

The Strong Lottery Ticket Hypothesis for Multi-Head Attention Mechanisms

Hikari Otsuka¹, Daiki Chijiwa², Yasuyuki Okoshi¹,
Daichi Fujiki¹, Susumu Takeuchi², Masato Motomura¹

¹Institute of Science Tokyo

²NTT, Inc.

otsuka.hikari@artic.iir.isct.ac.jp

Abstract

The strong lottery ticket hypothesis (SLTH) conjectures that high-performing subnetworks, called strong lottery tickets (SLTs), are hidden in randomly initialized neural networks. Although recent theoretical studies have established the SLTH across various neural architectures, the SLTH for transformer architectures still lacks theoretical understanding. In particular, the current theory of the SLTH does not yet account for the multi-head attention (MHA) mechanism, a core component of transformers. To address this gap, we introduce a theoretical analysis of the existence of SLTs within MHAs. We prove that, if a randomly initialized MHA of H heads and input dimension d has the hidden dimension $O(d \log(Hd^{3/2}))$ for the key and value, it contains an SLT that approximates an arbitrary MHA with the same input dimension with high probability. Furthermore, by leveraging this theory for MHAs, we extend the SLTH to transformers without normalization layers. We empirically validate our theoretical findings, demonstrating that the approximation error between the SLT within a source model (MHA and transformer) and an approximate target counterpart decreases exponentially by increasing the hidden dimension of the source model.

Code — <https://github.com/Hikari43/SLTH-for-MHAs>

Extended version — <https://arxiv.org/abs/2511.04217>

1 Introduction

The *lottery ticket hypothesis* (Frankle and Carbin 2019)—overparameterized networks contain subnetworks that achieve comparable accuracy to fully trained networks even if trained in isolation—presented new possibilities for compact and high-performing models inherent in recent deep neural networks. Later, a stronger claim, which is formally defined as the *strong lottery ticket hypothesis* (SLTH), was proposed (Ramanujan et al. 2020; Malach et al. 2020): overparameterized networks contain subnetworks (called *strong lottery tickets* (SLTs)) that achieve comparable accuracy to the trained dense network even without any training. Whether such subnetworks exist is a fascinating question in itself, and studying them can bring us closer to understanding the principles behind overparameterized models.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

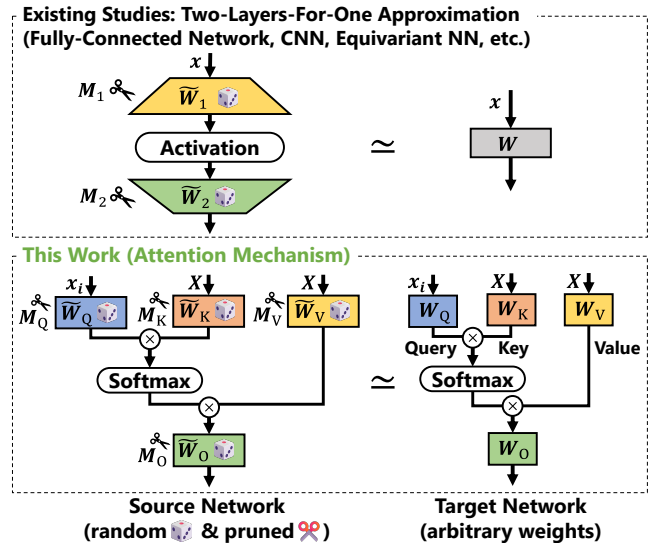


Figure 1: Comparison of the approximation techniques in conventional theories of the SLTH (top) and in our attention-specific approach (bottom). This work demonstrates that an arbitrary attention mechanism can be approximated by pruning a randomly initialized one.

The rigorous proof for the SLTH was firstly established in fully-connected networks. Early studies showed that a randomly-weighted fully-connected network of sufficient width (a *source network*) contains an SLT, which approximates an arbitrary fully-connected network with half the depth (a *target network*) (Malach et al. 2020; Orseau, Hutter, and Rivasplata 2020; Pensia et al. 2020). These theories are built on the foundational argument called a *two-layers-for-one approximation*: a two-layer source network with random weight matrices contains an SLT that approximates a single-layer target network with an arbitrary weight matrix (the top panel of Figure 1). Following this finding, subsequent studies have succeeded in proving the existence of SLTs in more complex networks, such as convolutional and equivariant networks (da Cunha, Natale, and Viennot 2022; Burkholz 2022a; Ferbach et al. 2023).

However, the theoretical foundation of the SLTH for *transformers*, which form the basis of modern language

models, remains unexplored—due to a transformer-specific component, *an attention mechanism*. As shown in the bottom panel of Figure 1 (right side), one of the distinctive structures in transformers is the inner product between two vectors called *query* and *key*, obtained as linear projections of given inputs. This structure fundamentally differs from the conventional components of non-transformer architectures for which the SLTH has been established (the top panel of Figure 1); thus, it remains a mystery whether transformers contain SLTs under existing theoretical insights. This gap motivates our key research question: *does an attention mechanism—an essential component of transformers—contain an SLT?*

In this work, we prove the existence of SLTs within attention mechanisms, extending the SLTH to transformers. More precisely, we prove a suitably pruned source attention mechanism with random weights can approximate any target attention mechanism with arbitrary weights:

Theorem 1 (informal). *Given inputs of length T , a suitably pruned randomly-initialized attention mechanism of the input dimension d and hidden dimension $n = O(d \log(d^{3/2}/\epsilon))$ can approximate an arbitrary attention mechanism of the same input dimension with an approximation error ϵ , with high probability.*

Our key idea is to reinterpret the inner product between the query and key vectors in the attention mechanism as a (linear) neural network weighted by the query and key projection matrices. Then, we can view the source and target inner products as neural networks with different numbers of layers: the source one has two layers with query and key projection matrices as its weights, while the target one has a single layer with a weight matrix obtained by merging these two projections. This reinterpretation makes it possible to apply a variant of the two-layers-for-one approximation, leading to the SLT existence within attention mechanisms (Theorem 3). Note that, as can be seen by comparing the top and bottom panels of Figure 1, our arguments do not require additional layers in the MHA for approximation, in contrast to the previous two-layers-for-one argument for fully-connected networks. By exploiting this theorem, we further establish the SLTH for transformers without normalization layers: a randomly-initialized transformer has an SLT that approximates an arbitrary transformer with similar structures (Theorem 6).

We also empirically validate our theory and confirm its implications. Specifically, we show that 1) the approximation error between the source and target attentions (or, more generally, source and target transformers) decays exponentially as the hidden dimension increases; and 2) this approximation error does not diverge even when the input length T increases. Also, based on our theoretical arguments, we derive a new, practical weight initialization scheme, leading to better SLTs in our experiments.

Our contributions are summarized as follows:

- We provide the first theoretical proof that SLTs exist within attention mechanisms and transformers by reinterpreting the inner product in attention mechanisms.
- We then empirically validate our theory under conditions

that are close to our theoretical assumptions. More precisely, we carefully designed a synthetic experiment to observe how the hidden dimension or input length affects the approximation error of SLTs.

- Furthermore, we demonstrate that our theory not only explains the empirical results, but also provides a new insight into a weight initialization for finding better SLTs in practical settings.

Notation: In this paper, scalars, vectors, and matrices are denoted by lowercase, bold lowercase, and bold uppercase letters, respectively. We use the norm of matrices and vectors $\|\cdot\|$ as the spectral norm unless otherwise specified by subscripts. We denote the uniform distribution on $[a, b]$ by $U[a, b]$. “ \odot ” represents an element-wise multiplication (i.e., the Hadamard product). The superscript (i) denotes the layer index, and we write $\{x^{(i)}\}_{i=1}^H$ to denote the set of elements $x^{(i)}$ indexed by i from 1 to H .

2 Preliminaries

This section reviews the prior theoretical studies on the strong lottery ticket hypothesis (SLTH) and the formulation of multi-head attention (MHA) mechanisms.

2.1 Strong Lottery Ticket Hypothesis

The strong lottery ticket hypothesis (SLTH) conjectured that a randomly-initialized network inherently contains subnetworks (strong lottery tickets (SLTs)) that achieve high accuracy comparable to trained dense networks, without any weight updates (Ramanujan et al. 2020; Malach et al. 2020). The first theoretical result of the SLTH was given by Malach et al. (2020). They proved the existence of SLTs in a fully-connected ReLU network. Subsequent studies relaxed the requirements for source networks to contain SLTs that approximate some target network (Orseau, Hutter, and Rivasplata 2020; Pensia et al. 2020; Burkholz 2022b). In particular, Pensia et al. (2020) introduced a subset-sum approximation technique (Lueker 1998) into the SLTH context and concluded that the logarithmic overparameterization of the source network to a given target is approximately optimal:

Lemma 2. *Given $\mathbf{x} \in \mathbb{R}^{d_1}$, $\mathbf{W} \in \mathbb{R}^{d_2 \times d_1}$, $\tilde{\mathbf{W}}_1 \in \mathbb{R}^{n \times d_1}$, and $\tilde{\mathbf{W}}_2 \in \mathbb{R}^{d_2 \times n}$, we define the target and pruned source fully-connected networks as*

$$F_T(\mathbf{x}) := \mathbf{W}\mathbf{x},$$

$$F_S(\mathbf{x}) := (\tilde{\mathbf{W}}_2 \odot \mathbf{M}_2) \text{ReLU}((\tilde{\mathbf{W}}_1 \odot \mathbf{M}_1)\mathbf{x}),$$

where $\mathbf{M}_1 \in \{0, 1\}^{n \times d_1}$ and $\mathbf{M}_2 \in \{0, 1\}^{d_2 \times n}$ are binary pruning masks. Assume that $\|\mathbf{W}\| \leq 1$, $\|\mathbf{x}\| \leq 1$, and each entry of $\tilde{\mathbf{W}}_1$ and $\tilde{\mathbf{W}}_2$ is drawn i.i.d. from $U[-1, 1]$. Also, for $0 < \epsilon < 1$, suppose that the hidden dimension n satisfies $n \geq d_1 C \log(2d_1 d_2 / \epsilon)$, where $C > 0$ is some universal constant. Then, with probability at least $1 - \epsilon$, there exists a choice of binary pruning masks \mathbf{M}_1 and \mathbf{M}_2 such that

$$\|F_T(\mathbf{x}) - F_S(\mathbf{x})\| \leq \epsilon.$$

This approach, which approximates a single weight matrix by pruning two randomly initialized matrices (the top

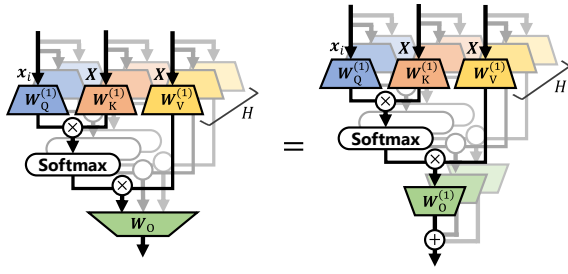


Figure 2: The structure of an MHA. By partitioning the output projection, the final result can be interpreted as the sum of outputs from all heads.

panel of Figure 1), is called the two-layers-for-one approximation, and is now the theoretical foundation of the SLTH for more complex architectures and problems (da Cunha, Natale, and Viennot 2022; Burkholz 2022a; Ferbach et al. 2023; Natale et al. 2024; Otsuka et al. 2025).

2.2 Multi-head Attention Mechanisms

Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T]^\top \in \mathbb{R}^{T \times d_1}$ be a sequence of T input vector embeddings. For each embedding \mathbf{x}_i , we define a binary attention mask $\mathbf{a}_i \in \{0, 1\}^\top$, where $\mathbf{a}_{i,j} = 1$ indicates that the i -th embedding attends to the j -th one. We assume that each embedding attends to at least one other (i.e., $\|\mathbf{a}_i\|_1 \geq 1$). Given such inputs, a multi-head attention (MHA) mechanism (Vaswani et al. 2017) is defined as a function that computes their pair-wise relationships at each of the H attention heads (the left panel of Figure 2). For the i -th embedding \mathbf{x}_i , the MHA is of the following form:

$$\begin{aligned} \text{Attn}(\mathbf{x}_i; \mathbf{X}, \{\mathbf{W}_Q^{(j)}, \mathbf{W}_K^{(j)}, \mathbf{W}_V^{(j)}\}_{j=1}^H, \mathbf{W}_O) \\ &:= [\text{head}_i^{(1)}, \dots, \text{head}_i^{(H)}] \mathbf{W}_O \in \mathbb{R}^{1 \times d_2}, \\ \text{head}_i^{(j)} &:= \sigma \left(\frac{\mathbf{q}_i^{(j)} \mathbf{K}^{(j)\top}}{\sqrt{d_K}}; \mathbf{a}_i \right) \mathbf{V}^{(j)}, \\ \mathbf{q}_i^{(j)} &:= \mathbf{x}_i^\top \mathbf{W}_Q^{(j)}, \quad \mathbf{K}^{(j)} := \mathbf{X} \mathbf{W}_K^{(j)}, \quad \mathbf{V}^{(j)} := \mathbf{X} \mathbf{W}_V^{(j)}, \\ \sigma(\mathbf{x}_i; \mathbf{a}_i)_j &:= \frac{a_{i,j} \exp(x_{i,j})}{\sum_{k=1}^T a_{i,k} \exp(x_{i,k})}. \end{aligned}$$

Here, we define $\mathbf{W}_Q^{(j)}, \mathbf{W}_K^{(j)} \in \mathbb{R}^{d_1 \times d_K}$, $\mathbf{W}_V^{(j)} \in \mathbb{R}^{d_1 \times d_V}$, and $\mathbf{W}_O \in \mathbb{R}^{H d_V \times d_2}$ as single-layer projections for the query $\mathbf{q}_i^{(j)}$, key $\mathbf{K}^{(j)}$, value $\mathbf{V}^{(j)}$, and output of the MHA, respectively. The softmax function with the attention mask is defined as $\sigma(\cdot)$. As shown in the right panel of Figure 2, by partitioning the output weight matrix \mathbf{W}_O into

$$\mathbf{W}_O := [\mathbf{W}_O^{(1)\top}, \dots, \mathbf{W}_O^{(H)\top}]^\top, \quad \mathbf{W}_O^{(j)} \in \mathbb{R}^{d_V \times d_2},$$

the form of $\text{Attn}(\cdot)$ can be represented as follows:

$$\begin{aligned} \text{Attn}(\mathbf{x}_i; \mathbf{X}, \{\mathbf{W}_Q^{(j)}, \mathbf{W}_K^{(j)}, \mathbf{W}_V^{(j)}\}_{j=1}^H, \mathbf{W}_O) \\ &= \text{Attn}(\mathbf{x}_i; \mathbf{X}, \mathbf{W}_{Q:O}^{(1:H)}) \\ &= \sum_{j=1}^H \text{head}_i^{(j)} \mathbf{W}_O^{(j)}. \end{aligned}$$

We denote the set of all weights as

$$\mathbf{W}_{Q:O}^{(1:H)} := \{\mathbf{W}_Q^{(j)}, \mathbf{W}_K^{(j)}, \mathbf{W}_V^{(j)}, \mathbf{W}_O^{(j)}\}_{j=1}^H.$$

3 Strong Lottery Ticket Hypothesis for Transformers

This section analyzes the existence of SLTs within multi-head attention (MHA) mechanisms and extends it to the transformer architecture without normalization layers. For a detailed proof, see Appendix A.

3.1 Setups

We consider two MHAs: a target MHA $\text{Attn}_T(\cdot)$ with arbitrary (tuned) weights, and a pruned source MHA $\text{Attn}_S(\cdot)$ with randomly-initialized weights, denoted as follows:

$$\text{Attn}_T(\mathbf{x}_i) = \text{Attn}(\mathbf{x}_i; \mathbf{X}, \mathbf{W}_{Q:O}^{(1:H)}), \quad (1)$$

$$\text{Attn}_S(\mathbf{x}_i) = \text{Attn}(\mathbf{x}_i; \mathbf{X}, (\tilde{\mathbf{W}} \odot \mathbf{M})_{Q:O}^{(1:H)}). \quad (2)$$

Here, similarly to the weight set $\mathbf{W}_{Q:O}^{(1:H)}$, we define the set of pruned random weights as

$$\begin{aligned} (\tilde{\mathbf{W}} \odot \mathbf{M})_{Q:O}^{(1:H)} &:= \{\tilde{\mathbf{W}}_Q^{(j)} \odot \mathbf{M}_Q^{(j)}, \tilde{\mathbf{W}}_K^{(j)} \odot \mathbf{M}_K^{(j)}, \\ &\quad \tilde{\mathbf{W}}_V^{(j)} \odot \mathbf{M}_V^{(j)}, \tilde{\mathbf{W}}_O^{(j)} \odot \mathbf{M}_O^{(j)}\}_{j=1}^H, \end{aligned}$$

where $\tilde{\mathbf{W}}_Q^{(j)}, \tilde{\mathbf{W}}_K^{(j)} \in \mathbb{R}^{d_1 \times n_K}$, $\tilde{\mathbf{W}}_V^{(j)} \in \mathbb{R}^{d_1 \times n_V}$, and $\tilde{\mathbf{W}}_O^{(j)} \in \mathbb{R}^{n_V \times d_2}$ are the randomly-weighted query, key, value, and output projections of the j -th head in $\text{Attn}_S(\cdot)$, respectively. Also, $\mathbf{M}_Q^{(j)}, \mathbf{M}_K^{(j)}, \mathbf{M}_V^{(j)}$, and $\mathbf{M}_O^{(j)}$ are their corresponding binary pruning masks. Note that the target and source MHAs have different key and value hidden dimensions: d_K and d_V for the target, and n_K and n_V for the source. We assume that $\alpha \geq \max(\sqrt{d_1}, \sqrt{d_2})$ for the inputs, and $\|\mathbf{W}_Q^{(j)}\|, \|\mathbf{W}_K^{(j)}\|, \|\mathbf{W}_V^{(j)}\|, \|\mathbf{W}_O^{(j)}\| \leq 1$ for the j -th head of the target MHA. The source MHA is initialized such that each entry of $\tilde{\mathbf{W}}_Q$ and $\tilde{\mathbf{W}}_K$ is drawn i.i.d. from $U[-n_K^{1/4}, n_K^{1/4}]$, and each entry of $\tilde{\mathbf{W}}_V$ and $\tilde{\mathbf{W}}_O$ is drawn i.i.d. from $U[-1, 1]$.

3.2 The Existence of SLTs Within an MHA

Now, we prove the following SLT existence theorem:

Theorem 3. *Let $\text{Attn}_T(\cdot)$ and $\text{Attn}_S(\cdot)$ be as defined in Equations (1) and (2). Then, with probability at least $1 - \epsilon$, there exists a choice of binary pruning masks $\mathbf{M}_Q^{(j)}, \mathbf{M}_K^{(j)}, \mathbf{M}_V^{(j)}, \mathbf{M}_O^{(j)}$ that satisfy*

$$\max_{i \in [T]} \|\text{Attn}_S(\mathbf{x}_i) - \text{Attn}_T(\mathbf{x}_i)\| \leq \epsilon,$$

if the source hidden dimensions satisfy

$$n_K \geq d_1 C \log \left(\frac{8H\alpha^3 d_1^{3/2}}{\epsilon} \right),$$

$$n_V \geq d_1 C \log \left(\frac{2H\alpha d_1 \sqrt{d_2}}{\epsilon} \right),$$

for some universal constant $C > 0$.

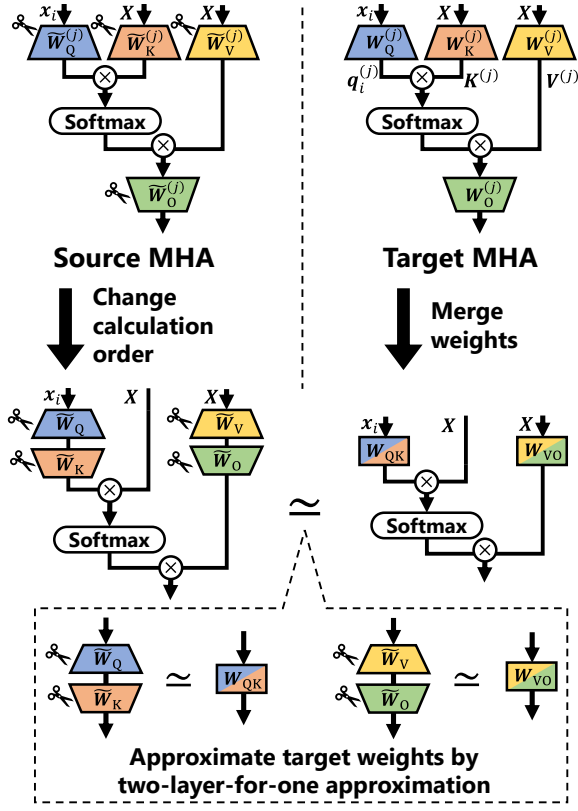


Figure 3: The diagram of our proof. By merging the target projections and changing the calculation order of the source MHA, we can apply a variant of the two-layers-for-one approximation technique and approximate the target MHA while keeping the original source and target structures.

Figure 3 shows an overview of our proof. To prove Theorem 3, we begin by focusing on the part before the softmax, the target and source inner products for the j -th head:

$$\frac{1}{\sqrt{d_K}} \mathbf{q}^{(j)} \mathbf{K}^{(j)\top} = \frac{1}{\sqrt{d_K}} (\mathbf{x}_i^\top \mathbf{W}_Q^{(j)}) (\mathbf{X} \mathbf{W}_K^{(j)})^\top, \quad (3)$$

$$\frac{1}{\sqrt{n_K}} (\mathbf{x}_i^\top (\tilde{\mathbf{W}}_Q^{(j)} \odot \mathbf{M}_Q^{(j)})) (\mathbf{X} (\tilde{\mathbf{W}}_K^{(j)} \odot \mathbf{M}_K^{(j)}))^\top. \quad (4)$$

Since the only difference lies in the projection matrices, we consider the problem of pruning the source projections $\tilde{\mathbf{W}}_Q^{(j)}$ and $\tilde{\mathbf{W}}_K^{(j)}$ to approximate the target projections $\mathbf{W}_Q^{(j)}$ and $\mathbf{W}_K^{(j)}$. A naive idea might be to approximate each target projection independently. In this case, a single source random matrix must approximate each target matrix. However, pruning a single random matrix cannot generally approximate arbitrary ones; thus, this approach is infeasible. To overcome this limitation, we revisit the structure of the target inner product. By closely examining the formulation of the target inner product (Equation (3)), we observe that the query and (transposed) key projections appear adjacently and can be merged into a single joint projection (the right

panel of Figure 3):

$$\frac{1}{\sqrt{d_K}} (\mathbf{x}_i^\top \mathbf{W}_Q^{(j)}) (\mathbf{X} \mathbf{W}_K^{(j)})^\top = \mathbf{x}_i^\top \mathbf{W}_{QK}^{(j)} \mathbf{X}^\top, \quad (5)$$

$$\mathbf{W}_{QK}^{(j)} := \frac{1}{\sqrt{d_K}} \mathbf{W}_Q^{(j)} (\mathbf{W}_K^{(j)})^\top.$$

This reformulation enables us to reinterpret the original problem—not as approximating two target matrices—but as approximating a single merged projection matrix.

We now approximate this merged matrix $\mathbf{W}_{QK}^{(j)}$ by pruning the two source projections. On the source side (Equation (4)) as well, the query and key projections are adjacent. Thus, the source inner product can be viewed as a computation that first calculates the query and key projections (the left panel of Figure 3):

$$\frac{1}{\sqrt{n_K}} (\mathbf{x}_i^\top (\tilde{\mathbf{W}}_Q^{(j)} \odot \mathbf{M}_Q^{(j)})) (\mathbf{X} (\tilde{\mathbf{W}}_K^{(j)} \odot \mathbf{M}_K^{(j)}))^\top$$

$$= \mathbf{x}_i^\top \left((\tilde{\mathbf{W}}_Q^{(j)} \odot \mathbf{M}_Q^{(j)}) (\tilde{\mathbf{W}}_K^{(j)} \odot \mathbf{M}_K^{(j)})^\top \right) \mathbf{X}^\top,$$

$$\tilde{\mathbf{W}}_Q^{\prime(j)} := \frac{1}{n_K^{1/4}} \tilde{\mathbf{W}}_Q^{(j)}, \quad \tilde{\mathbf{W}}_K^{\prime(j)} := \frac{1}{n_K^{1/4}} \tilde{\mathbf{W}}_K^{(j)},$$

where each entry of $\tilde{\mathbf{W}}_Q^{\prime(j)}$ and $\tilde{\mathbf{W}}_K^{\prime(j)}$ is drawn i.i.d. from $U[-1, 1]$ as per our assumption. Therefore, the task reduces to selecting masks $\mathbf{M}_Q^{(j)}$ and $\mathbf{M}_K^{(j)}$ such that the source matrix product $(\tilde{\mathbf{W}}_Q^{\prime(j)} \odot \mathbf{M}_Q^{(j)}) (\tilde{\mathbf{W}}_K^{\prime(j)} \odot \mathbf{M}_K^{(j)})^\top$ closely approximates the target $\mathbf{W}_{QK}^{(j)}$. This allows us to draw an analogy to the conventional theoretical results of the SLTH, particularly the two-layers-for-one approximation (Lemma 2). We therefore establish and apply a variant of Lemma 2, which guarantees the existence of binary pruning masks that achieve such an approximation (the bottom panel of Figure 3):

Lemma 4. *Let $\mathbf{W} \in \mathbb{R}^{d_2 \times d_1}$ be a target matrix with $\|\mathbf{W}\| \leq 1$, and $\tilde{\mathbf{W}}_1 \in \mathbb{R}^{n \times d_1}$ and $\tilde{\mathbf{W}}_2 \in \mathbb{R}^{d_2 \times n}$ be source matrices whose entries are drawn i.i.d. from $U[-1, 1]$. Suppose that $n \geq d_1 C \log(d_1 d_2 / \epsilon)$ for some universal constant $C > 0$. Then, with probability at least $1 - \epsilon$, there exists a choice of binary pruning masks \mathbf{M}_1 and \mathbf{M}_2 such that*

$$\left\| \mathbf{W} - (\tilde{\mathbf{W}}_2 \odot \mathbf{M}_2) (\tilde{\mathbf{W}}_1 \odot \mathbf{M}_1) \right\|_{\max} \leq \frac{\epsilon}{d_1 d_2}.$$

We now turn to the components after the softmax function: the value and output projections. Similar to the query and key case, the value and output projections appear adjacently and can also be merged into a single composite transformation. Thus, we aim to approximate the target merged matrix $\mathbf{W}_{VO}^{(j)} := \mathbf{W}_V^{(j)} \mathbf{W}_O^{(j)}$. This approximation follows the same principle as before: we leverage the matrix product on the source side $(\tilde{\mathbf{W}}_V^{(j)} \odot \mathbf{M}_V^{(j)}) (\tilde{\mathbf{W}}_O^{(j)} \odot \mathbf{M}_O^{(j)})$ to approximate the merged matrix $\mathbf{W}_{VO}^{(j)}$. Lemma 4 ensures that, with high probability, this approximation is successful via appropriately chosen binary pruning masks $\mathbf{M}_V^{(j)}$ and $\mathbf{M}_O^{(j)}$.

Assuming that all weights in the target MHA are approximated by the above procedure, we next analyze the error of the entire attention mechanism by investigating the behavior of the softmax. As a natural idea, one might consider exploiting the 1-Lipschitz continuity of the softmax (Gao and Pavel 2017), which enables internal errors to propagate linearly to the output. However, since the MHA subsequently multiplies the softmax output and the input matrix \mathbf{X} , applying Lipschitz continuity results in a loose upper bound of the error between MHAs: as $\|\mathbf{X}\|$ can grow with T in the worst case, the bound depends on the input length T .

In contrast to this general approach, we provide a more precise analysis. In our setting, thanks to the accurate weight approximation technique mentioned earlier, the internal error of softmax is guaranteed to be finite and small. Leveraging this property, we analyze the softmax output and \mathbf{X} simultaneously to obtain a T -independent bound as follows:

Lemma 5. *Let $\epsilon \in \mathbb{R}^{d_1}$ be an error vector with $\|\epsilon\|_{\max} \leq \epsilon_{\max}$ for some $0 \leq \epsilon_{\max} \leq 1/2$. Then,*

$$\max_{i \in [T]} \|\sigma(\mathbf{x}_i; \mathbf{a}_i)\mathbf{X} - \sigma(\mathbf{x}_i + \epsilon; \mathbf{a}_i)\mathbf{X}\| \leq 4\sqrt{d_1}\alpha\epsilon_{\max}.$$

Since this lemma provides a bound independent of \mathbf{a}_i , our theory holds for models with arbitrary attention masks, including encoder (Devlin et al. 2019) and decoder models (Radford et al. 2019). By applying these above analyses to each attention head, we complete the proof of Theorem 3. For the full proof, see Appendix A.3. We also empirically validate two main theoretical findings in Section 4.2: the accurate approximation of the target MHA becomes feasible with larger source hidden dimensions, and the approximation error remains independent of the input length T .

Proof Sketch of Theorem 3: First, for each attention head, we reformulate the problem by merging the four original target projection matrices into two merged matrices: one combining the query and key projections, and the other the value and output projections. Applying Lemma 4 to these merged matrices enables us to prune each source head to produce an inner product that closely approximates the target one. Next, using Lemma 5, we bound how errors in approximating the query and key matrices propagate through the softmax operation. Lemma 5 ensures that the approximation error of the softmax depends on the approximation accuracy of the query-key projections and does not scale with the input length T . Thus, provided the source hidden dimensions n_K and n_V are sufficiently large, there exists a choice of binary masks for the source MHA which approximate the target MHA within an error ϵ . Also, by suitably setting lower bounds on n_K and n_V , a union bound guarantees that the approximation succeeds across all heads with probability at least $1 - \epsilon$.

3.3 The Existence of SLTs Within a Transformer

By leveraging our main theorem, we now extend the SLTH to transformers. We consider a transformer without the normalization layers for the original definition (Vaswani et al. 2017). The target transformer of B blocks are of the follow-

ing form:

$$\begin{aligned} \text{Tf}_T(\mathbf{x}_i) &:= \text{Blk}_T^{(B)}(\text{Blk}_T^{(B-1)} \dots \text{Blk}_T^{(1)}(\mathbf{x}_i)), \\ \text{Blk}_T^{(b)}(\mathbf{x}_i^{(b)}) &:= \text{F}_T^{(b)}(\text{Attn}_T(\mathbf{x}_i^{(b)})^\top + \mathbf{x}_i^{(b)}) \\ &\quad + \text{Attn}_T(\mathbf{x}_i^{(b)})^\top + \mathbf{x}_i^{(b)}, \end{aligned}$$

where $\text{Blk}_T^{(b)}$ is a b -th target block, and $\mathbf{x}_i^{(b)} \in \mathbb{R}^d$ is the i -th input embedding of the b -th target block. We employ single-layer projection $\text{F}_T^{(b)}(\cdot)$ for the fully-connected network of each target block. Similarly, we define the pruned source transformer as follows:

$$\begin{aligned} \text{Tf}_S(\mathbf{x}_i) &:= \text{Blk}_S^{(B)}(\text{Blk}_S^{(B-1)} \dots \text{Blk}_S^{(1)}(\mathbf{x}_i)), \\ \text{Blk}_S^{(b)}(\mathbf{x}_i^{(b)}) &:= \text{F}_S^{(b)}(\text{Attn}_S(\mathbf{x}_i^{(b)})^\top + \mathbf{x}_i^{(b)}) \\ &\quad + \text{Attn}_S(\mathbf{x}_i^{(b)})^\top + \mathbf{x}_i^{(b)}, \end{aligned}$$

where $\text{Blk}_S^{(b)}(\cdot)$ is a b -th source block and $\mathbf{x}_i^{(b)}$ is the i -th input embedding of the b -th source block. We set the hidden dimension of $\text{F}_S^{(b)}(\cdot)$ as $n_{\text{FC}}^{(b)}$ and assume the hidden dimensions of the MHA in the b -th source block as a same value $n_{\text{MHA}}^{(b)}$ for simplicity. Then, we prove the following theorem:

Theorem 6. *Assume $B \geq 2$. Then, with probability at least $1 - \epsilon$ for $0 < \epsilon < 1$, there exists a choice of binary pruning masks that satisfies*

$$\|\text{Tf}_S(\mathbf{x}_i) - \text{Tf}_T(\mathbf{x}_i)\| \leq \epsilon,$$

if the hidden dimensions of b -th source MHA and fully-connected network satisfy

$$\begin{aligned} n_{\text{MHA}}^{(b)} &\geq d_1 C \log \left(\frac{c_1^{f_1(b,B)} H^{f_2(b,B)} d_1^{f_3(b,B)}}{\epsilon} \right), \\ n_{\text{FC}}^{(b)} &\geq d_1 C \log \left(\frac{c_2^{g_1(b,B)} H^{g_2(b,B)} d_1^{g_3(b,B)}}{\epsilon} \right), \end{aligned}$$

for universal constants $C > 0$ and $c_1, c_2 > 0$ including α . Here, $f_1, f_2, f_3, g_1, g_2, g_3$ are quadratic forms of b and B .

Proof Sketch of Theorem 6: From the existing work (Lemma 2) and Theorem 3, we already know that an MHA and FFN contain SLTs with high probability if each module has a large hidden dimension; thus, by determining the lower bound of the hidden dimension of each module based on the error propagation from the input to output, we can prove that there exists an SLT, which approximates the output of an target transformer to an error of ϵ , within a randomly initialized transformer. By the union bound, the probability that all approximations hold simultaneously is at least $1 - \epsilon$.

For simplicity, this theorem uses target and source fully-connected networks as a single-layer and two-layer ReLU networks F_T and F_S in Lemma 2. It can be generalized to an L -layer target fully-connected network by applying the multi-layer approximation by Pensia et al. (2020). We show that theorem and its proof in Appendix A.5.

4 Experimental Results

This section empirically validates our SLTH theorems.

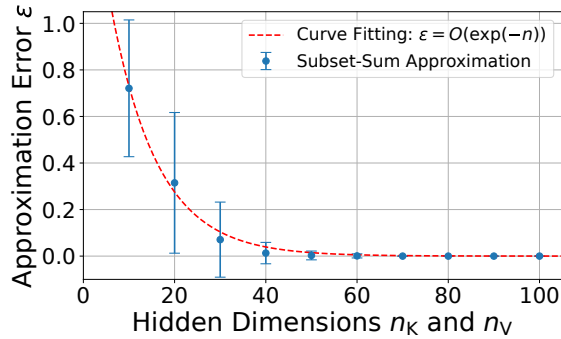


Figure 4: The approximation error ϵ of SLTs within a source MHA for the hidden dimensions $n_K = n_V$. This result shows that the error ϵ satisfies $\epsilon = O(\exp(-n))$, consistently with Theorem 3.

4.1 Experimental Settings

To empirically validate the approximation guarantees established by our SLTH theorems, we evaluate the approximation error on a synthetic dataset for angular velocity estimation. The input consists of a sequence of two-dimensional vectors arranged on the unit circle with a fixed angular velocity. A regression token is used to estimate this velocity, and the source model uses the same regression token as the target model to ensure input consistency. The source and target models are both implemented as either single-head attention mechanisms or single-head transformers as defined in Section 3.3. Both models are initialized according to our theoretical setup: the entries of the query and key projection weights are drawn i.i.d. from $U[-n_K^{1/4}, n_K^{1/4}]$, and those of the value and output projection weights from $U[-1, 1]$. To identify SLTs that approximate the target network, we implement the weight approximation technique described in Lemma 4, which is based on the subset-sum approximation of Pensa et al. (2020). The target MHA is approximated using 100 randomly initialized source MHAs, and we report the mean and standard deviation of the approximation error.

We also investigate whether our theoretical insights generalize to practical settings. In this setting, we search for SLTs by the `edge-popup` algorithm (Ramanujan et al. 2020), which finds accurate subnetworks by backpropagation, instead of learning weights. We train models from the GPT-2 family (mini¹, small, and medium) (Radford et al. 2019) on the WikiText-103 dataset (Merity et al. 2017). The weights of these models are initialized based on the GPT-2 initialization scheme. For each model, we repeat training three times with different random seeds and report the mean and standard deviation of the final performance. See Appendix B for further details on experimental settings.

4.2 Verification of Main Theorems

We empirically verify our theoretical results by pruning a source network to approximate the target network.

¹A 4-layer GPT-2 (Wolf et al. 2020). For details, see the extended version.

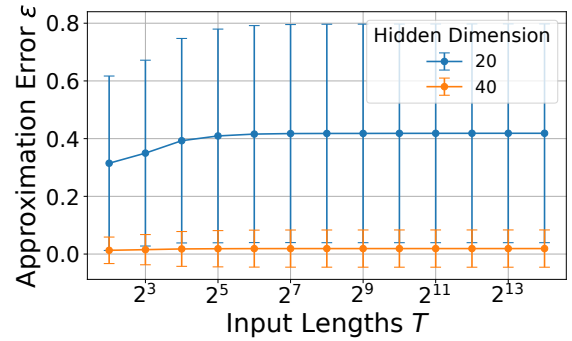


Figure 5: The approximation error ϵ of SLTs within an MHA for the sequence length T . This result suggests that the error ϵ does not diverge as T increases, as implied by Theorem 3.

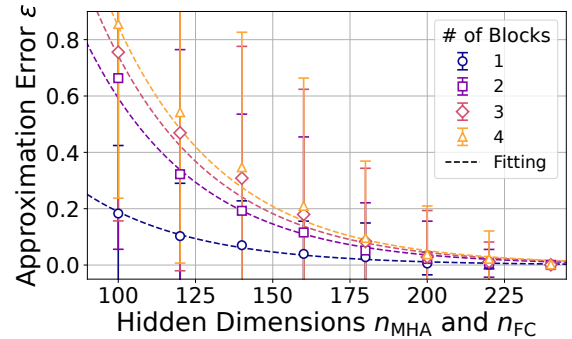


Figure 6: The approximation error ϵ of SLTs within a randomly initialized transformer for the source hidden dimensions $n_{\text{MHA}} = n_{\text{FC}}$. This result suggests that error accumulates as the number of blocks increases, while each error holds $\epsilon = O(\exp(-n_{\text{MHA}}))$, consistently with Theorem 3.

Varying the Hidden Dimensions: We validate Theorem 3 by showing that increasing the hidden dimensions leads to an exponential decrease in approximation error. When we fit the empirical results to the function $\epsilon = \gamma \exp(-\delta n_K)$, we obtain $\epsilon = 0.8 \exp(-0.06 n_K)$, which closely matches the observations. This finding supports our theoretical claim: given a target MHA, each source hidden dimension requires $O(\log(1/\epsilon))$ for the existence of SLTs.

Varying the Sequence Length: Theorem 3 also implies that the existence of SLTs in MHAs is independent of the input length T . In other words, with sufficiently large hidden dimensions, the approximation error has an upper bound that does not depend on T . Figure 5 empirically supports this argument: even as T increases, the error remains bounded, and the bound decreases with larger hidden dimensions.

Varying the Number of Blocks: To validate Theorem 6, we analyze how the approximation error behaves across different numbers of transformer blocks. We set $n_{\text{MHA}} = n_{\text{FFN}}$ and use an untrained target model to be close to our theoretical assumptions. As in the MHA experiment, we fit an exponential decay $\epsilon = \gamma \exp(-\delta n_K)$ to the error of each block, using the same decay rate δ obtained from the first

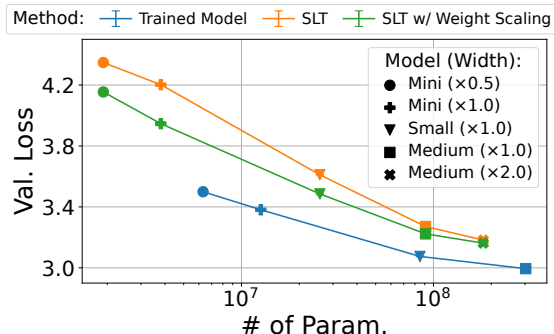


Figure 7: Loss comparison between SLTs with and without query and key weight scaling. By introducing a scale based on our theoretical assumptions, we can obtain better SLTs.

block, as predicted by Theorem 6. Figure 6 shows that, consistent with our theoretical implication, the approximation error decreases rapidly with increasing hidden dimensions for all numbers of blocks. Despite fitting only the coefficient γ per block, the shared δ provides curves that closely match the empirical results, supporting our theoretical claim that only the scale factor varies across blocks.

4.3 Behavior of SLTs in Practical Settings

In the theoretical analysis, we employ a non-conventional initialization strategy: the query and key projection weights are initialized from $U[-n_K^{1/4}, n_K^{1/4}]$, scaled by a factor of $n_K^{1/4}$ compared to the value and output weights, which are initialized from $U[-1, 1]$. This weight scaling was introduced to facilitate the application of the weight approximation lemma in our analysis, and played an important role in establishing our theory. Its theoretical contribution motivates the following question: *does this scaled initialization strategy also benefit SLTs in realistic scenarios?* We empirically evaluate SLTs using the GPT-2 architectures and the WikiText-103 dataset. Figure 7 compares the validation loss of SLTs with and without scaling the query and key weights by $n_K^{1/4} \simeq 2.8$, with respect to the number of nonzero parameters. We observe that SLTs with the weight scaling tend to exhibit lower loss, approaching the performance of trained models. Interestingly, this specific scaling factor $n_K^{1/4}$ is nearly optimal for finding better SLTs: as shown in Figure 8, increasing the scale from 1 gradually decreases the loss up to a certain point, but further increasing it beyond $n_K^{1/4}$ results in increased loss. In all models, the lowest loss is consistently achieved around this scaling factor $n_K^{1/4}$. These findings suggest that our initialization strategy actually helps to ensure the existence of better SLTs within the practical transformer models.

5 Related Work

Strong Lottery Tickets: Zhou et al. (2019) and Ramanujan et al. (2020) empirically found the subnetworks that achieve high accuracy without any weight training. The exist-

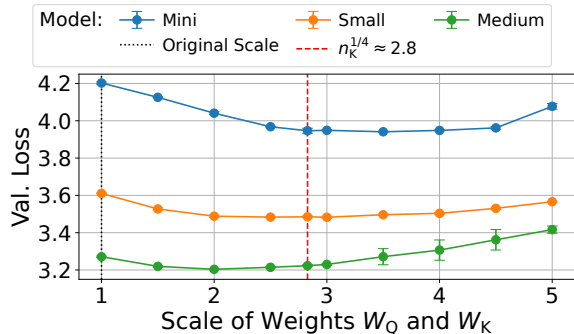


Figure 8: Loss comparison with respect to the weight scaling factor applied to query and key weights. Interestingly, in all models, the loss reaches its minimum near the weight scaling of our theoretical assumptions.

tence of such high-performing subnetworks has been called the strong lottery ticket hypothesis (SLTH), and its theoretical proof was firstly provided in fully-connected ReLU networks (Malach et al. 2020; Orseau, Hutter, and Rivasplata 2020; Pensia et al. 2020; Burkholz 2022b).

Based on these pioneering studies, the SLTH has been extended in three main directions. The first direction involves introducing additional flexibility for relaxing the overparameterization of the source network (Chijiwa et al. 2021; Xiong, Liao, and Kyrillidis 2023). The second direction, in contrast, imposes additional constraints on the source network (Gadhikar, Mukherjee, and Burkholz 2023; Otsuka et al. 2025; Natale et al. 2024). The third direction extends the SLTH to various architectures (Diffenderfer and Kailkhura 2021; Burkholz 2022b; Fischer and Burkholz 2021; da Cunha, Natale, and Viennet 2022; Da Cunha and d’Amore 2023; Burkholz 2022a; Ferbach et al. 2023). Our work contributes to this third direction by proving the SLTH for attention mechanisms and transformers.

Randomly Weighted Transformers: Several studies have empirically investigated the capabilities of randomly weighted transformers. Shen et al. (2021a) demonstrated that a transformer with a few randomly weighted layers achieves accuracy comparable to fully trained models on translation and language understanding tasks. Zhong and Andreas (2024) found that randomly weighted transformers can solve toy tasks with high accuracy as the hidden dimension increases. Some studies empirically showed the existence of SLTs within randomly weighted transformers (Shen et al. 2021b; Ito et al. 2025). Our analysis provides theoretical support for these empirical results about the SLT existence. Furthermore, it provides a theoretical explanation for the improved performance of randomly weighted transformers as the hidden dimension increases, particularly when the pruning is used for optimization.

6 Conclusion

This work investigated the existence of SLTs within a multi-head attention (MHA) mechanism. We extended the exist-

ing theory of the SLTH to MHAs and proved that, if the source MHA has logarithmically large hidden dimensions, it contains an SLT that approximates an arbitrary MHA with high probability. Our proof revealed that, for the SLTH in MHAs, additional layers are not required for approximation, in contrast to the existing theories that rely on approximating a single-layer structure by a two-layer one. Furthermore, by exploiting our findings, we established the theory of the SLTH for transformers without normalization layers. We empirically validated our theory and confirmed that the results are consistent with the theoretical implications. Interestingly, our theoretical implication, which provides an appropriate weight scale for initializing query and key projection weights, contributed to improving the performance of SLTs in practical settings. Our results not only extend SLTH to transformers, but also indicate a new research direction in the SLTH for practical transformer models. We hope these findings will lead to a fundamental understanding of overparameterized models.

Acknowledgments

This work was supported in part by JSPS KAKENHI Grant Number JP23H05489, JP25K03092, JP25KJ1236, and JST-ALCA-Next Japan Grant # JPMJAN24F3.

References

- Burkholz, R. 2022a. Convolutional and residual networks provably contain lottery tickets. In *International Conference on Machine Learning*, 2414–2433. PMLR.
- Burkholz, R. 2022b. Most activation functions can win the lottery without excessive depth. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Chijiwa, D.; Yamaguchi, S.; Ida, Y.; Umakoshi, K.; and Inoue, T. 2021. Pruning randomly initialized neural networks with iterative randomization. *Advances in neural information processing systems*, 34: 4503–4513.
- Da Cunha, A.; and d’Amore, F. 2023. Polynomially overparameterized convolutional neural networks contain structured strong winning lottery tickets. *Advances in Neural Information Processing Systems*, 36: 25929–25957.
- da Cunha, A.; Natale, E.; and Viennot, L. 2022. Proving the strong lottery ticket hypothesis for convolutional neural networks. In *ICLR 2022-10th International Conference on Learning Representations*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.
- Diffenderfer, J.; and Kailkhura, B. 2021. Multi-prize lottery ticket hypothesis: Finding accurate binary neural networks by pruning a randomly weighted network. In *International Conference on Learning Representations*.
- Ferbach, D.; Tsigotis, C.; Gidel, G.; and Bose, J. 2023. A general framework for proving the equivariant strong lottery ticket hypothesis. In *The Eleventh International Conference on Learning Representations*.
- Fischer, J.; and Burkholz, R. 2021. Towards strong pruning for lottery tickets with non-zero biases. *arXiv preprint arXiv:2110.11150*.
- Frankle, J.; and Carbin, M. 2019. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*.
- Gadhikar, A. H.; Mukherjee, S.; and Burkholz, R. 2023. Why random pruning is all we need to start sparse. In *International Conference on Machine Learning*, 10542–10570. PMLR.
- Gao, B.; and Pavel, L. 2017. On the properties of the softmax function with application in game theory and reinforcement learning. *arXiv preprint arXiv:1704.00805*.
- Ito, H.; Yan, J.; Otsuka, H.; Kawamura, K.; Motomura, M.; Chu, T. V.; and Fujiki, D. 2025. Uncovering strong lottery tickets in graph transformers: A path to memory efficient and robust graph learning. *Transactions on Machine Learning Research*.
- Lueker, G. S. 1998. Exponentially small bounds on the expected optimum of the partition and subset sum problems. *Random Structures & Algorithms*, 12(1): 51–62.
- Malach, E.; Yehudai, G.; Shalev-Schwartz, S.; and Shamir, O. 2020. Proving the lottery ticket hypothesis: Pruning is all you need. In *International Conference on Machine Learning*, 6682–6691. PMLR.
- Merity, S.; Xiong, C.; Bradbury, J.; and Socher, R. 2017. Pointer sentinel mixture models. In *International Conference on Learning Representations*.
- Natale, E.; Ferre’, D.; Giambartolomei, G.; Giroire, F.; and Mallmann-Trenn, F. 2024. On the sparsity of the strong lottery ticket hypothesis. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Orseau, L.; Hutter, M.; and Rivasplata, O. 2020. Logarithmic pruning is all you need. *Advances in Neural Information Processing Systems*, 33: 2925–2934.
- Otsuka, H.; Chijiwa, D.; García-Arias, Á. L.; Okoshi, Y.; Kawamura, K.; Chu, T. V.; Fujiki, D.; Takeuchi, S.; and Motomura, M. 2025. Partially frozen random networks contain compact strong lottery tickets. *Transactions on Machine Learning Research*.
- Pensia, A.; Rajput, S.; Nagle, A.; Vishwakarma, H.; and Pappaliopoulos, D. 2020. Optimal lottery tickets via subset sum: Logarithmic over-parameterization is sufficient. *Advances in neural information processing systems*, 33: 2599–2610.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Ramanujan, V.; Wortsman, M.; Kembhavi, A.; Farhadi, A.; and Rastegari, M. 2020. What’s hidden in a randomly weighted neural network? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11893–11902.

Shen, S.; Baevski, A.; Morcos, A.; Keutzer, K.; Auli, M.; and Kiela, D. 2021a. Reservoir transformers. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 4294–4309. Online: Association for Computational Linguistics.

Shen, S.; Yao, Z.; Kiela, D.; Keutzer, K.; and Mahoney, M. 2021b. What’s hidden in a one-layer randomly weighted transformer? In Moens, M.-F.; Huang, X.; Specia, L.; and Yih, S. W.-t., eds., *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2914–2921. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; Davison, J.; Shleifer, S.; von Platen, P.; Ma, C.; Jernite, Y.; Plu, J.; Xu, C.; Le Scao, T.; Gugger, S.; Drame, M.; Lhoest, Q.; and Rush, A. 2020. HuggingFace’s transformers: State-of-the-art natural language processing. In Liu, Q.; and Schlangen, D., eds., *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45. Online: Association for Computational Linguistics.

Xiong, Z.; Liao, F.; and Kyrillidis, A. 2023. Strong lottery ticket hypothesis with ε -perturbation. In *International Conference on Artificial Intelligence and Statistics*, 6879–6902. PMLR.

Zhong, Z.; and Andreas, J. 2024. Algorithmic capabilities of random transformers. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

Zhou, H.; Lan, J.; Liu, R.; and Yosinski, J. 2019. Deconstructing lottery tickets: Zeros, signs, and the supermask. *Advances in neural information processing systems*, 32.