Learnable Low-Rank Polynomial Sketch for Effective Linear Attention

Zongyue Qin, Haoran Deng, Jason Cong, Yizhou Sun
Department of Computer Science
University of California Los Angeles
CA, USA
{qinzongyue,denghaoran,cong,yzsun}@cs.ucla.edu

Abstract

Softmax attention in Transformers suffers from quadratic complexity in sequence length, making it impractical for long-context applications. Linear attention alleviates this issue by replacing the exponential kernel with alternative functions that enable linear-time computation. Among existing linear attention approaches, recent studies have shown that polynomial kernels are particularly effective, as they exhibit sparse and spiky behavior similar to softmax attention, which emphasizes large dot products while suppressing irrelevant interactions. However, the exact computation of high-degree polynomial kernels is infeasible for high-dimensional representations. As a result, prior work relies on approximate polynomial kernels. This introduces a non-negligible approximation error. In this paper, we show that, from the perspective of polynomial kernel approximation, existing linear attention methods are still suboptimal. We propose **L**earnable **Lo**w-Rank **P**olynomial **S**ketch (LLoPS), a principled and flexible framework for approximating polynomial kernels with linear attention. Our method learns a low-rank polynomial sketch that provably achieves a strictly smaller approximation error than existing approaches. Experiment results show that LLoPS achieves the highest performance across extensive benchmarks, comparing with various linear attention baselines.

1 Introduction

Transformers have become the dominant architecture for sequence modeling across natural language processing, vision, and multimodal tasks. Despite their empirical success, the quadratic time and memory complexity of the attention mechanism with respect to sequence length severely limits scalability to long-context settings. This bottleneck is particularly problematic in emerging applications such as long-document understanding, code modeling, and retrieval-augmented generation, where context lengths can easily reach tens or hundreds of thousands of tokens.

Linear attention methods [Xiong et al., 2021, Team et al., 2025, Schlag et al., 2021] address this challenge by retaining the core attention formulation while reducing its complexity from quadratic to linear in sequence length. By rewriting softmax attention as a kernelized dot-product and exploiting associativity, linear attention enables streaming computation with constant memory. Recent large-scale deployments [Team et al., 2025] demonstrate that linear attention can attain a strong empirical performance while supporting extremely long contexts, thus making it a promising direction for next-generation foundation models.

Existing work on linear attention can be broadly categorized into two complementary directions. The first focuses on improving the state update rule, designing better ways to accumulate and normalize the running key-value statistics over time [Team et al., 2025, Yang et al., 2024]. The second direction, which this paper focuses on, aims to improve the feature map of query and key vectors so that the

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

induced kernel more faithfully approximates softmax attention [Arora et al., 2024, Zhang et al., 2024]. Prior studies on feature map design have observed that high-order polynomial kernels can approximate the softmax kernel surprisingly well [Kacham et al., 2023, Arora et al., 2024, Zhang et al., 2024]. Intuitively, both exponential and polynomial kernels can exhibit sparse, spiky behavior that emphasizes large dot-products while suppressing irrelevant interactions. However, computing exact polynomial kernels is infeasible for high degree kernel and high-dimensional vectors, as the computation complexity grows exponentially with the polynomial degree.

To address this challenge, PolySketchFormer [Kacham et al., 2023] exploits a data-independent sketching technique to approximate exact polynomial kernels. More recently, BASED [Arora et al., 2024] realizes state-of-the-art performance in language modeling by combining linear attention and sliding window attention. In their linear attention formulation, queries and keys are first compressed into a low-dimensional space and then the exact polynomial kernel is computed in that space.

In this paper, we show that from the perspective of polynomial kernel approximation, the previous methods that approximate polynomial kernel are suboptimal. Our key insight is that, when the output feature dimension is fixed, there exists a theoretically optimal sketch that minimizes the approximation error to the target polynomial kernel. We prove that a linear projection for the p-th order outer products of input vectors can achieve this optimal approximation, whereas existing approaches such as BASED [Arora et al., 2024] and PolySketchFormer [Kacham et al., 2023] can be viewed as constrained or data-independent approximations to this optimal mapping. Building on this observation, we further propose a structured decomposition of the optimal linear mapping using low-rank tensor rank decompositions (also known as CP decompositions) and column-wise Kronecker products, which significantly reduces computational costs. We provide theoretical analysis showing that this structured approximation is strictly tighter than prior feature-map constructions under the same dimensional budget.

Our contributions are summarized as follows: (1) We provide a theoretical characterization of the optimal sketching strategy for approximating polynomial kernels in the context of linear attention. (2) Motivated by this analysis, we propose a novel **L**earnable **Lo**w-Rank **P**olynomial **S**ketch (LLoPS), a linear attention kernel that accurately approximates polynomial attention while remaining computationally efficient. (3) We further derive approximation error analysis for LLoPS and show that several existing linear attention kernels can be interpreted as special cases within our unified framework. (4) We demonstrate through experiments that LLoPS consistently improves the effectiveness of linear attention transformers across a broad range of benchmarks.[1]

2 Preliminaries

2.1 Softmax-based Full Attention

Given a sequence of queries $\boldsymbol{Q} \in \mathbb{R}^{n\times d}$, keys $\boldsymbol{K} \in \mathbb{R}^{n\times d}$, and values $\boldsymbol{V} \in \mathbb{R}^{n\times d_v}$, standard Transformers compute attention as $\mathrm{Attn}(\boldsymbol{Q},\boldsymbol{K},\boldsymbol{V}) = \mathrm{softmax}\Big(\frac{\boldsymbol{Q}\boldsymbol{K}^\top}{\sqrt{d}}\Big)\boldsymbol{V}$, where the softmax is applied row-wise. Because this formulation has $\mathcal{O}(n^2)$ time and memory complexity due to the explicit computation of the attention matrix $\boldsymbol{Q}\boldsymbol{K}^\top$, it is prohibitively expensive for long input sequences.

2.2 Linear Attention

Previous studies have observed that softmax attention can be interpreted as kernelized dot-product attention with an exponential kernel. By approximating the exponential function with a feature map $\phi : \mathbb{R}^d \to \mathbb{R}^m$ and exploiting associativity, attention can be rewritten as $\mathrm{Attn}(\boldsymbol{Q},\boldsymbol{K},\boldsymbol{V}) = \big(\phi(\boldsymbol{Q})\phi(\boldsymbol{K})^\top\big)\boldsymbol{V} = \phi(\boldsymbol{Q})\big(\phi(\boldsymbol{K})^\top\boldsymbol{V}\big)$, where $\phi(\boldsymbol{Q}), \phi(\boldsymbol{K}) \in \mathbb{R}^{n\times m}$ are mapped features with dimension m and $\phi(\boldsymbol{Q})\phi(\boldsymbol{K})^\top$ is the inner product in the new feature space that approximates $\mathrm{softmax}\Big(\frac{\boldsymbol{Q}\boldsymbol{K}^\top}{\sqrt{d}}\Big)$.[2] Crucially, the overall complexity is reduced from quadratic to linear in the sequence length n. This enables streaming computation and a constant memory state, which are key advantages of linear attention.

[1] Code available at `https://github.com/ZongyueQin/LLoPS`

[2] We use $\phi(\boldsymbol{Q})$ to denote applying $\phi(\cdot)$ to each row vector in $\boldsymbol{Q}$. Similar for $\phi(\boldsymbol{K})$ and later usages.

Recent work reveals that polynomial kernels provide a strong approximation to the exponential kernel underlying softmax attention. In particular, a degree-p polynomial kernel, denoted as $\kappa(\cdot,\cdot)$, takes the form $\kappa(\boldsymbol{q},\boldsymbol{k}) = (\boldsymbol{q}^\top \boldsymbol{k})^p$, where $\boldsymbol{q},\boldsymbol{k} \in \mathbb{R}^d$ are query and key vectors.

However, computing the exact polynomial feature map required by linear attention involves expanding all degree-p monomials, i.e., $\psi_p(\boldsymbol{x}) = \mathrm{vec}(\boldsymbol{x}^{\otimes p})$, where $\otimes p$ is p-th order outer product, whose dimensionality scales as $\mathcal{O}(d^p)$. This makes exact polynomial kernels impractical for real-world models. BASED addresses this issue by first compressing the query and key vectors into a low-dimensional space using a linear projection,

$$\tilde{\boldsymbol{q}} = \boldsymbol{W}_q \boldsymbol{q}, \quad \tilde{\boldsymbol{k}} = \boldsymbol{W}_k \boldsymbol{k}, \qquad \boldsymbol{W}_q, \boldsymbol{W}_k \in \mathbb{R}^{r \times d},\ r \ll d, \tag{1}$$

and then computing exact polynomial features in the compressed space. In practice, BASED utilizes a second-order polynomial kernel $\kappa(\boldsymbol{q},\boldsymbol{k}) = \psi_2(\boldsymbol{q})^\top \psi_2(\boldsymbol{k})$, which corresponds to the quadratic term in the Taylor expansion of the exponential kernel.[3] By combining this polynomial linear attention with sliding-window attention to capture local interactions, BASED achieves state-of-the-art language modeling performance while maintaining favorable computational efficiency.

2.3 Related Work

We broadly categorize prior studies on linear attention into two directions: (i) the design of feature maps for query and key vectors, and (ii) the design of state update rules.

Feature Map Design. This line of work focuses on constructing efficient and expressive feature maps $\phi(\cdot)$ such that the softmax kernel $\exp(\boldsymbol{q}^\top \boldsymbol{k})$ can be approximated by an inner product $\phi(\boldsymbol{q})^\top \phi(\boldsymbol{k})$ [Katharopoulos et al., 2020, Choromanski et al., 2020, Qin et al., 2022, Keles et al., 2023, Meng et al., 2025]. Among the many design choices, several recent studies have identified polynomial kernels as particularly effective approximations, as they share important inductive biases with softmax attention, such as sparsity and sharp selectivity [Han et al., 2023, Zhang et al., 2024, Kacham et al., 2023, Arora et al., 2024]. As discussed by [Arora et al., 2024], many existing feature maps are substantially less effective than polynomial-kernel-based approaches when evaluated in realistic settings. Moreover, for current methods that explicitly approximate polynomial kernels [Han et al., 2023, Kacham et al., 2023, Arora et al., 2024], their approximation accuracy is still limited. Our work finds that these approximations can be further improved through a more principled and flexible feature-map construction. Overall, there remains significant room for boosting the effectiveness of linear attention feature maps.

Additionally, despite the extensive exploration of feature map design, many prior works do not assess performance on large-scale, real-world language modeling tasks. In some cases, linear attention variants are reported to outperform softmax attention, which contrasts with empirical observations in large language models where linear attention usually has sub-optimal effectiveness.

State Update Rules. Another line of work focuses on designing improved state update rules for linear attention that aim to stabilize training and better capture long-range dependencies [Schlag et al., 2021, Yang et al., 2024, Team et al., 2025]. These methods often resemble state space models [Gu and Dao, 2024, Peng et al., 2023] in that they maintain and update a recurrent state, rather than strictly adhering to the standard attention formulation.

Importantly, these approaches are largely orthogonal to feature map design: their update rules can be combined with different feature maps by replacing $\boldsymbol{q}$ and $\boldsymbol{k}$ with $\phi(\boldsymbol{q})$ and $\phi(\boldsymbol{k})$. Since our work focuses exclusively on improving feature maps for polynomial kernel approximation, we do not directly compare with methods that primarily modify the state update mechanism.

Polynomial Sketches for Neural Networks. Polynomial sketches have also been incorporated into neural network modules in prior work. For example, PolySketchFormer [Kacham et al., 2023] uses polynomial sketches to approximate attention, and EUGens [Kim et al., 2026] applies polynomial sketches within feedforward layers. While these works demonstrate the practical usefulness of polynomial sketching, their theoretical analyses are based on data-independent sketches, where the

[3] BASED proposes to use the Taylor expansion $1 + x + x^2/2$. For simplicity, we omit the constant and linear terms, as they can be handled by concatenation and do not affect the core analysis.

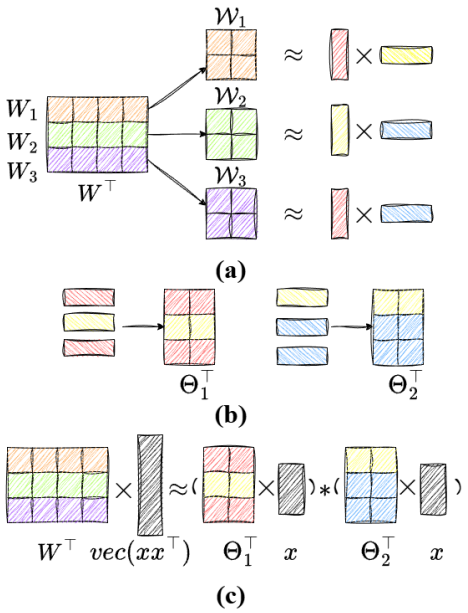


Figure 1: Illustration of LLoPS with an example where $p = 2, d = 2, m = 3$. (a) Given the optimal linear projection matrix $\boldsymbol{W}$, each column of $\boldsymbol{W}$ can be written as a 2 by 2 matrix, and be approximated by its rank-1 CP decomposition; (b) The rank-1 CP decomposition can be stacked into two weight matrices $\boldsymbol{\Theta_1}, \boldsymbol{\Theta_2}$; (c) In this way, the linear projection $\boldsymbol{W}^\top \text{vec}(\boldsymbol{x}\boldsymbol{x}^\top)$ can be approximated as $(\boldsymbol{\Theta}_\mathbf{1}^\top \boldsymbol{x}) * (\boldsymbol{\Theta}_\mathbf{2}^\top \boldsymbol{x})$ where $*$ is an element-wise product.

sketching map is fixed independently of the data distribution. This setting differs fundamentally from ours: LLoPS learns a data-dependent sketch tailored to the target representations, which can yield substantially better approximation under a fixed sketch dimension. To the best of our knowledge, our work is the first to provide theoretical analysis and approximation guarantees for such data-dependent polynomial sketches.

3 Methodology

Polynomial kernel is considered to be a good candidate to replace softmax kernel [Kacham et al., 2023, Zhang et al., 2024, Arora et al., 2024]. A key reason for the effectiveness of softmax attention is its sparse and spiky attention distribution: for a given query, only a small subset of keys receives large attention weights, while the majority of keys are assigned values close to zero. This behavior arises from the exponential function, which sharply amplifies differences in inner products, effectively separating keys that are well aligned with the query from those that are not.

Polynomial kernels exhibit a similar spikiness property, as higher-order powers increasingly emphasize large inner products while suppressing smaller ones. Besides, polynomial functions are often used to approximate exponential function (e.g., Taylor expansion) [Zhang et al., 2024, Arora et al., 2024]. Thus, a polynomial kernel is a natural and effective alternative to the exponential kernel in the context of linear attention. However, as discussed in Sec 2.2, an exact computation of polynomial kernels becomes impractical for high degrees or high-dimensional vectors due to the exponential growth of dimensionality.

Therefore, our work focuses on designing feature maps for queries and keys separately, denoted by $\phi_Q(\boldsymbol{Q}), \phi_K(\boldsymbol{K}) : \mathbb{R}^{n\times d} \to \mathbb{R}^{n\times m}$, that accurately approximate polynomial kernels while remaining computationally efficient. Specifically, we aim to minimize the following approximation error:

$$\left\|\phi_Q(\boldsymbol{Q})\phi_K(\boldsymbol{K})^\top - (\boldsymbol{Q}\boldsymbol{K}^\top)^{*p}\right\|_F^2, \tag{2}$$

where $*p$ is the element-wise p-th power and $\|\cdot\|_f$ is Frobenius norm.

To achieve this goal, we introduce **L**earnable **Lo**w-Rank **P**olynomial **S**ketch (LLoPS). We first demonstrate that, when the output dimension m is fixed, complicated feature maps are unnecessary to achieve

the best compression. Instead, linear projections of $\boldsymbol{q}^{\otimes p}$ and $\boldsymbol{k}^{\otimes p}$ are sufficient to attain the theoretically optimal approximation to the target polynomial kernel. Building on this, we further decompose the optimal linear projection and approximate it using a low-rank CP (CANDECOMP/PARAFAC) decomposition [Papalexakis, 2016], which is also known as tensor rank decomposition, that enables a practical computation with linear attention while preserving the approximation quality. Finally, we provide a detailed error analysis of the proposed feature maps and elucidate their connections to existing linear attention kernels.

3.1 Optimal Sketch of Polynomial Kernel

We first consider the optimal compression of polynomial kernel when the output dimension m is fixed. For clarity, let $\psi_p : \mathbb{R}^d \to \mathbb{R}^{d^p}$, where $\psi_p(\boldsymbol{x}) = \text{vec}(\boldsymbol{x}^{\otimes p})$; let $\boldsymbol{A} = (\boldsymbol{Q}\boldsymbol{K}^\top)^{*p} = \psi_p(\boldsymbol{Q})\psi_p(\boldsymbol{K})^\top$[4]; and let $\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^\top$ be the SVD decomposition of $\boldsymbol{A}$ where $\boldsymbol{\Sigma} = diag(\sigma_1, \cdots, \sigma_n)$.

Since $\phi_Q(\boldsymbol{Q})$ and $\phi_K(\boldsymbol{K})$ are both an $n \times m$ matrix ($n \gg m$), the rank of $\phi_Q(\boldsymbol{Q})\phi_K(\boldsymbol{K})^\top$ is at most m. By the property of SVD decomposition, the minimum possible loss of Eq. 2 is $\sum_{i=m+1}^{n} \sigma_i^2$ where σ_i is the i-th singular value. It is achieved when

$$\phi_Q(\boldsymbol{Q})\phi_K(\boldsymbol{K})^\top = \boldsymbol{U}\boldsymbol{\Sigma_m}\boldsymbol{V}^\top \tag{3}$$

where $\boldsymbol{\Sigma}_m = diag(\sigma_1, \cdots, \sigma_m, 0, \cdots, 0)$.

The following theorem shows how to achieve the minimum approximation error.

Theorem 3.1 (Optimal Polynomial Sketch)**.** *Let $\psi_p(\boldsymbol{x}) : \mathbb{R}^d \to \mathbb{R}^{d^p} = vec(\boldsymbol{x}^{\otimes p})$. There exists $\boldsymbol{W}_Q^*, \boldsymbol{W}_K^* \in \mathbb{R}^{d^p \times m}$ such that $\phi_Q^*(\boldsymbol{Q})\phi_K^*(\boldsymbol{K})^\top = \boldsymbol{U}\boldsymbol{\Sigma_m}\boldsymbol{V}^\top$ where*

$$\phi_Q^*(\boldsymbol{Q}) = \psi_p(\boldsymbol{Q})\boldsymbol{W}_Q^* \quad \phi_K^*(\boldsymbol{K}) = \psi_p(\boldsymbol{K})\boldsymbol{W}_K^* \tag{4}$$

The proof is illustrated in Appendix A.1.

Theorem 3.1 suggests that we could compress $\psi_p(\boldsymbol{Q}), \psi_p(\boldsymbol{K})$ into a m-dimensional space via linear mapping, while achieving the minimal approximation error stated in Eq 3.

3.2 Learnable Low-Rank Polynomial Sketch

Although Theorem 3.1 illustrates how to achieve the optimal approximation, since $\boldsymbol{W}_Q^*$ and $\boldsymbol{W}_K^*$ have d^p rows, the linear projection is still too computationally expensive. To address this problem, we propose to decompose and approximate the linear projection matrices utilizing their low-rank CP decompositions (also known as tensor rank decomposition [Papalexakis, 2016]). Figure 1 depicts an example of our decomposed sketch for a toy example where $p = 2, d = 2, m = 3$.

Let $\boldsymbol{W}$ denote $\boldsymbol{W}_Q^*$ or $\boldsymbol{W}_K^*$, and let $\boldsymbol{W}_i \in \mathbb{R}^{d^p}$ ($1 \le i \le m$) denote the i-th column of $\boldsymbol{W}$. We can write $\boldsymbol{W}_i$ as a (p)-way tensor $\mathcal{W}_i \in \mathbb{R}^{d \times \cdots \times d}$. According to CP decomposition, $\mathcal{W}_i$ can be represented, with a sufficiently large R_i, as a linear combination of R_i rank-1 tensors:

$$\mathcal{W}_i = \sum_{r=1}^{R_i} \lambda_r^{(i)} \boldsymbol{w}_{1,r}^{(i)} \otimes \boldsymbol{w}_{2,r}^{(i)} \otimes \cdots \otimes \boldsymbol{w}_{p,r}^{(i)} \tag{5}$$

where $\lambda_r^{(i)} \in \mathbb{R}$, $\boldsymbol{w}_{j,r}^{(i)} \in \mathbb{R}^d$ ($1 \le j \le p$), and $\otimes$ denotes outer product.

Inspired by the CP decomposition in Eq. 5, the key idea of LLoPS is as follows. Each row of $\boldsymbol{W}$ is expressed as a sum of rank-one outer products of learnable vectors, as illustrated in Figure 1(a). These vectors are trained implicitly to approximate the low-rank CP decomposition of each row vector $\boldsymbol{W}_i$. We then organize the learned factors into a collection of projection matrices, as shown in Figure 1(b), which enables an efficient computation of the feature map, avoiding explicitly constructing the high-dimensional polynomial features $\psi_p(\boldsymbol{Q})$ and $\psi_p(\boldsymbol{K})$ through a sequence of lightweight matrix multiplications and element-wise operations instead, as depicted in Figure 1(c).

Specifically, $\mathcal{W}_i$ can be approximated with its rank-t CP decomposition, i.e.,

$$\mathcal{W}_i \approx \sum_{r=1}^{t} \lambda_r^{(i)} \boldsymbol{w}_{1,r}^{(i)} \otimes \boldsymbol{w}_{2,r}^{(i)} \otimes \cdots \otimes \boldsymbol{w}_{p,r}^{(i)} \triangleq \hat{\mathcal{W}}_i \tag{6}$$

[4]Function ψ_p is applied to each row of the matrices.

By reshaping the tensor back to vector, we have an approximate column vector $\hat{\boldsymbol{W}}_i = \text{vec}(\hat{\mathcal{W}}_i)$ and thus the approximate linear projector $\hat{\boldsymbol{W}} = (\hat{\boldsymbol{W}}_1, \cdots, \hat{\boldsymbol{W}}_m) \in \mathbb{R}^{d^p \times m}$. Next, we demonstrate how to efficiently compute $\psi_p(\boldsymbol{Q})\hat{\boldsymbol{W}}$ and $\psi_p(\boldsymbol{K})\hat{\boldsymbol{W}}$ without explicitly computing $\psi_p(\boldsymbol{Q})$, $\psi_p(\boldsymbol{K})$, and $\hat{\boldsymbol{W}}$, which are with extremely high dimensions.

Notice that the inner product between the outer product of vectors $\{\boldsymbol{w}_{1,r}, \ldots, \boldsymbol{w}_{p,r}\}$ and $\text{vec}(\boldsymbol{x}^{\otimes p})$ equals the product of $\{\langle \boldsymbol{w}_{j,r}, \boldsymbol{x}\rangle : 1 \le j \le p\}$, i.e.,

$$\langle \text{vec}(\boldsymbol{w}_{1,r} \otimes \cdots \otimes \boldsymbol{w}_{p,r}), \text{vec}(\boldsymbol{x}^{\otimes p})\rangle = \prod_{j=1}^{p} \langle \boldsymbol{w}_{j,r}, \boldsymbol{x}\rangle \tag{7}$$

So we have $\langle \hat{\boldsymbol{W}}_i, \text{vec}(\boldsymbol{x}^{\otimes p})\rangle = \sum_{r=1}^{t} \lambda_r^{(i)} \prod_{j=1}^{p} \langle \boldsymbol{w}_{j,r}^{(i)}, \boldsymbol{x}\rangle$.

Thus

$$\begin{aligned}
&\psi_p(\boldsymbol{Q})\hat{\boldsymbol{W}} \\
=&(\psi_p(\boldsymbol{Q})\hat{\boldsymbol{W}}_1, \cdots, \psi_p(\boldsymbol{Q})\hat{\boldsymbol{W}}_m) \\
=&\big(\sum_{r=1}^{t} \lambda_r^{(1)} \prod_{j=1}^{p} \boldsymbol{Q}\boldsymbol{w}_{j,r}^{(1)}, \cdots, \sum_{r=1}^{t} \lambda_r^{(m)} \prod_{j=1}^{p} \boldsymbol{Q}\boldsymbol{w}_{j,r}^{(m)}\big) \\
=&\sum_{r=1}^{t} \big(\lambda_r^{(1)} \prod_{j=1}^{p} \boldsymbol{Q}\boldsymbol{w}_{j,r}^{(1)}, \cdots, \lambda_r^{(m)} \prod_{j=1}^{p} \boldsymbol{Q}\boldsymbol{w}_{j,r}^{(m)}\big)
\end{aligned} \tag{8}$$

With Eq 8, we have the following theorem to approximate the optimal projection matrix $\boldsymbol{W}$ with a set of matrices $\boldsymbol{\Theta}$.

Theorem 3.2. *Let $\boldsymbol{\Theta}_{j,r} \in \mathbb{R}^{d\times m}$ be the stack of $\boldsymbol{w}_{j,r}^{(i)}$ for $i = 1, \ldots, m$, i.e.,*

$$\boldsymbol{\Theta}_{j,r} = (\lambda_r^{(1)}\boldsymbol{w}_{j,r}^{(1)}, \cdots, \lambda_r^{(m)}\boldsymbol{w}_{j,r}^{(m)}) \tag{9}$$

where $1 \le j \le p$, $1 \le r \le t$. We have

$$\begin{aligned}
\psi_p(\boldsymbol{Q})\hat{\boldsymbol{W}} &= \sum_{r=1}^{t} \prod_{j=1}^{p} \boldsymbol{Q}\boldsymbol{\Theta}_{\boldsymbol{j},\boldsymbol{r}}, \\
\hat{\boldsymbol{W}} &= \sum_{r=1}^{t} \big(\boldsymbol{\Theta}_{\boldsymbol{1},\boldsymbol{r}} \odot_{col} \boldsymbol{\Theta}_{\boldsymbol{2},\boldsymbol{r}} \odot_{col} \cdots \odot_{col} \boldsymbol{\Theta}_{\boldsymbol{p},\boldsymbol{r}}\big)
\end{aligned} \tag{10}$$

where $\odot_{col}$ is the column-wise Kronecker product. Similarly, $\psi_p(\boldsymbol{K})\hat{\boldsymbol{W}} = \sum_{r=1}^{t} \prod_{j=1}^{p} \boldsymbol{K}\boldsymbol{\Theta}_{\boldsymbol{j},\boldsymbol{r}}$.

More importantly, this formulation avoids explicitly constructing the exponentially large polynomial feature map $\psi_p(\boldsymbol{Q})$: we first compute the projected matrices $\boldsymbol{Q}\boldsymbol{\Theta}_{j,r}$ for all $j = 1, \ldots, p$ and $r = 1, \ldots, t$ using standard matrix multiplications. Then, for each rank component r, we combine the p projected terms via element-wise multiplication, corresponding to the product over j. Finally, we sum the resulting features across all t components. Overall, the total computation cost is reduced from $O(nmd^p)$ to $O(ndmtp)$.

In practice, we find that the setting $t = 1$ already attains an accurate approximation of the polynomial kernel. Setting $t > 1$ will lead to too many learnable parameters. Therefore, we fix $t = 1$ and our proposed feature maps is formulated as below:

Definition 3.3 (**L**earnable **Lo**w-Rank **P**olynomial **S**ketch)**.**

$$\phi_Q(\boldsymbol{Q}) = \prod_{j=1}^{p} \boldsymbol{Q}\boldsymbol{\Theta}_{\boldsymbol{j}}^{(Q)}, \quad \phi_K(\boldsymbol{K}) = \prod_{j=1}^{p} \boldsymbol{K}\boldsymbol{\Theta}_{\boldsymbol{j}}^{(K)} \tag{11}$$

where $\boldsymbol{\Theta}_{\boldsymbol{j}}^{(Q)}$ and $\boldsymbol{\Theta}_{\boldsymbol{j}}^{(Q)}$ $(1 \le j \le p)$ are learnable.

The computation complexity is $O(ndmp)$, and the number of trainable parameters are $O(dmp)$. During training, $\boldsymbol{\Theta}$ are trained implicitly to approximate the low-rank CP decomposition of the optimal projecion matrix $\boldsymbol{W}_Q^*$ and $\boldsymbol{W}_K^*$.

The following theorem analyze the error of our proposed learnable polynomial sketch.

Theorem 3.4 (LLoPS Error)**.** *Let $\boldsymbol{W}^* \in \mathbb{R}^{d^p \times m}$ be the optimal matrix suggested in Theorem 3.1. Then*

$$\min_{\boldsymbol{\Theta},\ldots,\boldsymbol{\Theta}_p} \|\boldsymbol{W}^* - \boldsymbol{\Theta}_1 \odot \ldots \odot \boldsymbol{\Theta}_p\|_F^2 = \|\boldsymbol{W}^*\|_F^2 - \sum_{i=1}^{m} \tau_i^2 \tag{12}$$

where

$$\tau_i = \max_{\|\boldsymbol{u}_1\|=\ldots=\|\boldsymbol{u}_p\|=1} \langle \boldsymbol{W}^*[:,i], \boldsymbol{u}_1 \otimes \ldots \otimes \boldsymbol{u}_p \rangle \tag{13}$$

is the tensor spectral norm (i.e., the largest tensor singular value) of $\mathcal{W}_i$, and $\mathcal{W}_i \in \mathbb{R}^{d \times \ldots \times d}$ is the order p tensor by reshaping the ith column of $\boldsymbol{W}^* \in \mathbb{R}^{d^p}$.

The proof is given in Appendix A.2.

Some previous studies found that making the attention score of linear attention to be non-negative is helpful for training stability [Kacham et al., 2023, Han et al., 2023]. If non-negativity is required, we could adopt the approach used in [Kacham et al., 2023] to apply another square to the feature map, i.e., $\phi_Q(\boldsymbol{Q}) = \left(\prod_{j=1}^{p} \boldsymbol{Q}\boldsymbol{\Theta}_{\boldsymbol{j}}^{(Q)}\right)^{*2}$, and similarly for $\phi_K(\boldsymbol{K})$.

3.3 Relationship with Existing Feature Maps

In this subsection, we discuss the relationship between our proposed feature map and several existing approaches for approximating polynomial kernels in linear attention.

Element-wise Polynomial Feature Maps Some earlier linear attention methods [Han et al., 2023] adopt element-wise polynomial feature maps of the form $\phi(\boldsymbol{x}) = (\boldsymbol{\Theta}\boldsymbol{x})^{*p}$ where $*p$ denotes the element-wise p-th power. This construction can be viewed as a special case of our formulation in which all projection matrices are shared, i.e., $\boldsymbol{\Theta}_1 = \cdots = \boldsymbol{\Theta}_p$. Under this constraint, our kernel reduces exactly to an element-wise polynomial feature map.

BASED This method [Arora et al., 2024] approximates the second-order polynomial kernel by first projecting queries and keys into a low-dimensional space via a linear mapping, and then applying the quadratic feature map $\psi_2(\cdot)$. Specifically, for $\boldsymbol{W} \in \mathbb{R}^{m' \times d}$ and $\boldsymbol{x} \in \mathbb{R}^d$, BASED computes

$$\psi_2(\boldsymbol{W}\boldsymbol{x}) = \mathrm{vec}(\boldsymbol{W}\boldsymbol{x} \otimes \boldsymbol{W}\boldsymbol{x}) = (\boldsymbol{W} \otimes_{\mathrm{kron}} \boldsymbol{W})\,\mathrm{vec}(\boldsymbol{x}\boldsymbol{x}^\top), \tag{14}$$

where $\otimes_{\mathrm{kron}}$ is the full Kronecker product.

This construction can also be recovered as a special case of our kernel. In particular, let

$$\boldsymbol{\Theta}_1^\top = \boldsymbol{W} \odot_{\mathrm{col}} \mathbf{1}, \qquad \boldsymbol{\Theta}_2^\top = \mathbf{1} \odot_{\mathrm{col}} \boldsymbol{W},$$

where $\mathbf{1} \in \mathbb{R}^{m' \times d}$ is an all-ones matrix. Under this choice, we have $\boldsymbol{\Theta}_1 \odot_{\mathrm{col}} \boldsymbol{\Theta}_2 = (\boldsymbol{W} \otimes_{\mathrm{kron}} \boldsymbol{W})^\top$. Therefore, with these specific parameter settings, our feature map exactly recovers the BASED kernel.

PolySketchFormer This approach [Kacham et al., 2023] approximates polynomial kernels using randomized sketching techniques. Its feature map can be expressed as multiplication by a randomized sketching matrix $\boldsymbol{R} \in \mathbb{R}^{d^p \times m}$ satisfying the (ϵ, p)-AMM property [Woodruff et al., 2014]. Following prior constructions [Ahle et al., 2020], $\boldsymbol{R}$ can be formulated as a column-wise Kronecker product of p sparse randomized matrices, permitting efficient approximation of high-order polynomial kernels.

From our perspective, PolySketchFormer corresponds to a special case of our framework in which the projection matrices $\{\boldsymbol{\Theta}_i\}_{i=1}^p$ are fixed, randomized, and sparse, rather than learned from data.

Taken together, these observations reveal that several existing polynomial-kernel-based linear attention methods can be interpreted as special or constrained instances of our proposed feature map. By allowing more general and learnable projections, our formulation strictly generalizes these prior approaches. As a consequence, for a fixed output dimension, our method can achieve an approximation error that is no worse than these existing methods, and in practice often substantially better, as shown in the next section. This is equivalent to enforcing a constraint where half of the learnable matrices share parameters with the other half.

Table 1: Zero-shot accuracy (%) across diverse NLP benchmarks for pretrained models at two scales. Bold numbers indicate the best performance among linear-complexity methods within each model size. BASED+CT is BASED under the same continual training setting as LLoPS.

360M Models										
Method	Wino	PIQA	LAM	Hella	ARC-E	ARC-C	FDA	SWDE	SQuAD	Avg
Softmax Attn	51.78	63.71	29.32	33.60	42.85	24.32	57.89	56.08	27.88	43.05
Mamba	50.36	63.76	27.89	33.88	42.42	24.66	5.90	17.37	24.83	32.34
BASED	51.30	63.55	26.64	33.23	41.88	24.23	**17.42**	27.18	27.41	34.76
BASED+CT	50.91	65.40	27.73	36.85	43.10	**25.51**	**8.35**	30.87	26.27	35.00
BASED+**LLoPS**	**53.04**	**66.70**	**29.23**	**37.07**	**43.94**	**25.51**	7.53	**31.41**	**31.94**	**36.26**
1B Models										
Method	Wino	PIQA	LAM	Hella	ARC-E	ARC-C	FDA	SWDE	SQuAD	Avg
Softmax Attn	52.49	67.30	38.07	39.12	51.05	22.10	68.78	68.68	35.79	49.26
Mamba	51.30	66.81	38.87	39.34	46.80	26.11	11.07	28.17	29.46	37.55
BASED	51.62	66.70	37.86	39.10	45.54	25.60	24.05	42.84	34.58	40.88
BASED+CT	51.22	68.44	36.43	41.59	47.01	**27.30**	24.41	46.80	36.23	42.16
BASED+**LLoPS**	**51.62**	**68.72**	**38.11**	**42.69**	**47.31**	27.13	**26.32**	**47.97**	**36.36**	**42.91**

4 Experiments

In this section, we report the main experiment results comparing our method with state-of-the-art linear attention methods. Additional experiments and ablation studies are reported in Appendix B.

4.1 Zero-Shot NLP Benchmarks

Following prior work on alternative sequence modeling architectures [Gu and Dao, 2024, Arora et al., 2024], we evaluate our method on a suite of zero-shot NLP benchmarks. We compare against the standard softmax Transformer, Mamba, and BASED, using the publicly released 360M and 1B parameter checkpoints from Arora et al. [2024]. Given that these architectures require extensive pre-training from scratch, additional linear attention baselines were excluded from this evaluation.

Since the BASED attention kernel can be viewed as a special case of our proposed kernel, we initialize our model from the BASED checkpoint and continue training it to obtain our final model to reduce training time (BASED+LLoPS). For LLoPS-360M, the polynomial degree $p = 4$; for LLoPS-1B $p = 2$. Importantly, the original baselines in Arora et al. [2024] were pretrained on the Pile dataset [Gao et al., 2020], which is no longer publicly available. Therefore, we perform continued training utilizing the SlimPajama-6B dataset [Soboleva et al., 2023][5] with the standard next-token-prediction objective, which is 50 times smaller in size compared to Pile. For a fair comparison under this updated data setting, the BASED models are also post-trained with SlimPajama-6B for the same number of iterations. This allows us to isolate the benefits of our proposed kernel from differences arising purely from additional training data or compute. The evaluated benchmarks are implemented in `lm_eval` [Gao et al., 2023], including: WinoGrande [Sakaguchi et al., 2021], PIQA [Bisk et al., 2020], LAMBADA [Paperno et al., 2016], HellaSwag [Zellers et al., 2019], ARC-challenge [Clark et al., 2018], SQUAD [Rajpurkar et al., 2016], SWDE [Lockard et al., 2019, Arora et al., 2023], FDA [Arora et al., 2023]. The tasks include information extraction and commonsense reasoning QA.

The results are reported in Table 1. Notably, for 360M models, we observe that after continued training on SlimPajama-6B, both BASED and our method exhibit a substantial accuracy drop on the FDA benchmark. We hypothesize that this degradation arises from a mismatch between the SlimPajama-6B and the FDA task distribution, which may induce catastrophic forgetting of task-relevant knowledge. Since this phenomenon occurs for both BASED and our method, it is more likely due to the choice of training data rather than a limitation of our proposed kernel. Despite this issue, our method produces a higher average accuracy than BASED. These results further support the effectiveness of our approach in improving the performance of linear attention models across diverse zero-shot NLP benchmarks.

[5] `https://huggingface.co/datasets/DKYoon/SlimPajama-6B`

Table 2: Performance comparison across Long Range Arena (LRA) benchmarks. **Bold numbers** indicate the best performance in linear attention methods, underlined numbers indicate the second-best, and *italicized numbers* indicate the third-best.

Method	Image	Text	Retrieval	ListOps	Pathfinder	Avg
Softmax Attn	64.36	82.96	83.23	48.49	70.45	69.90
Nyström	**58.66**	66.19	**83.90**	*45.46*	69.71	64.78
Skyformer	58.54	65.25	79.55	46.77	72.73	*64.57*
Hedgehog	52.68	*66.59*	82.24	41.03	**74.16**	63.34
PolaFormer	47.74	73.06	80.50	38.86	70.53	62.14
Flatten	51.48	66.13	78.72	38.51	70.27	61.02
PolySketchFormer	51.40	64.60	81.10	43.00	66.20	61.26
BASED	51.13	65.73	80.76	43.55	*72.98*	62.83
LLoPS	*57.54*	**73.94**	*82.17*	**48.29**	73.84	**67.16**

4.2 Long Range Arena Benchmark

Following prior work [Zhang et al., 2024, Meng et al., 2025, Han et al., 2023], we evaluate LLoPS on the Long Range Arena (LRA) benchmark [Tay et al., 2020]. Instead of the original train-from-scratch protocol, we adopt the pretraining-and-finetuning protocol of Amos et al. [2023], which better reflects the behavior of pretrained transformers. We apply the same protocol to all reproduced baselines for a fair comparison, while reporting the best numbers from prior papers and our training results. All methods share the same architecture except for the attention layer, and we compare against softmax attention, Nyström [Xiong et al., 2021], Skyformer [Chen et al., 2021], Hedgehog [Zhang et al., 2024], PolaFormer [Meng et al., 2025], Flatten Transformer [Han et al., 2023], PolySketchFormer [Kacham et al., 2023], and BASED [Arora et al., 2024]. We set the polynomial degree of LLoPS to $p = 4$.

Table 2 summarizes the results. Under this protocol, softmax attention consistently outperforms linear attention variants, aligning with empirical observations from large-scale language modeling. Among linear attention methods, LLoPS achieves top-3 performance on every task and obtains the highest average accuracy, improving over the strongest linear attention baseline by 2.28%. These results show that LLoPS substantially improves linear attention effectiveness while preserving its efficiency advantages.

4.3 Polynomial Kernel Approximation Error

Furthermore, we directly assess the error of our proposed sketch in approximating polynomial kernels. To that end, we obtain a set of query and key vectors with the following steps: (1) we train a transformer with the exact polynomial kernel ($Attn(\boldsymbol{q}, \boldsymbol{k}, \boldsymbol{v}) = (\boldsymbol{q}^\top \boldsymbol{k})^p \boldsymbol{v}$, $p = 2, 3$) on the synthetic associative recall (AR) task introduced in Arora et al. [2024]; (2) we select the first (layer 0), middle (layer 8), and last attention layers (layer 16) of the trained model and sample the corresponding query and key vectors with synthetic AR data. After obtaining the query and key vectors, for each layer, we apply different sketching methods to approximate the polynomial kernel on these sampled vectors.

We compare our sketch against the data-independent sketching approach used in PolySketchFormer [Kacham et al., 2023] (PSF), as well as neural network-based sketches (NN). For the neural sketch baseline, the feature maps ϕ_q and ϕ_k are parameterized as two-layer MLPs and then trained to approximate the exact polynomial kernel. To quantify approximation quality, we report the relative mean absolute error (RMAE), defined as

$$\mathrm{RMAE} = \mathbf{E}_{\boldsymbol{q},\boldsymbol{k}} \left| \frac{(\boldsymbol{q}^\top \boldsymbol{k})^p - \phi_q(\boldsymbol{q})^\top \phi_k(\boldsymbol{k})}{(\boldsymbol{q}^\top \boldsymbol{k})^p} \right|. \tag{15}$$

The results are summarized in Table 3. We observe that our sketch consistently achieves substantially smaller approximation error than PSF. This aligns with the intuition that learnable methods are usually better than data-independent methods. Furthermore, our learnable sketch also outperforms the neural network-based feature maps. We attribute this to two main reasons. First, Theorem 3.1 suggests that a linear projection is sufficient to achieve optimal compression for polynomial kernels, making more complex nonlinear parameterizations unnecessary. Second, our structured and decomposed linear formulation is easier to optimize than MLP-based sketches, which may suffer from additional training instability. Overall, these results demonstrate that our proposed sketch provides an effective and accurate approximation of polynomial kernels in practical transformer query and key distributions.

Table 3: Relative mean absolute error of different sketches approximating polynomial kernels with degree $p = 2, 3$.

p	Method	Layer 0	Layer 8	Layer 15	Avg
2	PSF	0.186	0.247	0.445	0.293
	NN	0.035	0.070	0.058	0.054
	LLoPS	0.008	0.032	0.025	0.022
3	PSF	0.740	0.738	0.745	0.741
	NN	0.059	0.114	0.143	0.105
	LLoPS	0.018	0.040	0.058	0.038

4.4 Efficiency Comparison

We compare the forward-pass runtime of LLoPS with softmax attention, BASED, and PolySketchFormer (PSF). All methods are evaluated on the same hardware and in the same implementation environment, with an input dimension of $d = 128$ and sequence lengths ranging from 1,000 to 10,000. We report the total wall-clock time over 1,000 forward passes. For LLoPS, we use polynomial degree $p = 2$ and output feature dimension $m = 256$. For BASED, queries and keys are projected to 16 dimensions, resulting in $16^2 + 16 + 1 = 273$ features, including the constant and linear terms. PSF is evaluated with sketch dimension $m = 32$, the smaller of the 32- and 64-dimensional settings considered in the original PSF paper. This smaller setting favors PSF in terms of computational cost, whereas LLoPS uses a substantially larger sketch dimension.

Figure 2 shows that LLoPS is consistently faster than PSF across all evaluated sequence lengths, achieving 1.47–1.56× speedups despite using an eightfold larger sketch dimension (256 versus 32). At a sequence length of 10,000, LLoPS requires 25.55 seconds for 1,000 forward passes, compared with 39.94 seconds for PSF, 78.22 seconds for BASED, and 301.30 seconds for softmax attention. LLoPS also outperforms BASED at all evaluated sequence lengths except 1,000, where BASED is slightly faster. These results demonstrate that LLoPS retains strong practical efficiency while providing a more expressive learnable polynomial sketch.

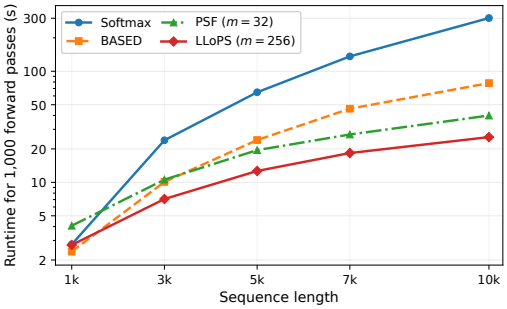


Figure 2: Forward-pass runtime of softmax attention, BASED, PSF, and LLoPS across sequence lengths (1,000 forward passes, log scale). PSF uses $m = 32$, while LLoPS uses $m = 256$. Lower is better.

5 Conclusion

We studied polynomial-kernel-based linear attention through the lens of sketching. We provided a theoretical characterization of the optimal sketch for approximating polynomial kernels in the linear attention setting, revealing that structured and learnable projections can significantly improve approximation quality over previous data-independent feature constructions [Kacham et al., 2023]. Building on this insight, we proposed LLoPS, a Learnable Decomposed Polynomial Sketch that offers an accurate and efficient approximation to polynomial attention while retaining the favorable computational properties of linear attention. Empirically, LLoPS consistently improves the effectiveness of linear attention Transformers across a diverse set of benchmarks.

Limitation. Since the original pre-training dataset used by BASED is no longer publicly available, we are unable to evaluate LLoPS under the same from-scratch training setting. Nevertheless, the results in Table 1 show that, when BASED and LLoPS are continually trained from the same initialization for the same number of iterations, LLoPS consistently achieves higher accuracy. We believe this controlled comparison is sufficient to demonstrate the effectiveness of LLoPS.

Acknowledgement. This work was partially supported by NSF 2303037, NSF 2312501, NSF 2531008, NSF 2553619, SRC JUMP 2.0 Center, CDSC Industry Partners[6], Amazon Research Awards, Snapchat, AMD AI/HPC Fund, and Google Gifts.

[6]https://cdsc.ucla.edu/partners/

References

Thomas D Ahle, Michael Kapralov, Jakob BT Knudsen, Rasmus Pagh, Ameya Velingker, David P Woodruff, and Amir Zandieh. Oblivious sketching of high-degree polynomial kernels. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 141–160. SIAM, 2020.

Ido Amos, Jonathan Berant, and Ankit Gupta. Never train from scratch: Fair comparison of long-sequence models requires data-driven priors. *arXiv preprint arXiv:2310.02980*, 2023.

Simran Arora, Brandon Yang, Sabri Eyuboglu, Avanika Narayan, Andrew Hojel, Immanuel Trummer, and Christopher Ré. Language models enable simple systems for generating structured views of heterogeneous data lakes. *arXiv preprint arXiv:2304.09433*, 2023.

Simran Arora, Sabri Eyuboglu, Michael Zhang, Aman Timalsina, Silas Alberti, Dylan Zinsley, James Zou, Atri Rudra, and Christopher Ré. Simple linear attention language models balance the recall-throughput tradeoff. *arXiv preprint arXiv:2402.18668*, 2024.

Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439, 2020.

Yifan Chen, Qi Zeng, Heng Ji, and Yun Yang. Skyformer: Remodel self-attention with gaussian kernel and nystr\" om method. *Advances in Neural Information Processing Systems*, 34:2122–2135, 2021.

Krzysztof Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.

Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.

Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac'h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for few-shot language model evaluation, 12 2023. URL `https://zenodo.org/records/10256836`.

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. In *First conference on language modeling*, 2024.

Dongchen Han, Xuran Pan, Yizeng Han, Shiji Song, and Gao Huang. Flatten transformer: Vision transformer using focused linear attention. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5961–5971, 2023.

Praneeth Kacham, Vahab Mirrokni, and Peilin Zhong. Polysketchformer: Fast transformers via sketching polynomial kernels. *arXiv preprint arXiv:2310.01655*, 2023.

Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pages 5156–5165. PMLR, 2020.

Feyza Duman Keles, Pruthuvi Mahesakya Wijewardena, and Chinmay Hegde. On the computational complexity of self-attention. In *International conference on algorithmic learning theory*, pages 597–619. PMLR, 2023.

Sang Min Kim, Byeongchan Kim, Arijit Sehanobish, Somnath Basu Roy Chowdhury, Rahul Kidambi, Dongseok Shim, Avinava Dubey, Snigdha Chaturvedi, Min-hwan Oh, and Krzysztof Choromanski. Eugens: Efficient, unified, and general dense layers. *arXiv preprint arXiv:2601.22563*, 2026.

Colin Lockard, Prashant Shiralkar, and Xin Luna Dong. Openceres: When open information extraction meets the semi-structured web. 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)*, pages 3047–3056, 2019.

Weikang Meng, Yadan Luo, Xin Li, Dongmei Jiang, and Zheng Zhang. Polaformer: Polarity-aware linear attention for vision transformers. *arXiv preprint arXiv:2501.15061*, 2025.

Evangelos E Papalexakis. Automatic unsupervised tensor mining with quality assessment. In *Proceedings of the 2016 SIAM International Conference on Data Mining*, pages 711–719. SIAM, 2016.

Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc-Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. The lambada dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 1525–1534, 2016.

Bo Peng, Eric Alcaide, Quentin Anthony, Alon Albalak, Samuel Arcadinho, Stella Biderman, Huanqi Cao, Xin Cheng, Michael Chung, Matteo Grella, et al. Rwkv: Reinventing rnns for the transformer era. *arXiv preprint arXiv:2305.13048*, 2023.

Zhen Qin, Weixuan Sun, Hui Deng, Dongxu Li, Yunshen Wei, Baohong Lv, Junjie Yan, Lingpeng Kong, and Yiran Zhong. cosformer: Rethinking softmax in attention. *arXiv preprint arXiv:2202.08791*, 2022.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.

Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.

Imanol Schlag, Kazuki Irie, and Jürgen Schmidhuber. Linear transformers are secretly fast weight programmers. In *International conference on machine learning*, pages 9355–9366. PMLR, 2021.

Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. SlimPajama: A 627B token cleaned and deduplicated version of RedPajama, June 2023. URL `https://huggingface.co/datasets/cerebras/SlimPajama-627B`.

Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. Long range arena: A benchmark for efficient transformers. *arXiv preprint arXiv:2011.04006*, 2020.

Kimi Team, Yu Zhang, Zongyu Lin, Xingcheng Yao, Jiaxi Hu, Fanqing Meng, Chengyin Liu, Xin Men, Songlin Yang, Zhiyuan Li, et al. Kimi linear: An expressive, efficient attention architecture. *arXiv preprint arXiv:2510.26692*, 2025.

David P Woodruff et al. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 10(1–2):1–157, 2014.

Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nyströmformer: A nyström-based algorithm for approximating self-attention. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 14138–14148, 2021.

Songlin Yang, Jan Kautz, and Ali Hatamizadeh. Gated delta networks: Improving mamba2 with delta rule. *arXiv preprint arXiv:2412.06464*, 2024.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.

Michael Zhang, Kush Bhatia, Hermann Kumbong, and Christopher Ré. The hedgehog & the porcupine: Expressive linear attentions with softmax mimicry. *arXiv preprint arXiv:2402.04347*, 2024.

A Proofs

A.1 Proof of Theorem 3.1

Proof. For simplicity, let $\tilde{\boldsymbol{Q}} = \psi_p(\boldsymbol{Q})$ and $\tilde{\boldsymbol{K}} = \psi_p(\boldsymbol{K})$. Notice that

$$\tilde{\boldsymbol{Q}}\tilde{\boldsymbol{K}}^\top = (\boldsymbol{Q}\boldsymbol{K}^\top)^{\circ p} \tag{16}$$

Let $\phi_Q(\boldsymbol{Q}) = \tilde{\boldsymbol{Q}}\boldsymbol{W}_1$, $\phi_K(\boldsymbol{K}) = \tilde{\boldsymbol{K}}\boldsymbol{W}_2$, where $\boldsymbol{W}_1, \boldsymbol{W}_2 \in \mathbb{R}^{d^p \times m}$. We need to prove that there exist $\boldsymbol{W}_1, \boldsymbol{W}_2$ such that

$$\tilde{\boldsymbol{Q}}\boldsymbol{W}_1\boldsymbol{W}_2^\top\tilde{\boldsymbol{K}}^\top = \boldsymbol{U_m}\boldsymbol{\Sigma_m}\boldsymbol{V_m}^\top \tag{17}$$

where $\boldsymbol{S}_m = \boldsymbol{U_m}\boldsymbol{\Sigma_m}\boldsymbol{V_m}^\top$ is the truncated SVD decomposition of $\boldsymbol{S} = \tilde{\boldsymbol{Q}}\tilde{\boldsymbol{K}}^\top$.

Notice that the column and row spaces of $\boldsymbol{S}_m$ are subspaces of the column and row spaces of $\boldsymbol{S}$.

Since $\boldsymbol{S} = \tilde{\boldsymbol{Q}}\tilde{\boldsymbol{K}}^\top$, the columns of $\boldsymbol{S}$ are linear combinations of the columns of $\tilde{\boldsymbol{Q}}$. Therefore, the range (column space) of $\boldsymbol{S}$ is a subspace of the range of $\tilde{\boldsymbol{Q}}$:

$$\text{range}(\boldsymbol{S}) \subseteq \text{range}(\tilde{\boldsymbol{Q}}) \tag{18}$$

Similarly, the rows of $\boldsymbol{S}$ are linear combinations of the rows of $\tilde{\boldsymbol{K}}^\top$:

$$\text{range}(\boldsymbol{S}^\top) \subseteq \text{range}(\tilde{\boldsymbol{K}}) \tag{19}$$

As a result, $\text{range}(\boldsymbol{U_m}) \subseteq \text{range}(\tilde{\boldsymbol{Q}})$ and $\text{range}(\boldsymbol{V_m}) \subseteq \text{range}(\tilde{\boldsymbol{K}})$

Therefore, there must exist $\boldsymbol{W}_1, \boldsymbol{W}_2$ such that $\tilde{\boldsymbol{Q}}\boldsymbol{W}_1 = \boldsymbol{U_m}\boldsymbol{\Sigma_m}^{1/2}$ and $\tilde{\boldsymbol{K}}\boldsymbol{W}_2 = \boldsymbol{V_m}\boldsymbol{\Sigma_m}^{1/2}$.

So $\tilde{\boldsymbol{Q}}\boldsymbol{W}_1(\tilde{\boldsymbol{K}}\boldsymbol{W}_2)^\top = \boldsymbol{S}_m$ □

A.2 Proof of Theorem 3.4

Proof.

$$\min_{\boldsymbol{\Theta}_1,\ldots,\boldsymbol{\Theta}_p} \|\boldsymbol{W}^* - \boldsymbol{\Theta}_1 \odot \ldots \odot \boldsymbol{\Theta}_p\|_F^2 = \sum_{i=1}^{m} \min_{\boldsymbol{\Theta}_{1,i},\ldots,\boldsymbol{\Theta}_{p,i}} \|\boldsymbol{\Theta}_{1,i} \odot \ldots \odot \boldsymbol{\Theta}_{p,i} - \boldsymbol{W}^*[:,i]\|_F^2 \tag{20}$$

where $\boldsymbol{\Theta}_{j,i}$ is the i-th row of $\boldsymbol{\Theta}_j$, $1 \le j \le p$, $1 \le i \le m$.

By the definition of tensor spectral norm and the best Rank-1 CP approximation, for each i,

$$\min_{\boldsymbol{w}_{1,i},\ldots,\boldsymbol{w}_{p,i}} \|\boldsymbol{w}_{1,i} \odot \ldots \odot \boldsymbol{w}_{p,i} - \boldsymbol{W}^*[:,i]\|_F^2 = \|\boldsymbol{W}^*[:,i]\|_2^2 - \tau_i^2 \tag{21}$$

□

A.3 Population-Level Optimality of Polynomial Sketches

Theorem 3.1 characterizes the optimal polynomial sketch for a fixed pair of query and key matrices. We extend this result to the population setting, where query and key vectors are sampled from underlying distributions.

Let $q \sim \mathcal{D}_Q$ and $k \sim \mathcal{D}_K$ be independently sampled query and key vectors, respectively. Define the polynomial feature map

$$\psi_p(x) = \text{vec}(x^{\otimes p}) \in \mathbb{R}^{d^p}. \tag{22}$$

For an output feature dimension m, we consider the population-level approximation objective

$$\mathcal{L}_{\text{pop}}(W_Q, W_K) = \mathbb{E}_{q,k}\left[\left(\psi_p(q)^\top\psi_p(k) - \psi_p(q)^\top W_Q W_K^\top \psi_p(k)\right)^2\right], \tag{23}$$

where $W_Q, W_K \in \mathbb{R}^{d^p \times m}$.

Corollary (Population-Level Optimality). Suppose that the query and key distributions have finite $2p$-th moments, i.e.,

$$\mathbb{E}_{q\sim\mathcal{D}_Q}\left[\|q\|_2^{2p}\right] < \infty, \qquad \mathbb{E}_{k\sim\mathcal{D}_K}\left[\|k\|_2^{2p}\right] < \infty. \tag{24}$$

Let Q and K contain n_Q and n_K independent samples from their respective distributions. Then, as $n_Q, n_K \to \infty$, the minimum normalized empirical approximation error in Theorem 3.1 converges almost surely to the minimum population approximation error:

$$\lim_{n_Q,n_K\to\infty} \frac{1}{n_Q n_K} \min_{W_Q,W_K} \left\|\psi_p(Q)\psi_p(K)^\top - \psi_p(Q)W_Q W_K^\top \psi_p(K)^\top\right\|_F^2 = \min_{W_Q,W_K} \mathcal{L}_{\mathrm{pop}}(W_Q, W_K). \tag{25}$$

Furthermore, if the population operator has a nonzero singular-value gap at rank m, the corresponding leading singular subspaces also converge to their population counterparts.

Proof. Let $z_q = \psi_p(q)$ and $z_k = \psi_p(k)$, and define the population second-moment matrices

$$C_Q = \mathbb{E}[z_q z_q^\top], \qquad C_K = \mathbb{E}[z_k z_k^\top]. \tag{26}$$

For any rank-m matrix $B = W_Q W_K^\top$, the population objective can be expressed as

$$\begin{aligned}\mathcal{L}_{\mathrm{pop}}(W_Q, W_K) &= \mathbb{E}_{q,k}\left[\left(z_q^\top (I - B) z_k\right)^2\right] \\ &= \left\|C_Q^{1/2}(I - B)C_K^{1/2}\right\|_F^2.\end{aligned} \tag{27}$$

Let $M = C_Q^{1/2} C_K^{1/2}$, and denote its singular values by $\sigma_1(M) \geq \sigma_2(M) \geq \cdots$. By the Eckart–Young–Mirsky theorem, the minimum population approximation error is

$$\min_{W_Q,W_K} \mathcal{L}_{\mathrm{pop}}(W_Q, W_K) = \sum_{i>m} \sigma_i(M)^2. \tag{28}$$

The optimal rank-m approximation of M can be realized by suitable linear projections W_Q and W_K on the supports of C_Q and C_K.

Next, define the empirical second-moment matrices

$$\widehat{C}_Q = \frac{1}{n_Q}\psi_p(Q)^\top \psi_p(Q), \qquad \widehat{C}_K = \frac{1}{n_K}\psi_p(K)^\top \psi_p(K). \tag{29}$$

Under the finite-moment assumptions, the strong law of large numbers implies

$$\widehat{C}_Q \xrightarrow{\text{a.s.}} C_Q, \qquad \widehat{C}_K \xrightarrow{\text{a.s.}} C_K. \tag{30}$$

Since the feature dimension d^p is fixed, this convergence also holds in operator norm. By continuity of the matrix square root,

$$\widehat{C}_Q^{1/2}\widehat{C}_K^{1/2} \xrightarrow{\text{a.s.}} C_Q^{1/2} C_K^{1/2} = M. \tag{31}$$

The nonzero singular values of $\psi_p(Q)\psi_p(K)^\top/\sqrt{n_Q n_K}$ coincide with those of $\widehat{C}_Q^{1/2}\widehat{C}_K^{1/2}$. Consequently, the optimal normalized empirical approximation error satisfies

$$\mathcal{L}_{\mathrm{emp}}^* = \sum_{i>m} \sigma_i\left(\widehat{C}_Q^{1/2}\widehat{C}_K^{1/2}\right)^2. \tag{32}$$

By continuity of singular values,

$$\mathcal{L}_{\mathrm{emp}}^* \xrightarrow{\text{a.s.}} \sum_{i>m} \sigma_i(M)^2 = \mathcal{L}_{\mathrm{pop}}^*. \tag{33}$$

Moreover, when $\sigma_m(M) > \sigma_{m+1}(M)$, standard singular-subspace perturbation results imply convergence of the leading rank-m singular subspaces. Without such a spectral gap, the optimal approximation error still converges, although the optimal subspaces need not be unique. This completes the proof.

Discussion. This corollary extends the finite-sample optimality established in Theorem 3.1 to the underlying query and key distributions. In particular, it shows that the optimal empirical approximation error converges to the population optimum as the sample sizes increase. The result characterizes the optimality of unrestricted linear polynomial sketches; the approximation error introduced by the practical low-CP-rank parameterization is analyzed separately below.

A.4 Approximation Error of Practical Low-CP-Rank Sketches

Although Theorem 3.1 establishes the optimality of unrestricted linear projections of polynomial features, the corresponding projection matrices have $d^p \times m$ parameters, making their direct implementation computationally expensive. LLoPS addresses this limitation through a structured low-CP-rank parameterization.

In practice, we adopt CP rank $t = 1$, which provides a favorable trade-off between approximation expressiveness, optimization difficulty, and computational efficiency. While increasing t enlarges the hypothesis class and can reduce the best-achievable projection approximation error, it also increases the number of learnable parameters and the computational cost of training and inference. Our experiments indicate that $t = 1$ achieves strong approximation quality and downstream performance while preserving the efficiency of the proposed kernel.

We next characterize how the approximation errors of the practical projection matrices propagate to the polynomial kernel matrix and attention outputs.

A.4.1 Polynomial Kernel Approximation Error

Let

$$Z_Q = \psi_p(Q), \qquad Z_K = \psi_p(K), \tag{34}$$

where ψ_p is applied row-wise. Denote the unrestricted optimal projections in Theorem 3.1 by W_Q^* and W_K^*. The resulting optimal rank-m kernel approximation is

$$A^* = Z_Q W_Q^* W_K^{*\top} Z_K^\top. \tag{35}$$

In contrast, the practical LLoPS parameterization with CP rank $t = 1$ produces

$$\widehat{A} = Z_Q \widehat{W}_Q \widehat{W}_K^\top Z_K^\top, \tag{36}$$

where $\widehat{W}_Q$ and $\widehat{W}_K$ are structured projection matrices.

Define the projection errors

$$E_Q = \widehat{W}_Q - W_Q^*, \qquad E_K = \widehat{W}_K - W_K^*. \tag{37}$$

Expanding the product gives

$$\begin{aligned} \widehat{W}_Q \widehat{W}_K^\top &= (W_Q^* + E_Q)(W_K^* + E_K)^\top \\ &= W_Q^* W_K^{*\top} + E_Q W_K^{*\top} + W_Q^* E_K^\top + E_Q E_K^\top. \end{aligned} \tag{38}$$

Therefore,

$$\widehat{A} - A^* = Z_Q \left(E_Q W_K^{*\top} + W_Q^* E_K^\top + E_Q E_K^\top \right) Z_K^\top. \tag{39}$$

Using the triangle inequality and the matrix norm inequality

$$\|XYZ\|_F \le \|X\|_2 \|Y\|_F \|Z\|_2, \tag{40}$$

we obtain

$$\begin{aligned} \|\widehat{A} - A^*\|_F &\le \|Z_Q\|_2 \|Z_K\|_2 \Big(\|E_Q W_K^{*\top}\|_F + \|W_Q^* E_K^\top\|_F + \|E_Q E_K^\top\|_F \Big) \\ &\le \|Z_Q\|_2 \|Z_K\|_2 \Big(\|E_Q\|_F \|W_K^*\|_2 + \|W_Q^*\|_2 \|E_K\|_F + \|E_Q\|_F \|E_K\|_F \Big). \end{aligned} \tag{41}$$

For notational convenience, define

$$\Gamma = \|E_Q\|_F \|W_K^*\|_2 + \|W_Q^*\|_2 \|E_K\|_F + \|E_Q\|_F \|E_K\|_F. \tag{42}$$

Then,

$$\|\widehat{A} - A^*\|_F \leq \|Z_Q\|_2 \|Z_K\|_2 \Gamma. \tag{43}$$

Theorem 3.4 characterizes the minimum projection error achievable by the CP-rank-one parameterization. Specifically, for an optimal rank-one CP approximation of each column of W_Q^*,

$$\|E_Q\|_F^2 = \|W_Q^*\|_F^2 - \sum_{i=1}^{m} (\tau_i^Q)^2, \tag{44}$$

and similarly,

$$\|E_K\|_F^2 = \|W_K^*\|_F^2 - \sum_{i=1}^{m} (\tau_i^K)^2, \tag{45}$$

where τ_i^Q and τ_i^K denote the tensor spectral norms defined in Theorem 3.4.

These expressions characterize the best-achievable projection errors under the CP-rank-one constraint. For learned projections obtained through numerical optimization, the actual errors may be larger, depending on optimization quality.

We further obtain an explicit bound under bounded query and key norms. Recall that

$$\|\psi_p(x)\|_2 = \|\operatorname{vec}(x^{\otimes p})\|_2 = \|x\|_2^p. \tag{46}$$

Assume that

$$\|q_i\|_2 \leq B_Q, \qquad \|k_j\|_2 \leq B_K \tag{47}$$

for all query and key vectors. Then,

$$\|Z_Q\|_2 \leq \|Z_Q\|_F = \left(\sum_{i=1}^{n_Q} \|\psi_p(q_i)\|_2^2 \right)^{1/2} \leq \sqrt{n_Q} B_Q^p, \tag{48}$$

$$\|Z_K\|_2 \leq \|Z_K\|_F \leq \sqrt{n_K} B_K^p. \tag{49}$$

Substituting into Eq. (41) yields

$$\boxed{\|\widehat{A} - A^*\|_F \leq \sqrt{n_Q n_K}\, B_Q^p B_K^p \Gamma.} \tag{50}$$

This bound explicitly connects the projection-space approximation errors characterized by Theorem 3.4 to the induced polynomial kernel approximation error.

A.4.2 Attention Output Error

We next extend the analysis from the kernel matrix to the corresponding attention outputs.

Consider the unnormalized polynomial attention outputs

$$Y^* = A^* V, \qquad \widehat{Y} = \widehat{A} V, \tag{51}$$

where $V \in \mathbb{R}^{n_K \times d_v}$ is the value matrix.

Their difference satisfies

$$\begin{aligned} \|\widehat{Y} - Y^*\|_F &= \|(\widehat{A} - A^*)V\|_F \\ &\leq \|\widehat{A} - A^*\|_F \|V\|_2. \end{aligned} \tag{52}$$

Combining this inequality with Eq. (50), we obtain

$$\boxed{\|\widehat{Y} - Y^*\|_F \leq \sqrt{n_Q n_K}\, B_Q^p B_K^p \Gamma \|V\|_2.} \tag{53}$$

Corollary (Attention Output Error). Under the bounded query and key assumptions above, the difference between the unnormalized attention outputs produced by the practical CP-rank-one LLoPS sketch and the unrestricted optimal rank-m polynomial sketch is bounded by Eq. (53).

Furthermore, let $A = Z_Q Z_K^\top$ denote the exact polynomial kernel matrix. By the triangle inequality,

$$\begin{aligned}\|AV - \widehat{A}V\|_F &\le \|(A - A^*)V\|_F + \|(A^* - \widehat{A})V\|_F \\ &\le \left(\|A - A^*\|_F + \sqrt{n_Q n_K}\, B_Q^p B_K^p \Gamma\right) \|V\|_2. \end{aligned} \tag{54}$$

By Theorem 3.1,

$$\|A - A^*\|_F^2 = \sum_{i>m} \sigma_i(A)^2, \tag{55}$$

where $\sigma_i(A)$ denotes the i-th largest singular value of A.

Thus, the total attention output error is bounded by two components: the irreducible rank-m approximation error characterized by Theorem 3.1, and the additional error arising from the structured low-CP-rank parameterization.

Discussion. The analysis establishes a direct connection between the projection approximation error in Theorem 3.4 and the resulting kernel and attention output errors. In particular, reducing the projection errors $\|E_Q\|_F$ and $\|E_K\|_F$ reduces the derived upper bounds.

These results also clarify the distinction between the unrestricted optimal polynomial sketch and its practical low-CP-rank realization. While LLoPS with $t = 1$ does not generally attain the unrestricted optimum, it provides a computationally efficient approximation whose error can be characterized through the quality of the structured projections.

We note that these bounds apply to unnormalized polynomial attention. Extending them to normalized attention additionally requires controlling the normalization denominators.

A.5 Comparison with Existing Polynomial Feature Maps

In Section 3.3, we show that several existing polynomial feature maps can be interpreted as special or constrained cases of LLoPS. We further clarify the implications of this relationship for approximation optimality.

Let $\mathcal{H}_{\mathrm{LLoPS}}$ denote the class of polynomial kernel approximations representable by LLoPS with a fixed output feature dimension. Similarly, let $\mathcal{H}_{\mathrm{prior}}$ denote the hypothesis class of a prior polynomial feature-map method under a comparable output dimension.

For the prior constructions identified as special cases of LLoPS in Section 3.3, we have

$$\mathcal{H}_{\mathrm{prior}} \subseteq \mathcal{H}_{\mathrm{LLoPS}}. \tag{56}$$

This inclusion follows from the fact that LLoPS relaxes structural restrictions imposed by existing feature maps, such as shared projection matrices, fixed randomized projections, or specific parameter-sharing patterns.

For example, element-wise polynomial feature maps impose the constraint $\Theta_1 = \cdots = \Theta_p$, whereas LLoPS allows these projection matrices to be learned independently. Likewise, the BASED polynomial feature map corresponds to a particular structured choice of the LLoPS projection matrices, as discussed in Section 3.3.

Let $\mathcal{L}(h)$ denote the approximation error of a kernel approximation h. The hypothesis-class inclusion immediately implies

$$\inf_{h \in \mathcal{H}_{\mathrm{LLoPS}}} \mathcal{L}(h) \le \inf_{h \in \mathcal{H}_{\mathrm{prior}}} \mathcal{L}(h). \tag{57}$$

Importantly, this inequality is not necessarily strict for every query and key distribution. Equality can occur when an optimal approximation within the larger LLoPS hypothesis class is already representable by the constrained prior method.

Nevertheless, the additional flexibility of LLoPS enlarges its representational capacity. For nondegenerate settings in which the prior structural constraints exclude a better approximation available to LLoPS, the inequality in Eq. (57) is strict.

Table 4: Comparion of Based and LLoPS on synthetic associative recall task. For LLoPS, $m = 256$, $p = 4$. For both methods $d = 128$.

d	128	192
BASED	91.9	98.1
LLoPS	93.2	99.0

Therefore, the theoretical guarantee is that LLoPS achieves a no-worse best-achievable approximation error than the included prior feature-map classes, with the possibility of strict improvement when the additional expressiveness is beneficial.

Whether this increased expressiveness translates into lower approximation error after numerical optimization, or improved generalization on downstream tasks, remains an empirical question.

Our polynomial kernel approximation experiments in Section 4.3 provide supporting evidence. Across the evaluated polynomial degrees and Transformer layers, LLoPS consistently achieves lower approximation error than the compared sketching baselines. Furthermore, the zero-shot language modeling and Long Range Arena results demonstrate improved downstream performance in the evaluated settings.

Taken together, the theoretical and empirical results establish that LLoPS generalizes several existing polynomial feature-map constructions, guarantees a no-worse optimal approximation error within the corresponding hypothesis classes, and empirically achieves more accurate kernel approximations and stronger downstream performance in the evaluated settings.

B Additional Experiment Results

B.1 Synthetic Associative Recall Task

We compare our method with Based on the synthetic associative reacall task introduced in Arora et al. [2024]. For both methods, the hidden dimension of the model is set to $d \in \{64, 128\}$. The query and key dimensions of BASED is 16. The sketch dimension of LLoPS $m = 256$ and $p = 4$. Table 4 shows the result. We can see that LLoPS achieves higher accuracy than BASED, further proving its effectiveness.

B.2 Polynomial Degree Ablation

Here, we study how the polynomial degree p affects the effectiveness of LLoPS. We consider the synthetic associative recall task Arora et al. [2024] and train three 2-layer transformers from scratch using polynomial kernels with degrees $p \in \{2, 4, 6\}$. All models use a hidden dimension of 128, and the output sketch dimension is fixed to 256 across settings. The results are summarized in Table 5.

We observe that increasing the degree from $p = 2$ to $p = 4$ leads to improved accuracy. This suggests that moderate higher-order interactions can better distinguish key vectors that align well with the queries from those that do not, thereby enhancing retrieval performance. However, when further increasing the degree to $p = 6$, accuracy declines. We hypothesize that this drop arises because the sketch dimension m remains fixed: as p grows, the polynomial kernel becomes more difficult to approximate accurately, and the resulting increase in approximation error can negatively impact model effectiveness. Another possible reason is that, since the number of parameters increase with p, a higher value of p make the model easier to overfit the training data, which would also cause the test accuracy to drop.

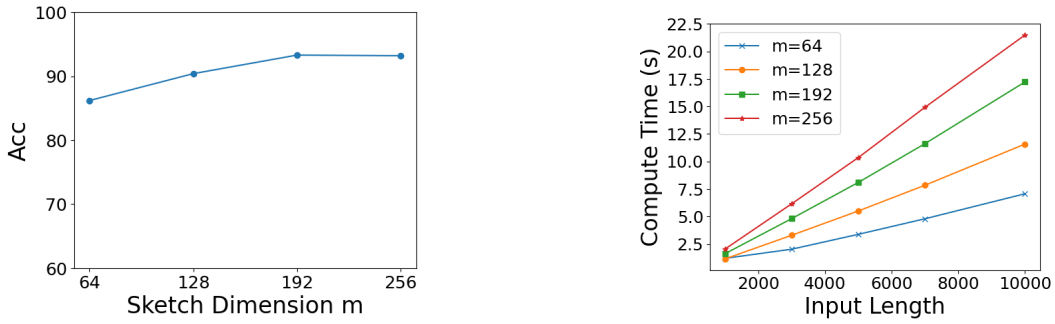
Overall, these results highlight an important trade-off: choosing p requires balancing representational power and approximation accuracy, and higher degree is not always better in practice.

B.3 Sketch Dimension Ablation

Here we demonstrate the effect of sketch dimension m to effectiveness and effciency. We fix the model dimension $d = 128$ and vary $m \in \{64, 128, 192, 256\}$. We measure the accuracy for synthetic association recall task [Arora et al., 2024] and the forward computation time. The results show

Table 5: Accuracy of synthetic association recall task under different polynomial degree p with $d = 128$.

p	2	4	6
Acc	88.7	93.2	86.6



(a) m vs Acc

(b) m vs efficency

Figure 3: Results of how accuracy and forward time changes with sketch dimension $m \in \{64, 128, 192, 256\}$ and $d = 128$, $p = 4$.

that when m increases, the effective improves but the efficiency drops, which fits the intuition. Nevertheless, we note that when the sketch dimension drops from 256 to 192, the efficiency improves while the accuracy does not drop. This suggests that our method has great potential in approximating polynomial kernels with limited budget of state space size.

B.4 FDA Character-Level Perplexity

We evaluated the original BASED-360M checkpoint, BASED after continued pretraining on SlimPajama-6B, and BASED+LLoPS after the same continued-pretraining budget. The results are summarized in Table 6.

Continued pretraining increases the BASED perplexity from 11.89 to 12.07, providing direct evidence of degradation on the FDA distribution without introducing LLoPS. Importantly, BASED+LLoPS achieves a perplexity of 11.94, which is lower than BASED after continued pretraining and close to the original BASED checkpoint. Thus, this controlled analysis does not support the hypothesis that the data-dependent LLoPS feature map causes additional FDA forgetting. Instead, it suggests that continued pretraining contributes to the regression, while LLoPS does not exacerbate it.

Although the increase from 11.89 to 12.07 appears numerically moderate, for sequence-level predictions, such differences can accumulate across tokens and substantially affect the relative likelihoods or rankings of candidate sequences. Therefore, a moderate character-level perplexity degradation may plausibly produce a larger change in the downstream FDA score. The remaining difference between perplexity and downstream accuracy also indicates that these metrics capture different aspects of model behavior.

Table 6: Character-level perplexity on the FDA dataset.

Model	FDA Perplexity
BASED before continued pretraining	**11.89**
BASED after continued pretraining	12.07
BASED+LLoPS	11.94

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [N/A].
- [N/A] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for [N/A]).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will also be asked to include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While [Yes] is generally preferable to [No], it is perfectly acceptable to answer [No] provided a proper justification is given (e.g., error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering [No] or [N/A] is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS Paper Checklist"**,
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers**.

1. **Claims**

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

   Answer: [Yes]

   Justification: The main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

2. **Limitations**

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

   Answer: [Yes]

   Justification: The limitation is discussed in conclusion.

3. **Theory assumptions and proofs**

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

Answer: [Yes]

Justification: The proofs are provided in the main paper and appendix.

4. **Experimental result reproducibility**

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

   Answer: [Yes]

   Justification: All information is disclosed.

5. **Open access to data and code**

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

   Answer: [No]

   Justification: The code will be released upon acceptance.

6. **Experimental setting/details**

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

   Answer: [Yes]

   Justification: All information is disclosed

7. **Experiment statistical significance**

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

   Answer: [Yes]

   Justification: We include error bars in Figure **??**. For other experiments, due to the high cost of pre-training, we cannot use different random seeds to train it multiple times.

8. **Experiments compute resources**

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

   Answer: [Yes]

   Justification: Information is disclosed

   Guidelines:

9. **Code of ethics**

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

   Answer: [Yes]

   Justification: This paper conform with the NeurIPS code of ethics.

10. **Broader impacts**

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

    Answer: [Yes]

    Justification: The impact of the paper is discussed in the main paper.

11. **Safeguards**

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

Answer: [N/A]

Justification: The paper has no such risk.

12. **Licenses for existing assets**

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

    Answer: [Yes]

    Justification: They are credited.

13. **New assets**

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

    Answer: [N/A]

    Justification: the paper does not release new assets.

14. **Crowdsourcing and research with human subjects**

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

    Answer: [N/A]

    Justification: the paper does not involve crowdsourcing nor research with human subjects.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

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

    Answer: [N/A]

    Justification: the paper does not involve crowdsourcing nor research with human subjects.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: the core method development in this research does not involve LLMs as any important, original, or non-standard components.