}^{\frac{1}{2}}(1-f_{k,\eta}(z))\int_{-1}^{1} f_{k,\eta}((1+2\varepsilon)(z-x))\mathbb{I}\left\{z-x\in\left[-\frac{1}{2},\frac{1}{2}\right]\right\}\,dxdz \\
&= p_k + \frac{1}{2}\int_{-\frac{1}{2}}^{\frac{1}{2}}(1-f_{k,\eta}(z))\int_{-\frac{1}{2}-\varepsilon}^{\frac{1}{2}+\varepsilon}(1+2\varepsilon)f_{k,\eta}(u)\,dudz \\
&= p_k + \left(\varepsilon+\frac{1}{2}\right)(1-p_k)p_k.
\end{aligned}$$

To facilitate the proof, we define the following auxiliary variable

$$\tilde{p}_{k+1} = p_k + \int_{-\frac{1}{2}}^{\frac{1}{2}}(1-f_{k,\eta}(z))f_{k,\eta}((1+2\varepsilon)(z-x_{k+1}))\mathbb{I}\left\{z-x_{k+1}\in\left[-\frac{1}{2},\frac{1}{2}\right]\right\}\,dz.$$

Note that

$$\tilde{p}_{k+1} \leq p_{k+1}.$$

A.2. Growth up to $\frac{1}{4}$

Following the work of [5] define K as below

$$K = \min\left\{k\geq 0 : p_k > \frac{1}{4}\right\}$$

Lemma 1 *Let $\beta=\frac{\varepsilon}{4}+\frac{3}{8}$. For all $0\leq k\leq K$, it holds that*

$$\mathbb{P}\left(p_{k+1}\geq(1+\beta)p_k \mid \{x_i\}_{i=1}^k\right) \geq \frac{4}{7}$$

Proof

As above, we can show that

$$\mathbb{E}\left[\tilde{p}_{k+1}\right] = p_k + \left(\varepsilon + \frac{1}{2}\right) p_k(1-p_k) \geq \left(\frac{3}{4}\varepsilon + \frac{11}{8}\right) p_k$$

Moreover, we have that $\tilde{p}_{k+1} \leq (2+2\varepsilon)p_k$. Thus, we can apply the reverse Markov's inequality

$$\begin{aligned}
\mathbb{P}\left(p_{k+1} \geq (1+\beta)p_k \mid \{x_i\}_{i=1}^k\right) &\geq \mathbb{P}\left(\tilde{p}_{k+1} \geq (1+\beta)p_k \mid \{x_i\}_{i=1}^k\right) \\
&\geq \frac{\mathbb{E}\left[\tilde{p}_{k+1}\right] - (1+\beta)p_k}{\tilde{p}_{k+1} - (1+\beta)p_k} \\
&\geq \frac{\frac{3}{4}\varepsilon + \frac{11}{8} - (1+\beta)}{2+2\varepsilon-(1+\beta)} \\
&= \frac{4\varepsilon}{7\varepsilon+5} \geq \frac{4}{7}
\end{aligned}$$

■

Lemma 2 *With probability at least $1 - \exp\left(-\frac{1}{7}K_1\right)$ we have that*

$$K \leq O\left(\frac{\log \eta^{-1}}{\log(1+\beta)}\right)$$

Proof *As in previous work, we define the following partition of the interval $(0, \frac{1}{4})$.*

$$\begin{aligned}
I_1 &= (0, \eta] \\
I_i &= \left((1+\beta)^{i-1}\eta, (1+\beta)^i\eta\right] \\
I_{i^*} &= \left(2(1+\beta)^{i^*-1}\eta, \frac{1}{4}\right]
\end{aligned}$$

where i^* is the smallest integer such that $(1+\beta)^{i^*-1}\eta \geq \frac{1}{4}$, that is

$$i^* = \left\lceil \frac{\log \frac{1}{4\eta}}{\log(1+\beta)} \right\rceil$$

For $i > 1$, let $\hat{k}_i$ be given such that $p_{\hat{k}_i} \geq (1+\beta)^{i-1}\eta$, let $\hat{Y}_i$ be the smallest number of steps such that $p_{\hat{k}_i+\hat{Y}_i} > (1+\beta)^i\eta$. Then we have that $\hat{Y}_i \leq Y_i \sim \texttt{Geom}\left(\frac{4}{7}\right)$, since, according to lemma (1), we have

$$\mathbb{P}\left(p_{k+1} \geq (1+\beta)p_k \mid \{x_i\}_{i=1}^k\right) \geq \frac{4}{7}$$

Therefore, for all K^*, we have that

$$\mathbb{P}\left(K \geq K^*\right) \leq \mathbb{P}\left(\sum_{i=1}^{i^*} Y_i \geq K^*\right) = \mathbb{P}\left(B_{K^*} \leq i^*\right)$$

where $B_{K^*} \sim \texttt{Bin}\left(K^*, \frac{4}{7}\right)$. *Given that* $\mathbb{E}[B_{K^*}] = \frac{4}{7}K^*$, *we can apply the Hoeffding's inequality for binomial distribution*

$$\mathbb{P}\left(B_{K^*} \leq \frac{4}{7}K^* - t\right) \leq \exp\left(-\frac{2t^2}{K^*}\right)$$

choose $K^* = \frac{7}{2}i^*$ *and* $t = \frac{2}{7}K^*$ *gives that*

$$\mathbb{P}\left(K \leq K^*\right) \geq \mathbb{P}\left(B_{K^*} \geq i^*\right) \geq 1 - \exp\left(-\frac{1}{7}K^*\right)$$

■

A.3. Growth Beyond $\frac{1}{4}$

Recall the recurrence

$$\mathbb{E}\left[\tilde{p}_{k+1}\right] = p_k + \left(\varepsilon + \frac{1}{2}\right)(1 - p_k)p_k.$$

We define

$$Z_{k+1} = \begin{cases} \frac{\tilde{p}_{k+1} - p_k}{(1-p_k)p_k} & \text{if } p_k < 1 \\ \frac{1}{2} + \varepsilon & \text{if } p_k = 1. \end{cases}$$

Then we have $\mathbb{E}[Z_{k+1}] = \frac{1}{2} + \varepsilon$. Let $Y_k = -\frac{k}{2} - k\varepsilon + \sum_{i=K_1+1}^{K_1+k+1} Z_i$, then Y_k is a submartingale. We bound Z_{k+1} as follows

Lemma 3

$$0 \leq Z_{k+1} \leq 2 + 2\varepsilon$$

Proof *We notice that* $p_k \leq \tilde{p}_{k+1} \leq \min\left\{(2+2\varepsilon)p_k, 1\right\}$. *Consider two cases of* p_k*:*

Case 1*:* $p_k \leq \frac{1}{2+2\varepsilon}$. *In this case, we have* $1 - p_k \geq \frac{1+2\varepsilon}{2+2\varepsilon}$

$$Z_{k+1} \leq \frac{(2+2\varepsilon)p_k - p_k}{p_k(1-p_k)} = \frac{1+2\varepsilon}{1-p_k} \leq 2 + 2\varepsilon.$$

Case 2*:* $p_k \geq \frac{1}{2+2\varepsilon}$. *In this case, we use* $\tilde{p}_{k+1} \leq 1$

$$Z_{k+1} \leq \frac{1-p_k}{p_k(1-p_k)} = \frac{1}{p_k} \leq 2 + 2\varepsilon.$$

■

Thus,

$$|Y_{k+1} - Y_k| = \left| -\frac{1}{2} - \varepsilon + Z_{K_1+k+2} \right| \leq \frac{3}{2} + \varepsilon.$$

Let $n = K_1 + K_2$. Therefore, we can apply Azuma's inequality to get that

$$\begin{aligned}
\mathbb{P}\left(\sum_{i=K_1+1}^{n} Z_i \geq \left(\frac{1}{2} + \varepsilon\right)(K_2 - 1) - t \right) &= \mathbb{P}\left(\left(\frac{1}{2} + \varepsilon\right) n + \sum_{i=K_1+1}^{n} Z_i \geq -t \right) \\
&= \mathbb{P}\left(Y_{K_2-1} - Y_0 \geq -t\right) \\
&\geq 1 - \exp\left(-\frac{t^2}{2(K_2-1)(\frac{3}{2}+\varepsilon)^2} \right)
\end{aligned}$$

Let $t = \frac{1}{4}\left(\frac{3}{2} + \varepsilon\right)(K_2 - 1)$ gives that

$$\mathbb{P}\left(\sum_{i=1}^{n} Z_i \geq \left(\frac{1}{8} + \frac{3\varepsilon}{4}\right)(K_2 - 1) \right) \geq 1 - \exp\left(-\frac{K_2 - 1}{32} \right).$$

We use the same function as Lueker to track the growth of p_k, but starting from $\frac{1}{4}$:

$$\psi(p) = \lg p - \ln(1-p) + \frac{p}{2}.$$

Lemma 4 *For all $p_k \geq \frac{1}{4}$, we have that*

$$\psi(p_{k+1}) \geq \psi(\tilde{p}_{k+1}) \geq \psi(p_k) + Z_{k+1}.$$

Proof *The first inequality follows from the fact that ψ is monotone increasing on $(0,1)$ and $p_{k+1} \geq \tilde{p}_{k+1}$. We now show the second inequality. We first notice that*

$$\psi(\tilde{p}_{k+1}) - \psi(p_k) = \int_{p_k}^{p_{k+1}} \psi'(p)dp \geq \min_{p \in [p_k, \tilde{p}_{k+1}]} \psi'(p)(\tilde{p}_{k+1} - p_k).$$

It suffice to show that

$$\min_{p \in [p_k, \tilde{p}_{k+1}]} \psi'(p) \geq \frac{1}{p_k(1-p_k)}.$$

The first- and second-order derivative of ψ are

$$\begin{aligned}
\psi'(p) &= \frac{\lg e}{p} + \frac{1}{1-p} + \frac{1}{2} \\
\psi''(p) &= \frac{1}{(1-p)^2} - \frac{\lg e}{p^2}.
\end{aligned}$$

Therefore, ψ' attains its minimum at $p^* = \frac{(\lg e)^{\frac{1}{2}}}{1+(\lg e)^{\frac{1}{2}}}$, and ψ' increases monotonically on $(\frac{1}{4}, p^*]$ and decreases monotonically on $[p^*, 1]$. Notice that $p^* \leq \frac{1}{2}$. We consider two cases of p_k:

Case 1: $p_k \in (\frac{1}{4}, \frac{1}{2}]$. *Notice that*

$$\min_{p\in[p_k,\tilde{p}_{k+1}]} \psi'(p) \geq \psi'(p^*) = \left(1 + (\lg e)^{\frac{1}{2}}\right)^2 + \frac{1}{2} \geq \frac{16}{3}.$$

Moreover, we notice that $\frac{1}{p_k(1-p_k)}$ *decreases monotonically on* $(\frac{1}{4}, \frac{1}{2}]$. *Therefore, it attains its maximum at* p_k *with value of* $\frac{16}{3}$. *Thus*

$$\min_{p\in[p_k,\tilde{p}_{k+1}]} \psi'(p) \geq \frac{1}{p_k(1-p_k)}$$

Case 2: $p_k \in (\frac{1}{2}, 1]$. *Notice that* ψ *increases monotonically on* $(\frac{1}{2}, 1]$. *Thus*

$$\min_{p\in[p_k,\tilde{p}_{k+1}]} \psi'(p) = \psi'(p_k) = \frac{\lg e}{p_k} + \frac{1}{1-p_k} + \frac{1}{2} \geq \frac{1}{p_k} + \frac{1}{1-p_k} = \frac{1}{p_k(1-p_k)}.$$

■

Therefore, we have

$$\psi(p_n) \geq \psi(p_{K_1}) + \sum_{i=K_1+1}^{n} Z_i.$$

Plugging in the value of $\psi(p_n)$ and $\psi(p_{K_1})$, and notice that ψ increases monotonically with p

$$\begin{aligned}
-\ln(1-p_n) &= \psi(p_n) - \lg p_n - \frac{p_n}{2} \\
&\geq \psi(p_{K_1}) + \sum_{i=K_1+1}^{n} Z_i - \frac{1}{2} \\
&\geq \psi\left(\frac{1}{4}\right) + \sum_{i=K_1+1}^{n} Z_i - \frac{1}{2} \\
&= -\lg 4 - \ln\frac{3}{4} - \frac{1}{4} + \sum_{i=K_1+1}^{n} Z_i \\
&\geq \sum_{i=K_1+1}^{n} Z_i - 2
\end{aligned}$$

Since $\sum_{i=K_1+1}^{n} Z_i \geq \left(\frac{1}{8} + \frac{3\varepsilon}{4}\right)(K_2 - 1)$ with probability at least $1 - \exp\left(-\frac{K_2-1}{32}\right)$, as long as $K_2 \geq \frac{8}{1+6\varepsilon}\left(\ln\eta^{-1} + 2\right) + 1 = O(\left((1+6\varepsilon)^{-1}\ln\eta^{-1}\right) + 1)$, we have that with high probability $\sum_{i=K_1+1}^{n} Z_i \geq \ln\eta^{-1} + 2$, which implies that

$$-\ln(1-p_n) \geq \ln\eta^{-1} = -\ln(\eta)$$

which implies that $p_n \geq 1 - \eta$. Thus, as long as $n = K_1 + K_2$ with

$$K_1 = O\left(\frac{\log\eta^{-1}}{\log(5+\varepsilon) - \log 4}\right) \qquad K_2 = O\left(1 + \frac{\log\eta^{-1}}{(1+6\varepsilon)}\right),$$

we have that with probability at least $1 - \exp\left(-\frac{K_2-1}{32}\right) - \exp\left(-\frac{K_1}{7}\right)$, each point in $[-\frac{1}{2}, \frac{1}{2}]$ either has η approximation or lies with η distance to a point with η approximation. Therefore, each point in $[-\frac{1}{2}, \frac{1}{2}]$ has an 2η approximation.

Appendix B. Using Projected Gradient Descent to Find ε-Perturbation

Algorithm 1: PGD+StrongLTH

Data: Perturbation scale ε, neural network loss $\mathcal{L}$, initial weight $\mathbf{W}_0$, learning rate $\{\alpha_t\}_{t=0}^{T-1}$
Result: Optimal loss ℓ^*, mask $\mathbf{M}^*$ and sparsity level s

---

$\Delta\mathbf{W} \leftarrow 0$;
foreach $t \in \{0, \ldots, T-1\}$ **do**
  $\Delta\mathbf{W} \leftarrow \texttt{sign}(\Delta\mathbf{W} - \alpha_t \nabla\mathcal{L}(\mathbf{W}_t)) \cdot \min\{\texttt{abs}(\Delta\mathbf{W} - \alpha_t \nabla\mathcal{L}(\mathbf{W}_t)), \varepsilon\}$;
  $\mathbf{W}_{t+1} \leftarrow \mathbf{W}_0 + \Delta\mathbf{W}$;
end
$\ell^* \leftarrow \infty$, $\mathcal{M}^* \leftarrow$ None;
foreach *pruning level* $s \in \{0.1, 0.2, \ldots, 0.9\}$ **do**
  $\ell, \mathcal{M} \leftarrow \texttt{Edge-Popup}(\mathcal{L}, \mathbf{W}_T, s)$;
  if $\ell \leq \ell^*$ **then**
    $\ell^* \leftarrow \ell, \mathcal{M}^* \leftarrow \mathcal{M}$
  end
end
return ℓ^*, $\mathbf{M}^*$, s