

BEYOND LINEAR APPROXIMATIONS: A NOVEL PRUNING APPROACH FOR ATTENTION MATRIX

Yingyu Liang* Jiangxuan Long† Zhenmei Shi‡ Zhao Song§ Yufa Zhou¶

ABSTRACT

Large Language Models (LLMs) have shown immense potential in enhancing various aspects of our daily lives, from conversational AI to search and AI assistants. However, their growing capabilities come at the cost of extremely large model sizes, making deployment on edge devices challenging due to memory and computational constraints. This paper introduces a novel approach to LLM weight pruning that directly optimizes for approximating the attention matrix, a core component of transformer architectures. Unlike existing methods that focus on linear approximations, our approach accounts for the non-linear nature of the Softmax attention mechanism. We provide theoretical guarantees for the convergence of our Gradient Descent-based optimization method to a near-optimal pruning mask solution. Our empirical results demonstrate the effectiveness of our non-linear pruning approach in maintaining model performance while significantly reducing computational costs, which is beyond the current state-of-the-art methods, i.e., SparseGPT and Wanda, by a large margin. This work establishes a new theoretical foundation for pruning algorithm design in LLMs, potentially paving the way for more efficient LLM inference on resource-constrained devices.

1 INTRODUCTION

Large Language Models (LLMs) based on the transformer architecture (Vaswani et al., 2017), including GPT-4o (OpenAI, 2024a), Claude (Anthropic, 2024), and OpenAI’s recent o1 (OpenAI, 2024b), have shown immense potential to enhance our daily lives. They revolutionize fields like conversational AI (Liu et al., 2024), AI agents (Xi et al., 2023; Chen et al., 2024b), search AI (OpenAI, 2024b), and AI assistants (Mahmood et al., 2023; Zhang et al., 2023a; Kuo et al., 2024; Feng et al., 2024). With their growing capabilities, LLMs are powerful tools shaping the future of technology. However, the current state-of-the-art LLM weights number is extremely large. For instance, the smallest version of Llama 3.1 (Llama Team, 2024) needs 8 billion parameters, which takes more than 16GB GPU memory with half float precision and requires significant inference time. Deploying such models on edge devices such as smartphones becomes challenging due to their large memory and high computational cost. To reduce the LLM model size, many studies work on pruning the LLMs model weights to relax the device memory constraint and minimize response latency. The classical pruning problem in LLMs can be formulated as follows. Given a weight matrix $W \in \mathbb{R}^{d \times d}$ and some calibration data $X \in \mathbb{R}^{n \times d}$, where n is input token length, and d is feature dimension, the goal is to find a matrix \widetilde{W} under some sparse constraint such that $\|XW - X\widetilde{W}\|$ being small under some norm function. The above formulation has been widely used in many state-of-the-art pruning methods, such as SparseGPT (Frantar & Alistarh, 2023) and Wanda (Sun et al., 2024).

However, the current object functions only focus on the approximation of a linear function XW . Their optimal solutions do not have a good approximation to the attention matrix (see Figure 2 for details). Note that the attention mechanism is the kernel module of the transformer architecture. The high-level insight of their bad performance is that the Softmax function is very sensitive to

* yingyul@hku.hk. The University of Hong Kong. yliang@cs.wisc.edu. University of Wisconsin-Madison.

† lungchianghsuan@gmail.com. South China University of Technology.

‡ zhmeishi@cs.wisc.edu. University of Wisconsin-Madison.

§ magic.linuxkde@gmail.com. The Simons Institute for the Theory of Computing at UC, Berkeley.

¶ yufazhou@seas.upenn.edu. University of Pennsylvania.

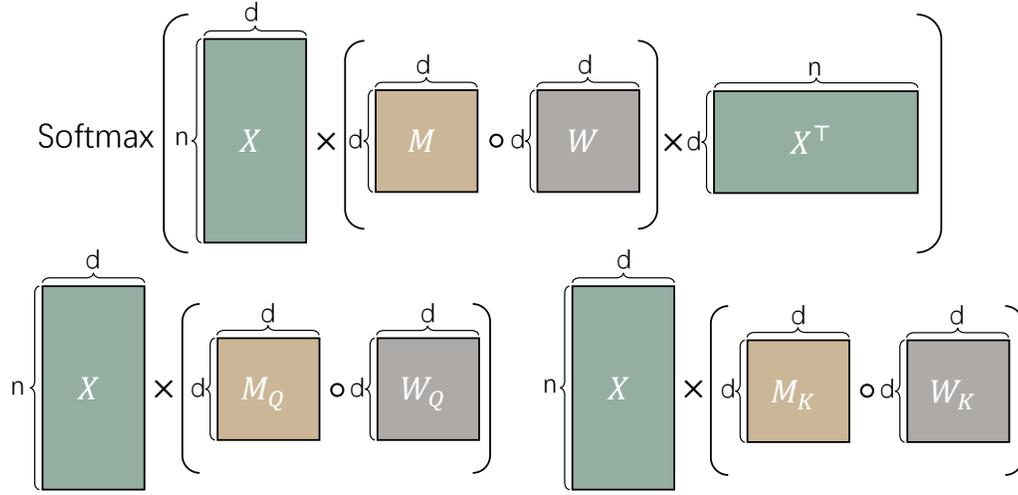


Figure 1: Comparison of our Attention Weights Pruning method and Linear Pruning method such as Wanda and SparseGPT. The top figure illustrates our proposed method of the attention matrix approximation, where pruning is applied directly to the fused attention weight matrix W , using only one pruning mask M . The bottom figure describes the Linear Pruning method of the linear function approximation, where pruning is applied separately to the query weight matrix W_Q and key weight matrix W_K , using two different pruning masks M_Q and M_K , respectively.

the large positive values of the input due to its \exp scaling effect, while pruning mask based on linear approximation cannot capture this sensitivity. Thus, in this work, we directly compute the pruning mask on weights to approximate the attention matrix, which is a highly non-linear function, $\text{Softmax}(XWX^\top) \in \mathbb{R}^{n \times n}$. To the best of our knowledge, this paper is the first work studying attention weight pruning to directly approximate the attention matrix. We provide a theoretical guarantee that optimization based on Gradient Descent (GD) on our loss function can converge to a good pruning mask solution (Theorem 1.3). Furthermore, we preliminarily verified the effectiveness of our method with empirical support (Section 6). Our theoretical foundation may pave the way for more efficient LLM inference on resource-constrained devices.

Key background. We introduce some key backgrounds. We define the attention matrix in the self-attention mechanism as below:

Definition 1.1 (Attention Matrix). *Let $X \in \mathbb{R}^{n \times d}$ be the input. Given query and key weights matrix $W_Q, W_K \in \mathbb{R}^{d \times d}$, we define $W := W_Q W_K^\top$. Then, we have the Softmax attention matrix being*

$$\text{Softmax}(XWX^\top) = D^{-1} \exp(XWX^\top),$$

where (1) $D := \text{diag}(\exp(XWX^\top) \cdot \mathbf{1}_n)$, (2) \exp denotes the exponential function and is applied entry-wisely, (3) $\text{diag}()$ operation takes a vector and outputs a diagonal matrix with the entries of that vector, and (4) $\mathbf{1}_n$ denotes the length- n all ones vector.

Further, we introduce the problem setup of our Attention Weights Pruning. By selectively reducing the number of non-zero elements in the attention weight matrix W in Definition 1.1, we can preserve model performance while lowering computational cost and GPU memory usage. Below, we formally define the Attention Weights Pruning problem and the corresponding loss function:

Definition 1.2 (Attention Weights Pruning). *Let $M \in [0, 1]^{d \times d}$ be the pruning mask. Let X, W be defined in Definition 1.1. Let $A := \exp(XWX^\top)$ and $\tilde{A} := \exp(X(M \circ W)X^\top)$, where \circ is the Hadamard product. Let $D := \text{diag}(A \cdot \mathbf{1}_n)$ and $\tilde{D} := \text{diag}(\tilde{A} \cdot \mathbf{1}_n)$. Let $\lambda \in \mathbb{R}_+$ be the regularization parameter. We define the Attention Weights Pruning loss function to be*

$$\mathcal{L}(M) := \frac{1}{2} \|D^{-1}A - \tilde{D}^{-1}\tilde{A}\|_F^2 + \frac{1}{2} \lambda \|M\|_F^2.$$

Thus, the Attention Weights Pruning optimization problem is $\min_M \mathcal{L}(M)$.

Our contributions. This is the first work studying the Attention Weights Pruning problem, which is an approximation problem to a non-linear function. We provide an algorithm for obtaining the near-optimal pruning mask based on Gradient Descent (GD) with a convergence guarantee.

Theorem 1.3 (Main result, informal version of Theorem 4.1). *For any $\epsilon > 0$, our Algorithm 1 can converge to the near-optimal pruning mask for the Attention Weights Pruning problem (Definition 1.2) in $O(d \text{ poly}(n)/\epsilon)$ time with $O(\xi + \epsilon)$ error, where ξ is a small term depending on intrinsic property of the data and weights.*

In the above theorem, ξ can be arbitrarily small as $\xi \rightarrow 0$ when the regularization coefficient $\lambda \rightarrow 0$. Thus, our analysis shows that although the objective function is highly non-linear, the GD training can converge to a near-optimal pruning mask solution.

Our experiments on synthetic clearly align and support our theoretical analysis (Section 6.1). Furthermore, we evaluate our non-linear pruning method and show it beyond SparseGPT and Wanda by a large margin on real data (C4 Dataset Raffel et al. (2020)) under real LLM (Llama 3.2-1B Meta (2024)) in Section 6.2 and Section 6.3.

Our contributions are as follows:

- This is the first work that analyzes the weights pruning problem based on Softmax attention, which is a non-linear function.
- We provide the closed form of the gradient of Attention Weights Pruning loss function (Theorem 5.3), and Lipschitz of that gradient (Theorem 5.4),
- We provide Gradient Descent based Algorithm 1 to obtain the near-optimal pruning mask and its convergence guarantee (Theorem 4.1).
- We conduct experiments to verify the effectiveness of our method (Section 6), showing our non-linear pruning method is beyond SparseGPT and Wanda by a large margin.

Roadmap. In Section 2, we review the related work. Section 3 introduces key concepts and definitions essential for the subsequent sections. In Section 4, we present our main result. Section 5 offers a technical overview of the methods employed. Experimental results are discussed in Section 6. Finally, Section 7 summarizes our findings and offers concluding remarks.

2 RELATED WORK

2.1 PRUNING AND COMPRESSION FOR LLMs

Model compression plays a critical role in improving the efficiency and deployment of large language models (LLMs) (Zhu et al., 2023) for its effectiveness in reducing computational overhead while preserving performance. Common compression techniques include quantization (Park et al., 2024; Xiao et al., 2023; Hooper et al., 2024), pruning (Chen et al., 2021; Hoefler et al., 2021; Hubara et al., 2021; Jin et al., 2022; Frantar & Alistarh, 2022; 2023; Sun et al., 2024; Li et al., 2024a; Zandieh et al., 2024; Zhang et al., 2024c; Xia et al., 2023; Ashkboos et al., 2024; Chen et al., 2025b), and knowledge distillation (Hsieh et al., 2023; Shridhar et al., 2023; Jiang et al., 2023; Wang et al., 2023). Specifically, pruning techniques have been developed extensively, such as unstructured pruning, which removes individual weights (Li et al., 2024a; Sun et al., 2024), and structured pruning, which eliminates entire components like neurons or attention heads (Michel et al., 2019; Ashkboos et al., 2024; Xia et al., 2024). Sun et al. (2024) proposed Wanda, a novel unstructured pruning technique that uses weight-activation products to induce up to 50% sparsity in LLMs without retraining, achieving competitive results with significantly lower computational cost. SparseGPT (Frantar & Alistarh, 2023) introduced a one-shot pruning method that achieves up to 60% sparsity in large GPT-family models with minimal impact on performance. A follow-up work (Li et al., 2024a) improved the complexity analysis of SparseGPT, reducing the running time from $O(d^3)$ to $O(d^{2.53})$, enabling faster pruning on LLMs. Together, these techniques contribute to more scalable and resource-efficient LLMs, maintaining competitive performance while substantially reducing computational resources.

2.2 ATTENTION ACCELERATION

The attention mechanism has faced criticism due to its quadratic time complexity with respect to context length (Vaswani et al., 2017). Addressing this criticism, a variety of approaches are employed, including sparse attention (Hubara et al., 2021; Kurtic et al., 2023; Frantar & Alistarh, 2023; Li et al., 2024a), low-rank approximations (Razenshteyn et al., 2016; Li et al., 2016; Hu et al., 2022; Zeng & Lee, 2024; Hu et al., 2024d; Li et al., 2025a), and kernel-based methods (Charikar et al., 2020; Liu & Zenke, 2020; Deng et al., 2025; Zandieh et al., 2023; Liang et al., 2024b), to reduce computational overhead and improve scalability. Aggarwal & Alman (2022) enable the derivation of a low-rank representation of the attention matrix, which accelerates both the training and inference processes of single attention layer, tensor attention, and multi-layer transformer, achieving almost linear time complexity (Alman & Song, 2023; 2024a;b; Liang et al., 2024g;c; Ke et al., 2025a; Hu et al., 2024h). Other approaches like Mamba (Gu & Dao, 2023; Dao & Gu, 2024), Linearizing Transformers (Zhang et al., 2024b; Mercat et al., 2024), Hopfield Models (Hu et al., 2023; Wu et al., 2024b; Hu et al., 2024c; Xu et al., 2024a; Wu et al., 2024a; Hu et al., 2024a;b;f), and PolySketchFormer (Kacham et al., 2023) focus on architectural modifications and implementation optimizations to enhance performance. System-level optimizations such as FlashAttention (Dao et al., 2022; Dao, 2023; Shah et al., 2024) and block-wise parallel decoding (Stern et al., 2018) further improve efficiency. Collectively, these innovations have significantly augmented transformer models’ ability to handle longer input sequences, unlocking broader applications across multiple sectors (Chen et al., 2023; Peng et al., 2023; Ding et al., 2024; Ma et al., 2024; Xu et al., 2024d; An et al., 2024; Jin et al., 2024; Li et al., 2024c; Liang et al., 2024d; Shi et al., 2024a).

3 PRELIMINARY

Notations. For any positive integer n , we use $[n]$ to denote set $\{1, 2, \dots, n\}$. For each $a, b \in \mathbb{R}^n$, we use $a \circ b \in \mathbb{R}^n$ to denote the Hadamard product, i.e., the i -th entry of $(a \circ b)$ is $a_i b_i$ for all $i \in [n]$. For $A \in \mathbb{R}^{m \times n}$, let $A_i \in \mathbb{R}^n$ denote the i -th row and $A_{*,j} \in \mathbb{R}^m$ denote the j -th column of A , where $i \in [m]$ and $j \in [n]$. We use $\exp(A)$ to denote a matrix where $\exp(A)_{i,j} := \exp(A_{i,j})$ for a matrix $A \in \mathbb{R}^{n \times d}$. We use $\|A\|_F$ to denote the Frobenius norm of a matrix $A \in \mathbb{R}^{n \times d}$, i.e., $\|A\|_F := \sqrt{\sum_{i \in [n]} \sum_{j \in [d]} |A_{i,j}|^2}$. For a symmetric matrix $A \in \mathbb{R}^{n \times n}$, $A \succeq 0$ means that A is positive semidefinite (PSD), i.e., for all $x \in \mathbb{R}^n$, we have $x^\top A x \geq 0$.

Attention weights pruning. We aim to determine a near-optimal pruning mask M for the attention weights in a self-attention mechanism. Furthermore, we incorporate causal attention masking¹ into our method to be more aligned with the current decoder-only LLM architecture, while our analysis can be applied to any general attention mask, e.g., block-wise attention mask. To formalize this, we provide the formal definition of causal attention mask and attention weights pruning in this section.

The causal attention mask ensures that each token in the sequence can attend only to itself and the preceding tokens. Here, we provide the formal definition of the causal attention mask:

Definition 3.1 (Causal attention mask, Liang et al. (2024c)). *We define the causal attention mask as $M_c \in \{0, 1\}^{n \times n}$, where $(M_c)_{i,j} = 1$ if $i \geq j$ and $(M_c)_{i,j} = 0$ otherwise.*

Now, we incorporate Attention Weights Pruning (see Definition 1.2) with causal attention mask M_c .

Definition 3.2 (Attention Weights Pruning with Causal Attention Mask). *Let $M_c \in \{0, 1\}^{n \times n}$ be the causal attention mask defined in Definition 3.1. Let $A := \exp(XWX^\top) \circ M_c$ and $\tilde{A} := \exp(X(M \circ W)X^\top) \circ M_c$. Let $D := \text{diag}(A \cdot \mathbf{1}_n)$ and $\tilde{D} := \text{diag}(\tilde{A} \cdot \mathbf{1}_n)$. We define Attention Weights Pruning with Causal Attention Mask loss function to be $\mathcal{L}(M) := \mathcal{L}_{\text{attn}}(M) + \mathcal{L}_{\text{reg}}(M)$ where $\mathcal{L}_{\text{attn}}(M) := \frac{1}{2} \|D^{-1}A - \tilde{D}^{-1}\tilde{A}\|_F^2$ and $\mathcal{L}_{\text{reg}}(M) := \frac{1}{2} \lambda \|M\|_F^2$.*

¹In this paper, we always use *pruning mask* to refer $M \in \mathbb{R}^{d \times d}$, which is our target, and use *causal attention mask* to refer $M_c \in \mathbb{R}^{n \times n}$, which is a fixed mask for standard self-attention.

Algorithm 1 Gradient Descent for Pruning Mask (Theorem 4.1). Let $X_1, X_2, \dots, X_k \in \mathbb{R}^{n \times d}$ be our calibration dataset of size k . We iteratively run our GD method on this dataset.

```

1: procedure PRUNEMASKGD( $X_1, X_2, \dots, X_k \in \mathbb{R}^{n \times d}, W \in \mathbb{R}^{d \times d}, M_c \in \{0, 1\}^{d \times d}, \rho \in [0, 1], \lambda \in [0, 1], \epsilon \in (0, 0.1)$ )
    $\triangleright$  Theorem 4.1
2:   Initialize  $M \in \{1\}^{d \times d}$ 
3:   Initialize  $\eta, T$  by Theorem 4.1
4:   for  $j = 1 \rightarrow k$  do  $\triangleright$  Iterate over the dataset
5:      $u_j \leftarrow \exp(X_j W X_j^\top)$ 
6:      $f_j \leftarrow \text{diag}((u_j \circ M_c) \cdot \mathbf{1}_n)^{-1}(u_j \circ M_c)$ 
7:   end for
8:   for  $i = 1 \rightarrow T$  do  $\triangleright$  Iterate over  $T$  steps
9:     for  $j = 1 \rightarrow k$  do  $\triangleright$  Iterate over the dataset
10:       $\tilde{u}_j \leftarrow \exp(X_j (M \circ W) X_j^\top)$ 
11:       $\tilde{f}_j \leftarrow \text{diag}((\tilde{u}_j \circ M_c) \cdot \mathbf{1}_n)^{-1}(\tilde{u}_j \circ M_c)$ 
12:       $c_j \leftarrow \tilde{f}_j - f_j$ 
13:       $p_{1,j} \leftarrow c_j \circ \tilde{f}_j$ 
14:       $p_{2,j} \leftarrow \text{diag}(p_{1,j} \cdot \mathbf{1}_n) \tilde{f}_j$ 
15:       $p_j \leftarrow p_{1,j} - p_{2,j}$ 
16:    end for
17:     $M \leftarrow M - (\eta/k) \cdot (W \circ (\sum_{j=1}^k X_j^\top p_j X_j) + \lambda M)$   $\triangleright$  Gradient Descent
18:  end for
19:   $m \leftarrow \text{vec}(M)$   $\triangleright$  Flatten  $M$  into a vector
20:   $m_{\text{sorted}} \leftarrow \text{sort}(m)$   $\triangleright$  Sort the elements of  $M$ 
21:   $\tau \leftarrow m_{\text{sorted}}[\lceil \rho \cdot d^2 \rceil]$   $\triangleright$  Get the  $\rho$ -th largest element
22:   $M_{ij} \leftarrow \begin{cases} 1 & \text{if } M_{ij} > \tau \\ 0 & \text{otherwise} \end{cases}$   $\triangleright$  Set the top  $\rho$  entries to 1, others to 0
23:  return  $M$ 
24: end procedure

```

4 MAIN RESULTS

We provide an Algorithm 1 for Attention Weights Pruning problem based on Gradient Descent (GD). We also prove the convergence for our GD algorithm in Theorem 4.1.

Theorem 4.1 (Main result, formal version of Theorem 1.3). *Let $M, X, W, \tilde{D}, \tilde{A}, \lambda, \mathcal{L}, \mathcal{L}_{\text{attn}}$ be defined in Definition 3.2. Assume $XX^\top \succeq \beta I$ and $\min_{i,j \in [n]} (\tilde{D}^{-1} \tilde{A})_{i,j} \geq \delta > 0$. Furthermore, Let $\mu = 2 \min_{i,j \in [d]} \{|W_{i,j}|\} \cdot \beta \cdot \delta$. Let $\xi = 12n^{-1.5} \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X\|_F^2 \cdot \lambda d / \mu$. Then, for any $\epsilon > 0$, provided $\eta < 1/L$ where L is the Lipschitz constant for $\nabla_M \mathcal{L}(M)$ (see Theorem 5.4), GD (Algorithm 1) with fixed step size η and run for $T = 4\mathcal{L}(M^{(0)}) / (\eta\mu\epsilon n^2)$ iterations results in the following guarantee,*

$$\frac{1}{n^2} \min_{t < T} \mathcal{L}_{\text{attn}}(M^{(t)}) + \frac{\lambda^2}{\mu n^2} \|M^{(t)}\|_F^2 \leq (\xi + \epsilon)/2.$$

Proof. Let $g(M) = 2\mathcal{L}_{\text{attn}}(M) + \frac{2\lambda^2}{\mu} \|M\|_F^2$. Note that $\mathcal{L}(M)$ satisfies the $(g(M), n^2\xi, 2, \mu)$ -proxy PL inequality (Lemma 5.5). Also, we have $\mathcal{L}(M)$ is non-negative and has L -Lipschitz gradients Theorem 5.4. Thus, we finish the proof using Theorem 5.2. \square

Remark 4.2. *The two assumptions in Theorem 4.1 are practical. The first assumption of the positive definite matrix is widely used in theoretical deep learning analysis, e.g., Li & Liang (2018); Du et al. (2019); Allen-Zhu et al. (2019b); Arora et al. (2019). The second assumption is natural, as $\tilde{D}^{-1} \tilde{A}$ is the pruned attention matrix, where each entry is*

$$\frac{\exp(X(M \circ W)X^\top)_{i,j}}{\sum_{j=1}^n \exp(X(M \circ W)X^\top)_{i,j}} > 0,$$

which has a natural lower bound.

Our error upper bound in Theorem 4.1 is $O(\xi + \epsilon)$, where ϵ can be arbitrarily small. For ξ , we can let it be small by choosing a proper λ , i.e., the ξ error term can be made arbitrarily small by choosing small λ . However, if we choose a very small λ , the algorithm’s run time gets larger as $T \propto 1/\eta \propto L \propto (\lambda + \text{other terms})$. Thus, although the objective function is highly non-linear, we can show that the Gradient Descent of our Algorithm 1 can converge to a good solution of the Attention Weights Pruning problem.

After solving the optimization problem, we obtain a pruning mask with real-valued entries. In practice, however, this pruning mask must be converted into a binary form, specifically $M \in \{0, 1\}^{d \times d}$. We define the pruning ratio ρ as the percentage of weights to be pruned. We apply this ratio by setting the pruning mask entries to zero for weights that fall below the ρ -th percentile and to one for those above. This ensures that only the specified proportion of weights is pruned.

5 TECHNIQUE OVERVIEW

In Section 5.1, we introduce some useful tools from previous work. In Section 5.2, we derive the close form of the gradient of Attention Weights Pruning. In Section 5.3, we calculate the Lipschitz constant of that gradient. In Section 5.4, we prove the PL inequality for our loss function.

5.1 PREVIOUS TOOLS ON CONVERGENCE OF GD

To analyze the convergence behavior of GD for our optimization problem (Definition 3.2), we first introduce the concept of g -proxy, ξ -optimal Polyak–Łojasiewicz(PL) inequality (Polyak, 1963; Łojasiewicz, 1963; Karimi et al., 2016), under which GD will converge:

Definition 5.1 (g -proxy, ξ -optimal PL inequality, Definition 1.2 in Frei & Gu (2021)). *We say that a function $f : \mathbb{R}^p \rightarrow \mathbb{R}$ satisfies a g -proxy, ξ -optimal Polyak–Łojasiewicz inequality with parameters $\alpha > 0$ and $\mu > 0$ (in short, f satisfies the (g, ξ, α, μ) -PL inequality) if there exists a function $g : \mathbb{R}^p \rightarrow \mathbb{R}$ and scalars $\xi \in \mathbb{R}$, $\mu > 0$ such that for all $w \in \mathbb{R}^p$, $\|\nabla f(w)\|^\alpha \geq \frac{1}{2}\mu(g(w) - \xi)$.*

PL inequality is a powerful tool for studying non-convex optimization, and it has been used in recent studies on provable guarantees for neural networks trained by gradient descent (Li & Liang, 2018; Allen-Zhu et al., 2019a;b;c; Frei et al., 2019; Cao & Gu, 2020; Ji & Telgarsky, 2019; Frei et al., 2021; Shi et al., 2021; 2024b). It provides a proxy convexity property, although the objective function is non-convex. In detail, for a function with good smoothness property, we can find some proxy functions and show the convergence by utilizing these proxy functions.

Leveraging this PL inequality, Frei & Gu (2021) derives the following GD convergence guarantees.

Theorem 5.2 (Theorem 3.1 in Frei & Gu (2021)). *Suppose $f(w)$ satisfies the $(g(\cdot), \xi, \alpha, \mu)$ -proxy PL inequality for some function $g(\cdot) : \mathbb{R}^p \rightarrow \mathbb{R}$. Assume that f is non-negative and has L -Lipschitz gradients. Then for any $\epsilon > 0$, provided $\eta < 1/L$, GD with fixed step size η and run for $T = 2\eta^{-1}(\mu\epsilon/2)^{-2/\alpha} f(w^{(0)})$ iterations results in the following guarantee, $\min_{t < T} g(w^{(t)}) \leq \xi + \epsilon$.*

The above theorem establishes that under the (g, ξ, α, μ) -PL inequality and Lipschitz continuity of the gradient, GD converges to a point where the proxy function $g(w)$ is within ϵ of ξ . To apply this result to our specific problem, we need to verify these conditions for our loss function $\mathcal{L}(M)$.

5.2 CLOSED FORM OF GRADIENT

As a first step, we compute the close form of the gradient $\nabla_M \mathcal{L}(M)$. The pruning mask M is inside a non-linear function Softmax, which complicates our calculation. We defer the proof to Section C.

Theorem 5.3 (Closed form of gradient, informal version of Theorem D.5). *Let $\mathcal{L}(M)$ be defined in Definition 3.2. Let p be defined in Definition D.1. Let $X \in \mathbb{R}^{n \times d}$, $M \in [0, 1]^{d \times d}$, $W \in \mathbb{R}^{d \times d}$. Then, we have*

$$\frac{d\mathcal{L}(M)}{dM} = W \circ (X^\top pX) + \lambda M.$$

Based on Theorem 5.3, we calculate the gradient of the pruning mask from Line 10 to Line 15 in our Algorithm 1.

5.3 LIPSCHITZ OF GRADIENT

Having obtained the close form of gradient, we proceed to investigate its Lipschitz continuity. We aim to show that the gradient $\nabla_M \mathcal{L}(M)$ is Lipschitz continuous with respect to M .

Theorem 5.4 (Lipschitz of the gradient, informal version of Theorem F.8). *Let R be some fixed constant that satisfies $R > 1$. Let $X \in \mathbb{R}^{n \times d}$, $W \in \mathbb{R}^{d \times d}$. We have $\|X\|_F \leq R$ and $\|W\|_F \leq R$. Let $\mathcal{L}(M)$ be defined in Definition 3.2. For $M, \tilde{M} \in \mathbb{R}^{d \times d}$, we have*

$$\|\nabla_M \mathcal{L}(M) - \nabla_M \mathcal{L}(\tilde{M})\|_F \leq (\lambda + 30dn^{7/2}R^6) \cdot \|M - \tilde{M}\|_F.$$

We defer the proof to Section F. Establishing the Lipschitz continuity of the gradient satisfies one of the necessary conditions for applying Theorem 5.2. The above theorem implicates that the gradient for M is upper bounded, providing a way to choose step size.

5.4 PL INEQUALITY OF GRADIENT

Next, we need to verify that our loss function satisfies the PL inequality with appropriate parameters. To complete the verification of the conditions required for convergence, we demonstrate that $\mathcal{L}(M)$ satisfies the PL inequality. We show that $\nabla_M \mathcal{L}(M)$ satisfies the PL inequality in this lemma:

Lemma 5.5 (PL inequality, informal version of Lemma G.10). *Let $M, X, W, \tilde{D}, \tilde{A}, \lambda, \mathcal{L}, \mathcal{L}_{\text{attn}}$ be defined in Definition 3.2. Assume that $XX^\top \succeq \beta I$ and $\min_{i,j \in [n]} (\tilde{D}^{-1} \tilde{A})_{i,j} \geq \delta > 0$. Also,*

- Let $\mu = 2 \min_{i,j \in [d]} \{|W_{i,j}|\} \cdot \beta \cdot \delta$.
- Let $\xi = 12\sqrt{n} \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X\|_F^2 \cdot \lambda d / \mu$.

We have $\|\nabla_M \mathcal{L}(M)\|_F^2 \geq \frac{1}{2}\mu(2\mathcal{L}_{\text{attn}}(M) + \frac{2\lambda^2}{\mu}\|M\|_F^2 - \xi)$.

We defer the proof to Section G. By confirming that $\mathcal{L}(M)$ satisfies the PL inequality and that its gradient is Lipschitz continuous, we then apply Theorem 5.2 to conclude that GD will converge to a solution within our desired error tolerance, and further prove Theorem 4.1.

To prove the PL inequality, we also need the following two key Lemmas, which introduce our two assumptions in our Theorem 4.1, $XX^\top \succeq \beta I$ and $\min_{i,j \in [n]} (\tilde{D}^{-1} \tilde{A})_{i,j} \geq \delta > 0$.

Lemma 5.6 (Informal version of Lemma G.4). *Let $B \in \mathbb{R}^{n \times n}$ and $X \in \mathbb{R}^{n \times d}$. Assume that $XX^\top \succeq \beta I$. Then, we have $\|X^\top BX\|_F \geq \beta \|B\|_F$.*

Lemma 5.7 (Informal version of Lemma G.7). *Let $B \in \mathbb{R}^{n \times n}$ and each row summation is zero, i.e., $B \cdot \mathbf{1}_n = \mathbf{0}_n$. Let $\tilde{B} \in [0, 1]^{n \times n}$ and each row summation is 1, i.e., $\tilde{B} \cdot \mathbf{1}_n = \mathbf{1}_n$. Assume that $\min_{i,j \in [n]} \tilde{B}_{i,j} \geq \delta > 0$. Then, we can show $\|B \circ \tilde{B} - \text{diag}((B \circ \tilde{B}) \cdot \mathbf{1}_n) \tilde{B}\|_F \geq \delta \cdot \|B\|_F$.*

6 EXPERIMENT

6.1 EVALUATION ON SYNTHETIC DATA

We discuss the synthetic experiments conducted to illustrate the effectiveness of our Algorithm 1.

Method and evaluation. We implement our method following the pseudocode in Algorithm 1, using NumPy and JAX for acceleration. We evaluate our method on unstructured sparsity, meaning that zeros can occur anywhere within the attention weight matrix W . Specifically, we use Definition 3.2 as our loss function, optimizing over the pruning mask M using gradient descent based on the closed-form expression derived in Theorem 5.3. To accelerate convergence, we leverage momentum into the optimization process and fix the momentum parameter at 0.9. After obtaining the optimal pruning mask, we convert M to a binary pruning mask to prune W , maintaining sparsity at

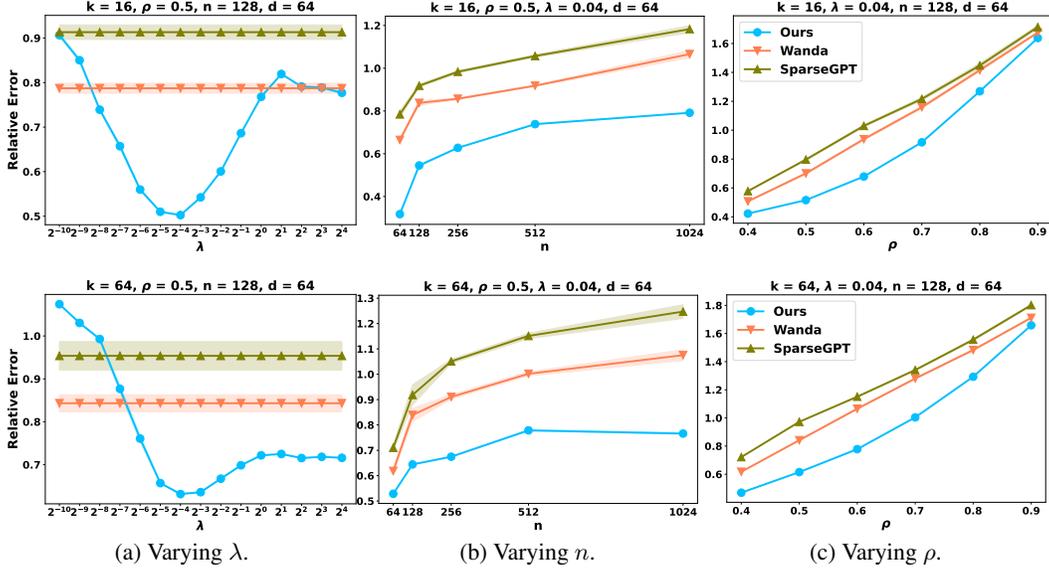


Figure 2: The comparison among our Algorithm 1, Wanda, and SparseGPT. The y -axis is a relative error, which is defined as $\frac{\|\tilde{D}^{-1}\tilde{A}-D^{-1}A\|_F^2}{\|D^{-1}A\|_F^2}$, where $D^{-1}A$ is original attention matrix and $\tilde{D}^{-1}\tilde{A}$ is approximated attention matrix based on three methods. We always use $d = 64$. We use $k = 16$ for the first row and $k = 64$ for the second row. The x -axis is (a) regularization coefficient λ for the left column; (b) input sequence length n for the middle column; (c) pruning ratio ρ for the right column.

the desired pruning ratio ρ . We use the relative error as our evaluation metric, which is defined as $\|\tilde{D}^{-1}\tilde{A}-D^{-1}A\|_F^2/\|D^{-1}A\|_F^2$, where \tilde{D} , \tilde{A} , D , A are defined in Definition 3.2.

Baselines. We compare our method with two linear pruning approaches, namely Wanda (Sun et al., 2024) and SparseGPT (Frantar & Alistarh, 2023). Wanda is a pruning method that removes weights with the smallest magnitudes multiplied by the corresponding input activations, achieving sparsity without requiring retraining or weight updates. SparseGPT is a second-order pruning method that utilizes the Hessian matrix to prune a portion of the weight matrix while simultaneously updating the remaining parameters. We implement Wanda and SparseGPT as described in their respective papers. Notably, since the settings of SparseGPT and Wanda are linear, we do not prune the fused weight matrix W directly; instead, we prune W_Q and W_K separately (see Figure 1).

Data. In order to assess the efficacy of different methods in approximating the attention matrix, we construct the data via a carefully defined generating process. Specifically, we create multiple independent random Gaussian matrices $G \in \mathbb{R}^{d \times d}$, where each entry of G drawn from a normal distribution, i.e., $G_{i,j} \sim \mathcal{N}(0, 1)$ for $i, j \in [d]$. Then, we perform singular value decomposition (SVD) on matrix G , i.e., $U, S, V^\top = \text{SVD}(G)$. We retain the first four singular values in S and set others to zero, constraining the rank to four. Our W_Q and W_K are then constructed as $U \text{diag}(S) V^\top$. The weight matrix W used in our setting is formed by taking the product $W = W_Q W_K^\top$. For $X \in \mathbb{R}^{n \times d}$, we generate it as a full-rank Gaussian random matrix.

Setup. In our experiments, the weight matrix dimension $d = 64$ is kept constant across all figures, and we simulate two datasets of size $k = 16$ and $k = 64$. We set the input sequence length $n = 128$ for experiments (a) and (c) in Figure 2. The pruning ratio $\rho = 0.5$ is set for experiments (a) and (b) in Figure 2. For our method, the regularization coefficient $\lambda := \tilde{\lambda}/n$ where we abuse the notation to denote $\tilde{\lambda}$ as the same used in Definition 3.2 and λ here is the parameter we really control in experiments. λ is set as 0.04 for experiments (b) and (c) in Figure 2 (intuition drawn from experiment (a)). The total number of epochs is set as $T = 100$. The step size is set as $\eta = 0.1/\lambda$ because Theorem 4.1 indicates that η is inversely proportional to λ with some constant, i.e., $\eta \propto 1/L \propto 1/(\lambda + \text{other terms})$.

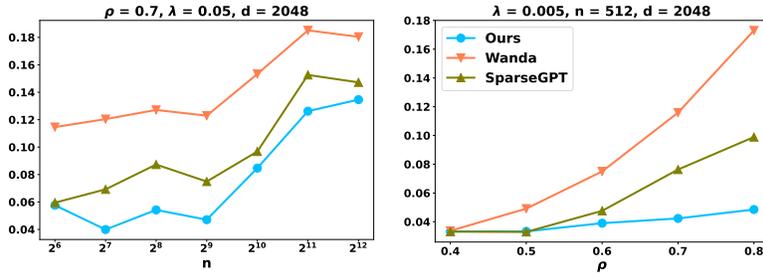


Figure 3: The comparison among our algorithm, Wanda, and SparseGPT on Llama 3.2-1B.

Results. Overall, the results in Figure 2 show that our Algorithm 1 outperforms Wanda and SparseGPT with a large margin, which supports our theoretical analysis in Theorem 4.1. In the following, we will discuss each setting in detail.

Relation with regularization coefficient λ . The leftmost column of Figure 2 investigates the impact of the regularization coefficient λ on relative error. As λ increases from very small values, the relative error initially decreases sharply for our algorithm, reaching a minimum before gradually rising again, which forms a U shape curve. This behavior indicates that there is an optimal λ where our algorithm achieves its best performance around 2^{-4} . The U-shape curve phenomena are well-known in most hyper-parameter choosing, e.g., regularization coefficient.

Relation with input sequence length n . The center column of Figure 2 explores how the relative error changes with respect to the input sequence length n . As n increases, the relative error for all three methods grows, though at different rates. Our method demonstrates a slower increase, maintaining a significant margin over both Wanda and SparseGPT, particularly for larger values of n . Wanda, while showing better performance than SparseGPT for larger sequence lengths, becomes comparable to SparseGPT as n is relatively small.

Relation with pruning ratio ρ . The rightmost column of Figure 2 illustrates the relationship between the relative error and the pruning ratio ρ for the three methods under comparison: our algorithm, Wanda, and SparseGPT. As the pruning ratio ρ increases, all methods exhibit a rise in relative error, indicating a degradation in approximation accuracy. However, our algorithm consistently outperforms both Wanda and SparseGPT across the range of ρ , with a lower relative error. SparseGPT and Wanda follow a similar trend, closely tracking each other.

6.2 EXPERIMENT ON REAL DATASET AND LLMs

In this subsection, we discuss the experiments conducted on the real dataset and LLMs to illustrate the effectiveness of our method.

Method and evaluation. We evaluate our method on unstructured sparsity, meaning that zeros can occur anywhere within the attention weight matrices W_Q and W_K . Specifically, we use the loss function defined below:

$$\mathcal{L}(M_Q, M_K) := \frac{1}{2} \|D^{-1}A - \tilde{D}^{-1}\tilde{A}\|_F^2 + \frac{1}{2} (\|M_Q\|_F^2 + \|M_K\|_F^2) \quad (1)$$

where D and A are the original attention matrix, M_Q, M_K are pruning masks, and

$$\tilde{A} := \exp(X(M_Q \circ W_Q)(M_K \circ W_K)^\top X^\top), \quad \tilde{D} := \text{diag}(\tilde{A} \cdot \mathbf{1}_n).$$

After obtaining the optimal pruning mask, we convert the pruning mask to a binary pruning mask to prune W_Q and W_K , maintaining sparsity at the desired pruning ratio ρ . We use the relative error as our evaluation metric.

Baselines. The pruning is performed on the last attention layer of the pretrained model Llama 3.2-1B Meta (2024). The baselines are Wanda and SparseGPT, the same as Section 6.2.

Data. To simulate real-world large language models, we utilize the Colossal Clean Crawled Corpus (C4 Dataset) Raffel et al. (2020), which is also used as the calibration dataset in our baselines Wanda and SparseGPT. Additionally, with a primary focus on pruning the attention matrix, we extract the

input hidden states corresponding to the target attention matrix from the pretrained Llama 3.2 model using a customized hook function and use these as our input X .

Setup. The weight matrix dimension is $d = 2048$ in Llama 3.2-1B. For all the experiments, we set λ as 0.05 and η as 0.005. For the varying n experiment, we set the pruning ratio ρ as 0.7. For the varying ρ experiment, we set the input sequence length n as 512.

Results. Overall, the results in Figure 3 show that our algorithm outperforms Wanda and SparseGPT in real-world LLMs, which supports our theoretical analysis in Theorem 4.1 and enhances our preliminary experiment in Section 6. In the following, we will discuss each setting in detail.

Relation with input sequence length n . The left column of Figure 3 shows that our algorithm outperforms the baselines in different sequence lengths continuously.

Relation with pruning ratio ρ . The right column of Figure 3 illustrates that as the pruning ratio ρ increases, all methods exhibit a rise in relative error. But our algorithm consistently outperforms both Wanda and SparseGPT across the range of ρ with a much lower increasing rate.

Assumptions verification. Notice that Theorem 4.1 relies on two assumptions: $XX^\top \succeq \beta I$ and $\min_{i,j \in [n]} (\tilde{D}^{-1} \tilde{A})_{i,j} \geq \delta > 0$. We verify these assumptions using the C4 dataset, obtaining $\beta \approx 0.034$ and $\delta \approx 0.0025$, thereby demonstrating the practicality of our theoretical framework.

6.3 EXPERIMENT ON END-TO-END PERPLEXITY

We present the experiment on end-to-end perplexity in this section.

Baselines. As our method focuses on pruning the attention matrix, it can be seamlessly combined with approaches that perform linear pruning on the MLP, such as Wanda and SparseGPT. Therefore, we conduct three groups of experiments: (1) using a dense MLP without pruning, (2) applying SparseGPT to prune the MLP layer, and (3) using Wanda to prune the MLP layer. Subsequently, we apply Wanda, SparseGPT, and our method to prune the attention weights.

Method and data. We use the same method and data as Section 6.2.

Setup. We use Llama 3.2 1B as the target model to prune, in which the hidden state dimension is 2048. We set the pruning ratio as 0.5 for MLP and 0.5 for Attention Layer, when pruning MLP and attention weights at the same time, we have pruning ratio ρ of the whole model as 0.5. We use 8 sentences in the C4 training dataset Raffel et al. (2020) as calibration data, and we use 128 sentences in the C4 validation dataset Raffel et al. (2020) to evaluate perplexity. We set 0.005 as the learning rate η and 0.05 as the regularization coefficient λ .

Results. As we can see in Table 1, our Attention pruning methods always outperform Wanda Attention pruning and SparseGPT Attention pruning by a large margin when combined with Dense/Wanda/SparseGPT MLP pruning methods. These empirical results support our theoretical analysis that our pruning method can converge. Furthermore, our method is broadly applicable in the real-world case and can be combined with many other pruning methods.

7 CONCLUSION

This paper introduces a novel approach to LLM weight pruning that directly optimizes for approximating the attention matrix. We provide theoretical guarantees for the convergence of our Gradient Descent-based algorithm to a near-optimal pruning mask solution. Experimental results demonstrated the method’s effectiveness in maintaining model performance while reducing computational costs. This work establishes a new theoretical foundation for pruning algorithm design in LLMs, potentially enabling more efficient inference on resource-constrained devices.

Table 1: Comparison of different methods based on Perplexity (PPL).

Method	PPL
Dense MLP + Dense Attn	12.487
Dense MLP + SparseGPT Attn	14.269
Dense MLP + Wanda Attn	14.912
Dense MLP + Our Attn	13.885
Wanda MLP + Wanda Attn	30.426
Wanda MLP + SparseGPT Attn	26.074
Wanda MLP + Our Attn	24.427
SparseGPT MLP + Wanda Attn	45.435
SparseGPT MLP + SparseGPT Attn	36.641
SparseGPT MLP + Our Attn	34.946

ACKNOWLEDGEMENT

Research is partially supported by the National Science Foundation (NSF) Grants 2023239-DMS, CCF-2046710, and Air Force Grant FA9550-18-1-0166.

REFERENCES

- Amol Aggarwal and Josh Alman. Optimal-degree polynomial approximations for exponentials and gaussian kernel density estimation. In *Proceedings of the 37th Computational Complexity Conference*, pp. 1–23, 2022.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. *Advances in neural information processing systems*, 32, 2019a.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via overparameterization. In *International Conference on Machine Learning*, 2019b.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. On the convergence rate of training recurrent neural networks. In *NeurIPS*, 2019c.
- Josh Alman and Zhao Song. Fast attention requires bounded entries. *Advances in Neural Information Processing Systems*, 36, 2023.
- Josh Alman and Zhao Song. The fine-grained complexity of gradient computation for training large language models. *arXiv preprint arXiv:2402.04497*, 2024a.
- Josh Alman and Zhao Song. How to capture higher-order correlations? generalizing matrix softmax attention to kronecker computation. In *The Twelfth International Conference on Learning Representations*, 2024b.
- Josh Alman and Zhao Song. Fast rope attention: Combining the polynomial method and fast fourier transform. *manuscript*, 2024c.
- Chenxin An, Fei Huang, Jun Zhang, Shansan Gong, Xipeng Qiu, Chang Zhou, and Lingpeng Kong. Training-free long-context scaling of large language models. *arXiv preprint arXiv:2402.17463*, 2024.
- Anthropic. The claude 3 model family: Opus, sonnet, haiku, 2024. https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pp. 322–332. PMLR, 2019.
- Saleh Ashkboos, Maximilian L Croci, Marcelo Gennari do Nascimento, Torsten Hoefler, and James Hensman. SliceGPT: Compress large language models by deleting rows and columns. In *The Twelfth International Conference on Learning Representations*, 2024.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- Yuan Cao and Quanquan Gu. Generalization error bounds of gradient descent for learning overparameterized deep relu networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 3349–3356, 2020.

- Ya-Ting Chang, Zhibo Hu, Xiaoyu Li, Shuiqiao Yang, Jiaojiao Jiang, and Nan Sun. Dihan: A novel dynamic hierarchical graph attention network for fake news detection. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, pp. 197–206, 2024.
- Moses Charikar, Michael Kapralov, Navid Nouri, and Paris Siminelakis. Kernel density estimation through density constrained near neighbor search. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 172–183. IEEE, 2020.
- Bo Chen, Xiaoyu Li, Yingyu Liang, Jiangxuan Long, Zhenmei Shi, and Zhao Song. Circuit complexity bounds for rope-based transformer architecture. *arXiv e-prints*, pp. arXiv–2411, 2024a.
- Bo Chen, Xiaoyu Li, Yingyu Liang, Zhenmei Shi, and Zhao Song. Bypassing the exponential dependency: Looped transformers efficiently learn in-context by multi-step gradient descent. In *International Conference on Artificial Intelligence and Statistics*, 2025a.
- Bo Chen, Yingyu Liang, Zhizhou Sha, Zhenmei Shi, and Zhao Song. Hsr-enhanced sparse attention acceleration. In *Conference on Parsimony and Learning*. PMLR, 2025b.
- Tianyi Chen, Bo Ji, Tianyu Ding, Biyi Fang, Guanyi Wang, Zhihui Zhu, Luming Liang, Yixin Shi, Sheng Yi, and Xiao Tu. Only train once: A one-shot neural network training and pruning framework. *Advances in Neural Information Processing Systems*, 34:19637–19651, 2021.
- Weize Chen, Ziming You, Ran Li, Yitong Guan, Chen Qian, Chenyang Zhao, Cheng Yang, Ruobing Xie, Zhiyuan Liu, and Maosong Sun. Internet of agents: Weaving a web of heterogeneous agents for collaborative intelligence. *arXiv preprint arXiv:2407.07061*, 2024b.
- Yifang Chen, Jiayan Huo, Xiaoyu Li, Yingyu Liang, Zhenmei Shi, and Zhao Song. Fast gradient computation for rope attention in almost linear time. *arXiv preprint arXiv:2412.17316*, 2024c.
- Yifang Chen, Xiaoyu Li, Yingyu Liang, Zhenmei Shi, and Zhao Song. The computational limits of state-space models and mamba via the lens of circuit complexity. In *Conference on Parsimony and Learning*. PMLR, 2025c.
- Yifang Chen, Xiaoyu Li, Yingyu Liang, Zhenmei Shi, and Zhao Song. Universal approximation of visual autoregressive transformers. *arXiv preprint arXiv:2502.06167*, 2025d.
- Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, and Jiaya Jia. Longlora: Efficient fine-tuning of long-context large language models. *arXiv preprint arXiv:2309.12307*, 2023.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022.
- Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. *arXiv preprint arXiv:2307.08691*, 2023.
- Tri Dao and Albert Gu. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. *arXiv preprint arXiv:2405.21060*, 2024.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359, 2022.
- Mehmet F Demirel, Shengchao Liu, Siddhant Garg, Zhenmei Shi, and Yingyu Liang. Attentive walk-aggregating graph neural networks. *Transactions on Machine Learning Research*, 2022.
- Yichuan Deng, Zhao Song, Yitan Wang, and Yuanyuan Yang. A nearly optimal size coresets algorithm with nearly linear time. *arXiv preprint arXiv:2210.08361*, 2022.

- Yichuan Deng, Zhao Song, Shenghao Xie, and Chiwun Yang. Unmasking transformers: A theoretical approach to data recovery via attention weights. *arXiv preprint arXiv:2310.12462*, 2023.
- Yichuan Deng, Jiangxuan Long, Zhao Song, Zifan Wang, and Han Zhang. Streaming kernel pca algorithm with small space. In *Conference on Parsimony and Learning*. PMLR, 2025.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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)*, pp. 4171–4186. Association for Computational Linguistics, 2019.
- Yiran Ding, Li Lyna Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. Longrope: Extending llm context window beyond 2 million tokens. *arXiv preprint arXiv:2402.13753*, 2024.
- Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *ICLR*. arXiv preprint arXiv:1810.02054, 2019.
- Tao Feng, Chuanyang Jin, Jingyu Liu, Kunlun Zhu, Haoqin Tu, Zirui Cheng, Guanyu Lin, and Jiakuan You. How far are we from agi. *arXiv preprint arXiv:2405.10313*, 2024.
- Elias Frantar and Dan Alistarh. Optimal brain compression: A framework for accurate post-training quantization and pruning. *Advances in Neural Information Processing Systems*, 35:4475–4488, 2022.
- Elias Frantar and Dan Alistarh. Sparseopt: Massive language models can be accurately pruned in one-shot. In *International Conference on Machine Learning*, pp. 10323–10337. PMLR, 2023.
- Spencer Frei and Quanquan Gu. Proxy convexity: A unified framework for the analysis of neural networks trained by gradient descent. *Advances in Neural Information Processing Systems*, 34: 7937–7949, 2021.
- Spencer Frei, Yuan Cao, and Quanquan Gu. Algorithm-dependent generalization bounds for over-parameterized deep residual networks. *Advances in neural information processing systems*, 32, 2019.
- Spencer Frei, Yuan Cao, and Quanquan Gu. Provable generalization of sgd-trained neural networks of any width in the presence of adversarial label noise. In *International Conference on Machine Learning*, pp. 3427–3438. PMLR, 2021.
- Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023a.
- Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, 2021.
- Yeqi Gao, Sridhar Mahadevan, and Zhao Song. An over-parameterized exponential regression. *arXiv preprint arXiv:2303.16504*, 2023b.
- Yeqi Gao, Zhao Song, Weixin Wang, and Junze Yin. A fast optimization view: Reformulating single layer attention in llm based on tensor and svm trick, and solving it in matrix multiplication time. *arXiv preprint arXiv:2309.07418*, 2023c.
- Yeqi Gao, Zhao Song, and Shenghao Xie. In-context learning for attention scheme: from single softmax regression to multiple softmax regression via a tensor trick. *arXiv preprint arXiv:2307.02419*, 2023d.
- Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.

- Torsten Hoefer, Dan Alistarh, Tal Ben-Nun, Nikoli Dryden, and Alexandra Peste. Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks. *Journal of Machine Learning Research*, 22(241):1–124, 2021.
- Coleman Hooper, Sehoon Kim, Hiva Mohammadzadeh, Michael W Mahoney, Yakun Sophia Shao, Kurt Keutzer, and Amir Gholami. Kvquant: Towards 10 million context length llm inference with kv cache quantization. *arXiv preprint arXiv:2401.18079*, 2024.
- Cheng-Yu Hsieh, Chun-Liang Li, Chih-kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 8003–8017, 2023.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.
- Jerry Yao-Chieh Hu, Donglin Yang, Dennis Wu, Chenwei Xu, Bo-Yu Chen, and Han Liu. On sparse modern hopfield model. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- Jerry Yao-Chieh Hu, Pei-Hsuan Chang, Haozheng Luo, Hong-Yu Chen, Weijian Li, Wei-Po Wang, and Han Liu. Outlier-efficient hopfield layers for large transformer-based models. In *Forty-first International Conference on Machine Learning (ICML)*, 2024a.
- Jerry Yao-Chieh Hu, Bo-Yu Chen, Dennis Wu, Feng Ruan, and Han Liu. Nonparametric modern hopfield models. *arXiv preprint arXiv:2404.03900*, 2024b.
- Jerry Yao-Chieh Hu, Thomas Lin, Zhao Song, and Han Liu. On computational limits of modern hopfield models: A fine-grained complexity analysis. In *Forty-first International Conference on Machine Learning (ICML)*, 2024c.
- Jerry Yao-Chieh Hu, Maojiang Su, En-Jui Kuo, Zhao Song, and Han Liu. Computational limits of low-rank adaptation (lora) for transformer-based models. *arXiv preprint arXiv:2406.03136*, 2024d.
- Jerry Yao-Chieh Hu, Wei-Po Wang, Ammar Gilani, Chenyang Li, Zhao Song, and Han Liu. Fundamental limits of prompt tuning transformers: Universality, capacity and efficiency. *arXiv preprint arXiv:2411.16525*, 2024e.
- Jerry Yao-Chieh Hu, Dennis Wu, and Han Liu. Provably optimal memory capacity for modern hopfield models: Tight analysis for transformer-compatible dense associative memories. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 37, 2024f.
- Jerry Yao-Chieh Hu, Weimin Wu, Yi-Chen Lee, Yu-Chao Huang, Minshuo Chen, and Han Liu. On statistical rates of conditional diffusion transformers: Approximation, estimation and minimax optimality. *arXiv preprint arXiv:2411.17522*, 2024g.
- Jerry Yao-Chieh Hu, Weimin Wu, Zhuoru Li, Sophia Pi, , Zhao Song, and Han Liu. On statistical rates and provably efficient criteria of latent diffusion transformers (dits). *Advances in Neural Information Processing Systems*, 38, 2024h.
- Itay Hubara, Brian Chmiel, Moshe Island, Ron Banner, Joseph Naor, and Daniel Soudry. Accelerated sparse neural training: A provable and efficient method to find n: m transposable masks. *Advances in neural information processing systems*, 34:21099–21111, 2021.
- Ziwei Ji and Matus Telgarsky. Polylogarithmic width suffices for gradient descent to achieve arbitrarily small test error with shallow relu networks. *arXiv preprint arXiv:1909.12292*, 2019.
- Yuxin Jiang, Chunkit Chan, Mingyang Chen, and Wei Wang. Lion: Adversarial distillation of proprietary large language models. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.

- Hongye Jin, Xiaotian Han, Jingfeng Yang, Zhimeng Jiang, Zirui Liu, Chia-Yuan Chang, Huiyuan Chen, and Xia Hu. Llm maybe longlm: Self-extend llm context window without tuning. *arXiv preprint arXiv:2401.01325*, 2024.
- Tian Jin, Michael Carbin, Dan Roy, Jonathan Frankle, and Gintare Karolina Dziugaite. Pruning’s effect on generalization through the lens of training and regularization. *Advances in Neural Information Processing Systems*, 35:37947–37961, 2022.
- Praneeth Kacham, Vahab Mirrokni, and Peilin Zhong. Polysketchformer: Fast transformers via sketches for polynomial kernels. *arXiv preprint arXiv:2310.01655*, 2023.
- Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2016, Riva del Garda, Italy, September 19-23, 2016, Proceedings, Part I 16*, pp. 795–811. Springer, 2016.
- Yekun Ke, Xiaoyu Li, Zhao Song, and Tianyi Zhou. Faster sampling algorithms for polytopes with small treewidth. In *2024 IEEE International Conference on Big Data (BigData)*. IEEE, 2024.
- Yekun Ke, Xiaoyu Li, Yingyu Liang, Zhizhou Sha, Zhenmei Shi, and Zhao Song. On computational limits and provably efficient criteria of visual autoregressive models: A fine-grained complexity analysis. *arXiv preprint arXiv:2501.04377*, 2025a.
- Yekun Ke, Yingyu Liang, Zhenmei Shi, Zhao Song, and Chiwun Yang. Curse of attention: A kernel-based perspective for why transformers fail to generalize on time series forecasting and beyond. In *Conference on Parsimony and Learning*. PMLR, 2025b.
- Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts, and Matei Zaharia. Demonstrate-search-predict: Composing retrieval and language models for knowledge-intensive nlp. *arXiv preprint arXiv:2212.14024*, 2022.
- Tzu-Sheng Kuo, Aaron Lee Halfaker, Zirui Cheng, Jiwoo Kim, Meng-Hsin Wu, Tongshuang Wu, Kenneth Holstein, and Haiyi Zhu. Wikibench: Community-driven data curation for ai evaluation on wikipedia. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pp. 1–24, 2024.
- Eldar Kurtic, Denis Kuznedelev, Elias Frantar, Michael Goin, and Dan Alistarh. Sparse finetuning for inference acceleration of large language models. *arXiv preprint arXiv:2310.06927*, 2023.
- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2021.
- Chenyang Li, Yingyu Liang, Zhenmei Shi, and Zhao Song. When can we solve the weighted low rank approximation problem in truly subquadratic time? In *International Conference on Artificial Intelligence and Statistics*, 2025a.
- Chenyang Li, Yingyu Liang, Zhenmei Shi, Zhao Song, and Tianyi Zhou. Fourier circuits in neural networks and transformers: A case study of modular arithmetic with multiple inputs. In *International Conference on Artificial Intelligence and Statistics*, 2025b.
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- Xiaoyu Li, Yingyu Liang, Zhenmei Shi, and Zhao Song. A tighter complexity analysis of sparsegpt. *arXiv preprint arXiv:2408.12151*, 2024a.
- Xiaoyu Li, Yingyu Liang, Zhenmei Shi, Zhao Song, and Junwei Yu. Fast john ellipsoid computation with differential privacy optimization. *arXiv preprint arXiv:2408.06395*, 2024b.
- Xiaoyu Li, Yingyu Liang, Zhenmei Shi, Zhao Song, and Yufa Zhou. Fine-grained attention i/o complexity: Comprehensive analysis for backward passes. *arXiv preprint arXiv:2410.09397*, 2024c.

- Xiaoyu Li, Jiangxuan Long, Zhao Song, and Tianyi Zhou. Fast second-order method for neural network under small treewidth setting. In *2024 IEEE International Conference on Big Data (BigData)*. IEEE, 2024d.
- Xiaoyu Li, Zhao Song, and Junwei Yu. Quantum speedups for approximating the john ellipsoid. *arXiv preprint arXiv:2408.14018*, 2024e.
- Xiaoyu Li, Yingyu Liang, Jiangxuan Long, Zhenmei Shi, Zhao Song, and Zhen Zhuang. Neural algorithmic reasoning for hypergraphs with looped transformers. *arXiv preprint arXiv:2501.10688*, 2025c.
- Xiaoyu Li, Yingyu Liang, Zhenmei Shi, Zhao Song, Wei Wang, and Jiahao Zhang. On the computational capability of graph neural networks: A circuit complexity bound perspective. *arXiv preprint arXiv:2501.06444*, 2025d.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. *Advances in neural information processing systems*, 31, 2018.
- Yuanzhi Li, Yingyu Liang, and Andrej Risteski. Recovery guarantee of weighted low-rank approximation via alternating minimization. In *International Conference on Machine Learning*, pp. 2358–2367. PMLR, 2016.
- Yingyu Liang, Heshan Liu, Zhenmei Shi, Zhao Song, and Junze Yin. Conv-basis: A new paradigm for efficient attention inference and gradient computation in transformers. *arXiv preprint arXiv:2405.05219*, 2024a.
- Yingyu Liang, Zhizhou Sha, Zhenmei Shi, and Zhao Song. Differential privacy mechanisms in neural tangent kernel regression. *arXiv preprint arXiv:2407.13621*, 2024b.
- Yingyu Liang, Zhizhou Sha, Zhenmei Shi, Zhao Song, and Yufa Zhou. Multi-layer transformers gradient can be approximated in almost linear time. *arXiv preprint arXiv:2408.13233*, 2024c.
- Yingyu Liang, Zhenmei Shi, Zhao Song, and Chiwun Yang. Toward infinite-long prefix in transformer. *arXiv preprint arXiv:2406.14036*, 2024d.
- Yingyu Liang, Zhenmei Shi, Zhao Song, and Yufa Zhou. Unraveling the smoothness properties of diffusion models: A gaussian mixture perspective. *arXiv preprint arXiv:2405.16418*, 2024e.
- Yingyu Liang, Zhenmei Shi, Zhao Song, and Yufa Zhou. Differential privacy of cross-attention with provable guarantee. *arXiv preprint arXiv:2407.14717*, 2024f.
- Yingyu Liang, Zhenmei Shi, Zhao Song, and Yufa Zhou. Tensor attention training: Provably efficient learning of higher-order transformers. *arXiv preprint arXiv:2405.16411*, 2024g.
- Yingyu Liang, Zhizhou Sha, Zhenmei Shi, Zhao Song, and Yufa Zhou. Looped relu mlps may be all you need as practical programmable computers. In *International Conference on Artificial Intelligence and Statistics*, 2025.
- Na Liu, Liangyu Chen, Xiaoyu Tian, Wei Zou, Kaijiang Chen, and Ming Cui. From llm to conversational agent: A memory enhanced architecture with fine-tuning of large language models. *arXiv preprint arXiv:2401.02777*, 2024.
- Tianlin Liu and Friedemann Zenke. Finding trainable sparse networks through neural tangent transfer. In *International Conference on Machine Learning*, pp. 6336–6347. PMLR, 2020.
- AI @ Meta Llama Team. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Stanislaw Lojasiewicz. A topological property of real analytic subsets. *Coll. du CNRS, Les équations aux dérivées partielles*, 117(87-89):2, 1963.
- Xuezhe Ma, Xiaomeng Yang, Wenhan Xiong, Beidi Chen, Lili Yu, Hao Zhang, Jonathan May, Luke Zettlemoyer, Omer Levy, and Chunting Zhou. Megalodon: Efficient llm pretraining and inference with unlimited context length. *arXiv preprint arXiv:2404.08801*, 2024.

- Amama Mahmood, Junxiang Wang, Bingsheng Yao, Dakuo Wang, and Chien-Ming Huang. Llm-powered conversational voice assistants: Interaction patterns, opportunities, challenges, and design guidelines. *arXiv preprint arXiv:2309.13879*, 2023.
- Jean Mercat, Igor Vasiljevic, Sedrick Keh, Kushal Arora, Achal Dave, Adrien Gaidon, and Thomas Kollar. Linearizing large language models. *arXiv preprint arXiv:2405.06640*, 2024.
- Meta. Llama 3.2: Revolutionizing edge ai and vision with open, customizable models. <https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/>, 2024. Accessed: 2024-11-21.
- Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one? *Advances in neural information processing systems*, 32, 2019.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. Cross-task generalization via natural language crowdsourcing instructions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 2022.
- Maxwell Nye, Anders Johan Andreassen, Gur AriGuy, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. Show your work: Scratchpads for intermediate computation with language models. *arXiv preprint arXiv:2112.00114*, 2021.
- OpenAI. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024a. Accessed: May 14.
- OpenAI. Introducing openai o1-preview. <https://openai.com/index/introducing-openai-o1-preview/>, 2024b. Accessed: September 12.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 2022.
- Gunho Park, Minsub Kim, Sungjae Lee, Jeonghoon Kim, Beomseok Kwon, Se Jung Kwon, Byeongwook Kim, Youngjoo Lee, Dongsoo Lee, et al. Lut-gemm: Quantized matrix multiplication based on luts for efficient inference in large-scale generative language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. Yarn: Efficient context window extension of large language models. *arXiv preprint arXiv:2309.00071*, 2023.
- Boris T Polyak. Gradient methods for the minimisation of functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4):864–878, 1963.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- Ilya Razenshteyn, Zhao Song, and David P Woodruff. Weighted low rank approximations with provable guarantees. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pp. 250–263, 2016.
- Clayton Sanford, Daniel J Hsu, and Matus Telgarsky. Representational strengths and limitations of transformers. *Advances in Neural Information Processing Systems*, 36, 2024.
- Jay Shah, Ganesh Bikshandi, Ying Zhang, Vijay Thakkar, Pradeep Ramani, and Tri Dao. Flashattention-3: Fast and accurate attention with asynchrony and low-precision. *arXiv preprint arXiv:2407.08608*, 2024.
- Xuan Shen, Zhao Song, Yufa Zhou, Bo Chen, Yanyu Li, Yifan Gong, Kai Zhang, Hao Tan, Jason Kuen, Henghui Ding, Zhihao Shu, Wei Niu, Pu Zhao, Yanzhi Wang, and Jiuxiang Gu. Lazydit: Lazy learning for the acceleration of diffusion transformers. *arXiv preprint arXiv:2412.12444*, 2024a.

- Xuan Shen, Zhao Song, Yufa Zhou, Bo Chen, Jing Liu, Ruiyi Zhang, Ryan A. Rossi, Hao Tan, Tong Yu, Xiang Chen, Yufan Zhou, Tong Sun, Pu Zhao, Yanzhi Wang, and Jiuxiang Gu. Numerical pruning for efficient autoregressive models. *arXiv preprint arXiv:2412.12441*, 2024b.
- Zhenmei Shi, Junyi Wei, and Yingyu Liang. A theoretical analysis on feature learning in neural networks: Emergence from inputs and advantage over fixed features. In *International Conference on Learning Representations*, 2021.
- Zhenmei Shi, Jiefeng Chen, Kunyang Li, Jayaram Raghuram, Xi Wu, Yingyu Liang, and Somesh Jha. The trade-off between universality and label efficiency of representations from contrastive learning. In *The Eleventh International Conference on Learning Representations*, 2023.
- Zhenmei Shi, Yifei Ming, Xuan-Phi Nguyen, Yingyu Liang, and Shafiq Joty. Discovering the gems in early layers: Accelerating long-context llms with 1000x input token reduction. *arXiv preprint arXiv:2409.17422*, 2024a.
- Zhenmei Shi, Junyi Wei, and Yingyu Liang. Provable guarantees for neural networks via gradient feature learning. *Advances in Neural Information Processing Systems*, 36, 2024b.
- Kumar Shridhar, Alessandro Stolfo, and Mrinmaya Sachan. Distilling reasoning capabilities into smaller language models. In *The 61st Annual Meeting Of The Association For Computational Linguistics*, 2023.
- Anshumali Shrivastava, Zhao Song, and Zhaozhuo Xu. A theoretical analysis of nearest neighbor search on approximate near neighbor graph. *arXiv preprint arXiv:2303.06210*, 2023.
- Ritwik Sinha, Zhao Song, and Tianyi Zhou. A mathematical abstraction for balancing the trade-off between creativity and reality in large language models. *arXiv preprint arXiv:2306.02295*, 2023.
- Zhao Song and Chiwun Yang. An automatic learning rate schedule algorithm for achieving faster convergence and steeper descent. *arXiv preprint arXiv:2310.11291*, 2023.
- Zhao Song, Lichen Zhang, and Ruizhe Zhang. Training multi-layer over-parametrized neural network in subquadratic time. In *Innovations in Theoretical Computer Science (ITCS)*, pp. 93:1–93:15, 2024.
- Mitchell Stern, Noam Shazeer, and Jakob Uszkoreit. Blockwise parallel decoding for deep autoregressive models. *Advances in Neural Information Processing Systems*, 31, 2018.
- Mingjie Sun, Zhuang Liu, Anna Bair, and J Zico Kolter. A simple and effective pruning approach for large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- Tianxiang Sun, Yunfan Shao, Hong Qian, Xuanjing Huang, and Xipeng Qiu. Black-box tuning for language-model-as-a-service. In *International Conference on Machine Learning*. PMLR, 2022.
- Xiaoyu Tan, Shaojie Shi, Xihe Qiu, Chao Qu, Zhenting Qi, Yinghui Xu, and Yuan Qi. Self-criticism: Aligning large language models with their understanding of helpfulness, honesty, and harmlessness. In *Proceedings of the 2023 conference on empirical methods in natural language processing: industry track*, pp. 650–662, 2023.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*. PMLR, 2023.

- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. In *The 61st Annual Meeting Of The Association For Computational Linguistics*, 2023.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- Jerry Wei, Le Hou, Andrew Kyle Lampinen, Xiangning Chen, Da Huang, Yi Tay, Xinyun Chen, Yifeng Lu, Denny Zhou, Tengyu Ma, and Quoc V Le. Symbol tuning improves in-context learning in language models. In *The 2023 Conference on Empirical Methods in Natural Language Processing*, 2023.
- Dennis Wu, Jerry Yao-Chieh Hu, Teng-Yun Hsiao, and Han Liu. Uniform memory retrieval with larger capacity for modern hopfield models. In *Forty-first International Conference on Machine Learning (ICML)*, 2024a.
- Dennis Wu, Jerry Yao-Chieh Hu, Weijian Li, Bo-Yu Chen, and Han Liu. STanhop: Sparse tandem hopfield model for memory-enhanced time series prediction. In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024b.
- Weimin Wu, Maojiang Su, Jerry Yao-Chieh Hu, Zhao Song, and Han Liu. Transformers are deep optimizers: Provable in-context learning for deep model training. *arXiv preprint arXiv:2411.16549*, 2024c.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*, 2023.
- Haojun Xia, Zhen Zheng, Yuchao Li, Donglin Zhuang, Zhongzhu Zhou, Xiafei Qiu, Yong Li, Wei Lin, and Shuaiwen Leon Song. Flash-llm: Enabling cost-effective and highly-efficient large generative model inference with unstructured sparsity. *Proceedings of the VLDB Endowment*, 17(2): 211–224, 2023.
- Mengzhou Xia, Tianyu Gao, Zhiyuan Zeng, and Danqi Chen. Sheared llama: Accelerating language model pre-training via structured pruning. In *The Twelfth International Conference on Learning Representations*, 2024.
- Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International Conference on Machine Learning*, pp. 38087–38099. PMLR, 2023.
- Sicong Xie, Qunwei Li, Weidi Xu, Kaiming Shen, Shaohu Chen, and Wenliang Zhong. Denoising time cycle modeling for recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1950–1955, 2022.
- Chenwei Xu, Yu-Chao Huang, Jerry Yao-Chieh Hu, Weijian Li, Ammar Gilani, Hsi-Sheng Goan, and Han Liu. Bishop: Bi-directional cellular learning for tabular data with generalized sparse modern hopfield model. In *Forty-first International Conference on Machine Learning (ICML)*, 2024a.
- Jun Xu, Weidi Xu, Mengshu Sun, Taifeng Wang, and Wei Chu. Extracting trigger-sharing events via an event matrix. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 1189–1201, 2022.
- Weidi Xu, Jingwei Wang, Lele Xie, Jianshan He, Hongting Zhou, Taifeng Wang, Xiaopei Wan, Jingdong Chen, Chao Qu, and Wei Chu. Logicmp: A neuro-symbolic approach for encoding first-order logic constraints. In *The Twelfth International Conference on Learning Representations*, 2024b.
- Zhuoyan Xu, Zhenmei Shi, Junyi Wei, Yin Li, and Yingyu Liang. Improving foundation models for few-shot learning via multitask finetuning. In *ICLR 2023 Workshop on Mathematical and Empirical Understanding of Foundation Models*, 2023.

- Zhuoyan Xu, Zhenmei Shi, and Yingyu Liang. Do large language models have compositional ability? an investigation into limitations and scalability. In *First Conference on Language Modeling*, 2024c.
- Zhuoyan Xu, Zhenmei Shi, Junyi Wei, Fangzhou Mu, Yin Li, and Yingyu Liang. Towards few-shot adaptation of foundation models via multitask finetuning. *arXiv preprint arXiv:2402.15017*, 2024d.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Amir Zandieh, Insu Han, Majid Daliri, and Amin Karbasi. Kdeformer: Accelerating transformers via kernel density estimation. In *ICML*. arXiv preprint arXiv:2302.02451, 2023.
- Amir Zandieh, Majid Daliri, and Insu Han. Qjl: 1-bit quantized jl transform for kv cache quantization with zero overhead. *arXiv preprint arXiv:2406.03482*, 2024.
- Yuchen Zeng and Kangwook Lee. The expressive power of low-rank adaptation. In *The Twelfth International Conference on Learning Representations*, 2024.
- Jiahao Zhang. Graph unlearning with efficient partial retraining. In *Companion Proceedings of the ACM on Web Conference 2024*, pp. 1218–1221, 2024.
- Jiahao Zhang, Rui Xue, Wenqi Fan, Xin Xu, Qing Li, Jian Pei, and Xiaorui Liu. Linear-time graph neural networks for scalable recommendations. In *Proceedings of the ACM on Web Conference 2024*, pp. 3533–3544, 2024a.
- Jieyu Zhang, Ranjay Krishna, Ahmed H Awadallah, and Chi Wang. Ecoassistant: Using llm assistant more affordably and accurately. *arXiv preprint arXiv:2310.03046*, 2023a.
- Michael Zhang, Kush Bhatia, Hermann Kumbong, and Christopher Ré. The hedgehog & the porcupine: Expressive linear attentions with softmax mimicry. *arXiv preprint arXiv:2402.04347*, 2024b.
- Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*, 2023b.
- Yifan Zhang, Yifeng Liu, Huizhuo Yuan, Zhen Qin, Yang Yuan, Quanquan Gu, and Andrew Chi-Chih Yao. Tensor product attention is all you need. *arXiv preprint arXiv:2501.06425*, 2025.
- Yuxin Zhang, Lirui Zhao, Mingbao Lin, Sun Yunyun, Yiwu Yao, Xingjia Han, Jared Tanner, Shiwei Liu, and Rongrong Ji. Dynamic sparse no training: Training-free fine-tuning for sparse llms. In *The Twelfth International Conference on Learning Representations*, 2024c.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*. PMLR, 2021.
- Huaixiu Steven Zheng, Swaroop Mishra, Xinyun Chen, Heng-Tze Cheng, Ed H Chi, Quoc V Le, and Denny Zhou. Step-back prompting enables reasoning via abstraction in large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srini Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, LILI YU, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. LIMA: Less is more for alignment. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- Xunyu Zhu, Jian Li, Yong Liu, Can Ma, and Weiping Wang. A survey on model compression for large language models. *arXiv preprint arXiv:2308.07633*, 2023.

Appendix

CONTENTS

1	Introduction	1
2	Related Work	3
2.1	Pruning and Compression for LLMs	3
2.2	Attention Acceleration	4
3	Preliminary	4
4	Main Results	5
5	Technique Overview	6
5.1	Previous Tools on Convergence of GD	6
5.2	Closed Form of Gradient	6
5.3	Lipschitz of Gradient	7
5.4	PL Inequality of Gradient	7
6	Experiment	7
6.1	Evaluation on Synthetic Data	7
6.2	Experiment on Real Dataset and LLMs	9
6.3	Experiment on End-to-end Perplexity	10
7	Conclusion	10
A	More Related Work	23
B	Preliminary	23
B.1	Notations	23
B.2	Facts	24
C	Gradient Calculation	26
C.1	Definitions	26
C.2	Gradient for Each Row of $X(M \circ W)X^\top$	28
C.3	Gradient for Each Row of $\tilde{u}(M)$	29
C.4	Gradient for Each Entry of $\tilde{\alpha}(M)$	29
C.5	Gradient for Each Entry of $\tilde{f}(M)$	30
C.6	Gradient for Each Entry of $c(M)$	31
C.7	Gradient for $\mathcal{L}_{\text{attn}}(M)$	32
C.8	Gradient for $\mathcal{L}_{\text{reg}}(M)$	32

C.9	Gradient for $\mathcal{L}(M)$	33
D	Matrix Form	34
D.1	Matrix Form of $B(M)$	34
D.2	Matrix Form of $\frac{d}{dM}\mathcal{L}(M)$	37
E	Bounds for Basic Functions	38
E.1	Basic Assumptions	38
E.2	Bounds for Basic Functions	38
E.3	Bounds for Gradient of $\tilde{f}(M)$	40
F	Lipschitz of Gradient	41
F.1	Useful Facts	41
F.2	Lipschitz of $\tilde{f}(M)$	42
F.3	Lipschitz of $c(M)$	42
F.4	Lipschitz of $\tilde{f}(M) \circ c(M)$	42
F.5	Lipschitz of $\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n)$	43
F.6	Lipschitz of $\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n)\tilde{f}(M)$	44
F.7	Lipschitz of Gradient	44
G	Convergence of Gradient Descent	45
G.1	Helpful Statements	45
G.2	Lower Bound on Frobenius Norm	46
G.3	Sandwich Lower Bound on Frobenius Norm	47
G.4	Lower Bound on Hadamard Product Between Two Matrices	48
G.5	Final Bound	48
G.6	PL Inequality	49

Roadmap. The appendix is organized as follows. In Section A, we reviewed more literature related to our paper. In Section B, we give the preliminary of our paper. In Section C, we provide a detailed gradient analysis of the loss function. In Section D, we provide details about how we integrate the gradient of loss function into matrix form. In Section E, we bound some basic functions to be used later. In Section F, we provide proof for the Lipschitz property of the gradient of the loss function. In Section G, we provide proof of convergence for GD.

A MORE RELATED WORK

Large Language Models. Transformer-based neural networks (Vaswani et al., 2017) have rapidly emerged as the dominant architecture for natural language processing in machine learning. When expanded to billions of parameters and trained on vast, diverse datasets, these systems are typically termed large language models (LLMs) or foundation models (Bommasani et al., 2021). Prominent LLM examples encompass BERT (Devlin et al., 2019), PaLM (Chowdhery et al., 2022), Llama (Touvron et al., 2023; Meta, 2024), and GPT4o (OpenAI, 2024a), which display adaptable competencies (Bubeck et al., 2023) across numerous downstream applications. To enhance LLMs for domain-specific uses, researchers have created multiple adaptation approaches. These include: adapter modules (Hu et al., 2022; Gao et al., 2023a; Zhang et al., 2023b; Shi et al., 2023); calibration mechanisms (Zhou et al., 2023; Zhao et al., 2021); multitask refinement (Gao et al., 2021; Von Oswald et al., 2023; Xu et al., 2024d; 2023); along with prompt engineering (Lester et al., 2021), scratchpad approaches (Nye et al., 2021), instruction optimization (Chung et al., 2022; Li & Liang, 2021; Mishra et al., 2022), symbolic adaptation (Wei et al., 2023; Xu et al., 2024b; 2022), black-box adjustments (Sun et al., 2022), human-aligned reinforcement learning (Ouyang et al., 2022), and structured reasoning techniques (Khattab et al., 2022; Yao et al., 2023; Wei et al., 2022; Zheng et al., 2024). Contemporary investigations cover tensor architecture innovations Alman & Song (2024b); Liang et al. (2024g); Sanford et al. (2024); Zhang et al. (2025), efficiency enhancements Alman & Song (2024c); Chen et al. (2024c; 2025b); Hu et al. (2024a;b;c;h; 2023); Ke et al. (2024); Liang et al. (2024a); Li et al. (2024b;c;d;a); Liang et al. (2024d); Shi et al. (2024a); Shen et al. (2024a;b); Song et al. (2024); Wu et al. (2024a;b); Xu et al. (2024a), plus ancillary studies Chen et al. (2024a); Demirel et al. (2022); Chang et al. (2024); Deng et al. (2022); Gao et al. (2023b;d); Li et al. (2025b;d); Liang et al. (2024f); Li et al. (2024e); Shrivastava et al. (2023); Sinha et al. (2023); Song & Yang (2023); Tan et al. (2023); Xie et al. (2022); Xu et al. (2024c); Zhang (2024); Zhang et al. (2024a); Chen et al. (2025d); Li et al. (2025c); Ke et al. (2025b); Liang et al. (2025); Chen et al. (2025a;c); Gao et al. (2023c); Hu et al. (2024g,e); Wu et al. (2024c).

B PRELIMINARY

In Section B.1, we introduce some notations we use in this paper. In Section B.2, we provide some basic facts.

B.1 NOTATIONS

For any positive integer n , we use $[n]$ to denote set $\{1, 2, \dots, n\}$. For two vectors $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^n$, we use $\langle x, y \rangle$ to denote the inner product between x, y , i.e., $\langle x, y \rangle = \sum_{i=1}^n x_i y_i$. For each $a, b \in \mathbb{R}^n$, we use $a \circ b \in \mathbb{R}^n$ to denote the Hadamard product, i.e. the i -th entry of $(a \circ b)$ is $a_i b_i$ for all $i \in [n]$. We use e_i to denote a vector where only i -th coordinate is 1, and other entries are 0. We use $\mathbf{1}_n$ to denote a length- n vector where all the entries are ones. We use $\|x\|_p$ to denote the ℓ_p norm of a vector $x \in \mathbb{R}^n$, i.e. $\|x\|_1 := \sum_{i=1}^n |x_i|$, $\|x\|_2 := (\sum_{i=1}^n x_i^2)^{1/2}$, and $\|x\|_\infty := \max_{i \in [n]} |x_i|$. For $A \in \mathbb{R}^{m \times n}$, let $A_i \in \mathbb{R}^n$ denote the i -th row and $A_{*,j} \in \mathbb{R}^m$ denote the j -th column of A , where $i \in [m]$ and $j \in [n]$. For a square matrix A , we use $\text{tr}[A]$ to denote the trace of A , i.e., $\text{tr}[A] = \sum_{i=1}^n A_{i,i}$. For two matrices $X, Y \in \mathbb{R}^{m \times n}$, the standard inner product between matrices is defined by $\langle X, Y \rangle := \text{tr}[X^\top Y]$. We use $\exp(A)$ to denote a matrix where $\exp(A)_{i,j} := \exp(A_{i,j})$ for a matrix $A \in \mathbb{R}^{n \times d}$. For $k > n$, for any matrix $A \in \mathbb{R}^{k \times n}$, we use $\|A\|$ to denote the spectral norm of A , i.e. $\|A\| := \sup_{x \in \mathbb{R}^n} \|Ax\|_2 / \|x\|_2$. We use $\|A\|_\infty$ to denote the ℓ_∞ norm of a matrix $A \in \mathbb{R}^{n \times d}$, i.e. $\|A\|_\infty := \max_{i \in [n], j \in [d]} |A_{i,j}|$. We use $\|A\|_F$ to denote the Frobenius norm of a matrix $A \in \mathbb{R}^{n \times d}$, i.e. $\|A\|_F := \sqrt{\sum_{i \in [n]} \sum_{j \in [d]} |A_{i,j}|^2}$. For a symmetric matrix $A \in \mathbb{R}^{n \times n}$,

we use $A \succeq 0$ (positive semidefinite (PSD)), if for all $x \in \mathbb{R}^n$, we have $x^\top A x \geq 0$. We use $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ to denote the minimum and the maximum eigenvalue of the square matrix A , respectively. Let $A \in \mathbb{R}^{n \times d}$. We use $a := \text{vec}(A)$ to denote a length nd vector. We stack rows of A into a column vector, i.e. $\text{vec}(A) := [a_1^\top, a_2^\top, \dots, a_n^\top]^\top$ where a_i^\top is the i -th row of A , or simply $\text{vec}(A)_{j+(i-1)d} := A_{i,j}$ for any $i \in [n], j \in [d]$.

B.2 FACTS

Fact B.1 (Indexing). *Suppose we have matrices $U \in \mathbb{R}^{n \times m}, V \in \mathbb{R}^{m \times d}$. We define*

$$\underbrace{X}_{n \times d} := \underbrace{U}_{n \times m} \underbrace{V}_{m \times d}.$$

Then, we have the following:

- Indexing for one row: $X_i = V^\top U_i \in \mathbb{R}^d$, i.e. $X_i^\top = U_i^\top V$, for $i \in [n]$.
- Indexing for one column: $X_{*,j} = UV_{*,j} \in \mathbb{R}^n$ for $j \in [d]$.

Fact B.2. *We have*

Part 1. *Suppose we have vectors $u \in \mathbb{R}^n, v \in \mathbb{R}^n$. For $i \in [n]$, we define*

$$x_i := u_i v_i.$$

Then we have the following:

$$\underbrace{x}_{n \times 1} = \underbrace{u \circ v}_{n \times 1} = \underbrace{\text{diag}(u)}_{n \times n} \underbrace{v}_{n \times 1} = \underbrace{\text{diag}(v)}_{n \times n} \underbrace{u}_{n \times 1}$$

Part 2. *Suppose we have matrix $W \in \mathbb{R}^{n \times n}$, vector $u \in \mathbb{R}^n$. For $i \in [n]$, we define*

$$X_{*,j} = W_{*,j} u_j.$$

Then we have the following:

- $X = W \text{diag}(u)$

Fact B.3 (Calculus). *We have*

Part 1. (Scalar calculus) *For any $t \in \mathbb{R}$, function $f : \mathbb{R} \rightarrow \mathbb{R}$, we have*

$$\bullet \frac{df^n(t)}{dt} = n f^{n-1}(t) \frac{df(t)}{dt}.$$

Part 2. (Vector calculus) *For any $x, y \in \mathbb{R}^n, t \in \mathbb{R}$, we have*

- $\frac{d(x \circ y)}{dt} = \frac{dx}{dt} \circ y + \frac{dy}{dt} \circ x$. (Product rule of vector Hadamard product)
- $\frac{d\langle x, y \rangle}{dt} = \langle \frac{dx}{dt}, y \rangle + \langle x, \frac{dy}{dt} \rangle$. (Product rule of inner product)
- $\frac{dx}{dx_i} = e_i$.

Part 3. (Matrix calculus) *For any $X, Y \in \mathbb{R}^{n \times m}, Z \in \mathbb{R}^{m \times d}, t \in \mathbb{R}$ which is independent of Z , function $f : \mathbb{R} \rightarrow \mathbb{R}^{n \times d}$, functions $f_1(t), f_2(t), \dots, f_n(t) : \mathbb{R} \rightarrow \mathbb{R}^{n \times d}$, we have*

- $\frac{d(X \circ Y)}{dt} = \frac{dX}{dt} \circ Y + \frac{dY}{dt} \circ X$. (Product rule of matrix Hadamard product)
- $\frac{d \exp(f(t))}{dt} = \exp(f(t)) \circ \frac{df(t)}{dt}$, where $\exp(\cdot)$ is applied entry-wise.
- $\frac{d(XZ)}{dt} = \frac{dX}{dt} Z$.
- $\frac{d(ZX^\top)}{dt} = Z \frac{dX^\top}{dt}$.

- $\frac{d}{dt} \sum_{i=1}^n f_i(t) = \sum_{i=1}^n \frac{df_i(t)}{dt}$.
- $\frac{dX}{dX_{i,j}} = \underbrace{e_i}_{n \times 1} \underbrace{e_j^\top}_{1 \times m}$.

Fact B.4 (Basic algebra). *Let $u \in \mathbb{R}^n$, $v \in \mathbb{R}^n$, $w \in \mathbb{R}^n$, $X \in \mathbb{R}^{n \times d}$, $Y \in \mathbb{R}^{n \times d}$, and $Z \in \mathbb{R}^{n \times n}$. Then, we have*

- $\langle u, v \rangle = \langle v, u \rangle = u^\top v = v^\top u$
- $u \circ v = v \circ u = \text{diag}(u)v = \text{diag}(v)u$
- $\langle u, v \rangle = \langle u \circ v, \mathbf{1}_n \rangle$
- $\langle u \circ v, w \rangle = \langle u \circ w, v \rangle = \langle w \circ v, u \rangle$
- $u^\top (v \circ w) = u^\top \text{diag}(v)w$
- $(X \circ Y)^\top = X^\top \circ Y^\top$
- $X \circ e_i e_j^\top = X_{i,j} e_i e_j^\top$
- $\text{diag}(u)Z \text{diag}(v) = (uv^\top) \circ Z$
- $XY^\top = \sum_{i \in [d]} X_{*,i} Y_{*,i}^\top$
- $X_{i,j} Y_{i,j} = (X \circ Y)_{i,j}$
- $\sum_{j \in [n]} u \circ A_{*,j} = u \circ \sum_{j \in [n]} A_{*,j}$
- $\|X\|_F^2 = \text{tr}[XX^\top]$
- $\text{tr}[XY^\top] = \text{tr}[Y^\top X]$
- $\|\text{diag}(u)\|_F = \|u\|_2$

Fact B.5 (Norm bounds). *For $a \in \mathbb{R}$, $u \in \mathbb{R}^d$, $X, Y \in \mathbb{R}^{n \times d}$, $Z \in \mathbb{R}^{d \times m}$ we have*

- $\|aX\|_F = |a| \|X\|_F$ (*absolute homogeneity*).
- $\|X + Y\|_F \leq \|X\|_F + \|Y\|_F$ (*triangle inequality*).
- $|\langle X, Y \rangle| \leq \|X\|_F \cdot \|Y\|_F$ (*Cauchy–Schwarz inequality*).
- $\|X^\top\|_F = \|X\|_F$.
- $\|Xu\|_2 \leq \|X\| \cdot \|u\|_2$
- $\|X \circ Y\|_F \leq \|X\|_F \cdot \|Y\|_F$.
- *For any $i \in [n]$, $j \in [d]$, we have $|X_{i,j}| \leq \|X\|_F$.*
- $\|X\| \leq \|X\|_F \leq \sqrt{k} \|X\|$ *where k is the rank of X .*
- $\|Y \cdot Z\|_F \leq \|Y\|_F \cdot \|Z\|_F$.

Fact B.6. *For matrices $A, B \in \mathbb{R}^{m \times n}$, we have*

$$\|A + B\|_F^2 = \|A\|_F^2 + \|B\|_F^2 + 2\langle A, B \rangle.$$

Proof. We can show

$$\begin{aligned} \|A + B\|_F^2 &= \text{tr}[(A + B)^\top (A + B)] \\ &= \text{tr}[A^\top A + A^\top B + B^\top A + B^\top B] \end{aligned}$$

$$\begin{aligned}
&= \text{tr}[A^\top A] + \text{tr}[B^\top B] + 2 \text{tr}[A^\top B] \\
&= \|A\|_F^2 + \|B\|_F^2 + 2 \text{tr}[A^\top B] \\
&= \|A\|_F^2 + \|B\|_F^2 + 2\langle A, B \rangle
\end{aligned}$$

where the first step follows from $\text{tr}[A^\top A] = \|A\|_F^2$ for matrix $A \in \mathbb{R}^{m \times n}$, the second step follows from the basic algebra, the third follows from $\text{tr}[X^\top Y] = \text{tr}[XY^\top]$ for matrices $X, Y \in \mathbb{R}^{m \times n}$, the fourth step follows from $\text{tr}[A^\top A] = \|A\|_F^2$ for matrix $A \in \mathbb{R}^{m \times n}$, and the last step follows from definition of inner product of matrices. \square

Lemma B.7. *Let $M \in \mathbb{R}^{n \times n}$. Let $X \in \mathbb{R}^{n \times n}$ be independent of M . We have*

$$\frac{d(M \circ X)}{dM_{i,j}} = X_{i,j} e_i e_j^\top$$

Proof. We can show

$$\begin{aligned}
\frac{d(M \circ X)}{dM_{i,j}} &= M \circ \frac{dX}{dM_{i,j}} + X \circ \frac{dM}{dM_{i,j}} \\
&= X \circ \frac{dM}{dM_{i,j}} \\
&= X \circ (e_i e_j^\top) \\
&= X_{i,j} e_i e_j^\top
\end{aligned}$$

where the first step, the second step, and the third step follow from Fact B.3, the fourth step follows from Fact B.4. \square

C GRADIENT CALCULATION

C.1 DEFINITIONS

In this section, we introduce some definitions we used to compute $\frac{d\mathcal{L}(M)}{dM}$. First, we introduce the exponential function.

Definition C.1 (Exponential function u, \tilde{u}). *If the following conditions hold*

- Let $X \in \mathbb{R}^{n \times d}$.
- Let $W \in \mathbb{R}^{d \times d}$.
- Let $M \in [0, 1]^{d \times d}$.
- Let $i_0 \in [n]$.

We define $u \in \mathbb{R}^{n \times n}$ as follows

$$u := \exp(XWX^\top).$$

We define $\tilde{u}(M) \in \mathbb{R}^{n \times n}$ as follows

$$\tilde{u}(M) := \exp(X(M \circ W)X^\top).$$

We define i_0 -th row of $\tilde{u}(M)$ as follows

$$\tilde{u}(M)_{i_0} := \exp(X(M \circ W)X^\top)_{i_0}.$$

Then, we introduce the sum function.

Definition C.2 (Sum function of softmax $\alpha, \tilde{\alpha}$). *If the following conditions hold*

- Let $M \in [0, 1]^{d \times d}$.

- Let $M_c \in \{0, 1\}^{n \times n}$ be the causal attention mask defined in Definition 3.1.
- Let $u, \tilde{u}(M)$ be defined as Definition C.1.
- Let $i_0 \in [n]$.

We define $\alpha \in \mathbb{R}^n$ as follows

$$\alpha := (u \circ M_c) \cdot \mathbf{1}_n.$$

We define $\tilde{\alpha}(M) \in \mathbb{R}^n$ as follows

$$\tilde{\alpha}(M) := (\tilde{u}(M) \circ M_c) \cdot \mathbf{1}_n.$$

We define i_0 -th entry of $\tilde{\alpha}(M)$ as follows

$$\tilde{\alpha}(M)_{i_0} := \langle (\tilde{u}(M) \circ M_c)_{i_0}, \mathbf{1}_n \rangle$$

Then, we introduce the Softmax probability function.

Definition C.3 (Softmax probability function f, \tilde{f}). *If the following conditions hold*

- Let $M \in [0, 1]^{d \times d}$.
- Let $M_c \in \{0, 1\}^{n \times n}$ be the causal attention mask defined in Definition 3.1.
- Let $u, \tilde{u}(M)$ be defined as Definition C.1.
- Let $\alpha, \tilde{\alpha}(M)$ be defined as Definition C.2.
- Let $i_0, j_0 \in [n]$.

We define $f \in \mathbb{R}^{n \times n}$ for each $j \in [n]$ as follows

$$f := \text{diag}(\alpha)^{-1} (u \circ M_c).$$

We define $\tilde{f}(M) \in \mathbb{R}^{n \times n}$ for each $j \in [n]$ as follows

$$\tilde{f}(M) := \text{diag}(\tilde{\alpha}(M))^{-1} (\tilde{u}(M) \circ M_c).$$

We define i_0 -th row of $\tilde{f}(M)$ as follows

$$\tilde{f}(M)_{i_0} := \tilde{\alpha}(M)_{i_0}^{-1} (\tilde{u}(M) \circ M_c)_{i_0}.$$

We define the entry in i_0 -th row, j_0 -th column of $\tilde{f}(M)$ as follows

$$\tilde{f}(M)_{i_0, j_0} := \tilde{\alpha}(M)_{i_0}^{-1} (\tilde{u}(M) \circ M_c)_{i_0, j_0}.$$

Then, we introduce the one-unit loss function.

Definition C.4 (One unit loss function c). *If the following conditions hold*

- Let f, \tilde{f} be defined in Definition C.3.
- Let $M \in [0, 1]^{d \times d}$.
- Let $i_0, j_0 \in [n]$.

We define $c(M) \in \mathbb{R}^{n \times n}$ as follows

$$c(M) := \tilde{f}(M) - f$$

We define i_0 -th row of $c(M)$ as follows

$$c(M)_{i_0} := \tilde{f}(M)_{i_0} - f_{i_0}$$

We define j_0 -th column of $c(M)$ as follows

$$c(M)_{*, j_0} := \tilde{f}(M)_{*, j_0} - f_{*, j_0}$$

We define the entry in i_0 -th row, j_0 -th column of $c(M)$ as follows

$$c(M)_{i_0, j_0} := \tilde{f}(M)_{i_0, j_0} - f_{i_0, j_0}$$

Then, we introduce the reconstruction error.

Definition C.5 (Reconstruction Error $\mathcal{L}_{\text{attn}}$). *If the following conditions hold*

- Let $M \in [0, 1]^{d \times d}$.
- Let $c(M)$ be defined in Definition C.4.

We define $\mathcal{L}_{\text{attn}}(M) \in \mathbb{R}$ as follows

$$\mathcal{L}_{\text{attn}}(M) := \frac{1}{2} \|c(M)\|_F^2 = \frac{1}{2} \sum_{i_0=1}^n \sum_{j_0=1}^n c(M)_{i_0, j_0}^2.$$

Then, we introduce the regularization term.

Definition C.6 (Regularization Term \mathcal{L}_{reg}). *If the following conditions hold*

- $M \in [0, 1]^{d \times d}$.

We define $\mathcal{L}_{\text{reg}}(M) \in \mathbb{R}$ as follows

$$\mathcal{L}_{\text{reg}}(M) := \frac{1}{2} \lambda \|M\|_F^2.$$

Finally, we introduce the overall loss function.

Definition C.7 (Overall loss function \mathcal{L}). *If the following conditions hold*

- Let $M \in [0, 1]^{d \times d}$.
- Let $\mathcal{L}_{\text{attn}}(M)$ be defined in Definition C.5.
- Let $\mathcal{L}_{\text{reg}}(M)$ be defined in Definition C.6.
- Let $\lambda \in \mathbb{R}_+$ be the regularization parameter.

We define $\mathcal{L}(M)$ as follows

$$\mathcal{L}(M) := \mathcal{L}_{\text{attn}}(M) + \mathcal{L}_{\text{reg}}(M)$$

C.2 GRADIENT FOR EACH ROW OF $X(M \circ W)X^\top$

We introduce the Lemma of gradient for each row of $X(M \circ W)X^\top$.

Lemma C.8. *Let $i_1 \in [d]$, $j_1 \in [d]$, $i_0 \in [n]$, we have*

$$\underbrace{\frac{d(X(M \circ W)X^\top)_{i_0}}{dM_{i_1, j_1}}}_{n \times 1} = \underbrace{W_{i_1, j_1}}_{\text{scalar}} \underbrace{X_{i_0, i_1}}_{\text{scalar}} \underbrace{X_{*, j_1}}_{n \times 1}$$

Proof. We can simplify the derivative expression

$$\begin{aligned} \frac{d(X(M \circ W)X^\top)_{i_0}}{dM_{i_1, j_1}} &= \frac{dX(X(M \circ W))_{i_0}}{dM_{i_1, j_1}} \\ &= \frac{dX(M \circ W)^\top X_{i_0}}{dM_{i_1, j_1}} \\ &= X \frac{d(M \circ W)^\top}{dM_{i_1, j_1}} X_{i_0} \end{aligned} \tag{2}$$

where the first and second step follows from Fact B.1, the third step follows from Fact B.3.

We further compute Eq. (2):

$$\frac{d(M \circ W)^\top}{dM_{i_1, j_1}} = \frac{dM^\top \circ W^\top}{dM_{i_1, j_1}}$$

$$\begin{aligned}
&= \frac{dM^\top \circ W^\top}{d(M^\top)_{j_1, i_1}} \\
&= (W^\top)_{j_1, i_1} e_{j_1} e_{i_1}^\top \\
&= W_{i_1, j_1} e_{j_1} e_{i_1}^\top
\end{aligned} \tag{3}$$

where the first follows from Fact B.4, the second step follows from for any matrix X , $X_{i,j} = (X^\top)_{j,i}$, the third step follows from Fact B.7, and the fourth step follows from for any matrix X , $X_{i,j} = (X^\top)_{j,i}$.

Finally, we have

$$\begin{aligned}
\frac{d(X(M \circ W)X^\top)_{i_0}}{dM_{i_1, j_1}} &= XW_{i_1, j_1} e_{j_1} e_{i_1}^\top X_{i_0} \\
&= W_{i_1, j_1} (X e_{j_1}) (e_{i_1}^\top X_{i_0}) \\
&= W_{i_1, j_1} X_{*, j_1} X_{i_0, i_1}
\end{aligned}$$

where the first step follows from Eq. (2) and Eq. (3), and the second step and the third step follow from basic algebra. \square

We introduce the Lemma of the gradient for each row of $\tilde{u}(M)$.

C.3 GRADIENT FOR EACH ROW OF $\tilde{u}(M)$

Lemma C.9. *If the following conditions hold:*

- Let $\tilde{u}(M)$ be defined in Definition C.1.

Let $i_1 \in [d]$, $j_1 \in [d]$, $i_0 \in [n]$, we have

$$\frac{d\tilde{u}(M)_{i_0}}{dM_{i_1, j_1}} = \underbrace{\tilde{u}(M)_{i_0}}_{n \times 1} \circ \underbrace{(W_{i_1, j_1})}_{\text{scalar}} \underbrace{(X_{i_0, i_1})}_{\text{scalar}} \underbrace{(X_{*, j_1})}_{n \times 1}$$

Proof. We have

$$\begin{aligned}
\frac{d\tilde{u}(M)_{i_0}}{dM_{i_1, j_1}} &= \frac{d \exp(X(M \circ W)X^\top)_{i_0}}{dM_{i_1, j_1}} \\
&= \exp(\underbrace{X}_{n \times d} \underbrace{(M \circ W)}_{d \times d} \underbrace{X^\top}_{d \times 1})_{i_0} \circ \underbrace{\frac{d(X(M \circ W)X^\top)_{i_0}}{dM_{i_1, j_1}}}_{n \times 1} \\
&= \underbrace{\tilde{u}(M)_{i_0}}_{n \times 1} \circ \underbrace{\frac{d(X(M \circ W)X^\top)_{i_0}}{dM_{i_1, j_1}}}_{n \times 1} \\
&= \underbrace{\tilde{u}(M)_{i_0}}_{n \times 1} \circ \underbrace{(W_{i_1, j_1})}_{\text{scalar}} \underbrace{(X_{i_0, i_1})}_{\text{scalar}} \underbrace{(X_{*, j_1})}_{n \times 1}
\end{aligned}$$

where the first step follows from Definition C.1, the second step follows from Fact B.3, the third step follows from Definition C.1, and the fourth step follows from Lemma C.8. \square

C.4 GRADIENT FOR EACH ENTRY OF $\tilde{\alpha}(M)$

We introduce the Lemma of gradient for each entry of $\tilde{\alpha}(M)$.

Lemma C.10. *If the following conditions hold:*

- Let $\tilde{u}(M)$ be defined in Definition C.1.
- Let $\tilde{\alpha}(M)$ be defined in Definition C.2.

Let $i_1 \in [d]$, $j_1 \in [d]$, $i_0 \in [n]$, we have

$$\underbrace{\frac{d\tilde{\alpha}(M)_{i_0}}{dM_{i_1, j_1}}}_{\text{scalar}} = \langle \tilde{u}(M)_{i_0} \circ (M_c)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle$$

Proof. We have

$$\begin{aligned} \underbrace{\frac{d\tilde{\alpha}(M)_{i_0}}{dM_{i_1, j_1}}}_{\text{scalar}} &= \underbrace{\frac{d\langle (\tilde{u}(M) \circ M_c)_{i_0}, \mathbf{1}_n \rangle}{dM_{i_1, j_1}}}_{\text{scalar}} \\ &= \left\langle \frac{d(\tilde{u}(M) \circ M_c)_{i_0}}{dM_{i_1, j_1}}, \mathbf{1}_n \right\rangle \\ &= \left\langle \frac{d\tilde{u}(M)_{i_0}}{dM_{i_1, j_1}} \circ (M_c)_{i_0}, \mathbf{1}_n \right\rangle \\ &= \langle \tilde{u}(M)_{i_0} \circ (W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1}) \circ (M_c)_{i_0}, \mathbf{1}_n \rangle \\ &= \langle \tilde{u}(M)_{i_0} \circ (M_c)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle \end{aligned}$$

where the first step follows from Definition C.2, the second step follows from the product rule of inner product in Fact B.3, the third step follows from the product rule of Hadamard product in Fact B.3, the fourth step follows from Lemma C.9, and the last step follows from Fact B.4. \square

C.5 GRADIENT FOR EACH ENTRY OF $\tilde{f}(M)$

We introduce the Lemma of the gradient for each entry of $\tilde{f}(M)$.

Lemma C.11. *If the following conditions hold:*

- Let $\tilde{u}(M)$ be defined in Definition C.1.
- Let $\tilde{\alpha}(M)$ be defined in Definition C.2.
- Let $\tilde{f}(M)$ be defined in Definition C.3.

Let $i_1 \in [d]$, $j_1 \in [d]$, $i_0 \in [n]$, $j_0 \in [n]$, we have

$$\frac{d\tilde{f}(M)_{i_0, j_0}}{dM_{i_1, j_1}} = \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} - \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle$$

Proof. We have

$$\begin{aligned} \frac{d\tilde{f}(M)_{i_0, j_0}}{dM_{i_1, j_1}} &= \frac{d\tilde{\alpha}(M)_{i_0}^{-1} (\tilde{u}(M) \circ M_c)_{i_0, j_0}}{dM_{i_1, j_1}} \\ &= \frac{d\tilde{\alpha}(M)_{i_0}^{-1}}{dM_{i_1, j_1}} (\tilde{u}(M) \circ M_c)_{i_0, j_0} + \frac{d(\tilde{u}(M) \circ M_c)_{i_0, j_0}}{dM_{i_1, j_1}} \tilde{\alpha}(M)_{i_0}^{-1} \end{aligned} \quad (4)$$

where the first step follows from Definition C.3, and the second step follows from Fact B.3.

In the following part, we compute the two terms separately.

For the first term above, we have

$$\frac{d\tilde{\alpha}(M)_{i_0}^{-1}}{dM_{i_1, j_1}} (\tilde{u}(M) \circ M_c)_{i_0, j_0}$$

$$\begin{aligned}
&= (\tilde{u}(M) \circ M_c)_{i_0, j_0} (-1) \tilde{\alpha}(M)_{i_0}^{-2} \frac{d\tilde{\alpha}(M)_{i_0}}{dM_{i_1, j_1}} \\
&= - (\tilde{u}(M) \circ M_c)_{i_0, j_0} \langle \tilde{u}(M)_{i_0} \circ (M_c)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle / \tilde{\alpha}(M)_{i_0}^2 \\
&= - (\tilde{\alpha}(M)_{i_0}^{-1} (M_c)_{i_0, j_0} \tilde{u}(M)_{i_0, j_0}) \langle \tilde{\alpha}(M)_{i_0}^{-1} \tilde{u}(M)_{i_0} \circ (M_c)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle \\
&= - \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle
\end{aligned} \tag{5}$$

where the first step follows from Fact B.3, the second step follows from Lemma C.10, the third step follows from basic algebra, and the fourth step follows from Definition C.3.

For the second term above, we have

$$\begin{aligned}
&\frac{d(\tilde{u}(M) \circ M_c)_{i_0, j_0} \tilde{\alpha}(M)_{i_0}^{-1}}{dM_{i_1, j_1}} \\
&= \frac{d\tilde{u}(M)_{i_0, j_0} (M_c)_{i_0, j_0} \tilde{\alpha}(M)_{i_0}^{-1}}{dM_{i_1, j_1}} \\
&= (M_c)_{i_0, j_0} \left(\frac{d\tilde{u}(M)_{i_0}}{dM_{i_1, j_1}} \right)_{j_0} \tilde{\alpha}(M)_{i_0}^{-1} \\
&= ((M_c)_{i_0, j_0} \tilde{u}(M)_{i_0, j_0} \tilde{\alpha}(M)_{i_0}^{-1}) W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} \\
&= \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1}
\end{aligned} \tag{6}$$

where the first step and the second step follow from basic algebra, the third step follows from Lemma C.9, and the fourth step follows from Definition C.3.

So, we have

$$\begin{aligned}
\frac{d\tilde{f}(M)_{i_0, j_0}}{dM_{i_1, j_1}} &= \frac{d\tilde{\alpha}(M)_{i_0}^{-1}}{dM_{i_1, j_1}} (\tilde{u}(M) \circ M_c)_{i_0, j_0} + \frac{d(\tilde{u}(M) \circ M_c)_{i_0, j_0} \tilde{\alpha}(M)_{i_0}^{-1}}{dM_{i_1, j_1}} \\
&= \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} - \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle
\end{aligned}$$

where the first step follows from Eq. (4), and the second step follows from Eq. (5) and Eq. (6). \square

C.6 GRADIENT FOR EACH ENTRY OF $c(M)$

We introduce the Lemma of gradient for each entry of $c(M)$.

Lemma C.12. *If the following conditions hold:*

- Let $\tilde{f}(M)$, f be defined in Definition C.3.
- Let $c(M)$ be defined in Definition C.4.

Let $i_1 \in [d]$, $j_1 \in [d]$, $i_0 \in [n]$, $j_0 \in [n]$, we have

$$\frac{dc(M)_{i_0, j_0}}{dM_{i_1, j_1}} = \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} - \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle$$

Proof. We have

$$\begin{aligned}
\frac{dc(M)_{i_0, j_0}}{dM_{i_1, j_1}} &= \frac{d(\tilde{f}(M)_{i_0, j_0} - f_{i_0, j_0})}{dM_{i_1, j_1}} \\
&= \frac{d\tilde{f}(M)_{i_0, j_0}}{dM_{i_1, j_1}} \\
&= \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} - \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle
\end{aligned}$$

where the first step follows from Definition C.4, the second step follows from Fact B.3, and the third step follows from Lemma C.11. \square

C.7 GRADIENT FOR $\mathcal{L}_{\text{attn}}(M)$

We introduce the Lemma of the gradient for $\mathcal{L}_{\text{attn}}(M)$.

Lemma C.13. *If the following conditions hold:*

- Let $\tilde{f}(M)$ be defined in Definition C.3.
- Let $c(M)$ be defined in Definition C.4.
- Let $\mathcal{L}_{\text{attn}}(M)$ be defined in Definition C.5.

Let $i_1 \in [d]$, $j_1 \in [d]$, $i_0 \in [n]$, $j_0 \in [n]$, we have

$$\frac{d\mathcal{L}_{\text{attn}}(M)}{dM_{i_1, j_1}} = \sum_{i_0=1}^n \sum_{j_0=1}^n B_1(M) + B_2(M)$$

where we have definitions:

- $B_1(M) := c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1}$
- $B_2(M) := -c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle$

Proof. We have

$$\begin{aligned} \frac{d\mathcal{L}_{\text{attn}}(M)}{dM_{i_1, j_1}} &= \frac{1}{2} \frac{d\|c(M)\|_F^2}{dM_{i_1, j_1}} \\ &= \frac{1}{2} \frac{d \sum_{i_0=1}^n \sum_{j_0=1}^n (c(M)_{i_0, j_0})^2}{dM_{i_1, j_1}} \\ &= \frac{1}{2} \sum_{i_0=1}^n \sum_{j_0=1}^n \frac{d(c(M)_{i_0, j_0})^2}{dM_{i_1, j_1}} \\ &= \sum_{i_0=1}^n \sum_{j_0=1}^n c(M)_{i_0, j_0} \frac{dc(M)_{i_0, j_0}}{dM_{i_1, j_1}} \end{aligned}$$

where the first step follows from Definition C.5, the second step follows from the definition of Frobenius's norm of the matrix, the third step follows from Fact B.3, and the fourth step follows from Fact B.3.

Following Lemma C.12, we have

$$\begin{aligned} &\sum_{i_0=1}^n \sum_{j_0=1}^n c(M)_{i_0, j_0} \frac{dc(M)_{i_0, j_0}}{dM_{i_1, j_1}} \\ &= \sum_{i_0=1}^n \sum_{j_0=1}^n c(M)_{i_0, j_0} (\tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} - \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle) \\ &= c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} - c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle \\ &:= \sum_{i_0=1}^n \sum_{j_0=1}^n B_1(M) + B_2(M) \end{aligned}$$

where the second step follows from basic algebra. \square

C.8 GRADIENT FOR $\mathcal{L}_{\text{reg}}(M)$

We introduce the Lemma of the gradient for $\mathcal{L}_{\text{reg}}(M)$.

Lemma C.14. *If the following conditions hold:*

- Let $\mathcal{L}_{\text{reg}}(M)$ be defined in Definition C.6.

Let $i_1 \in [d]$, $j_1 \in [d]$, we have

$$\frac{d\mathcal{L}_{\text{reg}}(M)}{dM_{i_1, j_1}} = B_3(M)$$

where we have the definition:

- $B_3(M) := \lambda M_{i_1, j_1}$

Proof. We have

$$\begin{aligned} \frac{d\mathcal{L}_{\text{reg}}(M)}{dM_{i_1, j_1}} &= \frac{1}{2} \lambda \frac{d\|M\|_F^2}{dM_{i_1, j_1}} \\ &= \frac{1}{2} \lambda \left(\frac{d}{dM_{i_1, j_1}} \sum_{i_0=1}^d \sum_{j_0=1}^d M_{i_0, j_0}^2 \right) \\ &= \lambda M_{i_1, j_1} \\ &:= B_3(M) \end{aligned}$$

where the first step follows from Definition C.6, the second step follows from the definition of Frobenius's norm of the matrix, and the third step follows from Fact B.3. \square

C.9 GRADIENT FOR $\mathcal{L}(M)$

We introduce the Lemma of the gradient for $\mathcal{L}(M)$.

Lemma C.15. *If the following conditions hold:*

- Let $\tilde{u}(M)$ be defined in Definition C.1.
- Let $\tilde{\alpha}(M)$ be defined in Definition C.2.
- Let $\tilde{f}(M)$ be defined in Definition C.3.
- Let $\mathcal{L}_{\text{attn}}(M)$ be defined in Definition C.5.
- Let $\mathcal{L}_{\text{reg}}(M)$ be defined in Definition C.6.
- Let $\mathcal{L}(M)$ be defined in Definition C.7.

Let $i_1 \in [d]$, $j_1 \in [d]$, $i_0 \in [n]$, $j_0 \in [n]$, we have

$$\frac{d\mathcal{L}(M)}{dM_{i_1, j_1}} = \sum_{i_0=1}^n \sum_{j_0=1}^n (B_1(M) + B_2(M)) + B_3(M)$$

where we have definitions:

- $B_1(M) := c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1}$
- $B_2(M) := -c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle$
- $B_3(M) := \lambda M_{i_1, j_1}$

Proof.

$$\begin{aligned} \frac{d\mathcal{L}(M)}{dM_{i_1, j_1}} &= \frac{d\mathcal{L}_{\text{attn}}(M) + \mathcal{L}_{\text{reg}}(M)}{dM_{i_1, j_1}} \\ &= \frac{d\mathcal{L}_{\text{attn}}(M)}{dM_{i_1, j_1}} + \frac{d\mathcal{L}_{\text{reg}}(M)}{dM_{i_1, j_1}} \end{aligned}$$

$$= \sum_{i_0=1}^n \sum_{j_0=1}^n (B_1(M) + B_2(M)) + B_3(M)$$

where the first step follows from Definition C.7, the second step follows from Fact B.3, and the third step follows from Lemma C.13 and Lemma C.14. \square

D MATRIX FORM

D.1 MATRIX FORM OF $B(M)$

Given the matrix form, we define p to simplify the calculation.

Definition D.1. *If the following conditions hold*

- Let $X \in \mathbb{R}^{n \times d}$.
- Let $M \in [0, 1]^{d \times d}$.
- Let $W \in \mathbb{R}^{d \times d}$.
- Let $c(M)$ be defined in Definition C.4.
- Let $\tilde{f}(M)$ be defined in Definition C.3.

We define p_1 as follows

$$p_1 := c(M) \circ \tilde{f}(M)$$

We define the j_0 -th column of p_1 as follows

$$(p_1)_{*,j_0} := (c(M) \circ \tilde{f}(M))_{*,j_0}$$

We define p_2 as follows

$$p_2 := \text{diag}(p_1 \cdot \mathbf{1}_n) \tilde{f}(M)$$

We define the i_0 -th row of p_2 as follows

$$(p_2)_{i_0} := \mathbf{1}_n^\top (p_1)_{i_0} \tilde{f}(M)_{i_0} = \tilde{f}(M)_{i_0} c(M)^\top \tilde{f}(M)_{i_0}$$

We define p as follows

$$p := p_1 - p_2 = c(M) \circ \tilde{f}(M) - \text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)$$

We introduce the matrix view of $B_1(M)$ and its summation.

Lemma D.2 (Matrix view of $B_1(M)$). *If we have the below conditions,*

- Let $B_1(M, i_1, j_1) := c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1}^\top$, which is defined in Lemma C.15
- We define $C_1(M) \in \mathbb{R}^{d \times d}$. For all $i_1, j_1 \in [d]$, let $C_1(i_1, j_1)$ denote the (i_1, j_1) -th entry of $C_1(M)$. We define $C_1(i_1, j_1) = B_1(M, i_1, j_1)$.

Then, we can show that

- Part 1. For $i_0, j_0 \in [n]$

$$C_1(M) = \underbrace{c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0}}_{\text{scalar}} (W \circ (X_{i_0} X_{j_0}^\top))$$

- Part 2.

$$\sum_{i_0=1}^n \sum_{j_0=1}^n C_1(M) = W \circ (X^\top p_1 X)$$

Proof. Part 1. We have

$$\begin{aligned}
C_1(i_1, j_1) &= c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} \\
&= \underbrace{c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0}}_{\text{scalar}} \underbrace{(X_{i_0})_{i_1}}_{d \times 1} \underbrace{(W_{*, j_1})_{i_1}}_{d \times 1} \underbrace{(X_{j_0})_{j_1}}_{\text{scalar}} \\
&= \underbrace{c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0}}_{\text{scalar}} \underbrace{(\text{diag}(X_{i_0}) W_{*, j_1})_{i_1}}_{d \times 1} \underbrace{(X_{j_0})_{j_1}}_{\text{scalar}}
\end{aligned}$$

where the first step follows from the definition of C_1 , the second step follows from Fact B.1, and the third step follows from Fact B.2.

Following from Fact B.1, we can get j_1 -th column of C_1

$$\begin{aligned}
C_1(*, j_1) &= \underbrace{c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0}}_{\text{scalar}} \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{W_{*, j_1}}_{d \times 1} \underbrace{(X_{j_0})_{j_1}}_{\text{scalar}} \\
&= \underbrace{c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0}}_{\text{scalar}} \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{(W \text{diag}(X_{j_0}))_{*, j_1}}_{d \times d}
\end{aligned}$$

where the second step follows from Fact B.2.

Following from Fact B.1, we can get $C_1(M)$

$$\begin{aligned}
C_1(M) &= \underbrace{c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0}}_{\text{scalar}} \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{W}_{d \times d} \underbrace{\text{diag}(X_{j_0})}_{d \times d} \\
&= \underbrace{c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0}}_{\text{scalar}} (W \circ (X_{i_0} X_{j_0}^\top)) \tag{7}
\end{aligned}$$

where the second step follows from Fact B.4.

Part 2. We further compute the summation of $C_1(M)$.

$$\begin{aligned}
\sum_{i_0=1}^n \sum_{j_0=1}^n C_1(M) &= \sum_{i_0=1}^n \sum_{j_0=1}^n \underbrace{c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0}}_{\text{scalar}} (W \circ X_{i_0} X_{j_0}^\top) \\
&= \sum_{i_0=1}^n \sum_{j_0=1}^n (W \circ (c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} X_{i_0} X_{j_0}^\top)) \\
&= W \circ \sum_{i_0=1}^n \sum_{j_0=1}^n c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} X_{i_0} X_{j_0}^\top \\
&= W \circ \sum_{j_0=1}^n \sum_{i_0=1}^n ((p_1)_{*, j_0})_{i_0} X_{i_0} X_{j_0}^\top
\end{aligned}$$

where the first step follows from Eq. (7), the second step follows from basic algebra, the third step follows from Fact B.4, and the fourth step follows from Definition D.1.

Then following from Fact B.2, we have

$$\begin{aligned}
&W \circ \sum_{j_0=1}^n \sum_{i_0=1}^n ((p_1)_{*, j_0})_{i_0} X_{i_0} X_{j_0}^\top \\
&= W \circ \sum_{j_0=1}^n X^\top (p_1)_{*, j_0} X_{j_0}^\top \\
&= W \circ (X^\top p_1 X)
\end{aligned}$$

□

We introduce the matrix view of $B_2(M)$ and its summation.

Lemma D.3 (Matrix view of $B_2(M)$). *If we have the below conditions,*

- Let $B_2(M, i_1, j_1) := -c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle$ be defined in Lemma C.15.
- We define $C_2(M) \in \mathbb{R}^{d \times d}$. For all $i_1, j_1 \in [d]$, let $C_2(i_1, j_1)$ denote the (i_1, j_1) -th entry of $C_2(M)$. We define $C_2(i_1, j_1) = B_2(M, i_1, j_1)$.

Then, we can show that

- Part 1. For $i_0, j_0 \in [n]$

$$C_2(M) = -c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{W}_{d \times d} \underbrace{\text{diag}(X^\top \tilde{f}(M)_{i_0})}_{d \times d}$$

- Part 2.

$$\sum_{i_0=1}^n \sum_{j_0=1}^n C_2(M) = -W \circ (X^\top p_2 X)$$

Proof. Part 1. We have

$$\begin{aligned} -C_2(i_1, j_1) &= c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle \\ &= c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \underbrace{\tilde{f}(M)_{i_0}^\top}_{1 \times n} \underbrace{X_{*, j_1}^\top W_{i_1, j_1} X_{i_0, i_1}}_{n \times 1} \\ &= c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \underbrace{\tilde{f}(M)_{i_0}^\top}_{1 \times n} \underbrace{X_{*, j_1}^\top (X_{i_0})_{i_1} (W_{*, j_1})_{i_1}}_{n \times 1} \\ &= c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \underbrace{\tilde{f}(M)_{i_0}^\top}_{1 \times n} \underbrace{X_{*, j_1}^\top (\text{diag}(X_{i_0}) W_{*, j_1})_{i_1}}_{n \times 1} \end{aligned}$$

where the first step follows from the definition of C_2 , the second step, the third step, and the fourth step follow from Fact B.4.

Following from Fact B.1, we can get j_1 -th column of C_2

$$\begin{aligned} -C_2(*, j_1) &= \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{W_{*, j_1}}_{d \times 1} \underbrace{c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \tilde{f}(M)_{i_0}^\top}_{1 \times n} \underbrace{X_{*, j_1}}_{n \times 1} \\ &= c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{W_{*, j_1}}_{d \times 1} \underbrace{\tilde{f}(M)_{i_0}^\top}_{1 \times n} \underbrace{X_{*, j_1}}_{n \times 1} \\ &= c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{W_{*, j_1}}_{d \times 1} \underbrace{(X^\top \tilde{f}(M)_{i_0})_{j_1}}_{\text{scalar}} \\ &= c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{(W \text{diag}(X^\top \tilde{f}(M)_{i_0}))_{*, j_1}}_{d \times d} \end{aligned}$$

where the second step and the fourth step follows from Fact B.4, and the third step follows from Fact B.1.

Following from Fact B.1, we can get C_2 .

$$-C_2(M) = c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{W}_{d \times d} \underbrace{\text{diag}(X^\top \tilde{f}(M)_{i_0})}_{d \times d} \quad (8)$$

Part 2. We further compute the summation of C_2

$$-\sum_{i_0=1}^n \sum_{j_0=1}^n C_2(M) = \sum_{i_0=1}^n \sum_{j_0=1}^n c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} \underbrace{\text{diag}(X_{i_0})}_{d \times d} \underbrace{W}_{d \times d} \underbrace{\text{diag}(X^\top \tilde{f}(M)_{i_0})}_{d \times d}$$

$$\begin{aligned}
&= \sum_{i_0=1}^n \sum_{j_0=1}^n c(M)_{i_0, j_0} \tilde{f}(M)_{i_0, j_0} ((X_{i_0} \tilde{f}(M)_{i_0}^\top X) \circ W) \\
&= W \circ \sum_{i_0=1}^n (X_{i_0} \tilde{f}(M)_{i_0}^\top X) \sum_{j_0=1}^n ((p_1)_{i_0})_{j_0}
\end{aligned}$$

where the first step follows from Eq. (8), the second step and the third step follow from Fact B.4.

Following from Fact B.2, we have

$$\begin{aligned}
&W \circ \sum_{i_0=1}^n (X_{i_0} \tilde{f}(M)_{i_0}^\top X) \sum_{j_0=1}^n ((p_1)_{i_0})_{j_0} \\
&= W \circ \sum_{i_0=1}^n (X_{i_0} \tilde{f}(M)_{i_0}^\top X) \mathbf{1}_n^\top (p_1)_{i_0} \\
&= W \circ (X^\top \text{diag}(p_1 \cdot \mathbf{1}_n) \tilde{f}(M) X) \\
&= W \circ (X^\top p_2 X)
\end{aligned}$$

where the third step follows from Definition D.1. \square

We introduce the matrix view of $B_3(M)$.

Lemma D.4 (Matrix view of $B_3(M)$). *If the following conditions hold*

- Let $B_3(M, i_1, j_1) := \lambda M_{i_1, j_1}$ be defined in Lemma C.15.
- We define $C_3(M) \in \mathbb{R}^{d \times d}$. For all $i_1, j_1 \in [d]$, let $C_3(i_1, j_1)$ denote the (i_1, j_1) -th entry of $C_3(M)$. We define $C_3(i_1, j_1) = B_3(M, i_1, j_1)$.

We can show that

$$C_3(M) = \lambda M.$$

Proof. The proof is straightforward. By the definition of $C_3(M)$, for all $i_1, j_1 \in [d]$, the (i_1, j_1) -th entry of $C_3(M)$ is given by $C_3(i_1, j_1) = B_3(M, i_1, j_1) = \lambda M_{i_1, j_1}$. Thus, the entire matrix $C_3(M)$ has entries that correspond to those of λM . Therefore, we can conclude that $C_3(M) = \lambda M$ as required. \square

D.2 MATRIX FORM OF $\frac{d}{dM} \mathcal{L}(M)$

We introduce the matrix form of the overall loss function.

Theorem D.5 (Close form of gradient, formal version of Theorem 5.3). *If the following conditions hold*

- Let $\mathcal{L}(M)$ be defined in Definition C.7.
- Let p be defined in Definition D.1.
- Let $X \in \mathbb{R}^{n \times d}$.
- Let $M \in [0, 1]^{d \times d}$.
- Let $W \in \mathbb{R}^{d \times d}$.

We can show that

$$\frac{d\mathcal{L}(M)}{dM} = W \circ (X^\top p X) + \lambda M.$$

Proof. We have

$$\begin{aligned} \frac{d\mathcal{L}(M)}{dM} &= \sum_{i_0=1}^n \sum_{j_0=1}^n (C_1(M) + C_2(M)) + C_3(M) \\ &= W \circ (X^\top p_1 X) - W \circ (X^\top p_2 X) + \lambda M \\ &= W \circ (X^\top (p_1 - p_2) X) + \lambda M \\ &= W \circ (X^\top p X) + \lambda M \end{aligned}$$

where the first step follows from Lemma C.15, the second step follows from Lemma D.2, Lemma D.3, and Lemma D.4, the third step follows from basic algebra, and the fourth step follows from Definition D.1. \square

E BOUNDS FOR BASIC FUNCTIONS

E.1 BASIC ASSUMPTIONS

Here, we introduce our bounded parameters assumption.

Assumption E.1 (Bounded parameters). *We assume the following conditions*

- Let R be some fixed constant satisfies $R > 1$.
- Let $X \in \mathbb{R}^{n \times d}$, $W \in \mathbb{R}^{d \times d}$. We have $\|X\|_F \leq R$ and $\|W\|_F \leq R$.

Here, we present the lemma of bounds for M and M_c .

Lemma E.2 (Bounds for M and M_c). *Let $M \in [0, 1]^{d \times d}$ and $M_c \in \{0, 1\}^{n \times n}$ be the causal attention mask defined in Definition 3.1. For M , we have*

$$\|M\|_F \leq d$$

For M_c , we have

$$\|M_c\|_F \leq n$$

Proof. This Lemma simply follows from the definition of the Frobenius norm, given that the max value of each entry in M and M_c is 1. \square

E.2 BOUNDS FOR BASIC FUNCTIONS

We first introduce the lemma of bounds for basic functions.

Lemma E.3. *Under Assumption E.1, for all $i_0 \in [n]$, $j_0 \in [n]$, $i_1 \in [d]$, $j_1 \in [d]$, we have the following bounds*

- Part 1.

$$\|\tilde{f}(M)\|_F \leq \sqrt{n}$$

- Part 2.

$$\|c(M)\|_F \leq 2\sqrt{n}$$

- Part 3.

$$\|(c(M) \circ \tilde{f}(M))\|_F \leq 2\sqrt{n}$$

- Part 4.

$$|\tilde{f}(M)_{i_0, j_0}| \leq 1$$

- *Part 5.*

$$|W_{i_1, j_1}| \leq R$$

- *Part 6.*

$$|X_{i_0, i_1}| \leq R$$

- *Part 7.*

$$\|\tilde{f}(M)_{i_0}\|_2 \leq 1$$

- *Part 8.*

$$|\tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1}| \leq R^3$$

- *Part 9.*

$$|\tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle| \leq R^3$$

- *Part 10.*

$$\|\text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n)\|_F \leq 2n$$

Proof. **Proof of Part 1.** Each entry in $\tilde{f}(M)$ present a probability, thus for $i_0 \in [n]$, $j_0 \in [n]$, we have

$$0 \leq \tilde{f}(M)_{i_0, j_0} \leq 1.$$

For any i_0 -th row of $\tilde{f}(M)$, following from the definition of Softmax function, we know

$$\sum_{j_0=1}^n \tilde{f}(M)_{i_0, j_0} = 1.$$

So we have

$$\sum_{j_0=1}^n \tilde{f}(M)_{i_0, j_0}^2 \leq 1$$

which follows from $\tilde{f}(M)_{i_0, j_0} \leq (\tilde{f}(M)_{i_0, j_0})^2$. Then, we can show

$$\|\tilde{f}(M)\|_F = \sqrt{\sum_{i_0=1}^n \sum_{j_0=1}^n \tilde{f}(M)_{i_0, j_0}^2} \leq \sqrt{n}$$

Proof of Part 2. Following from **Part 1.**, we can show

$$\|\tilde{f}(M)\|_F \leq \sqrt{n}$$

and

$$\|f\|_F \leq \sqrt{n}.$$

Then we have

$$\begin{aligned} \|c(M)\|_F &= \|\tilde{f}(M) - f\|_F \\ &\leq \|\tilde{f}(M)\|_F + \|f\|_F \\ &\leq 2\sqrt{n} \end{aligned}$$

where the first step follows from Definition C.4, the second step follows the triangle inequality.

Proof of Part 3. We have $0 \leq \tilde{f}(M)_{i_0, j_0} \leq 1$, so we have

$$\|(c(M) \circ \tilde{f}(M))\|_F \leq \|c(M)\|_F$$

$$\leq 2\sqrt{n}$$

where the second step follows from **Part 2.**

Proof of Part 4. See **Proof of Part 1.**

Proof of Part 5. The proof simply follows from Assumption E.1 and Fact B.5.

Proof of Part 6. The proof simply follows from Assumption E.1 and Fact B.5.

Proof of Part 7. See **Proof of Part 1.**

Proof of Part 8. The proof simply follows from **Part 4.**, **Part 5.**, **Part 6.** and **Part 7.**

Proof of Part 9. The proof simply follows from **Part 4.**, **Part 5.**, **Part 6.** and **Part 7.**

Proof of Part 10. We have

$$\begin{aligned} \|\text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n)\|_F &= \|(c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n\|_2 \\ &\leq \|\mathbf{1}_n\|_2 \|(c(M) \circ \tilde{f}(M))\|_F \\ &= \sqrt{n} \cdot 2\sqrt{n} \\ &= 2n \end{aligned}$$

where the first step follows from Fact B.4 the second step follows from Fact B.5, the third step follows from **Part 3.**, and the last step follows from simple algebra. \square

E.3 BOUNDS FOR GRADIENT OF $\tilde{f}(M)$

We introduce the lemma of bounds for the gradient of $\tilde{f}(M)$.

Lemma E.4. *If the following conditions hold*

- Let $\tilde{f}(M)$ be defined in Definition C.3.
- Assumption E.1 holds.

Then we have

$$\left\| \frac{\text{d vec}(\tilde{f}(M))}{\text{d vec}(M)} \right\|_F \leq 2dnR^3$$

Proof. We have

$$\begin{aligned} & \left| \frac{\text{d}f(M)_{i_0, j_0}}{\text{d}M_{i_1, j_1}} \right| \\ &= \left| \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} - \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle \right| \\ &\leq \left| \tilde{f}(M)_{i_0, j_0} W_{i_1, j_1} X_{i_0, i_1} X_{j_0, j_1} \right| + \left| \tilde{f}(M)_{i_0, j_0} \langle \tilde{f}(M)_{i_0}, W_{i_1, j_1} X_{i_0, i_1} X_{*, j_1} \rangle \right| \\ &\leq 2R^3 \end{aligned}$$

For $\frac{\text{d vec}(\tilde{f}(M))}{\text{d vec}(M)}$, we can show

$$\begin{aligned} \left\| \frac{\text{d vec}(\tilde{f}(M))}{\text{d vec}(M)} \right\|_F &= \sqrt{\sum_{i_2=1}^{n^2} \sum_{j_2=1}^{d^2} \left| \frac{\text{d vec}(\tilde{f}(M))_{i_0}}{\text{d vec}(M)_{j_0}} \right|^2} \\ &\leq 2ndR^3 \end{aligned}$$

\square

F LIPSCHITZ OF GRADIENT

F.1 USEFUL FACTS

Here, we introduce the fact of the mean value theorem for matrix function.

Fact F.1 (Mean value theorem for matrix function, Fact C.6 in Liang et al. (2024e)). *If the following conditions hold*

- Let $X, Y \in C \subset \mathbb{R}^{d \times d}$ where C is an open convex domain.
- Let $g(X) : C \rightarrow \mathbb{R}^{n \times n}$ be a differentiable matrix function on C .
- Let $\|\frac{d \text{vec}(g(X))}{d \text{vec}(X)}\|_F \leq R$ for all $x \in C$.

We have

$$\|g(Y) - g(X)\|_F \leq R\|Y - X\|_F.$$

Proof. For the convenience of proof, we define x and y as follows:

- $x := \text{vec}(X)$ and $y := \text{vec}(Y)$.
- $h(x) := \text{vec}(g(X))$ and $h(y) := \text{vec}(g(Y))$.
- $h'(a)$ denotes a matrix which its (i, j) -th term is $\frac{dh(a)_i}{da_j}$.

Assume we have 1-variable function $\gamma(c) = f(x + c(y - x))$, we can apply Mean Value Theorem:

$$f(y) - f(x) = \gamma(1) - \gamma(0) = \gamma'(t)(1 - 0) = \nabla f(x + t(y - x))^\top (y - x) \quad (9)$$

where $t \in [0, 1]$. Let $G(c) := (h(y) - h(x))^\top h(c)$, we have

$$\begin{aligned} \|g(Y) - g(X)\|_F^2 &= G(y) - G(x) \\ &= \nabla G(x + t(y - x))^\top (y - x) \\ &= \underbrace{(h'(x + t(y - x)))}_{d^2 \times n^2} \cdot \underbrace{(h(y) - h(x))^\top}_{n^2 \times 1} \cdot (y - x) \\ &\leq \|h'(x + t(y - x))\| \cdot \|h(y) - h(x)\|_2 \cdot \|y - x\|_2 \end{aligned}$$

where the second step follows from Eq. (9), the third step follows from the chain rule, the fourth step follows from the Cauchy-Schwartz inequality.

By definition of matrix Frobenius norm and vector ℓ_2 norm, we have

$$\|g(Y) - g(X)\|_F = \|h(y) - h(x)\|_2$$

and

$$\|Y - X\|_F = \|y - x\|_2$$

so we can show

$$\|g(Y) - g(X)\|_F \leq R\|Y - X\|_F$$

which follows from $\|\frac{d \text{vec}(g(X))}{d \text{vec}(X)}\|_F \leq R$ for all $x \in C$. \square

Here, we introduce the fact of Lipschitz for the product of functions.

Fact F.2 (Lipschitz for product of functions, Fact H.2 in Deng et al. (2023)). *Under following conditions*

- Let $\{f_i(x)\}_{i=1}^n$ be a sequence of functions with the same domain and range.

- For each $i \in [n]$, we have
 - $f_i(x)$ is bounded: $\forall x, \|f_i(x)\|_F \leq R_i$ with $R_i \geq 1$.
 - $f_i(x)$ is Lipschitz continuous: $\forall x, y, \|f_i(x) - f_i(y)\|_F \leq L_i \|x - y\|_F$.

Then we have

$$\left\| \prod_{i=1}^n f_i(x) - \prod_{i=1}^n f_i(y) \right\|_F \leq 2^{n-1} \cdot \max_{i \in [n]} \{L_i\} \cdot \left(\prod_{i=1}^n R_i \right) \cdot \|x - y\|_F$$

F.2 LIPSCHITZ OF $\tilde{f}(M)$

We introduce the lemma about Lipschitz of $\tilde{f}(M)$.

Lemma F.3 (Lipschitz of $\tilde{f}(M)$). *Under the following conditions*

- Assumption E.1 holds.
- Let $\tilde{f}(M)$ be defined as Definition C.3.

For $M, \tilde{M} \in \mathbb{R}^{d \times d}$, we have

$$\|\tilde{f}(M) - \tilde{f}(\tilde{M})\|_F \leq 2dnR^3 \|M - \tilde{M}\|_F$$

Proof. We have

$$\begin{aligned} \|\tilde{f}(M) - \tilde{f}(\tilde{M})\|_F &\leq \|\nabla \tilde{f}(M)\|_F \cdot \|M - \tilde{M}\|_F \\ &\leq 2dnR^3 \cdot \|M - \tilde{M}\|_F \end{aligned}$$

where the first step follows from Fact F.1, the second step follows from Lemma E.4. \square

F.3 LIPSCHITZ OF $c(M)$

We introduce the lemma about Lipschitz of $c(M)$.

Lemma F.4 (Lipschitz of $c(M)$). *Under the following conditions*

- Assumption E.1 holds.
- Let $c(M)$ be defined as Definition C.4.

For $M, \tilde{M} \in \mathbb{R}^{d \times d}$, we have

$$\|c(M) - c(\tilde{M})\|_F \leq 2dnR^3 \|M - \tilde{M}\|_F$$

Proof. We have

$$\begin{aligned} \|c(M) - c(\tilde{M})\|_F &\leq \|\nabla c(M)\|_F \cdot \|M - \tilde{M}\|_F \\ &= \|\nabla \tilde{f}(M)\|_F \cdot \|M - \tilde{M}\|_F \\ &\leq 2dnR^3 \cdot \|M - \tilde{M}\|_F \end{aligned}$$

where the first step follows from Fact F.1, the second step follows from Lemma C.12, the third step follows from Lemma E.4. \square

F.4 LIPSCHITZ OF $\tilde{f}(M) \circ c(M)$

We introduce the lemma about Lipschitz of $\tilde{f}(M) \circ c(M)$.

Lemma F.5 (Lipschitz of $\tilde{f}(M) \circ c(M)$). *Under the following conditions*

- Assumption E.1 holds.
- Let $c(M)$ be defined as Definition C.4.
- Let $\tilde{f}(M)$ be defined as Definition C.3.

For $M, \tilde{M} \in \mathbb{R}^{d \times d}$, we have

$$\|\tilde{f}(M) \circ c(M) - \tilde{f}(\tilde{M}) \circ c(\tilde{M})\|_F \leq 6dn^{3/2}R^3\|M - \tilde{M}\|_F$$

Proof. We have

$$\begin{aligned} \text{LHS} &\leq \|\tilde{f}(M) \circ c(M) - \tilde{f}(M) \circ c(\tilde{M})\|_F + \|\tilde{f}(M) \circ c(\tilde{M}) - \tilde{f}(\tilde{M}) \circ c(\tilde{M})\|_F \\ &\leq \|\tilde{f}(M)\|_F \cdot \|c(M) - c(\tilde{M})\|_F + \|c(\tilde{M})\|_F \cdot \|\tilde{f}(M) - \tilde{f}(\tilde{M})\|_F \\ &\leq \sqrt{n} \cdot \|c(M) - c(\tilde{M})\|_F + 2\sqrt{n} \cdot \|\tilde{f}(M) - \tilde{f}(\tilde{M})\|_F \\ &\leq \sqrt{n} \cdot 2dnR^3\|M - \tilde{M}\|_F + 2\sqrt{n} \cdot 2dnR^3\|M - \tilde{M}\|_F \end{aligned}$$

where the first step follows from triangle inequality, the second step follows from Fact B.5, the third step follows from Lemma E.3, the fourth step follows from Lemma F.4 and Lemma F.3.

So we have

$$\|\tilde{f}(M) \circ c(M) - \tilde{f}(\tilde{M}) \circ c(\tilde{M})\|_F \leq 6dn^{3/2}R^3\|M - \tilde{M}\|_F$$

□

F.5 LIPSCHITZ OF $\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n)$

We introduce the lemma about Lipschitz of $\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n)$.

Lemma F.6 (Lipschitz of $\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n)$). *If the following conditions hold*

- Assumption E.1 holds.
- Let $c(M)$ be defined as Definition C.4.
- Let $\tilde{f}(M)$ be defined as Definition C.3.

For $M, \tilde{M} \in \mathbb{R}^{d \times d}$, we have

$$\|\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n) - \text{diag}((\tilde{f}(\tilde{M}) \circ c(\tilde{M})) \cdot \mathbf{1}_n)\|_F \leq 6dn^2R^3\|M - \tilde{M}\|_F$$

Proof. We have

$$\begin{aligned} \text{LHS} &= \|(\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n - (\tilde{f}(\tilde{M}) \circ c(\tilde{M})) \cdot \mathbf{1}_n\|_2 \\ &= \|((\tilde{f}(M) \circ c(M)) - (\tilde{f}(\tilde{M}) \circ c(\tilde{M}))) \cdot \mathbf{1}_n\|_2 \\ &\leq \|\tilde{f}(M) \circ c(M) - \tilde{f}(\tilde{M}) \circ c(\tilde{M})\| \cdot \|\mathbf{1}_n\|_2 \\ &= \sqrt{n} \|\tilde{f}(M) \circ c(M) - \tilde{f}(\tilde{M}) \circ c(\tilde{M})\| \end{aligned} \tag{10}$$

where the first step follows from Fact B.4, the second step follows from basic algebra, the third step follows from Fact B.5, and the fourth step follows from $\|\mathbf{1}_n\|_2 = \sqrt{n}$.

Then we have

$$\|\tilde{f}(M) \circ c(M) - \tilde{f}(\tilde{M}) \circ c(\tilde{M})\| \leq \|\tilde{f}(M) \circ c(M) - \tilde{f}(\tilde{M}) \circ c(\tilde{M})\|_F \tag{11}$$

which follows from Fact B.5.

Following Eq. (10), Eq. (11) and Lemma F.5, we have

$$\text{LHS} \leq \sqrt{n} \cdot 6dn^{3/2}R^3\|M - \tilde{M}\|_F = 6dn^2R^3\|M - \tilde{M}\|_F$$

□

F.6 LIPSCHITZ OF $\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n) \tilde{f}(M)$

We introduce the lemma about Lipschitz of $\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n) \tilde{f}(M)$.

Lemma F.7 (Lipschitz of $\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n) \tilde{f}(M)$). *If the following conditions hold*

- Assumption E.1 holds.
- Let $c(M)$ be defined as Definition C.4.
- Let $\tilde{f}(M)$ be defined as Definition C.3.

For $M, \tilde{M} \in \mathbb{R}^{d \times d}$, we have

$$\|\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n) \tilde{f}(M) - \text{diag}((\tilde{f}(\tilde{M}) \circ c(\tilde{M})) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M})\|_F \leq 24dn^{7/2}R^3 \|M - \tilde{M}\|_F$$

Proof. Following Fact F.2, we have

$$\begin{aligned} & \|\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n) \tilde{f}(M) - \text{diag}((\tilde{f}(\tilde{M}) \circ c(\tilde{M})) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M})\|_F \\ & \leq 2^1 \cdot \max\{6dn^2R^3, 6dn^{3/2}R^3\} \cdot (\sqrt{n} \cdot 2n) \|M - \tilde{M}\|_F \\ & = 24dn^{7/2}R^3 \|M - \tilde{M}\|_F \end{aligned}$$

where we have the upper bound in Lemma E.3, the Lipschitz of $\text{diag}((\tilde{f}(M) \circ c(M)) \cdot \mathbf{1}_n) \tilde{f}(M)$ in Lemma F.3 and Lemma F.6. \square

F.7 LIPSCHITZ OF GRADIENT

We introduce the lemma about Lipschitz of the gradient.

Theorem F.8 (Lipschitz of the gradient, formal version of Theorem 5.4). *We can show $\nabla_M \mathcal{L}(M)$ is L -Lipschitz.*

If the following conditions hold

- Assumption E.1 holds.
- Let $c(M)$ be defined as Definition C.4.
- Let $\tilde{f}(M)$ be defined as Definition C.3.

For $M, \tilde{M} \in \mathbb{R}^{d \times d}$, we have

$$\|\nabla_M \mathcal{L}(M) - \nabla_M \mathcal{L}(\tilde{M})\|_F \leq (\lambda + 30dn^{7/2}R^6) \cdot \|M - \tilde{M}\|_F$$

Proof. We have

$$\begin{aligned} & \|\nabla_M \mathcal{L}(M) - \nabla_M \mathcal{L}(\tilde{M})\|_F \\ & = \|W \circ (X^\top (c(M) \circ \tilde{f}(M) - \text{diag}(c(M) \circ \tilde{f}(M) \cdot \mathbf{1}_n) \tilde{f}(M) \\ & \quad - c(\tilde{M}) \circ \tilde{f}(\tilde{M}) + \text{diag}(c(\tilde{M}) \circ \tilde{f}(\tilde{M}) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M}))X) + \lambda M - \lambda \tilde{M}\|_F \\ & \leq \|W \circ (X^\top (c(M) \circ \tilde{f}(M) - \text{diag}(c(M) \circ \tilde{f}(M) \cdot \mathbf{1}_n) \tilde{f}(M) \\ & \quad - c(\tilde{M}) \circ \tilde{f}(\tilde{M}) + \text{diag}(c(\tilde{M}) \circ \tilde{f}(\tilde{M}) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M}))X)\|_F + \|\lambda(M - \tilde{M})\|_F \quad (12) \end{aligned}$$

where the first step follows from Theorem D.5, and the second step follows the triangle inequality. Now, we proof these two terms separately.

For the first term, we have

$$\begin{aligned} & \|W \circ (X^\top (c(M) \circ \tilde{f}(M) - \text{diag}(c(M) \circ \tilde{f}(M) \cdot \mathbf{1}_n) \tilde{f}(M) \\ & \quad - c(\tilde{M}) \circ \tilde{f}(\tilde{M}) + \text{diag}(c(\tilde{M}) \circ \tilde{f}(\tilde{M}) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M}))X)\|_F \end{aligned}$$

$$\begin{aligned}
&\leq \|W\|_F \cdot \|X^\top (c(M) \circ \tilde{f}(M) - \text{diag}(c(M) \circ \tilde{f}(M) \cdot \mathbf{1}_n) \tilde{f}(M) \\
&\quad - c(\tilde{M}) \circ \tilde{f}(\tilde{M}) + \text{diag}(c(\tilde{M}) \circ \tilde{f}(\tilde{M}) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M}))X\|_F \\
&\leq \|W\|_F \cdot \|X\|_F^2 \cdot \|c(M) \circ \tilde{f}(M) - \text{diag}(c(M) \circ \tilde{f}(M) \cdot \mathbf{1}_n) \tilde{f}(M) \\
&\quad - c(\tilde{M}) \circ \tilde{f}(\tilde{M}) + \text{diag}(c(\tilde{M}) \circ \tilde{f}(\tilde{M}) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M})\|_F \\
&\leq \|W\|_F \cdot \|X\|_F^2 \cdot (\|c(M) \circ \tilde{f}(M) - c(\tilde{M}) \circ \tilde{f}(\tilde{M})\|_F \\
&\quad + \|\text{diag}(c(M) \circ \tilde{f}(M) \cdot \mathbf{1}_n) \tilde{f}(M) - \text{diag}(c(\tilde{M}) \circ \tilde{f}(\tilde{M}) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M})\|_F) \\
&= R^3 \cdot (\|c(M) \circ \tilde{f}(M) - c(\tilde{M}) \circ \tilde{f}(\tilde{M})\|_F \\
&\quad + \|\text{diag}(c(M) \circ \tilde{f}(M) \cdot \mathbf{1}_n) \tilde{f}(M) - \text{diag}(c(\tilde{M}) \circ \tilde{f}(\tilde{M}) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M})\|_F)
\end{aligned}$$

where the first step and the second step follow from Fact B.5, the third step follows from triangle inequality, and the fourth step follows from Assumption E.1.

Then we have

$$\begin{aligned}
&R^3 \cdot (\|c(M) \circ \tilde{f}(M) - c(\tilde{M}) \circ \tilde{f}(\tilde{M})\|_F \\
&\quad + \|\text{diag}(c(M) \circ \tilde{f}(M) \cdot \mathbf{1}_n) \tilde{f}(M) - \text{diag}(c(\tilde{M}) \circ \tilde{f}(\tilde{M}) \cdot \mathbf{1}_n) \tilde{f}(\tilde{M})\|_F) \\
&\leq R^3 \cdot (24dn^{7/2}R^3\|M - \tilde{M}\|_F + 6dn^{3/2}R^3\|M - \tilde{M}\|_F) \\
&\leq R^3 \cdot (30dn^{7/2}R^3\|M - \tilde{M}\|_F) \\
&= 30dn^{7/2}R^6\|M - \tilde{M}\|_F \tag{13}
\end{aligned}$$

where the first step follows from Lemma F.5 and Lemma F.7, the second step follows from $n \geq 1$.

For the second term, we have

$$\|\lambda(M - \tilde{M})\|_F = \lambda\|M - \tilde{M}\|_F \tag{14}$$

which follows from Fact B.5.

Finally, we have

$$\|\nabla_M \mathcal{L}(M) - \nabla_M \mathcal{L}(\tilde{M})\|_F \leq (\lambda + 30dn^{7/2}R^6) \cdot \|M - \tilde{M}\|_F$$

which follows from Eq. (12), Eq. (13), and Eq. (14). \square

G CONVERGENCE OF GRADIENT DESCENT

G.1 HELPFUL STATEMENTS

Here, we present useful facts that we use to prove our convergence result.

Fact G.1. *We can show that for $a, b \in \mathbb{R}$*

- *Part 1.*

$$\sqrt{a^2 + b^2} \geq \frac{|a| + |b|}{\sqrt{2}}$$

- *Part 2. Suppose $|a| > |b|$*

$$\sqrt{|a| - |b|} \geq \sqrt{|a|} - \sqrt{|b|}$$

Proof. **Proof of Part 1.** Square both side of the inequality in **Part 1.**, we have

$$\text{LHS} = a^2 + b^2$$

and

$$\text{RHS} = \frac{a^2 + 2|a| \cdot |b| + b^2}{2}.$$

So we just need to prove

$$\begin{aligned} \text{LHS} - \text{RHS} &= a^2 + b^2 - \frac{a^2 + 2|a| \cdot |b| + b^2}{2} \\ &= \frac{a^2 + b^2 - 2|a| \cdot |b|}{2} \\ &= \frac{(|a| - |b|)^2}{2} \\ &\geq 0 \end{aligned}$$

which is hold because for any $x \in \mathbb{R}$, $x^2 \geq 0$.

Proof of Part 2. Square both side of the inequality in **Part 2.**, we have

$$\text{LHS} = |a| - |b|$$

and

$$\text{RHS} = |a| + |b| - 2\sqrt{|a||b|}$$

So we just need to prove

$$\begin{aligned} \text{LHS} - \text{RHS} &= |a| - |b| - |a| - |b| + 2\sqrt{|a||b|} \\ &= 2\sqrt{|a||b|} - 2|b| \\ &= 2\sqrt{|b|}(\sqrt{|a|} - \sqrt{|b|}) \\ &\geq 0 \end{aligned}$$

which is hold because $|a| > |b|$ and $|b| \geq 0$. □

G.2 LOWER BOUND ON FROBENIUS NORM

In this section, we present the lemma for the lower bound on the Frobenius norm.

Lemma G.2. *If the following conditions hold*

- *Let $B \in \mathbb{R}^{d \times d}$.*
- *Let $M \in [0, 1]^{d \times d}$.*
- *Let $\lambda \in [0, 1]$ be some constant.*
- *Suppose that $\|B\|_F \leq R$.*

Then, we can show

- *Part 1.*

$$\|B + \lambda M\|_F^2 \geq \|B\|_F^2 + \lambda^2 \|M\|_F^2 - 2R\lambda d$$

- *Part 2.*

$$\|B + \lambda M\|_F \geq \frac{1}{\sqrt{2}}(\|B\|_F + \lambda \|M\|_F) - \sqrt{2R\lambda d}$$

Proof. **Proof of Part 1.** We can show that

$$\begin{aligned} \|B + \lambda M\|_F^2 &= \|B\|_F^2 + \lambda^2 \|M\|_F^2 + 2\langle B, \lambda M \rangle \\ &\geq \|B\|_F^2 + \lambda^2 \|M\|_F^2 - 2\|B\|_F \cdot \|\lambda M\|_F \\ &\geq \|B\|_F^2 + \lambda^2 \|M\|_F^2 - 2R\lambda d \end{aligned} \tag{15}$$

where the first step follows from Fact B.6, the second step follows from Fact B.5, the third step follows from the upper bound of $\|B\|_F$ and $\|M\|_F$.

Proof of Part 2. Taking the square root on both sides, we get

$$\begin{aligned}\|B + \lambda M\|_F &\geq \sqrt{\|B\|_F^2 + \lambda^2 \|M\|_F^2 - 2R\lambda d} \\ &\geq \sqrt{\|B\|_F^2 + \lambda^2 \|M\|_F^2} - \sqrt{2R\lambda d} \\ &\geq \frac{1}{\sqrt{2}}(\|B\|_F + \lambda \|M\|_F) - \sqrt{2R\lambda d}\end{aligned}$$

where the first step follows from Eq. (15), the second step follows from **Part 2.** of Fact G.1, and the third step follows from **Part 1.** of Fact G.1. \square

G.3 SANDWICH LOWER BOUND ON FROBENIUS NORM

Here, we introduce a sandwich trace fact.

Fact G.3. *If $A \succeq \beta I$, then $\text{tr}[B^\top AB] \geq \beta \text{tr}[B^\top B]$.*

Proof. As $A \succeq \beta I$, we have $A - \beta I \succeq 0$. Multiplying both sides by B^\top on the left and B on the right (noting that these operations preserve the positive semidefiniteness), we have

$$B^\top(A - \beta I)B \succeq 0.$$

Taking the trace and utilizing the property that the trace of a positive semidefinite matrix is non-negative, we have

$$\text{tr}[B^\top AB - \beta B^\top B] \geq 0,$$

which simplifies to

$$\text{tr}[B^\top AB] - \beta \text{tr}[B^\top B] \geq 0.$$

This concludes the proof. \square

We establish a sandwich lower bound on the Frobenius norm.

Lemma G.4 (Formal version of Lemma 5.6). *If the following conditions hold*

- Let $B \in \mathbb{R}^{n \times n}$ and $X \in \mathbb{R}^{n \times d}$.
- Assume that $XX^\top \succeq \beta I$.

Then, we have

$$\|X^\top BX\|_F \geq \beta \|B\|_F$$

Proof. We can show that

$$\begin{aligned}\|X^\top BX\|_F^2 &= \text{tr}[X^\top BXX^\top B^\top X] \\ &\geq \beta \cdot \text{tr}[X^\top BB^\top X] \\ &= \beta \cdot \text{tr}[B^\top XX^\top B] \\ &\geq \beta^2 \cdot \text{tr}[B^\top B] \\ &= \beta^2 \cdot \|B\|_F^2\end{aligned}$$

where the first step, the third step, and the fifth step follow from Fact B.4, the second step and the fourth step follows from Fact G.3 and $XX^\top \succeq \beta I$.

Taking the square root of both sides, we finish the proof. \square

G.4 LOWER BOUND ON HADAMARD PRODUCT BETWEEN TWO MATRICES

We present the lemma for the lower bound on the Hadamard product between two matrices in this section.

Lemma G.5. *If the following conditions hold*

- Let $B, W \in \mathbb{R}^{d \times d}$.

Then, we have

$$\max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|B\|_F \geq \|W \circ B\|_F \geq \min_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|B\|_F.$$

Proof. The proof directly follows from the definition of the Frobenius norm. \square

G.5 FINAL BOUND

We introduce some useful lemmas that we use to prove the final bound.

Lemma G.6. *If the following conditions hold*

- Let $b \in \mathbb{R}^n$ and $\langle b, \mathbf{1}_n \rangle = 0$.
- Let $f \in [\delta, 1]^n$ and $\langle f, \mathbf{1}_n \rangle = 1$.

Then we have

$$\|(b - \langle b, f \rangle \mathbf{1}_n) \circ f\|_2 \geq \delta \|b\|_2.$$

Proof. Note that $\langle b, \mathbf{1}_n \rangle = 0$ so that b and $\mathbf{1}_n$ are orthogonal with each other. Then, we have

$$\begin{aligned} \|(b - \langle b, f \rangle \mathbf{1}_n) \circ f\|_2 &\geq \delta \|b - \langle b, f \rangle \mathbf{1}_n\|_2 \\ &= \delta \sqrt{\|b\|_2^2 + \|\langle b, f \rangle \mathbf{1}_n\|_2^2} \\ &\geq \delta \|b\|_2, \end{aligned}$$

where the second step is from the Pythagorean theorem. \square

We present our final bound for proving the PL inequality.

Lemma G.7 (Formal version of Lemma 5.7). *If the following conditions hold*

- Let $B \in \mathbb{R}^{n \times n}$ and each row summation is zero, i.e., $B \cdot \mathbf{1}_n = \mathbf{0}_n$.
- Let $\tilde{f}(M) \in [0, 1]^{n \times n}$ and each row summation is 1, i.e., $\tilde{f}(M) \cdot \mathbf{1}_n = \mathbf{1}_n$.
- Assume that $\min_{i,j \in [n]} \tilde{f}(M)_{i,j} \geq \delta > 0$.

Then, we can show

$$\|B \circ \tilde{f}(M) - \text{diag}((B \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)\|_F \geq \delta \cdot \|B\|_F$$

Proof. For any $i \in [n]$, let $B_i \in \mathbb{R}^n$ be the i -th row of B , and we have $\langle B_i, \mathbf{1}_n \rangle = 0$ by the first condition.

For any $i \in [n]$, let $\tilde{f}(M)_i \in \mathbb{R}^n$ be the i -th row of $\tilde{f}(M)$, and we have $\langle \tilde{f}(M)_i, \mathbf{1}_n \rangle = 1$ by the second condition and $\tilde{f}(M)_{i,j} \in [\delta, 1]$ by the third condition.

By Lemma G.6, for any $i \in [n]$, we have

$$\|(B_i - \langle B_i, \tilde{f}(M)_i \rangle \mathbf{1}_n) \circ \tilde{f}(M)_i\|_2 \geq \delta \|B_i\|_2.$$

Then, we have

$$\|B \circ \tilde{f}(M) - \text{diag}((B \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)\|_F^2$$

$$\begin{aligned}
&= \sum_{i \in [n]} \|(B_i - \langle B_i, \tilde{f}(M)_i \rangle \mathbf{1}_n) \circ \tilde{f}(M)_i\|_2^2 \\
&\geq \sum_{i \in [n]} \delta^2 \|B_i\|_2^2 \\
&= \delta^2 \|B\|_F^2.
\end{aligned}$$

□

G.6 PL INEQUALITY

Here, we present the bound for one unit loss function.

Lemma G.8. *If the following conditions hold*

- Let $c(M)$ be defined in Definition C.4.

We have

$$\|c(M)\|_F \leq 2\sqrt{n}.$$

Proof. We have

$$\begin{aligned}
\|c(M)\|_F &\leq \|\tilde{f}(M)\|_F + \|f\|_F \\
&\leq 2\sqrt{n},
\end{aligned}$$

where the first step follows from Definition C.4 and triangle inequality, the second step follows $x_1^2 + \dots + x_n^2 \leq (x_1 + \dots + x_n)^2$ when $x_i \geq 0$ for any $i \in [n]$. □

We present the lemma to prove the PL inequality.

Lemma G.9. *If the following conditions hold*

- Let $\tilde{f}(M)$ be defined in Definition C.3.
- Let $c(M)$ be defined in Definition C.4.

We have

$$\|\text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)\|_F \leq \sqrt{n}.$$

Proof. We have

$$\begin{aligned}
\|\text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)\|_F &\leq \max_{i \in [n]} \{|(c(M)_i \circ \tilde{f}(M)_i) \cdot \mathbf{1}_n|\} \cdot \|\tilde{f}(M)\|_F \\
&\leq \|\tilde{f}(M)\|_F \\
&\leq \sqrt{n},
\end{aligned}$$

where the first step is by Frobenius norm definition and the second step follows from $\langle \tilde{f}(M)_i, \mathbf{1}_n \rangle = 1$ and $c(M)_i \in [-1, 1]^n$ for any $i \in [n]$. □

Finally, we can show the lemma for PL inequality.

Lemma G.10 (PL inequality, formal version of 5.5). *If the following conditions hold,*

- Let $M \in [0, 1]^{d \times d}$.
- Let $\lambda \in [0, 1]$ be some constant.
- Assume that $XX^\top \succeq \beta I$.
- Assume that $\min_{i,j \in [n]} \tilde{f}(M)_{i,j} \geq \delta > 0$.

- Let $\mathcal{L}(M)$ be defined in Definition C.7.

Furthermore,

- Let $\alpha = 2$.
- Let $\mu = 2 \min_{i,j \in [d]} \{|W_{i,j}|\} \cdot \beta \cdot \delta$.
- Let $\xi = 12\sqrt{n} \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X\|_F^2 \cdot \lambda d / \mu$.

We have

$$\|\nabla_M \mathcal{L}(M)\|_F^\alpha \geq \frac{1}{2} \mu (\|c(M)\|_F^2 + \frac{2\lambda^2}{\mu} \|M\|_F^2 - \xi).$$

Proof. We have $\tilde{f}(M) \cdot \mathbf{1}_n = \mathbf{1}_n$ and $f \cdot \mathbf{1}_n = \mathbf{1}_n$ by Definition C.3. Note that $c(M) = \tilde{f}(M) - f$ by Definition C.4. Thus, we have $c(M) \cdot \mathbf{1}_n = \mathbf{0}_n$.

On the other hand, we have

$$\begin{aligned} & \|W \circ (X^\top (c(M) \circ \tilde{f}(M) - \text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)) X)\|_F \\ & \leq \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X^\top (c(M) \circ \tilde{f}(M) - \text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)) X\|_F \\ & \leq \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X\|_F^2 \cdot \|c(M) \circ \tilde{f}(M) - \text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)\|_F \\ & \leq \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X\|_F^2 \cdot (\|c(M) \circ \tilde{f}(M)\|_F + \|\text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)\|_F) \\ & \leq \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X\|_F^2 \cdot (\|c(M)\|_F + \|\text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)\|_F) \\ & \leq \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X\|_F^2 \cdot (2\sqrt{n} + \sqrt{n}) \\ & = \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X\|_F^2 \cdot 3\sqrt{n} \end{aligned}$$

where the first and fourth steps follow Lemma G.5, the second step follows from Frobenius norm property, the third step follows from triangle inequality, the fifth step follows from Lemma G.8 and Lemma G.9.

Let $\alpha = 2$. We have the following

$$\begin{aligned} & \|\nabla_M \mathcal{L}(M)\|_F^2 \\ & = \|W \circ (X^\top (c(M) \circ \tilde{f}(M) - \text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)) X) + \lambda M\|_F^2 \\ & \geq \|W \circ (X^\top (c(M) \circ \tilde{f}(M) - \text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)) X)\|_F^2 + \lambda^2 \|M\|_F^2 - \alpha_1 \\ & \geq \alpha_2 \cdot \|X^\top (c(M) \circ \tilde{f}(M) - \text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)) X\|_F^2 + \lambda^2 \|M\|_F^2 - \alpha_1 \\ & \geq \alpha_2 \cdot \alpha_3 \cdot \|c(M) \circ \tilde{f}(M) - \text{diag}((c(M) \circ \tilde{f}(M)) \cdot \mathbf{1}_n) \tilde{f}(M)\|_F^2 + \lambda^2 \|M\|_F^2 - \alpha_1 \\ & \geq \alpha_2 \cdot \alpha_3 \cdot \alpha_4 \cdot \|c(M)\|_F^2 + \lambda^2 \|M\|_F^2 - \alpha_1 \\ & = \frac{1}{2} \mu (\|c(M)\|_F^2 + \frac{2\lambda^2}{\mu} \|M\|_F^2 - \xi), \end{aligned}$$

where the second step follows from Lemma G.2 and $\alpha_1 = 6\sqrt{n} \max_{i,j \in [d]} \{|W_{i,j}|\} \cdot \|X\|_F^2 \cdot \lambda d$, the third step follows from Lemma G.5 and $\alpha_2 = \min_{i,j \in [d]} \{|W_{i,j}|\}$, the fourth step follows from Lemma G.4 and $\alpha_3 = \beta$, the fifth step follows from Lemma G.7 and $\alpha_4 = \delta$, and the last step follows from $\mu = 2\alpha_2 \cdot \alpha_3 \cdot \alpha_4$ and $\xi = 2\alpha_1 / \mu$. \square