
On the Sparsity of the Strong Lottery Ticket Hypothesis

Emanuele Natale
Université Côte d’Azur,
CNRS, Inria, I3S, France

Daive Ferré
Université Côte d’Azur,
CNRS, Inria, I3S, France

Giordano Giambartolomei
Department of Informatics,
King’s College London

Frédéric Giroire
Université Côte d’Azur, CNRS,
Inria, I3S, France

Frederik Mallmann-Trenn
Department of Informatics,
King’s College London

Abstract

Considerable research efforts have recently been made to show that a random neural network N contains subnetworks capable of accurately approximating any given neural network that is sufficiently smaller than N , without any training. This line of research, known as the Strong Lottery Ticket Hypothesis (SLTH), was originally motivated by the weaker Lottery Ticket Hypothesis, which states that a sufficiently large random neural network N contains *sparse* subnetworks that can be trained efficiently to achieve performance comparable to that of training the entire network N . Despite its original motivation, results on the SLTH have so far not provided any guarantee on the size of subnetworks. Such limitation is due to the nature of the main technical tool leveraged by these results, the Random Subset Sum (RSS) Problem. Informally, the RSS Problem asks how large a random i.i.d. sample Ω should be so that we are able to approximate any number in $[-1, 1]$, up to an error of ε , as the sum of a suitable subset of Ω .

We provide the first proof of the SLTH in classical settings, such as dense and equivariant networks, with guarantees on the sparsity of the subnetworks. Central to our results, is the proof of an essentially tight bound on the Random Fixed-Size Subset Sum Problem (RFSS), a variant of the RSS Problem in which we only ask for subsets of a given size, which is of independent interest.

1 Introduction

The Lottery Ticket Hypothesis (LTH) is a research direction that has attracted considerable attention over the years, stemming from the empirical contrast between the fact that, while large neural networks can be successfully trained to achieve good performance on a given task and successively pruned to a great level of sparsity without compromising their performance, researchers have struggled to train sparse neural networks from scratch. The authors of [12] observed that, using a simple pruning strategy (namely Iterative Magnitude Pruning while rewinding the original weights of the remaining edges to their value at initialization), *starting from a sufficiently large random neural networks, it is possible to identify sparse subnetworks that can be trained to achieve the performance achievable by the starting network* (see Figure 2 in the appendix for an illustration). The previous statement, namely the LTH, soon gave rise to an even stronger one, corroborated by empirical works [29, 26] which proposed “training-by-pruning” algorithms (see Section 2 for details), providing evidence that *starting from a sufficiently large random neural networks, it is possible to identify sparse subnetworks that exhibit good performance as they are, without changing the original weights* (see Figure 3 in the appendix for an illustration). By removing the need to analyze the dynamics of

training, the last statement, namely the Strong Lottery Ticket Hypothesis (SLTH), allowed a fruitful series of rigorous proofs for increasingly more general architectures (see Section 2 for an overview). Such rigorous results can informally be stated as follows:

Theorem 1 (Informal statement of previous SLTH results). *With high probability, a random artificial neural network N_Ω with m parameters can be pruned so that the resulting subnetwork N_S ε -approximates (i.e., approximates up to an error ε) any target artificial neural network N_t with $O(m/\log_2(1/\varepsilon))$ parameters.*

It is important to note that, to this day, we only have proofs on the existence of such subnetworks, also called winning tickets, but it remains an open question how to find them reliably.

All theoretical results on the SLTH however have so far not investigated the interplay between the sparsity of the winning ticket N_S and the size of the random neural network N_Ω . This is in contrast to the original motivation of the LTH and to the practical application of the aforementioned training-by-pruning algorithms that motivated the SLTH, such as [15, 14]. In fact, to approximate target networks with $O(m/\log_2(1/\varepsilon))$ parameters, essentially all winning tickets N_S have $\Theta(m)$ parameters (see Appendix A), thus being roughly of the same size of the original network N_Ω . We thus ask the following natural question:

If we want to ε -approximate a family of target artificial neural networks with m_t parameters by pruning a fraction α , called sparsity, of the m parameters of a random artificial neural network N_Ω , how big should m be?

We are particularly interested in the regime in which the density parameter $\gamma = 1 - \alpha$ vanishes as the size of the network increases, so that the size of the winning ticket N_S is $\gamma m = o(m)$.

The above question has so far remained unanswered as a consequence of the limitation inherited from the core technical tool that has been leveraged so far to prove SLTH results, namely the Random Subset Sum (RSS) Problem [17]. Informally, the RSS asks how large a random i.i.d. sample Ω should be so that we are able to approximate any number in $[-1, 1]$ as the sum of a suitable subset of Ω . The applicability of RSS to the SLTH was first recognized by [24] within the proof strategy previously developed in [20].

1.1 Our Contribution

We answer the aforementioned question by introducing and proving a refined variant of the RSS Problem, namely the Random Fixed-Size Subset Sum Problem (RFSS), in which the approximation of the target values should be achieved by only considering subsets of fixed size k from a set of n samples (Theorem 2). We focus on subsets of fixed size k rather than subsets of size up to k for two main reasons. From a theoretical point of view, it is a stronger requirement, and practically speaking, using fixed-size subsets enables us to achieve SLTH results where the layers of the lottery ticket exhibit a uniform structure, potentially offering a computational advantage in their implementation.

In Section 4, we show how the density γ impacts the *overparameterization*, i.e., the ratio (m/m_t) between the number of parameters of the original network N_Ω and that of the class of target networks N_t that can be ε -approximated by pruning N_Ω down to a subnetwork N_S with γm parameters. In our analysis, we also compare and recover as special cases previous SLTH results such as [24, 20, 8, 3, 9]. For instance, when $\gamma m = \Theta(m)$, we recover up to a logarithmic factor the result of [24], which states that the overparameterization needed is $O(\log_2(m_t^2/\varepsilon^2))$. In the case of Dense Neural Networks, Theorem 3 thus bridges the gap between the two extreme cases of $\gamma m = \Theta(m_t)$ and $\gamma m = \Theta(m)$ considered in [20] and [24], respectively. It is worth noting that [24] is often considered an improvement over [20], as it exponentially reduces the overparameterization, albeit at the cost of a trivial sparsity level. Finally, we prove that our bounds on the overparameterization as a function of the subnetwork sparsity are essentially tight.

Organization of the paper. After reviewing the literature on the SLTH in Section 2, we introduce the Random Fixed-Size Subset Sum Problem in Section 3. In Section 4, we explore some applications of the RFSS Problem to the SLTH, and finally draw our conclusions in Section 5. Some limitations of our work, along with its potential impact, are discussed in Section 6.

2 Related Work

The SLTH is named after the LTH, which was introduced by Frankle and Carbin in [12]. At the time of writing, this paper has received over 3,300 citations, attesting to the significance and impact of the research topic. Surveying the LTH is thus besides the scope of this work, and we defer the reader to dedicated surveys such as [16].

The SLTH was empirically motivated by work investigating training-by-pruning algorithms such as [29, 26], namely algorithms that leverage the gradient of the network parameters to learn a good *mask* of the edges to be retained (i.e., a good subnetwork, called the winning ticket). [29] achieves this by learning a probability associated to each edge, which is then used to sample the edges that should be included in the subnetwork. [26] gets rid of the stochasticity involved in the aforementioned strategies by learning a score associated to each edge; the subnetwork is then determined by including the edges with the highest score. Such strategies are leveraged in [15, 14] in a federated learning setting, in order to improve the communication cost of distributed training by communicating the sampled masks of a fixed shared network, rather than the entire weights. However, these training-by-pruning algorithms are generally not computationally less expensive than classical training, since they also make use of backpropagation to update scores and are applied to a sufficiently large network to find a winning ticket. To reduce the computational cost of finding a good subnetwork, [13] shows, both theoretically and experimentally, that randomly pre-pruning the source network before looking for a winning ticket can be an effective approach. In [23], on top of randomly pruning the source network, some parameters are also frozen. Frozen parameters are forced to be part of the winning ticket and they do not have an associated score, which effectively reduces the search space for the training-by-pruning algorithms.

The first rigorous proof of the SLTH in the case of dense neural networks has been provided by [20], which establishes a framework that was inherited by the subsequent works. [24] crucially shows that the framework in [20] allows the application of the RSS analysis in [17], proving that, with no constraint on the size of the subnetworks, a random network with m -parameters can be pruned to approximate target networks with $m/\log(1/\varepsilon)$ parameters (we defer the reader to Theorem 3 for details on further constraints on the parameters). An alternative proof of the result in [24] was simultaneously shown in [22]. [8] and [3] successively extended [24] and [22] to convolutional neural networks (CNNs). By leveraging multidimensional generalizations of RSS [7, 2], [6] further extended the SLTH to structured pruning of CNNs and, as a special case, dense networks. Finally, [9] provided a general framework that proves the SLTH for equivariant networks.

As for refinements and generalizations of the above results, [4] shows that, at the cost of a quadratic overhead in the overparameterization w.r.t. [24], the number of layers of the random network N_Ω can be reduced to $\ell + 1$, where ℓ is the number of layers of the target networks N_z ; furthermore, while previous results only considered networks with ReLU activation, [4] shows how to extend the proof in [24] to a more general class of activations functions. [5] introduces the notion of universal lottery ticket, and show that it is possible to prune a sufficiently overparameterized random network so that the resulting subnetwork (the lottery ticket) can approximate certain class of functions up to an affine transformation of the output of the subnetwork (in this sense being universal). [11] shows how to extend the proof in [24] when neurons have random biases, and adapts the training-by-pruning algorithm of [26] to find a strong lottery ticket with a desired sparsity level. Motivated by theoretical insights on the existence of sparse strong lottery tickets, [10] develops a framework to plant the latter in large random network and investigates training-by-pruning algorithms, providing evidence that sparse strong lottery tickets typically exists for common machine learning tasks, and the difficulty to find them is of algorithmic nature.

Our proof of the RFSS Problem in Section 3 is based on the second moment method approach first explored by [18], and which has recently been refined to prove multidimensional generalizations of RSS by [7] and [2].

3 Fixed-Size Random Subset Sum

In this section we present our technical contributions on the RFSS, which are the foundation of our proofs regarding the sparsity of the SLTH.

Let us start by introducing some notation. We denote by $[n]$ the set $\{1, \dots, n\}$, for $n \in \mathbb{N}$. Given a set $\Omega = \{X_1, \dots, X_n\}$ and a set of indices $S \subseteq [n]$ we define $\Sigma_S^\Omega = \sum_{i \in S} X_i$, and we omit Ω when clear from the context. We now define a class of distributions for which our RFSS result holds.

Definition 1 (sum-bounded). *We say that a probability density function f is sum-bounded if there exist positive constants c_l and c_u such that, for all $k \in \mathbb{N}$, given k independent samples X_1, \dots, X_k with density f , the density of their sum $f_{\Sigma_{[k]}}$ satisfies*

$$\frac{c_l}{\sqrt{k}} \leq f_{\Sigma_{[k]}}(x) \leq \frac{c_u}{\sqrt{k}},$$

with the lower bound holding for all $x \in [-\sqrt{k}, \sqrt{k}]$ and the upper bound holding for all $x \in \mathbb{R}$.

At first, our definition of sum-bounded could look as a weaker version of a classical local limit theorem on the sum of random variables (e.g., see [25, Chapter VII, Theorem 7]). However, that is not the case, since we require a lower bound on the sum for any k , which is needed to prove our main result.

Denote, for all $x \in [0, 1]$, the binary entropy as

$$H_2(x) = -x \log_2 x - (1-x) \log_2(1-x).$$

Our main technical result is the following proof of a fixed-size subset variant of the RSS Problem.

Theorem 2. *Let $0 < \varepsilon < 1$, $c_{hyp} \geq 1$, k, n be integers with $1 \leq k \leq \frac{n}{2}$, and let $\Omega = \{X_1, \dots, X_n\}$ where the X_i 's are i.i.d. random variables with sum-bounded density. There exists a constant c_{thm} such that, if*

$$n \geq c_{hyp} \frac{\log_2 \frac{k}{\varepsilon}}{H_2\left(\frac{k}{n}\right)}, \quad (1)$$

then for every fixed $z \in [-\sqrt{k}, \sqrt{k}]$ it holds that

$$\Pr(\exists S \subset [n], |S| = k : |\Sigma_S - z| < \varepsilon) \geq c_{thm}.$$

Remark. The proof of Theorem 2 is given in Section 3.1, and it actually holds for any $1 \leq k \leq \lambda n$, for an arbitrary $\lambda \in [1/n, 1)$. We state the theorem this way for readability and because we are primarily interested in high-sparsity settings (i.e., small size k of the subsets), so considering values of $k \geq \frac{n}{2}$ does not add much to our analysis. The same remark also holds for Corollary 1.

The sum-bounded condition of Definition 1 is easily verified for distributions such as the Gaussian distribution. Previous SLTH results rely on a classical resampling argument by [17, Corollary 3.3], which shows how RSS results for Uniform $[-1, 1]$ independent random variables naturally extend to independent random variables that *contains* a uniform distribution, in the sense that they can be expressed as the mixture of distributions one of which is Uniform $[-1, 1]$ with constant probability.¹ The next lemma thus proves that the Uniform $[-1, 1]$ distribution is sum-bounded². A detailed proof of is provided in Appendix C.

Lemma 1. *The Uniform $[-1, 1]$ probability density function is sum-bounded, i.e., given a set $\mathcal{U}_n = \{U_i\}_{i \in [n]}$ of i.i.d. variables U_i with Uniform $[-1, 1]$ probability density function, there exist constants c_l and c_u such that the probability density function $f(x, n)$ of the sum $\Sigma_{[n]}^{\mathcal{U}_n}$ of these variables, for all $n \in \mathbb{N}$,*

$$\frac{c_l}{\sqrt{n}} \leq f(x, n) \leq \frac{c_u}{\sqrt{n}}, \quad (2)$$

with the lower bound holding for all $x \in [-\sqrt{n}, \sqrt{n}]$, and the upper bound holding for all $x \in \mathbb{R}$.

Finally, in our proofs on the Sparse SLTH in Section 4, we make use of the following corollary of Theorem 2, which ensures a uniform high probability of hitting any target $z \in [-\sqrt{k}, \sqrt{k}]$, considering independent random variables that contain a uniform distribution.

¹The definition in [17, Corollary 3.3] is actually more general, since it concerns a different problem.

²We believe that Lemma 1 is known, but we could not find a reference.

Corollary 1. Let $0 < p \leq 1$ and $\varepsilon \in (0, 1/2)$ be constants, k, n with $1 \leq k \leq \frac{n}{2}$, and let $\Omega = \{X_1, \dots, X_n\}$ be i.i.d. random variables whose density is a mixture of a Uniform($[-1, 1]$) with probability p , and some other density otherwise. There exists a positive constant c_{amp} that only depends on p such that, if

$$n \geq c_{amp} \frac{\log_2^2 \frac{k}{\varepsilon}}{H_2\left(\frac{k}{n}\right)}, \quad (3)$$

then

$$\Pr\left(\forall z \in \left[-\sqrt{k}, \sqrt{k}\right], \exists S_z \subset [n], |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon\right) \geq 1 - \varepsilon.$$

Proof Idea. The corollary follows from three arguments. First, by a standard sampling argument, we can assume that a constant fraction of the sample follows a Uniform $[-1, 1]$ distribution. Secondly, by Lemma 1, the uniform probability density function is sum-bounded. We can thus apply Theorem 2, which guarantees a success probability of c_{thm} for approximating a given target. Finally, by a standard probability amplification argument and a union bound applied to Theorem 2, by paying an extra factor $\log_2(k/\varepsilon)$ in Eq. 1, the constant c_{thm} can be assumed to be $1 - \varepsilon$, and the existence of a suitable subset S_z holds simultaneously for all $z \in [-\sqrt{k}, \sqrt{k}]$. Details are given in Appendix D. \square

For k big enough, we can get rid of the squared logarithmic dependency on k in the right hand side of Equation 3, as shown in the following Corollary, whose proof can be found in Appendix E.

Corollary 2. Let $0 < p \leq 1$ and $\varepsilon \in (0, 1/2)$ be constants, k, n be integers with $1 \leq k \leq \frac{n}{2}$ and $k \geq 2c_{amp} (\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon})$. Let $\Omega = \{X_1, \dots, X_n\}$ be i.i.d. random variables whose density is a mixture of a Uniform($[-1, 1]$) with probability p , and some other density otherwise. There exists a positive constant c_{amp} that only depends on p such that, if

$$n \geq 2c_{amp} \frac{\log_2^2 \frac{1}{\varepsilon}}{H_2\left(\frac{k}{n}\right)}, \quad (4)$$

then

$$\Pr\left(\forall z \in \left[-\sqrt{k}, \sqrt{k}\right], \exists S_z \subset [n], |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon\right) \geq 1 - \varepsilon.$$

As customary in conference versions of papers, our proofs adopt the convention of taking ceilings and floors as suitable for non integer fractional terms. This is done in the interest of the reader (and ours), and does not impact the results in any significant way.

3.1 Proof of Theorem 2

Proof of Theorem 2. For simplicity, throughout the proof we will often use c to denote any positive constant. Let $\mathcal{S}_k = \{S \subset [n] \mid |S| = k\}$ and define, for a fixed $z \in [-\sqrt{k}, \sqrt{k}]$,

$$Y = Y(z) = \sum_{S \in \mathcal{S}_k} Z_S$$

where $Z_S = Z_S(z) = \mathbf{1}_{\{|\Sigma_S - z| < \varepsilon\}}$. Following [18], we exploit the second moment method for RFSS, generalising it to arbitrary k .

$$\Pr(Y > 0) \geq \frac{(\mathbb{E}[Y])^2}{\mathbb{E}[Y^2]}, \quad (5)$$

it thus suffices to prove that

$$\mathbb{E}[Y^2] \leq c (\mathbb{E}[Y])^2. \quad (6)$$

We first rewrite Eq. 5 in a more convenient form. Let \tilde{S} and \tilde{S}' be two independently and uniformly at random chosen subsets of $[n]$ of size k , and denote $H_S(z)$ as the event that Σ_S ε -approximates z , namely

$$H_S = H_S(z) = \{|\Sigma_S - z| < \varepsilon\}.$$

We have

$$\mathbb{E}[Y] = \sum_{S \in \mathcal{S}_k} \mathbb{E}[Z_S] = \sum_{S \in \mathcal{S}_k} \Pr(H_S) = \binom{n}{k} \Pr(H_{\tilde{S}}) \quad (7)$$

and

$$\begin{aligned} \mathbb{E}[Y^2] &= \mathbb{E} \left[\left(\sum_{S \in \mathcal{S}_k} Z_S \right) \left(\sum_{S' \in \mathcal{S}_k} Z_{S'} \right) \right] = \sum_{S, S' \in \mathcal{S}_k} \mathbb{E}[Z_S Z_{S'}] \\ &= \sum_{S, S' \in \mathcal{S}_k} \Pr(H_S \wedge H_{S'}) = \binom{n}{k}^2 \Pr(H_{\tilde{S}} \wedge H_{\tilde{S}'}). \end{aligned} \quad (8)$$

Using Eqs. 7 and 8 we can rewrite the r.h.s. of Eq. 5 as follows

$$\frac{(\mathbb{E}[Y])^2}{\mathbb{E}[Y^2]} = \frac{[\Pr(H_{\tilde{S}})]^2}{\Pr(H_{\tilde{S}} \wedge H_{\tilde{S}'})} = \frac{\Pr(H_{\tilde{S}})}{\Pr(H_{\tilde{S}'} | H_{\tilde{S}})}.$$

Eq. 6 thus becomes

$$\Pr(H_{\tilde{S}'} | H_{\tilde{S}}) \leq c \Pr(H_{\tilde{S}}). \quad (9)$$

Let I_i denote the event $\{|\tilde{S} \cap \tilde{S}'| = i\}$ and $I_{a,b}$ the event $\bigcup_{a \leq i \leq b} I_i$. Fix $\mu \in (\lambda, 1)$. By the law of total probability and independence of I_i and $H_{\tilde{S}}$, we rewrite the l.h.s. of Eq. 9 as follows:

$$\begin{aligned} &\Pr(H_{\tilde{S}'} | H_{\tilde{S}}) \\ &= \Pr(H_{\tilde{S}'} \wedge I_k | H_{\tilde{S}}) + \Pr(H_{\tilde{S}'} \wedge I_{\mu k, k-1} | H_{\tilde{S}}) + \Pr(H_{\tilde{S}'} \wedge I_{0, \mu k-1} | H_{\tilde{S}}) \\ &= \Pr(I_k) \cdot \Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_k) \end{aligned} \quad (10)$$

$$+ \Pr(I_{\mu k, k-1}) \cdot \Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_{\mu k, k-1}) \quad (11)$$

$$+ \sum_{i=0}^{\mu k-1} (\Pr(I_i) \cdot \Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_i)). \quad (12)$$

To conclude the proof, it suffices to show that each addendum in Eqs. 10, 11 and 12 are upper-bounded by some constant multiple of ε/\sqrt{k} , since the lower bound in Definition 1 ensures that

$$\frac{\varepsilon}{\sqrt{k}} \leq c \Pr(H_{\tilde{S}}). \quad (13)$$

As for the first addendum (Eq. 10), since $\Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_k) = 1$, then

$$\begin{aligned} \Pr(I_k) \cdot \Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_k) &= \Pr(I_k) = \frac{1}{\binom{n}{k}} \stackrel{(a)}{\leq} \sqrt{\frac{8k(n-k)}{n}} 2^{-nH_2(\frac{k}{n})} \\ &\stackrel{(b)}{\leq} \sqrt{\frac{8k(n-k)}{n}} 2^{-c_{\text{hyp}} \log_2 \frac{k}{\varepsilon}} \stackrel{(c)}{\leq} 2\sqrt{2} \frac{\varepsilon}{\sqrt{k}}, \end{aligned} \quad (14)$$

where inequality (a) in Eq. 14 is a standard lower bound on $\binom{n}{k}$ holding for all $k \leq n-1$; in inequality (b) in Eq. 14 we used Eq. 1, namely $nH_2(\frac{k}{n}) \geq c_{\text{hyp}} \log_2 \frac{k}{\varepsilon}$; in inequality (c) in Eq. 14 we used that $c_{\text{hyp}} \geq 1$.

As for the second addendum (Eq. 11), we next show that

$$\Pr(I_{\mu k, k-1}) \Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_{\mu k, k-1}) \leq c \frac{\varepsilon}{\sqrt{k}} \quad (15)$$

by proving that

$$\Pr(I_{\mu k, k-1}) \leq \frac{c}{\sqrt{k}} \quad (16)$$

and

$$\Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_{\mu k, k-1}) \leq c\varepsilon. \quad (17)$$

First, observe that $I = |\tilde{S} \cap \tilde{S}'|$ follows a Hypergeometric(n, k, k) distribution, thus by Chebyshev's inequality

$$\begin{aligned} \Pr(I_{\mu k, k-1}) &\leq \Pr(I \geq \mu k) = \Pr\left(I - \frac{k^2}{n} \geq \mu k - \frac{k^2}{n}\right) \leq \frac{\text{Var}[I]}{\mu^2 k^2 \left(1 - \frac{k}{\mu n}\right)^2} \\ &\leq c' \frac{\frac{k^2}{n} \frac{n-k}{n} \frac{n-k}{n-1}}{k^2} \leq \frac{c}{\sqrt{k}}, \end{aligned}$$

having set $c' = \mu^2(1 - \lambda/\mu)^2 > 0$, thus proving Eq. 16. The proof of Eq. 17 is given in Appendix F, concluding the proof of Eq. 15.

As for the third addendum (Eq. 12), in Appendix F we show that

$$\sum_{i=0}^{\mu k-1} \Pr(I_i) \cdot \Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_i) \leq c \frac{\varepsilon}{\sqrt{k}}, \quad (18)$$

The three bounds on the addenda in Eqs. 10, 11, and 12 (Eqs. 14, 15, and 18, respectively), combined with Eq. 13, conclude the proof. \square

4 Sparse Strong Lottery Ticket Hypothesis (SSLTH)

We now apply our results on the RFSS problem to the SLTH and obtain guarantees on the sparsity of winning tickets for Dense Neural Networks (DNNs, Theorem 3) and Equivariant NNs (Theorem 4).

The next theorem essentially interpolates between the two extremes of [20][Theorem 2.1] (where $\gamma m = \Theta(m_t)$) and [24][Theorem 1] (where $\gamma m = \Theta(m)$), where we recall that m and m_t represent the number of parameters of the overparameterized and the target networks, respectively, and γ is the density of the winning ticket.

We use $\sigma(\cdot)$ to denote the ReLU activation function, i.e., $\sigma(x) = x \cdot \mathbf{1}_{x \geq 0}$, and $\|\mathbf{W}\|$ to denote the spectral norm of the matrix \mathbf{W} . Let \mathcal{F} to be a set of target ReLU neural networks $f : \mathbf{R}^{d_0} \rightarrow \mathbf{R}^{d_l}$ of depth l such that

$$\mathcal{F} = \{f : f(\mathbf{x}) = \mathbf{W}_l \sigma(\mathbf{W}_{l-1} \dots \sigma(\mathbf{W}_1 \mathbf{x})), \forall i \ \mathbf{W}_i \in \mathbb{R}^{d_i \times d_{i-1}} \text{ and } \|\mathbf{W}_i\| \leq 1\} \quad (19)$$

For a given $f \in \mathcal{F}$, for all $i \in [\ell]$, let $\rho_i = \max\{d_{i-1}/d_i, d_i/d_{i-1}\}$, and $\rho = \max_i \rho_i$. Then, recalling that c_{amp} is the constant defined in Corollary 1, we have the following result.

Theorem 3 (SSLTH for DNNs). *Let g be a randomly initialized feed-forward 2ℓ -layer neural network, in which each weight is drawn from a Uniform $[-1, 1]$ distribution, of the following form:*

$$g(\mathbf{x}) = \mathbf{M}_{2\ell} \sigma(\mathbf{M}_{2\ell-1} \dots \sigma(\mathbf{M}_1 \mathbf{x})).$$

Let $\gamma' = \gamma'(\varepsilon) \in (0, 1)$, $\mathbf{M}_{2i} \in \mathbb{R}^{d_i \times 2d_{i-1} n_i^*}$ and $\mathbf{M}_{2i-1} \in \mathbb{R}^{2d_{i-1} n_i^* \times d_{i-1}}$, with n_i^* satisfying

$$n_i^* = c_{\text{amp}} \frac{\log_2^2 \left(\frac{2\ell d_{i-1} d_i \gamma' n_i^*}{\varepsilon} \right)}{H_2(\gamma')}. \quad (20)$$

With probability at least $1 - \varepsilon$, for every $f \in \mathcal{F}$, where \mathcal{F} is defined as in Eq. 19, g can be pruned to obtain a subnetwork of sparsity at least $\alpha = 1 - \gamma$ that approximates f up to an error ε , having defined $\gamma = \rho \gamma'$.

Proof Idea. The theorem follows from a slight variation of the same approach detailed in [24], in which we use our Corollary 1 instead of [17][Corollary 2.5] when pruning g , allowing us to have control over the size of the pruned network. A detailed proof is provided in Appendix G. \square

To illustrate a simple example of how Theorem 3 addresses the main question asked in the introduction, consider the case where we want to approximate a target network with m_t parameters and ℓ layers, each of width d (so $\rho = 1$ and $\gamma' = \gamma$), by pruning an overparameterized network to achieve a desired sparsity level of $\alpha = 1 - \gamma$. The condition expressed by Equation 20 in Theorem 3




Figure 1: A qualitative plot showing the relationship between the density γ of a winning ticket and the overparameterization required by Theorem 3 for a target network with m_t parameters. Earlier results from Pensia et al. [24] and Malach et al. [20] are shown for comparison.

comes from the use of Corollary 1 when pruning network g , as shown in the proof. If, instead of Corollary 1, we use its simplified variant Corollary 2, it is easy to observe that Equation 20 would become

$$n_i^* = c_{\text{amp}} \frac{\log_2^2 \left(\frac{2\ell d_i - 1 d_i}{\epsilon} \right)}{H_2(\gamma')}. \quad (21)$$

Using this condition, Theorem 3 then tells us that we need to prune a randomly initialized network with twice as many layers and a number of parameters of the order of $d^2 \frac{\log^2 \frac{\ell d^2}{H(\gamma)}}{H(\gamma)}$.

We will now clarify the connection between Theorem 3 and the earlier results from [20] and [24]. Figure 1 provides a quick visual comparison.

Malach et al.[20]. When all layers have the same width d , [20] showed that any target network with l layers and a total of $m_t = d^2 l$ parameters can be ϵ -approximated by pruning a randomly initialized network with $2l$ layers. The overparameterization of this network, relative to the target network, is $O\left(\frac{m_t^2}{\epsilon^2} \log_2 \frac{m_t}{\epsilon}\right) = \tilde{O}\left(\frac{m_t^2}{\epsilon^2}\right)$. More specifically, the winning ticket found after pruning has a parameter count of the same order as the target network, resulting in a density of $\gamma = \tilde{O}\left(\frac{\epsilon^2}{m_t^2}\right)$. Notably, this density γ is the inverse of the overparameterization, as the size of the winning ticket matches that of the target network.

Next, we show that Theorem 3 also yields a density that is polynomial in $\frac{\epsilon}{m_t}$, when using an overparameterization of $\Theta\left(\frac{m_t^2}{\epsilon^2}\right)$. Let $z = \left(\frac{m_t}{\epsilon}\right)$, and note that $\gamma' = \gamma$ in Theorem 3, since all layers have the same width. As n_i^* in Theorem 3 represents the overparameterization with respect to the target network, let us set $n_i^* = cz^2$, for some constant c . Equation 20 then becomes

$$cz^2 \geq c_{\text{amp}} \frac{\log_2^2(cz^3\gamma)}{H(\gamma)} \quad (22)$$

We show that the inequality $cz^2 \geq c_{\text{amp}} \frac{\log_2^2(cz^3\gamma)}{\gamma \log_2(1/\gamma)}$ holds for some big enough constant c when setting $\gamma = \frac{\epsilon}{m_t} = \frac{1}{z}$, which implies that Equation 22 is also satisfied. We get $cz \geq c_{\text{amp}} \frac{\log_2^2(cz^2)}{\log_2(z)}$, which is satisfied for a big enough constant c (see Appendix I). Overall, when using an overparameterization $\Theta\left(\frac{m_t^2}{\epsilon^2}\right)$, we find a winning ticket with density $\frac{\epsilon}{m_t}$, as shown in Figure 1.

Pensia et al.[24]. For simplicity, let us still consider target networks where all layers have the same width d , and we apply Theorem 3 using the simplified condition from Equation 21. When

$\gamma m = \Theta(m)$, i.e. the density γ is a constant as in [24] (see Appendix A), the entropy term $H_2(\gamma')$ in the right-side of Equation 21 also becomes a constant. In this setting, we indeed recover the result shown in [24][Theorem 1], up to a logarithmic factor, as shown in Figure 1.

Quite similarly to Theorem 3, the next result essentially generalizes [9] up to a factor $\log_2 \frac{1}{\varepsilon}$. The theorem is stated with the understanding that for G -equivariant networks, in order to preserve G -equivariance, pruning is best done not with respect to the parameters expressing the network in the canonical basis (i.e. directly on the weights of the network), but with respect to the *equivariant parameters*, that is those coefficients expressing the linear layers of the network as a linear combination of the elements of the corresponding equivariant basis [9]. For simplicity, due to the technical set-up, we assume all feature spaces being $\mathbb{F} = (\mathbb{R}^d, \sigma)$, with σ the linear representation of the group G , and the same number n of such feature spaces being stacked in each layer. A G -equivariant linear map from the i th feature space to the $i + 1$ st can be decomposed in a corresponding equivariant basis denoted $\mathcal{B}_{i \rightarrow i+1} = \mathcal{B}$. Since all feature spaces are the same, we omit the layers' indices. When stacking n feature spaces in the input and output of the i th layer, the full equivariant basis is denoted $k_{n \rightarrow n}$, and finally the basis of the G -equivariant maps from \mathbb{F}^n to \mathbb{F}^n can be written as the Kronecker product $k_{n \rightarrow n} \otimes \mathcal{B}$. For any basis $\mathcal{B} = \{b_1, \dots, b_p\}$, we denote its cardinality $p = |\mathcal{B}|$ and define $\|\mathcal{B}\| = \max_{\|\beta\|_\infty} \|\sum_{k=1}^p \beta_k b_k\|$, with $\|\cdot\|$ in the r.h.s. being the operator norm inherited from the ℓ_p norm.

Theorem 4 (SSLTH for Equivariant Networks). *Let h be a random 2ℓ -layer G -equivariant network where all equivariant parameters are drawn from a Uniform $[-1, 1]$ distribution, every odd layer expressed in the associated equivariant basis $k_{\tilde{n} \rightarrow \tilde{n}} \otimes \mathcal{B}$ and every even layer expressed in the associated equivariant basis $k_{n \rightarrow \tilde{n}} \otimes \mathcal{B}$. Let $\gamma = \gamma(\varepsilon) \in (0, 1)$, with \tilde{n} satisfying*

$$\tilde{n} = c_{amp} \frac{\log_2^2 \left(\frac{2\ell n^2 \max\{|\mathcal{B}|, \|\mathcal{B}\|\} \gamma \tilde{n}}{\varepsilon} \right)}{H_2(\gamma)}.$$

With probability at least $1 - \varepsilon$, for every ℓ -layer G -equivariant neural network f , with all layers expressed in the associated equivariant basis $k_{n \rightarrow n} \otimes \mathcal{B}$, h can be pruned to obtain a G -equivariant subnetwork of sparsity at least $\alpha = 1 - \gamma$ that approximates f up to an error ε .

The proof, which we omit, is analogous to that of Theorem 3, since [9][Theorem 1] exploits the exact same pruning strategy of [24], except for the fact that it is applied not to the original parameters of the equivariant network, but to the network expressed in terms of its equivariant basis (the sparsity α is here also intended with respect to the equivariant parameters count). This allows the construction to apply without losing the property of equivariance in the pruned approximating subnetwork obtained. The crucial step is when Corollary 1 is applied in [9][Lemma 1], instead of [17][Corollary 2.5]. This is done in parallel, multiple times, across non-overlapping coefficients of the equivariant basis. Thanks to the careful preprocessing devised by the authors, this preserves equivariance and at the same time ensures that each application of Corollary 1 is independent of the others.

To conclude the section, we mention that Theorem 4 applies in particular to vanilla CNNs, which are a special case of equivariant neural networks where the group is $G = (\mathbb{Z}^2, +)$, recovering previous SLTH results on CNN [8, 3]. Furthermore, we remark that Theorem 3 can be revisited through the improvement upon the 2ℓ -depth overparameterization devised in [4], i.e., it is possible to provide sparsity guarantees also for overparameterizations requiring depth $\ell + 1$ only. The analysis is more technical and we omit it, but the ideas are analogous to what shown in [4]. An analogous improvement is suggested as future work in [9].

4.1 Lower bound on the required overparameterization

We now adapt the lower bound of [24] in order to almost match the required overparameterization of our Theorem 3, considering the simple scenario in which we want to approximate the family \mathcal{F} of all linear networks with weights forming a matrix having spectral norm less than \sqrt{k} ; more formally

$$\mathcal{F} := \{h_W : W \in \mathbb{R}^{d \times d}, \|W\| \leq \sqrt{k}\}, \quad \text{where } h_W(x) = Wx. \quad (23)$$

The formal claim states that, if a network with n parameters can approximate every $h_W \in \mathcal{F}$ with probability at least $1/2$ (after it is pruned down to k parameters), then the hypothesis of Theorem 2 in Eq. 1 must hold.³

³Equivalently, the hypothesis of Corollary 1 must hold up to a factor $\Theta(\log_2 \frac{k}{\varepsilon})$.

Theorem 5. Let $n, k \in \mathbb{N}$, with $1 \leq k \leq \lambda n$, having set $\lambda = 1 - 1/2\pi \approx 0.84$. Consider a neural network g with n parameters, and let \mathcal{G}_k be the set of neural networks that can be formed by pruning g down to k parameters. Let \mathcal{F} be as defined in Eq. 23. If it holds that, for some $\varepsilon < 1/16$,

$$\forall h_W \in \mathcal{F}, \mathbb{P} \left(\exists g' \in \mathcal{G}_k : \max_{\mathbf{x}: \|\mathbf{x}\| \leq 1} \|h_W(\mathbf{x}) - g'(\mathbf{x})\| < \varepsilon \right) \geq \frac{1}{2}, \quad (24)$$

then it holds that

$$n \geq \frac{d^2 \log_2 \frac{k}{\varepsilon}}{2 H_2 \left(\frac{k}{n} \right)}.$$

The theorem follows by adapting the packing argument of [24]. A detailed proof is provided in Appendix H.

5 Conclusions

In this work, we have extended previous results on the Strong Lottery Ticket Hypothesis by quantifying the required overparameterization as a function of the sparsity of the subnetworks. Central to our results is a proof of the Random Fixed-size Subset Sum (RFSS) Problem, a refinement of the seminal Random Subset Sum (RSS) Problem in which the subsets have a required fixed size.

A challenging open problem is to extend our analysis of RFSS to the multidimensional case, in which the random samples and targets are vectors in \mathbb{R}^d . Previous extension of RSS to the Multidimensional RSS have indeed allowed to prove structured-pruning version of the SLTH [8]. A Multidimensional RFSS result would then allow to quantify, in the structured pruning case, the dependency of the overparameterization w.r.t. the sparsity of the (structured) subnetworks.

Another future direction is to refine our analysis of the RFSS in Theorem 2 in order to improve the probability of success to $1 - \varepsilon$ rather than constant, thus allowing to avoid shaving off the extra factor $\log_2(1/\varepsilon)$ in our corollaries w.r.t. our lower bound, which is due to the amplification done in Corollary 1 to get to probability $1 - \varepsilon$.

Finally, an important future direction is to improve training-by-pruning methods such as [29, 26, 11, 10, 23] or to develop new ones, in order to allow to efficiently find strong lottery tickets of a desired sparsity, thus empirically validating our theoretical predictions.

6 Limitations and Impact

Limitations Similar to all the research conducted on the LTH and the SLTH, this work only proves the existence of lottery tickets. To this date, it is not clear if these subnetworks can be found reliably (no formal proof exists) in an efficient manner - however, empirical evidence suggests that efficient algorithms exist (e.g., [29, 26]).

Impact The contribution of this work is primarily theoretical and not confined to a specific domain. Its potential societal impact would, therefore, be closely tied to the particular scenarios to which it is applied. It could be interesting to compare the environmental impact of finding lottery tickets inside overparameterized networks. We also believe that our work has the potential to have a strong environmental impact as sparse NNs have massively reduced inference costs.

Acknowledgments and Disclosure of Funding

This research is supported by the EPSRC grant EP/W005573/1, and by the France 2030 program, managed by the French National Research Agency under grant agreements No. ANR-23-PECL-0003 and ANR-22-PEFT-0002. It was also funded in part by the European Network of Excellence dAIEDGE under Grant Agreement Nr. 101120726, by SmartNet and LearnNet, and by the French government National Research Agency (ANR) through the UCA JEDI (ANR-15-IDEX-01), EUR DS4H (ANR-17-EURE-004), and the 3IA Côte d’Azur Investments in the Future project with the reference number ANR-19-P3IA-0002.

References

- [1] Giampietro Allasia. “Approximation of the normal distribution functions by means of a spline function”. In: *Statistica* 41.2 (1981), pp. 325–332.
- [2] Sander Borst et al. “On the Integrality Gap of Binary Integer Programs with Gaussian Data”. In: *Mathematical Programming* 197.2 (Feb. 2023), pp. 1221–1263. ISSN: 1436-4646. DOI: [10.1007/s10107-022-01828-1](https://doi.org/10.1007/s10107-022-01828-1). (Visited on 06/05/2023).
- [3] Rebekka Burkholz. “Convolutional and Residual Networks Provably Contain Lottery Tickets”. In: *Proceedings of the 39th International Conference on Machine Learning*. Baltimore: PMLR, July 2022, pp. 2414–2433. (Visited on 02/27/2023).
- [4] Rebekka Burkholz. “Most Activation Functions Can Win the Lottery Without Excessive Depth”. In: *Thirty-Sixth Conference on Neural Information Processing Systems*. Dec. 2022. (Visited on 02/27/2023).
- [5] Rebekka Burkholz et al. “On the Existence of Universal Lottery Tickets”. In: *International Conference on Learning Representations*. virtual, Apr. 2022. (Visited on 02/27/2023).
- [6] Arthur da Cunha, Francesco D’Amore, and Emanuele Natale. “Polynomially Over-Parameterized Convolutional Neural Networks Contain Structured Strong Winning Lottery Tickets”. In: *Thirty-Seventh Conference on Neural Information Processing Systems*. Nov. 2023. (Visited on 04/23/2024).
- [7] Arthur Carvalho Walraven da Cunha et al. “Revisiting the Random Subset Sum Problem”. In: *31st Annual European Symposium on Algorithms, ESA 2023, September 4-6, 2023, Amsterdam, The Netherlands*. Ed. by Inge Li Gørtz et al. Vol. 274. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 37:1–37:11. DOI: [10.4230/LIPICs.ESA.2023.37](https://doi.org/10.4230/LIPICs.ESA.2023.37). URL: <https://doi.org/10.4230/LIPICs.ESA.2023.37>.
- [8] Arthur da Cunha, Emanuele Natale, and Laurent Viennot. “Proving the Strong Lottery Ticket Hypothesis for Convolutional Neural Networks”. In: *ICLR 2022 - 10th International Conference on Learning Representations*. Virtual, France, Apr. 2022. (Visited on 08/04/2022).
- [9] Damien Ferbach et al. “A General Framework For Proving The Equivariant Strong Lottery Ticket Hypothesis”. In: *The Eleventh International Conference on Learning Representations*. Sept. 2022. (Visited on 03/04/2024).
- [10] Jonas Fischer and Rebekka Burkholz. “Plant ’n’ Seek: Can You Find the Winning Ticket?” In: *International Conference on Learning Representations*. Apr. 2022. (Visited on 02/27/2023).
- [11] Jonas Fischer, Advait Gadhikar, and Rebekka Burkholz. *Lottery Tickets with Nonzero Biases*. June 2022. arXiv: [2110.11150](https://arxiv.org/abs/2110.11150) [cs]. (Visited on 05/13/2024).
- [12] Jonathan Frankle and Michael Carbin. “The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks”. In: *International Conference on Learning Representations*. Sept. 2018. (Visited on 10/20/2023).
- [13] Advait Harshal Gadhikar, Sohom Mukherjee, and Rebekka Burkholz. “Why Random Pruning Is All We Need to Start Sparse”. In: *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*. Ed. by Andreas Krause et al. Vol. 202. Proceedings of Machine Learning Research. PMLR, 2023, pp. 10542–10570. URL: <https://proceedings.mlr.press/v202/gadhikar23a.html>.
- [14] Berivan Isik et al. “Adaptive Compression in Federated Learning via Side Information”. In: *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*. PMLR, Apr. 2024, pp. 487–495. (Visited on 05/16/2024).

- [15] Berivan Isik et al. “Sparse Random Networks for Communication-Efficient Federated Learning”. In: *The Eleventh International Conference on Learning Representations*. Sept. 2022. (Visited on 01/18/2024).
- [16] Bohan Liu et al. *A Survey of Lottery Ticket Hypothesis*. Mar. 2024. DOI: [10.48550/arXiv.2403.04861](https://doi.org/10.48550/arXiv.2403.04861). arXiv: [2403.04861](https://arxiv.org/abs/2403.04861) [cs]. (Visited on 05/16/2024).
- [17] George S. Lueker. “Exponentially small bounds on the expected optimum of the partition and subset sum problem”. In: *Random Structures and Algorithms* 12 (1998), pp. 51–62.
- [18] George S. Lueker. “On the average difference between the solutions to linear and integer knapsack problems”. In: *Applied Probability - Computer Science, The Interface*. Vol. 1. Birkhäuser, 1982.
- [19] Florence Jessie MacWilliams and Neil James Alexander Sloane. *The Theory of Error-Correcting Codes*. Vol. 16. North-Holland Mathematical Library. North-Holland Publishing Company, 1977.
- [20] Eran Malach et al. “Proving the Lottery Ticket Hypothesis: Pruning Is All You Need”. In: *Proceedings of the 37th International Conference on Machine Learning*. ICML’20. JMLR.org, July 2020, pp. 6682–6691. (Visited on 03/26/2023).
- [21] James E Marengo, David L Farnsworth, Lucas Stefanic, et al. “A geometric derivation of the Irwin-Hall distribution”. In: *International Journal of Mathematics and Mathematical Sciences* 2017 (2017).
- [22] Laurent Orseau, Marcus Hutter, and Omar Rivasplata. “Logarithmic Pruning Is All You Need”. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. NIPS’20. Red Hook, NY, USA: Curran Associates Inc., Dec. 2020, pp. 2925–2934. ISBN: 978-1-71382-954-6. (Visited on 03/26/2023).
- [23] Hikari Otsuka et al. *Partial Search in a Frozen Network Is Enough to Find a Strong Lottery Ticket*. Feb. 2024. DOI: [10.48550/arXiv.2402.14029](https://doi.org/10.48550/arXiv.2402.14029). arXiv: [2402.14029](https://arxiv.org/abs/2402.14029) [cs, stat]. (Visited on 05/15/2024).
- [24] Ankit Pensia et al. “Optimal lottery tickets via SUBSETSUM: logarithmic overparameterization is sufficient”. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. NIPS ’20. , Vancouver, BC, Canada, Curran Associates Inc., 2020. ISBN: 9781713829546.
- [25] Valentin V. Petrov. *Sums of Independent Random Variables*. Ergebnisse der Mathematik und ihrer Grenzgebiete. 2. Folge. Springer Berlin, Heidelberg, 1975, 348 pages. DOI: <https://doi.org/10.1007/978-3-642-65809-9>.
- [26] Vivek Ramanujan et al. “What’s Hidden in a Randomly Weighted Neural Network?” In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2020, pp. 11890–11899. DOI: [10.1109/CVPR42600.2020.01191](https://doi.org/10.1109/CVPR42600.2020.01191).
- [27] N. Shakhaidarova. “Uniform local and global theorems for densities”. In: *Izv. Akad. Nauk UzSSR Ser. Fiz-Mat. Nauk* 5 (1966), pp. 90–91.
- [28] Irina Shevtsova. *On the absolute constants in the Berry Esseen type inequalities for identically distributed summands*. 2011. arXiv: [1111.6554](https://arxiv.org/abs/1111.6554) [math.PR].
- [29] Hattie Zhou et al. “Deconstructing Lottery Tickets: Zeros, Signs, and the Supermask”. In: *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019 (NIPS 2019)*. 2019, pp. 3592–3602.

A Lower Bound on the Ticket Size in [24]

The claim is a direct consequence of the proof of [24, Theorem 2] (Appendix B). There, in Step 3, it is shown that

$$|\mathcal{G}| \geq \frac{1}{2} \left(\frac{1}{2\varepsilon} \right)^{d^2},$$

where \mathcal{G} is the set of subnetworks that can be formed. Let m be the number of parameters of the original network. If we consider subnetworks of size at most γm ($0 \leq \gamma \leq 1$), we have⁴

$$|\mathcal{G}| \leq \sum_{i=1}^{\gamma m} \binom{m}{\gamma m} \leq 2^{\gamma m \log_2 \left(\frac{m}{\gamma m} e \right)},$$

which combined with the previous inequality implies

$$\gamma m \log_2 \left(\frac{e}{\gamma} \right) \geq d^2 \log_2 \left(\frac{1}{2\varepsilon} \right) - 1$$

If we have an overparameterized network of size $m = \mathcal{O}(d^2 \log_2 \left(\frac{1}{2\varepsilon} \right))$, as in [24], we need $\gamma m = \Theta(m)$ for the last inequality to be satisfied (note that $\log_2 \left(\frac{e}{\gamma} \right) \leq 1$, as $0 \leq \gamma \leq 1$).

B Visualizations




Figure 2: **Simplified representation of the procedure for finding Lottery Tickets (LTH).** A large random neural network (step 1) is trained by iterative pruning with rewind: when the loss reaches a local minimum (step 2), some weights with smallest absolute value are pruned (step 3) and the value of the remaining edges is then reset to that of the initialization (step 4); finally, training is resumed and the final network is obtained (step 5). Remarkably, the sparser subnetwork is consistently able to reach a loss not larger than that right after pruning.



Figure 3: **Simplified representation of the procedure for finding Strongly Lottery Tickets (SLTH) / Training by pruning.** Previous work has shown that it is possible to sparsify large random neural network in order to obtain subnetworks that achieve good performance for a task under consideration, motivating the *Strong Lottery Ticket Hypothesis*. No training is required.

C Proof of Uniform $[-1, 1]$ being Sum-Bounded

In this section we provide a detailed proof of Lemma 1, which states that the uniform distribution in $[-1, 1]$ is sum-bounded, as stated in Definition 1. We remark that, while the proof is written for uniform random variables, it should be possible to extend it to a family of densities which are unimodal, with bounded variance, and bounded third moment.

⁴follows from the upper bound $\sum_{i=1}^k \binom{n}{i} \leq \left(\frac{en}{k} \right)^k$ on the partial sum of binomial coefficients.

Proof of Lemma 1. Note first that the distribution of the sum of n i.i.d. variables in $[0, 1]$ is known as the Irwin–Hall distribution I_n .⁵ We will use that $\text{Var}[I_n] = \frac{n}{12}$, where $\text{Var}[X]$ denotes the variance of the random variable X .

For $n \geq 2$, $f(x, n)$ can be defined as the convolution of $f(x) = f(x, 1)$ and $f(x, n-1)$, i.e.,

$$f(x, n) = \int_{-\infty}^{+\infty} f(x - \tau, n-1) f(\tau) d\tau.$$

It is straightforward to show, by induction and an elementary substitution in the integral above, which is relied upon in the inductive step, that $f(x, n)$ is symmetric about 0, that is $f(x, n) = f(-x, n)$.

Let us now prove by induction that $f(x, n)$ is nondecreasing on the interval $[-n, 0]$ and nonincreasing over $[0, n]$ (for simplicity, since it vanishes outside $[-n, n]$, we can consider directly the negative half and positive half of the real line, respectively, in the argument that follows).

The claims hold trivially for $f(x)$; also note that

$$f(\tau) = \begin{cases} \frac{1}{2} & \text{if } -1 \leq \tau \leq 1 \\ 0 & \text{otherwise} \end{cases} \implies f(x, n) = \frac{1}{2} \int_{-1}^{+1} f(x - \tau, n-1) d\tau.$$

If $x \leq x' \leq -1$. Since $x - \tau \leq x' - \tau \leq 0$, by inductive hypothesis we have that $f(x - \tau, n-1) \leq f(x' - \tau, n-1)$ over the whole interval $\tau \in [-1, 1]$. Taking integrals yields $f(x, n) \leq f(x', n)$.

Now, consider the case when $x \leq -1 \leq x' \leq 0$. If $x + 1 \leq -x' - 1$, $x - \tau \leq x + 1 \leq -x' - 1 \leq -x' + \tau \leq -x + \tau$. By the symmetry about the origin, the inductive hypothesis is $f(x - \tau, n-1) = f(-x + \tau, n-1) \leq f(-x' + \tau, n-1) = f(x' - \tau, n-1)$ over the whole interval $\tau \in [-1, 1]$, since $-1 \leq -x' + \tau \leq -x + \tau$. Taking integrals yields $f(x, n) \leq f(x', n)$. Otherwise, there exists τ_0 such that $x - \tau_0 = -x' - 1$, $x - \tau > -x' - 1$ for all $\tau \in [-1, \tau_0]$ and $x - \tau < -x' - 1$ for all $\tau \in (\tau_0, 1]$. By symmetry, using $-x = x' + 1 - \tau_0$, $f(x - \tau, n-1) = f(-x + \tau, n-1) = f(x' + 1 + \tau - \tau_0, n-1)$. Thus, for all $\tau \in [-1, \tau_0]$, via the change of variable $\sigma = -(1 + \tau - \tau_0)$ in the middle integral below, we obtain that

$$\int_{-1}^{\tau_0} f(x - \tau, n-1) d\tau = \int_{-1}^{\tau_0} f(x' + 1 + \tau - \tau_0, n-1) d\tau = \int_{-1}^{\tau_0} f(x' - \sigma, n-1) d\sigma. \quad (25)$$

For all $\tau \in (\tau_0, 1]$, $x - \tau < -x' - 1 \leq -x' + \tau \leq -x + \tau$, by symmetry about the origin we have that $f(x - \tau, n-1) \leq f(x' - \tau, n-1)$ by the inductive hypothesis with the same reasoning of the case $x + 1 \leq -x' - 1$. Taking integrals over the range $[\tau_0, 1]$ for each term of the inductive hypothesis yields

$$\int_{\tau_0}^1 f(x - \tau, n-1) d\tau \leq \int_{\tau_0}^1 f(x' - \tau, n-1) d\tau \quad (26)$$

Eqs. 25 and 26 imply that $f(x, n) \leq f(x', n)$.

Trivially, if $-1 \leq x \leq x' \leq 0$, analogous ideas are put in place as for the previous case, therefore we omit the details. We have thus shown the nondecreasing monotonicity of $f(x, n)$ on the negative half of the real line. By the symmetry of $f(x, n)$ about the origin, on the positive half of the real line the nondecreasing monotonicity turns into nonincreasing monotonicity, and the proof is complete.

Lower bound (first inequality in Eq. 2). The variance of $\Sigma_{[n]}^{\mathcal{U}_n}$ is $n/3$ since $\Sigma_{[n]}^{\mathcal{U}_n} = 2(I_n(n) - n/2)$ and $\text{Var}[I_n(n)] = n/12$. We define $Z_n^u = \frac{\Sigma_{[n]}^{\mathcal{U}_n}}{\sqrt{n/3}}$ and we note with F_n its cumulative distribution function. Z_n^u has expectation 0 and standard deviation 1. Consider the probability

$$P_L(n) = \Pr(\sqrt{n} \leq \Sigma_{[n]}^{\mathcal{U}_n} \leq 2\sqrt{n}) = \Pr(\sqrt{3} \leq Z_n^u \leq 2\sqrt{3}).$$

⁵It should be known that I_n is unimodal with a mode in $n/2$, but we were not able to find a reference. It is instructive to note, assuming that I_n is unimodal with a mode in $n/2$, it directly follows that its probability density function is increasing on the interval $[0, n/2]$, and then decreasing over $[n/2, n]$. This implies that $f(x, n)$ (the density of $\Sigma_{[n]}^{\mathcal{U}_n}$), is non decreasing in the interval $[-n, 0]$, has maximum at 0, and non increasing $[0, n]$ for all $n \geq 2$.

Now, we use the following form of Berry–Esseen inequality, discussed in [21][p.2)].⁶

Theorem 6 (Allasia [1]). *For all $n \geq 1$,*

$$|F_n(z) - \Phi(z)| \leq \frac{\sqrt{3}}{20\sqrt{n}},$$

where $\Phi(z)$ is the cumulative distribution function of the standard normal distribution.

Theorem 6 implies

$$P_L(n) \geq \Phi(2\sqrt{3}) - \Phi(\sqrt{3}) - 2 \cdot \frac{\sqrt{3}}{20\sqrt{n}}.$$

When $n \geq 18$,

$$\Phi(2\sqrt{3}) - \Phi(\sqrt{3}) - 2 \cdot \frac{\sqrt{3}}{20\sqrt{n}} \geq \Phi(2\sqrt{3}) - \Phi(\sqrt{3}) - 2 \cdot \frac{\sqrt{3}}{20\sqrt{18}} = C_{18} > 0.$$

That is $P_L(n) \geq C_{18} > 0$. When $2 \leq n < 18$, $P_L(n) = F_n(2\sqrt{3}) - F_n(\sqrt{3}) = c_n > 0$. We thus have

$$P_L(n) \geq \min\{C_i, \text{ for } 2 \leq i \leq 18\} = c'_i > 0.$$

Recall that $P_L(n) = \Pr(\sqrt{n} \leq \Sigma_{[n]}^{\mathcal{U}_n} \leq 2\sqrt{n})$. As the density $f(x, n)$ is decreasing on \mathbb{R}^+ , we have

$$P_L(n) \leq f(\sqrt{n}, n)\sqrt{n}.$$

Thus,

$$f(\sqrt{n}, n) \geq \frac{P_L(n)}{\sqrt{n}}.$$

Since $P_L(n) \geq c'_i$ then for all $n \geq 2$

$$f(\sqrt{n}, n) \geq \frac{c'_i}{\sqrt{n}}.$$

When $n = 1$, the density $f(1, 1) = \frac{1}{2}$. So, by setting $c_l = \min(c'_i, \frac{1}{2})$, we get that, for all $n \geq 1$, for all $0 \leq x \leq \sqrt{n}$:

$$f(x, n) \geq f(\sqrt{n}, n) \geq \frac{c_l}{\sqrt{n}}.$$

By a symmetric argument, we also have for all $n \geq 1$, for all $-\sqrt{n} \leq x \leq 0$:

$$f(x, n) \geq f(-\sqrt{n}, n) \geq \frac{c_l}{\sqrt{n}}.$$

Upper bound (second inequality in Eq. 2). Here, we bound the probability distribution function $f(x, n)$ of $\Sigma_{[n]}^{\mathcal{U}_n} = \sqrt{n/3}Z_n$, where we recall that $Z_n^u = \frac{\Sigma_{[n]}^{\mathcal{U}_n}}{\sqrt{n/3}}$. Denoting f_z the probability distribution function of Z_n^u , we have

$$f_z(x, n) = f\left(\sqrt{\frac{n}{3}}x, n\right)\sqrt{\frac{n}{3}}.$$

We use the following local limit theorem, discussed in [25][p.214].

Theorem 7 (Sahaidarova [27]). *Let $\{X_n\}$ be a sequence of independent random variables with a common density $p(x)$, such that $E[|X_1|^3] < \infty$, $E[X_1] = 0$, $E[X_1^2] = 1$ and $\sup p(x) \leq C$. Let $p_n(x)$ be the density of the random variable $\frac{1}{\sqrt{n}} \sum_{j=1}^n X_j$. Then*

$$\sup_x |p_n(x) - \phi(x)| \leq \frac{A\beta_3}{\sqrt{n}} \max(1, C^3),$$

where ϕ is the probability distribution function of a standard gaussian, A is an absolute constant, and $\beta_3 = E[|X_1|^3]$.

⁶It is also possible to obtain our result via classical Berry-Esseen inequality, due to the improved upper bound of 0.4748 on the absolute constant, provided in [28]. This would require replacing with 900 the cut-off value for n , which is 18 in the current version of the argument.

The theorem can be applied to a uniform continuous distribution with density $p^u(x) = \frac{1}{2\sqrt{3}}$ in the interval $[-\sqrt{3}, \sqrt{3}]$, which has mean 0 and variance 1. We thus get, for every $x \in \mathbb{R}$,

$$f_z(x, n) = p_n^u(x) \leq \phi(0) + \frac{A\beta_3}{\sqrt{n}} = \frac{1}{2\pi} + \frac{A\beta_3}{\sqrt{n}} \leq \frac{1}{2\pi} + A\frac{3\sqrt{3}}{4} = c'_u.$$

In conclusion, setting $c_u = \sqrt{3}c'_u$, for every $x \in \mathbb{R}$ it holds that

$$f(x, n) = \sqrt{\frac{3}{n}} f_z\left(\sqrt{\frac{3}{n}}x, n\right) \leq \frac{\sqrt{3}c'_u}{\sqrt{n}} = \frac{c_u}{\sqrt{n}}.$$

□

D Proof of Corollary 1

Proof of Corollary 1. As anticipated, we proceed in three steps.

Step 1: Hoeffding bound. We start by showing, following the idea at the base of [17, Corollary 3.3], that if n' is large enough, a standard Hoeffding bound ensures that with high probability a constant fraction of the sample follows a Uniform $[-1, 1]$ distribution. Since we assumed that every X_i is a mixture of a Uniform $[-1, 1]$ distribution with probability p , and another distribution with density g (given by the factors G_i), we can rewrite $X_i = B_i \cdot U_i + (1 - B_i) \cdot G_i$, with U_i being the uniform random variable, G_i being the random variable with density g , B_i being independent Bernoulli random variables with probability p .

Fix $\alpha = \alpha(p) \neq p$, and assume, for now, that n' satisfies Eq. 1, and therefore, since $\varepsilon < 1/2$, choosing $c_{\text{hyp}} = c_{\text{hyp}}(p) \geq (\alpha - p)^{-2}$, ensures that, defining $\varepsilon' = \varepsilon/2$,

$$n' \geq c_{\text{hyp}} \log_2 \frac{1}{\varepsilon} \geq \frac{1}{2(\alpha - p)^2} \ln \frac{1}{\varepsilon'}$$

and therefore

$$\Pr\left(\sum_i^{n'} B_i \leq \alpha n'\right) \leq e^{-2(\alpha - p)^2 n'} \leq e^{-\ln \frac{1}{\varepsilon'}} = \varepsilon'.$$

Thus

$$\Pr\left(\sum_i^{n'} B_i > \alpha n'\right) \geq 1 - \varepsilon',$$

that is, with high probability, there is a set of indices $I \subseteq [n']$ of size $|I| \geq \alpha n'$, such that for each $i \in I$ it holds $B_i = 1$, i.e. X_i is uniformly distributed.

Step 2: Application of Theorem 2 via rejection-sampling. Lemma 1 ensures that the uniform distribution of the $|I|$ random variables selected in *Step 1* is sum-bounded. Conditionally on the event $\{\sum_i^{n'} B_i > \alpha n'\}$, we can discard all random variables indexed outside I and apply directly Theorem 2 to $\alpha n'$ of the remaining ones, for any fixed k and $z \in [-\sqrt{k}, \sqrt{k}]$, since $\alpha c_{\text{hyp}} \geq 1$ by construction. This guarantees a success probability of c'_{thm} for approximating the given target z ; thus,

$$\begin{aligned} & \Pr(\exists S_z \subset [n], |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon') \geq \\ & \Pr\left(\exists S_z \subset [n], |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon' \mid \sum_i^{n'} B_i > \alpha n'\right) \Pr\left(\sum_i^{n'} B_i > \alpha n'\right) \geq \\ & \Pr\left(\exists S_z \subset I, |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon' \mid |I| > \alpha n'\right) (1 - \varepsilon') \geq c'_{\text{thm}}(1 - \varepsilon') \geq \frac{3}{4}c'_{\text{thm}} = c_{\text{thm}}. \end{aligned}$$

Step 3: Amplification. Finally, by a standard probability amplification argument and a union bound applied to Theorem 2, by paying an extra factor $\log_2(k/\varepsilon)$ in Eq. 1, the constant c_{thm} can be amplified to $1 - \varepsilon$, and the existence of a suitable subset S_z holds simultaneously for all $z \in [-\sqrt{k}, \sqrt{k}]$. We now give more details on this amplification.

Recall that $\varepsilon' = \frac{\varepsilon}{2}$, and let $c_{\text{amp}} = c_{\text{amp}}(p) = 8 \frac{c_{\text{hyp}}}{c_{\text{thm}}}$ and $r = \frac{4}{c_{\text{thm}}} \ln \frac{k}{\varepsilon}$. By assumption,

$$n \geq c_{\text{amp}} \frac{\log_2^2 \frac{k}{\varepsilon}}{H_2\left(\frac{k}{n}\right)} \geq 2rc_{\text{hyp}} \frac{\log_2 \frac{k}{\varepsilon}}{H_2\left(\frac{k}{n}\right)} \geq rc_{\text{hyp}} \frac{\log_2 \frac{k}{\varepsilon'}}{H_2\left(\frac{k}{n}\right)},$$

where the last inequality is ensured by $\varepsilon < 1/2$. By *Step 2*, we can apply Theorem 2, with ε' and $n' \geq c_{\text{hyp}} \frac{\log_2 \frac{k}{\varepsilon'}}{H_2\left(\frac{k}{n}\right)} = n^*$, allowing us to prove that we can ε' -approximate any target z with probability at least c_{thm} . The probability of failing to approximate some given z is then at most $1 - c_{\text{thm}}$. From the sample Ω of sum-bounded random variables take r subsamples (without replacement) of cardinality n^* each, $\Omega_1, \dots, \Omega_r$. The probability of failing to approximate some given z with subsetsums from Ω is less than that of failing to approximate it with subsetsums from within every Ω_i 's, and the latter probability is at most $(1 - c_{\text{thm}})^r$; thus, for every $z \in [-\sqrt{k}, \sqrt{k}]$,

$$\Pr(\nexists S_z \subset [n], |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon') \leq (1 - c_{\text{thm}})^r.$$

By an union bound, we also have that

$$\begin{aligned} & \Pr\left(\forall z \in [-\sqrt{k}, \sqrt{k}], \exists S_z \subset [n], |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon\right) \\ & \geq \Pr\left(\forall z \in \left\{-\sqrt{k} + i\varepsilon' : i \in \left[\frac{2}{\varepsilon'}\sqrt{k}\right]\right\}, \exists S_z \subset [n], |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon'\right) \\ & = 1 - \Pr\left(\exists z \in \left\{-\sqrt{k} + i\varepsilon' : i \in \left[\frac{2}{\varepsilon'}\sqrt{k}\right]\right\}, \nexists S_z \subset [n], |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon'\right) \\ & \geq 1 - \sum_{z \in \{-\sqrt{k} + i\varepsilon' : i \in [\frac{2}{\varepsilon'}\sqrt{k}]\}} \Pr(\nexists S_z \subset [n], |S_z| = k : |\Sigma_{S_z} - z| < \varepsilon') \\ & \geq 1 - \frac{2}{\varepsilon'}\sqrt{k}(1 - c_{\text{thm}})^r = 1 - \frac{2}{\varepsilon'}\sqrt{k} \exp\left(\frac{4}{c_{\text{thm}}} \ln\left(\frac{k}{\varepsilon}\right) \cdot \ln(1 - c_{\text{thm}})\right) \\ & \geq 1 - \frac{2}{\varepsilon'}\sqrt{k} \exp\left(-4 \ln \frac{k}{\varepsilon}\right) = 1 - \frac{2}{\varepsilon'}\sqrt{k} \frac{\varepsilon^4}{k^4} \geq 1 - 4\varepsilon^3 \geq 1 - \varepsilon, \end{aligned}$$

where the last inequality is ensured by $\varepsilon < 1/2$. This completes the proof. \square

E Proof of Corollary 2

Proof of Corollary 2. By definition of binary entropy, we have

$$H_2\left(\frac{k}{n}\right) = \frac{k}{n} \log_2\left(\frac{n}{k}\right) + \left(1 - \frac{k}{n}\right) \log_2 \frac{n}{n-k} \quad (27)$$

In particular, since both terms in the previous equation are positive, we get

$$H_2\left(\frac{k}{n}\right) \geq \frac{k}{n} \log_2\left(\frac{n}{k}\right) \quad (28)$$

We now use eq. (28) to derive an upper bound for the quantity $\frac{c_{\text{amp}}}{H_2\left(\frac{k}{n}\right)} \frac{\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon}}{n}$, which will be used later:

$$\frac{c_{\text{amp}}}{H_2\left(\frac{k}{n}\right)} \frac{\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon}}{n} \leq \frac{c_{\text{amp}}}{\frac{k}{n} \log_2\left(\frac{n}{k}\right)} \frac{\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon}}{n}$$

$$= c_{\text{camp}} \frac{\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon}}{k} \frac{1}{\log_2 \left(\frac{n}{k}\right)} \quad (29)$$

$$\leq c_{\text{camp}} \frac{\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon}}{k} \quad (30)$$

$$\leq \frac{1}{2}, \quad (31)$$

where from eq. (29) to eq. (30) we used that $\log_2^{n/k} \geq 1$ for $k \leq n/2$, and then the hypothesis $k \geq 2c_{\text{camp}} (\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon})$ directly gives eq. (31). Let us now rewrite eq. (3) in a more convenient form:

$$\begin{aligned} n \frac{H_2 \left(\frac{k}{n}\right)}{c_{\text{camp}}} &\geq \log_2^2 \frac{k}{\varepsilon} \\ n \frac{H_2 \left(\frac{k}{n}\right)}{c_{\text{camp}}} &\geq \log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon} + \log_2^2 \frac{1}{\varepsilon} \\ n \left(\frac{H_2 \left(\frac{k}{n}\right)}{c_{\text{camp}}} - \frac{\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon}}{n} \right) &\geq \log_2^2 \frac{1}{\varepsilon} \\ n \left(1 - \frac{c_{\text{camp}}}{H_2 \left(\frac{k}{n}\right)} \frac{\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon}}{n} \right) &\geq \frac{c_{\text{camp}}}{H_2 \left(\frac{k}{n}\right)} \log_2^2 \frac{1}{\varepsilon} \end{aligned} \quad (32)$$

$$n \geq \frac{c_{\text{camp}}}{\left(1 - \frac{c_{\text{camp}}}{H_2 \left(\frac{k}{n}\right)} \frac{\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon}}{n} \right)} \log_2^2 \frac{1}{\varepsilon} \quad (33)$$

Using eq. (31) we get

$$\frac{c_{\text{camp}}}{\left(1 - \frac{c_{\text{camp}}}{H_2 \left(\frac{k}{n}\right)} \frac{\log_2^2 k + 2\log_2 k \cdot \log_2 \frac{1}{\varepsilon}}{n} \right)} \leq 2c_{\text{camp}} \quad (34)$$

To satisfy eq. (33), we can then choose n such that

$$n \geq 2c_{\text{camp}} \frac{\log_2^2 \frac{1}{\varepsilon}}{H_2 \left(\frac{k}{n}\right)}, \quad (35)$$

and then apply Corollary 1 to end the proof. \square

F Details for the proof of Theorem 2

Proof of Eq. 17 Define $A = \tilde{S}' \setminus \tilde{S}$, and observe that

$$\Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_{\mu k, k-1}) \quad (36)$$

$$= \sum_{i=\mu k}^{k-1} \Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_i) \Pr(I_i | H_{\tilde{S}}, I_{\mu k, k-1}) \quad (37)$$

$$\begin{aligned} &= \sum_{i=\mu k}^{k-1} \int_{-\infty}^{\infty} \Pr(|\Sigma_A - (z - y)| < \varepsilon | \Sigma_I = y, I_i, H_{\tilde{S}}) \Pr(\Sigma_I = y | H_{\tilde{S}}, I_i) dy \\ &\quad \cdot \Pr(I_i | H_{\tilde{S}}, I_{\mu k, k-1}) \end{aligned} \quad (38)$$

$$\begin{aligned} &= \sum_{i=\mu k}^{k-1} \int_{-\infty}^{\infty} \Pr(|\Sigma_A - (z - y)| < \varepsilon | \Sigma_I = y, I_i) \Pr(\Sigma_I = y | H_{\tilde{S}}, I_i) dy \\ &\quad \cdot \Pr(I_i | H_{\tilde{S}}, I_{\mu k, k-1}) \end{aligned} \quad (39)$$

$$\leq c\varepsilon \sum_{i=\mu k}^{k-1} \int_{-\infty}^{\infty} \Pr(\Sigma_I = y | H_{\tilde{S}}, I_i) dy \Pr(I_i | H_{\tilde{S}}, I_{\mu k, k-1}) \quad (40)$$

$$\leq c\varepsilon$$

where from Eq. 36 to Eq. 37 and from Eq. 37 to Eq. 38 we used the law of total probability;⁷ from Eq. 38 to Eq. 39 we dropped the redundant event $H_{\tilde{S}}$ in the conditioning, due to conditional independence; finally, from Eq. 39 to Eq. 40 we used Definition 1 which implies that for any $i \in \{\mu k, \dots, k-1\}$ it holds

$$\Pr(|\Sigma_A - (z - y)| < \varepsilon | \Sigma_I = y, I_i) = \Pr(|\Sigma_{[k-i]} - (z - y)| < \varepsilon) \leq c\varepsilon.$$

Proof of Eq. 18 Let $A = \tilde{S}' \setminus \tilde{S}$. Analogously to the calculations from Eq. 37 to Eq. 39, by the law of total probability we have

$$\begin{aligned} & \sum_{i=0}^{\mu k-1} \Pr(I_i) \cdot \Pr(H_{\tilde{S}'} | H_{\tilde{S}}, I_i) \\ &= \sum_{i=0}^{\mu k-1} \Pr(I_i) \cdot \int_{-\infty}^{\infty} \Pr(|\Sigma_A - (z - y)| < \varepsilon | \Sigma_I = y, I_i) \Pr(\Sigma_I = y | H_{\tilde{S}}, I_i) dy \\ &= \sum_{i=0}^{\mu k-1} \Pr(I_i) \cdot \int_{-\infty}^{\infty} \Pr(|\Sigma_{[k-i]} - (z - y)| < \varepsilon) \Pr(\Sigma_I = y | H_{\tilde{S}}, I_i) dy \end{aligned} \quad (41)$$

$$\begin{aligned} & \leq c \frac{\varepsilon}{\sqrt{k}} \sum_{i=0}^{\mu k-1} \Pr(I_i) \cdot \int_{-\infty}^{\infty} \Pr(\Sigma_I = y | H_{\tilde{S}}, I_i) dy \\ & \leq c \frac{\varepsilon}{\sqrt{k}} \end{aligned} \quad (42)$$

where from Eq. 41 to Eq. 42 we used Definition 1, which implies that for any $i \in \{0, \dots, \frac{9}{10}k - 1\}$ it holds

$$\Pr(|\Sigma_{[k-i]} - (z - y)| < \varepsilon) \leq c' \frac{\varepsilon}{\sqrt{k-i}} \leq c \frac{\varepsilon}{\sqrt{k}}.$$

G Proof of Theorem 3

In the proof we will refer to the following results, upon which [24][Theorem 1] relies (the statement below slightly differ as we fix two small typos in their notation and mixing coefficients). With the understanding that by a mixture D of a distribution D_1 and D_2 with probability p it is meant that the pdf (we adopt the convention that this term includes generalised functions, such as Dirac deltas for point masses) of D can be written as a convex combination of the pdf of D_1 and that of D_2 , that is $f_D = pf_{D_1} + (1-p)f_{D_2}$. For the unfamiliar reader, we note that in the literature this is often stated in short as $D = pD_1 + (1-p)D_2$.

Lemma 2 ([24][Corollary 1]). *Let $X \sim \text{Uniform}[0, 1]$ (or $X \sim \text{Uniform}[-1, 0]$) and $Y \sim \text{Uniform}[-1, 1]$ be independent random variables. Let P be the distribution of the XY and δ_0 the Dirac delta at 0. Let D be the distribution obtained as mixture of δ_0 and P with probability $1/2$. Then D is the mixture of a $\text{Uniform}[-1/2, 1/2]$ and some distribution Q with probability $\ln(2)/4$.*

Corollary 3 ([24][Corollary 2]). *Let X_1, \dots, X_n be iid with distribution D as defined in Lemma 2, where $n \geq C \ln(2/\varepsilon)$ for some universal constant C . Then*

$$\Pr\left(\forall z \in [-1, 1], \exists S \subset [n] : |z - \sum_{i \in S} X_i| \leq \varepsilon\right) \geq 1 - \varepsilon.$$

Proof of Theorem 3. The key idea is exploiting Corollary 1 at each step of the pruning strategy established in [24][Theorem 1], where Corollary 3 is used instead. Without loss of generality, we replace their $\min\{\varepsilon, \delta\}$ with ε . For the sake of easily following the approach adopted in [24], let us define $n^*(x)$ as the function

$$n^*(x) = c_{\text{camp}} \frac{\log_2^2(kx)}{H_2\left(\frac{k}{n^*(x)}\right)} \quad (43)$$

⁷For simplicity, we denote the density of Σ_I conditional on $H_{\tilde{S}} \cap I_i$ as $\Pr(\Sigma_I = y | H_{\tilde{S}}, I_i)$.

where $k = \gamma' n^*(x)$. In the following, we use n^* as short for $n^*(1/\varepsilon)$, and we will only explicitly provide an argument for n^* when it is different than $1/\varepsilon$. For instance, in the last step of the proof, we will use $n^*(2\ell d_i d_{i-1}/\varepsilon)$, which matches the definition of n_i^* given in Eq. 20.

Consider [24][Lemma 1]. When approximating a single link (that is, a weight), after the overparameterization (which creates an additional layer of width $2n^*$ in between the input and the output node) via $4n^*$ links, instead of pruning via Corollary 3, we prune via Corollary 1 twice in the second layer, that is we ensure that only $k = \gamma' n^*$ edges yield the desired approximation, both in the edges corresponding to the positive part of the input weights and in those corresponding to the negative part. Thus we obtain at most $4k$ surviving edges, after the preprocessing step and the pruning mask is applied. This yields a sparsity of at least $\alpha' = 1 - \gamma'$. Note that it is because of the preprocessing step that we go from distributions Uniform $[-1, 1]$ to distributions D , as defined in Lemma 2, which are shown to be a mixture with Uniform $[-1, 1]$ and therefore can be also handled via Corollary 1.

Consider [24][Lemma 2]. When approximating a real-valued multivariate linear function, after the overparameterization (which creates an additional layer of width $2dn^*(d/\varepsilon)$ in between the d input nodes and the output node) one simply iterates the ideas of the previous case d times. For each input node, the overparameterization surviving the preprocessing step on the weights of the input layer is $4n^*(d/\varepsilon)$. Pruning the second layer of the overparameterized link for each input via Corollary 1 with $k = \gamma' n^*(d/\varepsilon)$ (again, performing this both on the edges corresponding to the positive part of the input weights and in those corresponding to the negative part), instead of exploiting Corollary 3, yields that at most $4dk$ edges survive after the pruning mask is applied. This yields a sparsity of at least $\alpha' = 1 - \gamma'$.

Finally, consider [24][Lemma 3]. When approximating a layer with input dimension d_1 and output dimension d_2 , after the overparameterization (which creates an additional layer of width $2d_1 n^*(d_1 d_2/\varepsilon)$ in between the input nodes and the output nodes) one iterates the ideas of the previous case d_1 times in the input layer through the same preprocessing step, and d_2 times in the output layer, one for each of the d_1 blocks created by the preprocessing (essentially the weights in the input layer are *re-used* d_2 times). For each input node, the overparameterization surviving the preprocessing step is at most $2(d_2 + 1)n^*(d_1 d_2/\varepsilon)$. Overall, after the preprocessing step, we have at most $2d_1(d_2 + 1)n^*(d_1 d_2/\varepsilon)$ parameters. We then use Corollary 1 (with $k = \gamma' n^*(d_1 d_2/\varepsilon)$) to prune the number of parameters between the introduced additional layer and the d_2 outputs down to $2d_1 d_2 \gamma' n^*(d_1 d_2/\varepsilon)$. As for the edges between the d_1 inputs and the additional layer, only those that reach a neuron in the additional layer, from which there is at least one outgoing edge towards the d_2 outputs, are used; since for each of the d_1 blocks of $2n^*(d_1 d_2/\varepsilon)$ neurons in the additional layer we only kept $2\gamma' n^*(d_1 d_2/\varepsilon)$ outgoing edges to each of the d_2 output neurons, in the worst case (all the nodes involved in the subsetsums are disjoint) we keep $2d_2 \gamma' n^*(d_1 d_2/\varepsilon)$ of them for each of the d_1 neurons. Globally, we are left with a total of at most $2d_1 d_2 \gamma' n^*(d_1 d_2/\varepsilon)$ edges both in the input layer and in the output layer, thus a total of $4d_1 d_2 \gamma' n^*(d_1 d_2/\varepsilon)$ edges survive the pruning. The density of the surviving edges is then less than

$$\frac{4d_1 d_2 \gamma' n^*(d_1 d_2/\varepsilon)}{2d_1^2 n^*(d_1 d_2/\varepsilon) + 2d_1 d_2 n^*(d_1 d_2/\varepsilon)} = \frac{2d_2 \gamma'}{d_1 + d_2} = \frac{(d_1 \frac{d_2}{d_1} + d_2) \gamma'}{d_1 + d_2} \leq \rho_1 \gamma',$$

where $\rho_1 = \max\{d_1/d_2, d_2/d_1\}$ and in the last inequality we used that $d_1 d_2/d_1 + d_2 \leq \rho_1(d_1 + d_2)$ since $\rho_1 \geq 1$. This ensures a sparsity $\alpha' \geq 1 - \rho_1 \gamma'$.

[24][Theorem 1] consists of performing, for every $i \in [\ell]$, the previous step on layer i with input dimension d_{i-1} and output dimension d_i . The overparameterization creates an additional layer of nodes of width $2d_{i-1} n^*(2\ell d_{i-1} d_i/\varepsilon)$ in between the d_{i-1} input nodes and the d_i output nodes. Since the construction is stacked ℓ times, this generates 2ℓ layers for the overparameterized network, which will therefore have a starting number of parameters

$$m = \sum_{i=1}^{\ell} 2d_{i-1}^2 n^*(2\ell d_{i-1} d_i/\varepsilon) + 2d_{i-1} d_i n^*(2\ell d_{i-1} d_i/\varepsilon).$$

Corollary 1 applied to each stacked overparameterized layer instead of Corollary 3 as in the previous step yields that the total number of parameters left after the pruning is

$$m_t \leq \sum_{i=1}^{\ell} 4d_{i-1} d_i k_i,$$

where $k_i = \gamma' n^{*(2\ell d_{i-1} d_i / \varepsilon)}$. Recall that $\rho = \max_i \rho_i$, where $\rho_i = \max\{d_i/d_{i-1}, d_{i-1}/d_i\} \geq 1$. Recall that $\gamma = \rho\gamma'$. We obtain that

$$\begin{aligned}
m_t &\leq \sum_{i=1}^{\ell} 2d_{i-1} \frac{d_i}{d_{i-1}} d_{i-1} k_i + 2d_{i-1} d_i k_i \\
&\leq \sum_{i=1}^{\ell} 2d_{i-1} \rho_i d_{i-1} k_i + 2d_{i-1} d_i \rho_i k_i \\
&\leq \rho \sum_{i=1}^{\ell} 2d_{i-1}^2 k_i + 2d_{i-1} d_i k_i \\
&= \rho\gamma' \sum_{i=1}^{\ell} 2d_{i-1}^2 n^{*(2\ell d_{i-1} d_i / \varepsilon)} + 2d_{i-1} d_i n^{*(2\ell d_{i-1} d_i / \varepsilon)} = \gamma m
\end{aligned}$$

We then get that the density of the edges surviving the pruning is $m_t/m \leq \gamma$, which implies a sparsity of at least $\alpha = 1 - \gamma$. □

H Proof of Theorem 5

Proof of Theorem 5. Consider the space $\mathcal{W}_k = \{W \in \mathbb{R}^{d \times d} : \|W\| \leq \sqrt{k}\}$, and let \mathcal{P}_k be a 2ε -separated set of \mathcal{W}_k , i.e. a subset $\mathcal{P}_k \subset \mathcal{W}_k$ such that for all distinct $W, W' \in \mathcal{P}_k$ it holds $\|W - W'\| > 2\varepsilon$. We denote $\mathcal{W} = \mathcal{W}_1$, $\mathcal{P} = \mathcal{P}_1$, and the set of all possible subnetworks of g as \mathcal{G} (note that this does not denote \mathcal{G}_1 , the set of all subnetworks of size 1).

Step 1: Packing argument. In [24][Theorem 2, Step 1], it is shown that any function g' can only approximate at most one member of \mathcal{P} for bounded input x (say, $\|x\| \leq 1$). In particular, this also applies to functions g' representing the elements of \mathcal{G}_k .

Step 2: Relation between $|\mathcal{G}_k|$ and $|\mathcal{P}_k|$. By *Step 1*, in [24][Theorem 2, Step 2] it is shown that $|\mathcal{P}| \leq 2|\mathcal{G}|$, under the assumption of Eq. 24, with \mathcal{G}_k replaced by \mathcal{G}). Therefore, also by *Step 1*, replacing \mathcal{P} with \mathcal{P}_k and \mathcal{G} with \mathcal{G}_k in [24][Theorem 2, Step 2], it holds that $|\mathcal{P}_k| \leq 2|\mathcal{G}_k|$. Note that $|\mathcal{G}_k| = \binom{n}{k}$, the number of different ways in which we can select k parameters out of n , so we actually get

$$\binom{n}{k} > \frac{|\mathcal{P}_k|}{2}. \quad (44)$$

Step 3: Lower bound on $|\mathcal{P}_k|$. Let us now consider a 2ε -separated set \mathcal{P}_k^{\max} of maximal cardinality. In [24][Theorem 2, Step 3] it is shown that

$$|\mathcal{P}_k^{\max}| \geq \frac{\text{Vol}(\mathcal{W})}{\text{Vol}(\{W \in \mathcal{W} : \|W\| \leq 2\varepsilon\})} = \left(\frac{1}{2\varepsilon}\right)^{d^2}.$$

Here Vol is the Lebesgue measure in $\mathbb{R}^{d \times d}$ identified with \mathbb{R}^{d^2} . By the exact same argument, replacing \mathcal{W} with \mathcal{W}_k and thus \mathcal{P}_k^{\max} with \mathcal{P}_k^{\max} , it holds that

$$|\mathcal{P}_k^{\max}| \geq \frac{\text{Vol}(\mathcal{W}_k)}{\text{Vol}(\{W \in \mathcal{W}_k : \|W\| \leq 2\varepsilon\})} = \left(\frac{\sqrt{k}}{2\varepsilon}\right)^{d^2}.$$

Combining this fact with Eq. 44 applied to \mathcal{P}_k^{\max} implies that

$$\binom{n}{k} > \frac{1}{2} \left(\frac{\sqrt{k}}{2\varepsilon}\right)^{d^2}. \quad (45)$$

Step 4: Lower bound on n . Consider the standard bound found in [19]

$$\binom{n}{k} \leq \sqrt{\frac{n}{2\pi k(n-k)}} 2^{nH_2(k/n)}.$$

and combine it with with Eq. 45. It follows that

$$2^{nH_2(k/n)} \geq \frac{1}{2} \sqrt{\frac{2\pi k(n-k)}{n}} \left(\frac{\sqrt{k}}{2\varepsilon}\right)^{d^2}$$

and taking the logarithm of both sides yields the sought lower bound on n :

$$nH_2\left(\frac{k}{n}\right) \geq \frac{1}{2} \log_2\left(\frac{2\pi k(n-k)}{n}\right) + d^2 \log_2 \frac{\sqrt{k}}{2\varepsilon} - 1 \quad (46)$$

$$\geq d^2 \left(\frac{1}{2} \log_2 k + \log_2 \frac{1}{\varepsilon} - 1\right) - 1 \quad (47)$$

$$\geq \frac{d^2}{2} \log_2 \frac{k}{\varepsilon}, \quad (48)$$

where from Eq. 46 to Eq. 47 we exploited the definition of λ , which ensures that the first term in the r.h.s. of Eq. 46 is nonnegative;⁸ from Eq. 47 to Eq. 48 we used that for all $\varepsilon < 1/16$ it holds that

$$d^2 \left(\log_2 \frac{1}{\varepsilon} - 1\right) \geq 1.$$

□

I Details of comparison with Malach et al. [20]

We show that $cz \geq c_{\text{amp}} \frac{\log_2^2(cz^2)}{\log_2(z)}$ holds for a big enough constant c . Recall that $z = \frac{m_i}{\varepsilon}$, so we can always assume $\log_2(z) \geq 1$. We have

$$\log_2^2(cz^2) = (\log_2(c) + 2\log_2(z))^2 \quad (49)$$

$$= \log_2^2(c) + 4\log_2(c)\log_2(z) + 4\log_2^2(z) \quad (50)$$

$$\stackrel{(a)}{\leq} 6(\log_2^2(c) + \log_2^2(z)) \quad (51)$$

$$\stackrel{(b)}{\leq} 12\log_2^2(c)\log_2^2(z), \quad (52)$$

where in (a) we used that $2ab \leq a^2 + b^2$, and in (b) that $a + b \leq 2ab$ for a and b greater than 1.

We can then focus on showing that there is a big enough constant c such that $cz \geq 12c_{\text{amp}} \log_2^2(c)\log_2^2(z)$. We get $c \geq 12c_{\text{amp}} \log_2^2(c) \frac{\log_2^2(z)}{z}$, and we have

$$\log_2^2(c) \frac{\log_2^2(z)}{z} \leq \log_2^2(c) \quad (53)$$

$$\leq \sqrt{c}. \quad (54)$$

We can then focus on $c \geq 12c_{\text{amp}}\sqrt{c}$, which is satisfied for $c \geq 144c_{\text{amp}}^2$.

⁸This term being nonnegative is equivalent to $k(1 - k/n) \geq 1/2\pi$, and since $1 \leq k \leq \lambda n$, any $\lambda \leq 1 - 1/2\pi$ ensures it.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and the introduction precisely reflect our contribution.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have added a discussion on the limitations.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Assumptions for all results are accurately stated in our statements.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification: The paper does not include experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics [https://neurips.cc/public/EthicsGuidelines?](https://neurips.cc/public/EthicsGuidelines)

Answer: [Yes]

Justification: The research conducted confirms in every respect with the NeurIPS Code of Ehtics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: A small section discussing the potential impacts of our work has been added to the paper.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We not releasing any data or models (apart from mathematical models) hence there is no risk of misues.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: Everything is cited properly and no assests are required.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No assets are introduced.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No such research was conducted.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: There are not study participants.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.